
Online Learning of Pure States is as Hard as Mixed States

Maxime Meyer

Department of Mathematics & IPAL, IRL2955
National University of Singapore
Singapore
maxime.meyer@u.nus.edu

Soumik Adhikary

Centre for Quantum Technologies
National University of Singapore
Singapore
soumik@nus.edu.sg

Naixu Guo

Centre for Quantum Technologies
National University of Singapore
Singapore
naixug@u.nus.edu

Patrick Rebentrost

Centre for Quantum Technologies & School of Computing
National University of Singapore
Singapore
cqtfpr@nus.edu.sg

Abstract

Quantum state tomography, the task of learning an unknown quantum state, is a fundamental problem in quantum information. In standard settings, the complexity of this problem depends significantly on the type of quantum state that one is trying to learn, with pure states being substantially easier to learn than general mixed states. A natural question is whether this separation holds for any quantum state learning setting. In this work, we consider the online learning framework and prove the surprising result that learning pure states in this setting is as hard as learning mixed states. More specifically, we show that both classes share almost the same sequential fat-shattering dimension, leading to identical regret scaling. We also generalize previous results on full quantum state tomography in the online setting to (i) the ϵ -realizable setting and (ii) learning the density matrix only partially, using *smoothed analysis*.

1 Introduction

Learning information from an unknown quantum state is a fundamental task in quantum physics. An N -dimensional quantum state ρ is represented as an $N \times N$ positive semi-definite Hermitian matrix with unit trace. Given several perfect copies of ρ , full quantum state tomography seeks to reconstruct the complete matrix representation of ρ via measurements. It has a wide range of practical applications, including tasks such as characterizing qubit states for superconducting circuits [Lucero et al., 2008], nitrogen-vacancy (NV) centers in diamond [Neumann et al., 2010], and verifying successful quantum teleportation [Bouwmeester et al., 1997]. The ease of learning a quantum state is often characterized by the sample complexity, *i.e.* the number of independent copies of the quantum state required for an accurate reconstruction. For a general state ρ , the sample complexity scales as $\tilde{\Theta}(N^3)$ for incoherent measurements. However, if the state is known to be pure, the sample complexity can be improved to $\tilde{\Theta}(N)$ [Kueng et al., 2017, Haah et al., 2016, Chen et al., 2023]. Nevertheless, for n -qubit states, where $N = 2^n$, this scaling implies an exponential dependence on the number of qubits.

However, for many practical problems, such as certification of a quantum device [Gross et al., 2010, Flammia and Liu, 2011, Eisert et al., 2020] and property estimation for quantum chemistry [Huang et al., 2021, Wu et al., 2023, Guo et al., 2024, Raza et al., 2024, Miller et al., 2024], only partial

information about the quantum state is needed. To address such cases, quantum PAC learning, also known as pretty good tomography, was introduced by Aaronson [2007]. In this framework, a fixed distribution $\mathcal{D}$ governs the selection of two-outcome measurements E , which are represented by $N \times N$ Hermitian matrices with eigenvalues in $[0, 1]$. The learner then tries to output a hypothesis state ω such that $\text{Tr}(E\rho) \approx_{E \sim \mathcal{D}} \text{Tr}(E\omega)$ with high probability. The number of measurements required to output this hypothesis n -qubit state was shown to scale only linearly with the number of qubits n , yielding an exponential improvement over full-state tomography. However, a key limitation of both these approaches is that they do not account for adversarial environments, where the set of realizable measurements may evolve over time.

This limitation can be circumvented by generalizing to the online learning setting [Aaronson et al., 2018, Chen et al., 2024, Bansal et al., 2024], where learning a quantum state ρ is posed as a T -round repeated two-player game. In each round $t \in [T]$, Nature—also called the adversary—chooses a measurement E_t from the set of two-outcome measurements. The task of the learner is to predict the value $\text{Tr}(E_t\rho)$ by selecting a hypothesis ω_t and computing $\text{Tr}(E_t\omega_t)$ based on previous results. Thereafter, Nature returns the loss, a metric quantifying the difference between the prediction and the true value. The most adversarial scenario arises when the measurement at each round is chosen from the set of all two-outcome measurements without any constraints, *i.e.*, it can be chosen adversarially and adaptively. It has been shown that in such cases, a learner can output a hypothesis state which incurs an additional $O(\sqrt{nT})$ loss compared to the best possible hypothesis state after T rounds of the game [Aaronson et al., 2018]. This measure of how much worse a learner performs compared to the best possible strategy in hindsight is called regret and serves as a fundamental measure of performance for any online learning problem.

Given the clear separation in sample complexity between pure and mixed state tomography [Kuang et al., 2017, Haah et al., 2016], we ask the central question that this work aims to address:

Is there any separation between online learning of pure and mixed quantum states?

In this work we show that the answer is **No**. Our proof is based on the analysis of the sequential fat-shattering dimension of pure and mixed states. Informally, it can be seen as the minimum number of mistakes—defined as errors exceeding a threshold δ —that a learner must make before successfully learning a quantum state ρ against a perfect adversary. We have shown that both pure and mixed states share almost the same sequential fat-shattering dimension of order $\Theta(\frac{\delta^2}{n})$, and hence the same regret. Indeed, it has been demonstrated that both upper and lower bounds on regret can be expressed in terms of this dimension [Rakhlin et al., 2015b].

We believe that this result is surprising. Indeed, not only is there a significant difference in sample complexities between learning pure and mixed states in the standard tomographic settings, but this distinction also holds in specific online learning scenarios. For instance, under bandit feedback with adaptive measurements, Lumbreras et al. [2022, 2024] shows that the regret for mixed states grows exponentially faster than for pure states.

A prior work [Aaronson et al., 2018] established tight bounds on the sequential fat-shattering dimension of general mixed states. In this work, we establish lower bounds on the sequential fat-shattering dimension – and consequently on the regret – of several subclasses of quantum states, the most important of which is the class of pure states. Our approach employs a distinct proof strategy, providing new insights and extending naturally to various subsettings of the online quantum state learning problem. As a result, we show that the regret for online learning of both pure and mixed quantum state scales as $\Theta(\sqrt{nT})$.

A key feature of the online learning setting considered here is the adversary’s ability to select measurements in a completely unconstrained manner. As a result, this setting may effectively incur full state tomography, even in cases where only partial information about the state is needed. We therefore introduce the concept of smoothness for online learning of quantum states. Smoothed analysis was first introduced in Spielman and Teng [2004] as a tool that allows for interpolation between the worst and the average case analysis. Later, Haghtalab et al. [2024] extended this concept to the online setting, where the degree of adversariality is quantified by a smoothness parameter $\sigma \in [0, 1]$. The particular value $\sigma = 1$ corresponds to independent and identically distributed (i.i.d.) inputs, while the limit $\sigma \rightarrow 0$ corresponds to fully adversarial inputs. In this work, we establish an

upper bound on regret for smoothed online learning of quantum states, providing insights into the effect of adversariality on regret scaling.

1.1 Main Contributions

Equivalence between pure and mixed state learning: We show that pure and mixed states share almost the same sequential fat-shattering dimension, leading to identical regret scaling under the L_1 loss in Section 4. We also prove a novel dependence of minimax regret with L_2 loss on the Rademacher complexity. This allows us to extend the tightness of regret to the L_2 loss setting for both pure and mixed states.

Extensions of Online Learning of Quantum States to more realistic settings: We obtain new regret bounds for two settings of interest, which capture more accurately the settings encountered by experimentalists, in Section 5. The first is the ϵ -realizable setting, where the learner is able to measure $\text{Tr}(E_t \rho)$ with error ϵ at every round. The second is the smooth setting, where the learner is only interested about learning specific properties of the quantum state ρ .

1.2 Related Works

Pure and mixed state tomography: In full state tomography of an N dimensional quantum state ρ , the goal is to reconstruct its complete classical representation given several independent copies. The associated sample complexity refers to the number of copies required to obtain a classical description of ρ up to an accuracy ϵ . It has been shown that the sample complexity up to trace distance ϵ is $\tilde{\Theta}(Nr/\epsilon^2)$ for incoherent measurements [Haah et al., 2016, Kueng et al., 2017, Chen et al., 2023], where r is the rank of the density matrix ρ . This result highlights a fundamental separation in the sample complexities between pure and mixed state tomography, as the rank of pure states is $r = 1$. Given the fundamental nature of this separation in quantum information science, one might expect it to persist in the online learning setting. In fact, Lumbreras et al. [2022] studies the online learning of properties of quantum states under bandit feedback with adaptive measurements, obtaining a regret scaling of $\Theta(\sqrt{T})$ for mixed states. Subsequently, Lumbreras et al. [2024] improved this result to $\Theta(\text{polylog } T)$ for pure states with rank-1 projective measurements, showing an exponential separation. In contrast, our results demonstrate that, in the general online learning setting, the regret scaling for pure and mixed states is identical.

Existing bounds: Aaronson et al. [2018] showed that the sequential fat-shattering dimension of quantum states with parameter δ is tight, of order $\Theta(\frac{n}{\delta^2})$. They also use a result from Arora et al. [2012] to show that the regret is tight for the L_1 loss in the non-realizable case (that is when the data isn't assumed to come from an actual quantum state), being of order $\Theta(\sqrt{nT})$. In this paper, we generalize this result to several new settings of interest. We also introduce a new proof technique that allows us to provide lower bounds for sequential fat-shattering dimension for restricted settings, notably proving near-tightness for pure states in Theorem 4.1.

2 Preliminaries

2.1 Quantum States and Measurements

A general mixed n -qubit (quantum bit) quantum state ρ corresponds to a Hermitian positive semi-definite matrix of size 2^n and of trace 1. It will be called a **pure** state if and only if it is of rank 1 (or equivalently if and only if $\text{Tr}(\rho^2) = 1$). We can thus identify any pure state ρ to a vector in the 2^n dimensional Hilbert space $\mathbb{C}^{2^n}$, typically expressed in the Dirac notation as $|\psi\rangle$. Under this notation, any pure state vector can be expressed as $|\psi\rangle = \sum_{i=1}^{2^n} c_i |i\rangle$ with $c_i \in \mathbb{C}$ and $\sum_{i=1}^{2^n} |c_i|^2 = 1$. Here, $\{|i\rangle\}_{i=1}^{2^n-1}$ corresponds to the canonical basis of the vector space. For a pure state $|\psi\rangle$ the corresponding density matrix representation is given by $\rho = |\psi\rangle\langle\psi|$, where $\langle\psi| = |\psi\rangle^\dagger$ is the complex conjugate of $|\psi\rangle$.

Information from quantum states can be obtained via quantum measurements. In our work we will focus on two outcome measurements. They are represented by two-element positive operator-valued measure (POVM) $\{E, 1 - E\}$, where $E \in \text{Herm}_{\mathbb{C}}(2^n)$ and $\text{Spec}(E) \subset [0, 1]$. Since the second element of the POVM is uniquely determined by the first, a two-outcome measurement can

effectively be represented by a single operator E . A measurement E is said to *accept* a quantum state ρ with probability $\text{Tr}(E\rho)$ and *reject* it with probability $1 - \text{Tr}(E\rho)$. For a given quantum state ρ , predicting its acceptance probabilities for all measurements E is tantamount to characterizing it completely. Hence learning a quantum state ρ is equivalent to learning the function Tr_ρ , defined as $\text{Tr}_\rho(E) = \text{Tr}(E\rho)$.

2.2 Online Learning of Quantum States

Online learning, or the sequential prediction model, is a T round repeated two-player game [Cesa-Bianchi and Lugosi, 2006]. In each round $t \in [T]$ of the game, the learner is presented with an input from the sample space $\mathcal{X}$. Without any loss of generality, we can assume that x_t is sampled from a distribution $\mathcal{D}_t(\mathcal{X})$ where $\mathcal{D}_t$ may be chosen adversarially. The learner's goal is to learn an unknown function $h : \mathcal{X} \rightarrow \mathcal{Y}$ from the data they receive. Here, $\mathcal{Y}$ denotes the space of possible labels for each input x_t . The learning proceeds by designing an algorithm that outputs a sequence of functions $h_t : \mathcal{X} \rightarrow \mathcal{Y}$ chosen from a hypothesis class $\mathcal{H}$. After each round, the learner incurs a loss and aims to minimize the cumulative regret—which corresponds to the difference with the loss incurred by the best hypothesis in hindsight—at the end of all T rounds of the game. The main figure of merit used in online learning, which represents the regret of the best strategy from the learner, when presented with the hardest possible inputs at every round is called the minimax regret [Rakhlin et al., 2015a].

Definition 2.1 (Minimax regret). Let $\mathcal{X}$ denote the sample space, $\mathcal{Y}$ its associated label space, and $\mathcal{H}$ the hypothesis class. Let $\mathcal{P}$ and $\Delta(\mathcal{H})$ be sets of probability measures defined on $\mathcal{X}$ and $\mathcal{H}$ respectively. In each round $t \in [T]$ of the online learning process, the learner incurs a loss $\ell_t(h_t(x_t), y_t)$, where $y_t \in \mathcal{Y}$ is the true label associated to x_t . The minimax regret is then defined as:

$$\mathcal{V}_T = \left\langle \inf_{Q \in \Delta(\mathcal{H})} \sup_{y_t} \sup_{\mathcal{D}_t \in \mathcal{P}} \mathbb{E}_{h_t \sim Q} \mathbb{E}_{x_t \sim \mathcal{D}_t} \right\rangle_{t=1}^T \left[\sum_{t=1}^T \ell_t(h_t(x_t), y_t) - \inf_{h \in \mathcal{H}} \sum_{t=1}^T \ell_t(h(x_t), y_t) \right], \quad (1)$$

where $\langle \cdot \rangle_{t=1}^T$ denotes iterated application of the enclosed operators.

The question of learnability of an online learning problem can then be reduced to the study of $\mathcal{V}_T$. Given a pair $(\mathcal{H}, \mathcal{X})$, a problem is said to be online learnable if and only if $\lim_{T \rightarrow \infty} \mathcal{V}_T/T = 0$.

In the context of quantum state learning, we define the sample space as $\mathcal{X} \subset \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$. Denoting $\mathcal{C}_n$ as the set of all n -qubit quantum states, we set the hypothesis class to be $\mathcal{H}_n = \{\text{Tr}_\omega, \omega \in \mathcal{C}_n\}$, and $\mathcal{Q}$ can be seen as a distribution over $\mathcal{C}_n$. Here, the learner receives a sequence of measurements $(E_t)_{t \in [T]}$, each drawn from a distribution $\mathcal{D}_t$ (chosen adversarially) one at a time. Upon receiving each measurement, the learner selects a hypothesis $\omega_t \in \mathcal{C}_n$ and thereby incurs a loss of $\ell_t(\text{Tr}_{\omega_t}(E_t), y_t) = \ell_t(\text{Tr}(E_t \omega_t), y_t)$. In practice, for quantum tomography, ℓ_t is taken to be either the L_1 or the L_2 loss (that is $\ell_t(a, b) \in \{|a - b|, (a - b)^2\}$). Lastly, the label $y_t \in \mathcal{Y} = [0, 1]$ is revealed to the learner. It may be an approximation of $\text{Tr}(E_t \rho)$, but it is allowed to be arbitrary in general.

2.3 Sequential fat-shattering dimension

The main notion we will be studying in this paper is that of sequential fat-shattering dimension.

Definition 2.2 (Sequential fat-shattering dimension). A $\mathcal{X}$ -valued complete binary tree $\mathbf{x}$ of depth T is deemed to be δ -shattered by a hypothesis class $\mathcal{H}$ if there exists a real-valued complete binary tree $\mathbf{v}$ of same depth T such that for all paths $\epsilon \in \{\pm 1\}^{T-1}$,

$$\exists h \in \mathcal{H} : \forall t \in [T] \quad \epsilon_t[h(\mathbf{x}_t(\epsilon)) - \mathbf{v}_t(\epsilon)] \geq \frac{\delta}{2}.$$

The sequential fat-shattering dimension at scale δ , $\text{sfat}_\delta(\mathcal{H}, \mathcal{X})$, is defined to be the largest T for which $\mathcal{H}$ δ -shatters a $\mathcal{X}$ -valued tree of depth T .

Recall that a $\mathcal{X}$ -valued complete binary tree of depth T , $\mathbf{x}$, is defined as a sequence of T mappings $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$, where $\mathbf{x}_t : \{\pm 1\}^{t-1} \rightarrow \mathcal{X}$, with a constant function $\mathbf{x}_1 \in \mathcal{X}$ as the root. In simpler terms the tree can be seen as a collection of T length paths $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_{T-1}) \in \{\pm 1\}^{T-1}$ (+1

indicating right and -1 indicating left from any given node) and $\mathbf{x}_t(\epsilon) \equiv \mathbf{x}_t(\epsilon_1, \dots, \epsilon_{t-1}) \in \mathcal{X}$ denoting the label of the t -th node on the corresponding path ϵ .

This dimension is a fundamental property in online learning, as it both upper and lower bounds regret [Rakhlin et al., 2015a,b], as shown in Equations (2) and (3).

$$\mathcal{V}_T \leq \inf_{\alpha > 0} \left\{ 4\alpha TL - 12L\sqrt{T} \int_{\alpha}^1 \sqrt{\text{sfat}_{\delta}(\mathcal{H}, \mathcal{X}) \log\left(\frac{2eT}{\delta}\right)} d\delta \right\}. \quad (2)$$

Note that this upper bound also holds for $\bar{\mathcal{V}}_T$. The bound in Equation (2) was used in Aaronson et al. [2018] to derive the regret upper bounds for online quantum state learning. Similarly, the minimax regret can also be lower bounded by the sequential fat-shattering dimension, provided that $\ell_t(h_t(x_t), y_t) = |h_t(x_t) - y_t|$ and that $\mathcal{P}$ is taken to be the whole set of all distributions on $\mathcal{X}$.

$$\mathcal{V}_T \geq \frac{1}{4\sqrt{2}} \sup_{\delta > 0} \left\{ \sqrt{\delta^2 T \min\{\text{sfat}_{\delta}(\mathcal{H}, \mathcal{X}), T\}} \right\}. \quad (3)$$

3 Lower Bounds for Online Learning of Quantum States

In this section, we employ a distinct proof strategy than Aaronson et al. [2018] to obtain lower bounds on sequential fat-shattering dimension—recall that they proved $\text{sfat}_{\delta}(\mathcal{H}_n, \mathcal{X}) = \Theta(\frac{n}{\delta^2})$. This allows us to extend those bounds naturally to various subsettings of the online quantum state learning problem, characterized by a restricted hypothesis class $\mathcal{H} \subset \mathcal{H}_n$ and a constrained sample space $\mathcal{X}$. Such settings frequently arise in practical applications, where the focus is on characterizing specific subsets of quantum states. Furthermore, experimental constraints often limit the implementable set of measurement operations. We derive bounds on $\text{sfat}(\mathcal{H}, \mathcal{X})$ for several such practically relevant subsettings, leading up to the most general formulation of the learning problem.

Intuition of proofs: Our proofs are based on the construction of $\mathcal{X} \times [0, 1]$ -valued trees, which can be looked upon as a means to extract information about a quantum state ρ at every layer. After a certain number of layers, full information about the state is recovered. The number of layers achieved serves as a lower bound to the sequential fat-shattering dimension. Therefore, the goal is to construct the largest tree before full information is recovered.

In our construction, we define gaining information simply as approximating the coefficients of the density matrix ρ . At each layer, we approximate a different coefficient as seen in Section 3.2. This is how we construct the $\mathcal{X}$ -valued part of the tree. We can then approximate each coefficient to an error ϵ using the Halving Tree defined in Section 3.1 (Section 3.3 shows how to combine both trees). The intuition behind the choice of the coefficients we approximate comes from the study of chordal graphs in Section 3.4. An easier way to understand it is that we consider only the first row of the matrix, as it is of rank 1. Finally, we find the optimal error ϵ to obtain the lower bound.

To set the stage for the following sections, we first establish a few notations: since we will frequently consider pure states, the quantum state $\omega(\epsilon)$ will be denoted by its associated vector $|\psi(\epsilon)\rangle$, where $\omega(\epsilon) = |\psi(\epsilon)\rangle\langle\psi(\epsilon)|$. Furthermore, we define $N = 2^n$ to be the dimension of the Hilbert space under consideration. In addition, we will denote the pair of binary trees $(\mathbf{x}, \mathbf{v})$ by a single tree $\mathbf{T}$. We call $\mathbf{x}$ the $\mathcal{X}$ -valued part of $\mathbf{T}$, and $\mathbf{v}$ the real-valued part of $\mathbf{T}$.

We highlight that while many of the theorems in this section can be recovered from existing results, our contribution lies in introducing a unified proof framework. This scheme not only underpins our main result in Section 4, but also offers a versatile foundation that can be adapted to various sub-settings of interest.

3.1 Learning with respect to a single measurement

We start with the learning problem of estimating the expectation value of an unknown n -qubit quantum state with respect to a fixed measurement operator E . This learning problem is key to practical tasks such as quantum state discrimination and hypothesis testing [Barnett and Croke, 2009, Bae and Kwek, 2015]. Formally, we take $\mathcal{X} = \{E\}$ to be the sample space and keep $\mathcal{H}_n$ as the hypothesis class. As mentioned previously, we are focused on providing a lower bound on $\text{sfat}_{\delta}(\mathcal{H}_n, \mathcal{X})$, which could then be used to lower bound $\mathcal{V}_T$. We achieve this by constructing what

we call the Halving tree $\mathbf{T}_h$, which has a constant $\mathcal{X}$ -valued part, and a real-valued part as shown in Figure 1. Every halving tree $\mathbf{T}_h[i, T] = (\mathbf{x}, \mathbf{v})$ will thus be entirely determined by its depth T and the constant measurement $|i\rangle\langle i|$. The name follows from the distinctive structure exhibited by the real part of the tree $\mathbf{T}_h[i, T]$, as shown in Figure 1. This construction will prove useful as a crucial building block for establishing the regret bounds in more general settings (see Theorems 3.4, 3.6 and 4.1).




Figure 1: Real-valued part of the halving tree $\mathbf{T}_h$, up to $\frac{1}{N}$ factor.




Figure 2: $\mathcal{X}$ -valued part of the Von Neumann tree $\mathbf{T}_{vn}$.

Theorem 3.1. Let $E \in \text{Herm}_{\mathbb{C}}(2^n)$, $\text{Spec}(E) \subset [0, 1]$ be a fixed measurement, and $\mathcal{X} = \{E\}$ be the sample space. Let $\mathcal{H}_n = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n\}$ be the hypothesis class, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Then we have $\text{sfat}_{\delta}(\mathcal{H}_n, \mathcal{X}) = \Omega(\log_2(\frac{1}{\delta}))$.

We prove this result in Appendix A.

Remark 3.2. Note that the lower bound on the fat-shattering dimension obtained above is independent of n , and therefore still holds if the hypothesis class is induced by 1-qubit pure states.

3.2 Learning uniform superposition states

We now shift our attention to a harder setting. Consider the sample space $\mathcal{X}$ consisting of the N measurements corresponding to an orthogonal basis of $\mathcal{C}_n$. The hypothesis class will be induced by the uniform superpositions of basis states. Such states play an important role in fundamental quantum algorithms [Simon, 1994, Shor, 1997, Grover, 1997], and for quantum random number generators [Mannalatha et al., 2023]. We now provide a lower bound to the sequential fat-shattering dimension for this specific setting. We accomplish this by constructing what we call the Von Neumann tree $\mathbf{T}_{vn}$. The name follows from the fact that, while its real-valued part is constant, all nodes in the $\mathcal{X}$ -valued part of $\mathbf{T}_{vn}$ are labelled by Von Neumann measurements as shown in Figure 2.

Theorem 3.3. Let $(|0\rangle, \dots, |N-1\rangle)$ be an orthogonal basis of $\mathcal{C}_n$, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Denote the sample space as $\mathcal{X} = \{|i\rangle\langle i|, i \in \llbracket 0, N-1 \rrbracket\}$. Let $\mathcal{H} = \{\text{Tr}_{\omega}, \omega = \frac{1}{\sqrt{|I|}} \sum_{i \in I} |i\rangle, I \subset \llbracket 0, N-1 \rrbracket\}$ be the hypothesis class. Then we have $\text{sfat}_{\delta}(\mathcal{H}, \mathcal{X}) = \Omega(\min(\frac{1}{\delta}, 2^n))$.

We prove this result in Appendix B.

3.3 Learning general states using Von Neumann measurements

Building on the setting established in the previous section, we consider the sample space $\mathcal{X}$ consisting of the N measurements corresponding to an orthogonal basis of $\mathcal{C}_n$. The hypothesis class in the present setting is however induced by the set of all pure quantum states. The corresponding learning problem involves a learner estimating the expectation values of an unknown n -qubit quantum state with respect to N measurement operators, where the hypothesis is chosen from the set of all n -qubit pure quantum states. Related problems have been considered, for example, in Zhao et al. [2023].

We now provide a lower bound to the sequential fat-shattering dimension for this specific setting. We accomplish this by constructing what we call the Von Neumann Halving tree $\mathbf{T}_{vnh}$, which is constructed by combining $\mathbf{T}_h$ and $\mathbf{T}_{vn}$ as shown in Figure 3. Our construction illustrates the utility of $\mathbf{T}_h$, demonstrating its effectiveness as a tool to multiply existing lower bounds by a factor of n .

$$|0\rangle\langle 0| \xrightarrow{\mathbf{T}_h[0, T]} |1\rangle\langle 1| \xrightarrow{\mathbf{T}_h[1, T]} \dots \xrightarrow{\mathbf{T}_h[N-3, T]} |N-2\rangle\langle N-2| \xrightarrow{\mathbf{T}_h[N-2, T]}$$

Figure 3: Von Neumann Halving Tree (Horizontal)

Theorem 3.4. *Let $(|0\rangle, \dots, |N-1\rangle)$ be an orthogonal basis of $\mathcal{C}_n$ and $\mathcal{X} = \{|i\rangle\langle i|, i \in \llbracket 0, N-1 \rrbracket\}$ be the sample space. Let $\mathcal{H} = \{\text{Tr}_\omega, \omega \in \mathcal{C}_n, \text{Tr}(\omega^2) = 1\}$ be the hypothesis class. Here $\mathcal{C}_n$ is the set of all n -qubit quantum states. Then, for $\delta = 2^{-\frac{n}{\eta}}$, we have $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X}) = \Omega(\frac{n}{\delta\eta}) \forall \eta < 1$.*

We prove this result in Appendix C.

3.4 Tightness of sequential fat-shattering dimension of quantum states

Having considered several restricted settings in the previous sections, we now turn to the most general formulation of the online quantum state learning problem. Formally, we define the sample space to be the space of all 2-outcome measurements $\mathcal{X}$, while the Hypothesis class is given as $\mathcal{H}_n$. Our objective is to derive a tight lower bound on $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X})$. To accomplish this, we will use the techniques developed in the previous sections. In addition, we will establish new results on completion of partial matrices, which are essential for our analysis. We start with the latter.

Let us first introduce a few necessary definitions on partial matrices and their completions. A partial matrix is a matrix in which certain entries are specified while the other entries are free to be chosen. It is called partial symmetric if it is symmetric on the specified entries. And a completion of a partial matrix refers to a specific assignment of values to its unspecified entries. The main result we will use is the following, proved in Appendix D.

Lemma 3.5. *Any real partial symmetric matrix ω satisfying the following conditions*

1. $w_{11} = \frac{1}{2}$, and $w_{ii} = \frac{1}{2(N-1)} \forall i \in \llbracket 2, N \rrbracket$,
2. *Elements are specified on the set $\{w_{1i}, w_{i1}\}$, where $i \in [N]$,*
3. $\forall i \in \llbracket 2, N \rrbracket, |w_{1i}| \leq \frac{1}{2\sqrt{N-1}}$,

can be completed to a density matrix. We will denote $\text{part}(w_{12}, w_{13}, \dots, w_{1N})$ such a matrix.

We now derive the lower bound for $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X})$. The key idea is to construct a new tree analogous to the Von Neumann halving tree, with a crucial distinction that it accommodates more general measurements, extending beyond the $|i\rangle\langle i|$ type measurements that have been considered thus far. Furthermore, given this tree, we apply Lemma 3.5 to ensure that all paths on the tree can be associated to a valid density matrix. We show that this allows us to obtain the quadratic dependence on $\frac{1}{\delta}$ in the lower bound of $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X})$, thus establishing the almost tightness.

Theorem 3.6. *Let $\mathcal{X} = \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$ be the sample space. Define $\mathcal{H}_n = \{\text{Tr}_\omega, \omega \in \mathcal{C}_n\}$ as the hypothesis class, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Then, for $\delta = 2^{-\frac{n}{\eta}}$, we have $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X}) = \Omega(\frac{n}{\delta\eta}), \forall \eta < 2$.*

We prove this result in Appendix E. Note that this result directly implies tightness of regret for L_1 loss [Aaronson et al., 2018], as well as for L_2 loss as shown in Appendix G.

Remark 3.7. As mentioned in the beginning of this section, our method for lower bounding the sequential fat-shattering dimension differs from that employed in Aaronson [2007], and is adaptable to various restricted online quantum state learning settings. For cases involving further restrictions, whether on the sample space, the hypothesis class or the relationship between δ and n (Recall that Aaronson [2007] required $\delta \geq \sqrt{n}2^{-(n-5)/35}/8$) we conjecture that the techniques from Appendix E involving matrix completion would serve as a valuable foundation.

4 Online Learning of Pure States is as Hard as Mixed States

While the tightness of the sequential fat-shattering dimension, and hence of the minimax regret were already known, we emphasize that the results were derived for a hypothesis class being induced by

the set of all n -qubit quantum states. In this section, we consider a more restrictive setting where the hypothesis class is induced solely by the set of all n -qubit pure states. The derivation closely follows the techniques developed in the previous section with a notable distinction: whereas the previous sections relied on matrix completion results (without defining the states $\omega(\epsilon)$ explicitly), here we shall construct the states $\omega(\epsilon) = |\psi(\epsilon)\rangle\langle\psi(\epsilon)|$ explicitly. We show that the bounds on the sequential fat-shattering dimension and consequently the minimax regret remain tight in this setting.

Theorem 4.1. *Let $\mathcal{X} = \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$ be the sample space. Define $\mathcal{H} = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n, \text{Tr}[\omega^2] = 1\}$ as the hypothesis class, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Then, for $\delta = 2^{-\frac{n}{\eta}}$, we have $\text{sfat}_{\delta}(\mathcal{H}, \mathcal{X}) = \Omega(\frac{n}{\delta^{\eta}}), \forall \eta < 2$.*

Now, building on the result in Theorem 4.1, we proceed to demonstrate the tightness of the minimax regret $\mathcal{V}_T$, proving that it is asymptotically the same as for mixed states:

Corollary 4.2. *Let $\mathcal{X} = \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$ be the sample space. Define $\mathcal{H} = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n, \text{Tr}[\omega^2] = 1\}$ as the hypothesis class, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Then we have $\mathcal{V}_T = \Omega(\sqrt{nT})$, assuming the loss function under consideration is the L_1 -loss. This result still holds for the L_2 loss, as long as $T \leq 4^n$.*

We prove these results in Appendices F and G respectively.

An important aspect to notice here is that the significance of the sequential fat-shattering dimension extends beyond regret analysis. Several corollaries follow from Theorem 4.1. One such example is the fact that Theorem 1 from Aaronson et al. [2018], which has been shown to be optimal for mixed states, is also optimal for pure states.

Corollary 4.3. *Let ρ be an n -qubit mixed state, and let $E_1, E_2, \dots$ be a sequence of 2-outcome measurements that are revealed to the learner one by one, each followed by a value $b_t \in [0, 1]$ such that $|\text{Tr}(E_t \rho) - b_t| \leq \varepsilon/3$. Then there is an explicit strategy for outputting hypothesis states $\omega_1, \omega_2, \dots$ such that $|\text{Tr}(E_t \omega_t) - \text{Tr}(E_t \rho)| > \varepsilon$ for at most $O(\frac{n}{\varepsilon^2})$ values of t . This mistake bound is almost asymptotically optimal, even if we restrict ρ to be a pure state.*

5 Extensions of Online Learning of Quantum States

5.1 The ϵ -realizable setting

In practice, an experimenter does not have access to the exact value of $\text{Tr}(E_t \rho)$. Assuming completely adversarial feedback is overly pessimistic and renders the learning problem particularly challenging. This motivates the consideration of the ϵ -realizable setting, where the learner incurs the L_1 loss with respect to a fixed quantum state ρ , and receives noisy feedback y_t at each round. Specifically, the feedback y_t is an approximation of $\text{Tr}(E_t \rho)$ with error parametrized by ϵ . For instance, y_t may be drawn from a uniform distribution over the interval $[\text{Tr}(E_t \rho) - \epsilon, \text{Tr}(E_t \rho) + \epsilon]$, or from a Gaussian distribution $\mathcal{N}(\text{Tr}(E_t \rho), \epsilon)$. More generally, any non-deterministic distribution μ centered at $\text{Tr}(E_t \rho)$ is permissible.

We show that even in this more favorable setting, the regret admits the same lower bound as in the fully adversarial case—*asymptotically independent of ϵ* .

$$\bar{\mathcal{V}}_T = \left\langle \inf_{\mathcal{Q} \in \Delta(\mathcal{C}_n)} \sup_{\mathcal{D}_t \in \mathcal{P}} \mathbb{E}_{\omega_t \sim \mathcal{Q}} \mathbb{E}_{E_t \sim \mathcal{D}_t} \mathbb{E}_{y_t \sim \mu} \right\rangle_{t=1}^T \sup_{\rho \in \mathcal{C}_n} \left[\sum_{t=1}^T |\text{Tr}(E_t \omega_t) - \text{Tr}(E_t \rho)| \right]. \quad (4)$$

Theorem 5.1. *The minimax regret in the ϵ -realizable setting satisfies $\bar{\mathcal{V}}_T = \Theta(\sqrt{nT})$. In particular, these bounds are independent of ϵ .*

We prove this result in Appendix H.

5.2 Smoothed Online Learning of Quantum States

The online learning framework discussed thus far is fully adversarial, allowing the adversary to select arbitrary two-outcome measurements at each round t . However, as highlighted in Section 1, practical

learning tasks often target specific properties of the quantum state ρ , rather than aiming for full state tomography. Within the PAC learning paradigm, such tasks are typically modelled by fixing a distribution $\mathcal{D}$ over the space of measurements, reflecting the learner's interest in a particular subset of observables.

To bridge the adversarial and distributional regimes, we adopt the lens of *smoothed analysis*, which interpolates between worst-case and average-case models. Specifically, we require that at each round t , the adversary selects a distribution $\mathcal{D}_t$ over measurements that remains close to the target distribution $\mathcal{D}$. This models an adversary with limited power to perturb the measurement process.

This smoothed perspective is not merely a theoretical convenience — it captures an important aspect of experimental practice. Indeed, quantum devices rarely implement measurements with perfect fidelity. Instead, small temporal fluctuations in control parameters, thermal drift, or calibration errors can cause the realized measurement to deviate slightly from the intended one Leibfried et al. [2003]. From a learning-theoretic standpoint, these deviations can be modelled as stochastic perturbations: at each round, the performed measurement is sampled from a distribution that is ε -close to the nominal one. Thus, smoothed analysis offers a principled framework for reasoning about online learning in the presence of structured noise in the measurement process.

Definition 5.2 (Smooth distributions). A distribution μ is said to be σ -smooth with respect to a fixed distribution $\mathcal{D}$ for a $\sigma \in (0, 1]$ if and only if [Haghtalab et al., 2020]:

1. μ is absolutely continuous with respect to $\mathcal{D}$, i.e. every measurable set A such that $\mathcal{D}(A) = 0$ satisfies $\mu(A) = 0$
2. The Radon Nikodym derivative $d\mu/d\mathcal{D}$ satisfies the following relation:

$$\text{ess sup } \frac{d\mu}{d\mathcal{D}} \leq \frac{1}{\sigma}. \quad (5)$$

Let $\mathcal{B}(\sigma, \mathcal{D})$ be the set of all σ -smooth distributions with respect to $\mathcal{D}$. In smoothed online learning, the adversary is restricted by the condition $\mathcal{D}_t \in \mathcal{B}(\sigma, \mathcal{D})$. Note that in this setting we recover the case of an oblivious adversary for $\sigma = 1$, while we get the completely adversarial case for $\sigma \rightarrow 0$.

Recall the expression of minimax regret in Equation (1). In the smoothed setting, the expression gets slightly modified accounting for the restriction imposed on the adversary:

$$\mathcal{V}_T = \left\langle \inf_{\mathcal{Q} \in \Delta(\mathcal{H})} \sup_{\mathcal{D}_t \in \mathcal{B}(\sigma, \mathcal{D})} \mathbb{E}_{h_t \sim \mathcal{Q}} \mathbb{E}_{x_t \sim \mathcal{D}_t} \right\rangle_{t=1}^T \left[\sum_{t=1}^T \ell_t(h_t(x_t)) - \inf_{h \in \mathcal{H}} \sum_{t=1}^T \ell_t(h(x_t)) \right]. \quad (6)$$

Here, the key difference with Equation (1) is that $\mathcal{D}_t$ is now restricted to the set of all σ -smooth distributions with respect to $\mathcal{D}$ instead of all possible distributions on $\mathcal{X}$. We recall that for quantum state learning, we have $\mathcal{X} \subset \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$, $\mathcal{H}_n = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n\}$ and a target state ρ . For the sake of brevity, we will continue using the notation x_t to indicate input data and h to indicate the hypothesis. The derivation here closely follows the approach in Block et al. [2022] (which derives the regret bounds for classical smoothed online supervised learning) with one important difference; in the original derivation, for a given input $x_t \in \mathcal{X}$, the authors distinguish between a predicted label $\hat{y}_t \in \mathcal{Y}$ and $h_t(x_t) \in \mathcal{Y}$ where $\mathcal{Y}$ is the label space. We do not make this distinction as our labels are always related to our inputs via the hypothesis.

Theorem 5.3. Let $\mathcal{X} = \{E \in \text{Herm}_{\mathbb{C}}(2^n), \text{Spec}(E) \subset [0, 1]\}$ be the sample space. Define the hypothesis class $\mathcal{H} = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n, \text{Tr}[\omega^2] = 1\}$ as the set of n -qubit pure states, where $\mathcal{C}_n$ is the set of all n -qubit quantum states. Furthermore let $\sigma \in (0, 1]$ be the smoothness parameter. Then we

$$\text{have } \mathcal{V}_T = O\left(\sqrt{\frac{nT \log T}{\sigma}}\right).$$

We prove this result in Appendix K.

6 Conclusion and Limitations

Lower bounds on fat-shattering dimension: In this work, we established lower bounds on sequential fat-shattering dimension of various subproblems within the quantum state learning framework.

Crucially, we showed that pure and mixed states almost share the same asymptotical dimension. Note that, although our construction directly implies that the regular δ -fat-shattering dimension of pure states scales as $\Omega(\frac{1}{\delta^2})$, whether we can recover $\Omega(\frac{n}{\delta^2})$ for pure n -qubits in the offline setting remains an open question.

Consequences on regret: This lower bound on sequential fat-shattering dimension has several implications, including the key result that learning pure and mixed states in the online setting will incur the same asymptotical regret for the L_1 -loss. However, the lower bound for the L_2 loss might leave room for improvement. Additionally, sequential fat-shattering dimension may serve as a fundamental tool for deriving bounds on other key complexity measures in quantum state learning.

Smoothed online learning: Finally, we extend our analysis from standard online learning of quantum states to the smoothed online learning setting. To our knowledge, this work represents the first application of smoothed analysis to quantum state learning. In this setting, we establish an upper bound on the regret. However a key open question is whether this bound is tight.

Possible extensions to classical learning theory: While a general (mixed) quantum state can be represented as a positive semi-definite Hermitian matrix of rank (up to) 2^n , a pure state can be represented as a positive semi-definite Hermitian matrix of rank 1. The fact that this restriction on the rank of a matrix has no impact on the sequential fat-shattering dimension lower bound (and hence the regret lower bound) is highly non-intuitive and does not seem to apply in varied settings of classical learning: Alon et al. [2016] shows that rank of sign matrices impact VC dimension, Bartlett et al. [2022] shows that the ϵ -fat shattering dimension of learned SCW matrices of rank k is $O(n \cdot (m + k \log(\frac{n}{k}) + \log(\frac{1}{\epsilon})))$, and Srebro and Shraibman [2005] shows that the pseudodimension of matrices scale linearly with the rank.

The reason why this scaling of dimension with rank vanishes in quantum is the additional assumptions on the density matrices, which are semi-definite positive, Hermitian, and of trace 1. Such matrices have however been extensively studied in Learning Theory [Tsuda et al., 2004, Shen et al., 2008, 2009]. In addition, common matrices satisfying those conditions are normalized covariance and kernel matrices, and Laplacian matrices differ only for the fact that their trace is constant equal to $2m$, where m is the number of edges of the associated graph. We leave as future work to see how our results translate to these other settings of interest.

Acknowledgements

The work of Maxime Meyer, Soumik Adhikary, Naixu Guo, and Patrick Rebentrost was supported by the National Research Foundation, Singapore and A*STAR under its CQT Bridging Grant, its AI Singapore Programme (AISG Award No: AISG3-PhD-2026-01-068T) and its Quantum Engineering Programme under grant NRF2021-QEP2-02-P05. Authors thank Josep Lumberras, Aadil Oufkir, and Jonathan Allcock for their insightful discussions and kind suggestions.

References

- Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, September 2007. ISSN 1471-2946. doi: 10.1098/rspa.2007.0113. URL <http://dx.doi.org/10.1098/rspa.2007.0113>.
- Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. *Advances in neural information processing systems*, 31, 2018.
- Noga Alon, Shay Moran, and Amir Yehudayoff. Sign rank versus VC dimension. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 47–80, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v49/alon16.html>.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012. doi: 10.4086/toc.2012.v008a006. URL <https://theoryofcomputing.org/articles/v008a006>.
- Joonwoo Bae and Leong-Chuan Kwek. Quantum state discrimination and its applications. *Journal of Physics A: Mathematical and Theoretical*, 48(8):083001, January 2015. ISSN 1751-8121. doi: 10.1088/1751-8121/48/8/083001. URL <http://dx.doi.org/10.1088/1751-8121/48/8/083001>.
- Akshay Bansal, Ian George, Soumik Ghosh, Jamie Sikora, and Alice Zheng. Online learning of a panoply of quantum objects, 2024. URL <https://arxiv.org/abs/2406.04245>.
- Stephen M. Barnett and Sarah Croke. Quantum state discrimination. *Adv. Opt. Photon.*, 1(2):238–278, 2009. doi: 10.1364/AOP.1.000238.
- P. Bartlett, P. Indyk, and T. Wagner. Generalization bounds for data-driven numerical linear algebra. *Conference on Learning Theory*, 2022. URL <https://par.nsf.gov/biblio/10341774>.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, 2003. URL <https://api.semanticscholar.org/CorpusID:463216>.
- Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. Smoothed online learning is as easy as statistical learning. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 1716–1786. PMLR, 02–05 Jul 2022. URL <https://proceedings.mlr.press/v178/block22a.html>.
- Dik Bouwmeester, Jian-Wei Pan, Klaus Mattle, Manfred Eibl, Harald Weinfurter, and Anton Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, December 1997. ISSN 1476-4687. doi: 10.1038/37539. URL <http://dx.doi.org/10.1038/37539>.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 01 2006. ISBN 978-0-521-84108-5. doi: 10.1017/CBO9780511546921.
- Sitan Chen, Brice Huang, Jerry Li, Allen Liu, and Mark Sellke. When does adaptivity help for quantum state learning? In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 391–404, 2023. doi: 10.1109/FOCS57990.2023.00029.

- Xinyi Chen, Elad Hazan, Tongyang Li, Zhou Lu, Xinzhao Wang, and Rui Yang. Adaptive online learning of quantum states. *Quantum*, 8:1471, September 2024. ISSN 2521-327X. doi: 10.22331/q-2024-09-12-1471. URL <http://dx.doi.org/10.22331/q-2024-09-12-1471>.
- Jens Eisert, Dominik Hangleiter, Nathan Walk, Ingo Roth, Damian Markham, Rhea Parekh, Ulysse Chabaud, and Elham Kashefi. Quantum certification and benchmarking. *Nature Reviews Physics*, 2(7):382–390, June 2020. ISSN 2522-5820. doi: 10.1038/s42254-020-0186-4. URL <http://dx.doi.org/10.1038/s42254-020-0186-4>.
- Steven T. Flammia and Yi-Kai Liu. Direct fidelity estimation from few pauli measurements. *Physical Review Letters*, 106(23), June 2011. ISSN 1079-7114. doi: 10.1103/physrevlett.106.230501. URL <http://dx.doi.org/10.1103/PhysRevLett.106.230501>.
- Robert Grone, Charles R. Johnson, Eduardo M. Sá, and Henry Wolkowicz. Positive definite completions of partial hermitian matrices. *Linear Algebra and its Applications*, 58:109–124, 1984. ISSN 0024-3795. doi: [https://doi.org/10.1016/0024-3795\(84\)90207-6](https://doi.org/10.1016/0024-3795(84)90207-6). URL <https://www.sciencedirect.com/science/article/pii/0024379584902076>.
- David Gross, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Physical Review Letters*, 105(15), October 2010. ISSN 1079-7114. doi: 10.1103/physrevlett.105.150401. URL <http://dx.doi.org/10.1103/PhysRevLett.105.150401>.
- Lov K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79(2):325–328, July 1997. ISSN 1079-7114. doi: 10.1103/physrevlett.79.325. URL <http://dx.doi.org/10.1103/PhysRevLett.79.325>.
- Naixu Guo, Feng Pan, and Patrick Rebentrost. Estimating properties of a quantum state by importance-sampled operator shadows, 2024. URL <https://arxiv.org/abs/2305.09374>.
- Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 913–925, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341325. doi: 10.1145/2897518.2897585. URL <https://doi.org/10.1145/2897518.2897585>.
- Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis of online and differentially private learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9203–9215. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/685bfde03eb646c27ed565881917c71c-Paper.pdf.
- Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *Journal of the ACM*, 71(3):1–34, 2024.
- Hsin-Yuan Huang, Richard Kueng, and John Preskill. Efficient estimation of pauli observables by derandomization. *Physical Review Letters*, 127(3), July 2021. ISSN 1079-7114. doi: 10.1103/physrevlett.127.030503. URL <http://dx.doi.org/10.1103/PhysRevLett.127.030503>.
- Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low rank matrix recovery from rank one measurements. *Applied and Computational Harmonic Analysis*, 42(1):88–116, 2017.
- D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003. doi: 10.1103/RevModPhys.75.281. URL <https://link.aps.org/doi/10.1103/RevModPhys.75.281>.
- Erik Lucero, M. Hofheinz, M. Ansmann, Radoslaw C. Bialczak, N. Katz, Matthew Neeley, A. D. O’Connell, H. Wang, A. N. Cleland, and John M. Martinis. High-fidelity gates in a single josephson qubit. *Phys. Rev. Lett.*, 100:247001, Jun 2008. doi: 10.1103/PhysRevLett.100.247001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.247001>.

- Josep Lumbreras, Erkka Haapasalo, and Marco Tomamichel. Multi-armed quantum bandits: Exploration versus exploitation when learning properties of quantum states. *Quantum*, 6:749, June 2022. ISSN 2521-327X. doi: 10.22331/q-2022-06-29-749. URL <http://dx.doi.org/10.22331/q-2022-06-29-749>.
- Josep Lumbreras, Mikhail Terekhov, and Marco Tomamichel. Learning pure quantum states (almost) without regret, 2024. URL <https://arxiv.org/abs/2406.18370>.
- Vaisakh Mannalatha, Sandeep Mishra, and Anirban Pathak. A comprehensive review of quantum random number generators: concepts, classification and the origin of randomness. *Quantum Information Processing*, 22(12), December 2023. ISSN 1573-1332. doi: 10.1007/s11128-023-04175-y. URL <http://dx.doi.org/10.1007/s11128-023-04175-y>.
- Daniel Miller, Laurin E. Fischer, Kyano Levi, Eric J. Kuehnke, Igor O. Sokolov, Panagiotis Kl. Barkoutsos, Jens Eisert, and Ivano Tavernelli. Hardware-tailored diagonalization circuits. *npj Quantum Information*, 10(1), November 2024. ISSN 2056-6387. doi: 10.1038/s41534-024-00901-1. URL <http://dx.doi.org/10.1038/s41534-024-00901-1>.
- Philipp Neumann, Johannes Beck, Matthias Steiner, Florian Rempp, Helmut Fedder, Philip R. Hemmer, Jörg Wrachtrup, and Fedor Jelezko. Single-shot readout of a single nuclear spin. *Science*, 329(5991):542–544, 2010. doi: 10.1126/science.1189075. URL <https://www.science.org/doi/abs/10.1126/science.1189075>.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning via sequential complexities. *J. Mach. Learn. Res.*, 16(1):155–186, 2015a.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Sequential complexities and uniform martingale laws of large numbers. *Probability theory and related fields*, 161:111–153, 2015b.
- Asad Raza, Matthias C. Caro, Jens Eisert, and Sumeet Khatri. Online learning of quantum processes, 2024. URL <https://arxiv.org/abs/2406.04250>.
- Chunhua Shen, Alan Welsh, and Lei Wang. Psdboost: Matrix-generation linear programming for positive semidefinite matrices learning. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008. URL https://proceedings.neurips.cc/paper_files/paper/2008/file/605ff764c617d3cd28dbbdd72be8f9a2-Paper.pdf.
- Chunhua Shen, Junae Kim, Lei Wang, and Anton Hengel. Positive semidefinite metric learning with boosting. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009. URL https://proceedings.neurips.cc/paper_files/paper/2009/file/051e4e127b92f5d98d3c79b195f2b291-Paper.pdf.
- Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. ISSN 1095-7111. doi: 10.1137/s0097539795293172. URL <http://dx.doi.org/10.1137/S0097539795293172>.
- D.R. Simon. On the power of quantum computation. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 116–123, 1994. doi: 10.1109/SFCS.1994.365701.
- Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In Peter Auer and Ron Meir, editors, *Learning Theory*, pages 545–560, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. ISBN 978-3-540-31892-7.
- Koji Tsuda, Gunnar Rätsch, and Manfred K. K Warmuth. Matrix exponential gradient updates for on-line learning and bregman projection. In L. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 17. MIT Press, 2004. URL https://proceedings.neurips.cc/paper_files/paper/2004/file/bd70364a8fcb02366697df66f50b4d4-Paper.pdf.

Bujiao Wu, Jinzhao Sun, Qi Huang, and Xiao Yuan. Overlapped grouping measurement: A unified framework for measuring quantum states. *Quantum*, 7:896, January 2023. ISSN 2521-327X. doi: 10.22331/q-2023-01-13-896. URL <https://doi.org/10.22331/q-2023-01-13-896>.

Liming Zhao, Naixu Guo, Ming-Xing Luo, and Patrick Rebentrost. Provable learning of quantum states with graphical models, 2023. URL <https://arxiv.org/abs/2309.09235>.

A Proof of Theorem 3.1

Proof. Without loss of generality, we set $\mathcal{X} = \{|i\rangle\langle i|\}$, with $i \in \llbracket 0, N-2 \rrbracket$. Define $\mathbf{x}$ as the complete binary tree of depth T such that $\forall t \in [T], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{x}_t(\epsilon) = |i\rangle\langle i|. \quad (7)$$

Furthermore, define $\mathbf{v}$ (Figure 1) as the complete binary tree of depth T such that $\forall t \in [T], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{v}_t(\epsilon) = \frac{1}{2^t} \sum_{k=0}^{t-1} \epsilon_k 2^{t-k-1} = \sum_{k=0}^{t-1} \epsilon_k 2^{-k-1}. \quad (8)$$

Here, we set $\epsilon_0 = 1$.

Given $(\mathbf{x}, \mathbf{v})$, we now set:

$$|\psi(\epsilon)\rangle = \sqrt{\sum_{k=0}^{T-1} \epsilon_k 2^{-k-1}} |i\rangle + \sqrt{1 - \sum_{k=0}^{T-1} \epsilon_k 2^{-k-1}} |\perp\rangle. \quad (9)$$

Then, for $T = \lfloor \log_2(\frac{1}{\delta}) \rfloor$ (which implies $\delta \leq \frac{1}{2^T}$), $\forall \epsilon \in \{\pm 1\}^{T-1}, \forall t \in [T]$, we have:

$$\begin{aligned} \epsilon_t [\text{Tr}_{\omega(\epsilon)}(\mathbf{x}_t(\epsilon)) - \mathbf{v}_t(\epsilon)] &= \epsilon_t \left[\sum_{k=0}^{T-1} \epsilon_k 2^{-k-1} - \sum_{k=0}^{t-1} \epsilon_k 2^{-k-1} \right] \\ &= \epsilon_t \sum_{k=t}^{T-1} \epsilon_k 2^{-k-1} \\ &= 2^{-t-1} + \epsilon_t \sum_{k=t+1}^{T-1} \epsilon_k 2^{-k-1} \\ &\geq 2^{-t-1} - \sum_{k=t+1}^{T-1} 2^{-k-1} \\ &= 2^{-t-1} - (2^{-t-1} - 2^{-T}) \\ &\geq \delta, \end{aligned} \quad (10)$$

where the first inequality follows from the minimum value that the term $\epsilon_t \sum_{k=t+1}^{T-1} \epsilon_k 2^{-k-1}$ can take, and the last inequality is by direct computation and the assumption $T = \lfloor \log_2(\frac{1}{\delta}) \rfloor$. Thus, we show that the set $\mathcal{X}$ is δ -shattered by the hypothesis class $\mathcal{H}_n$ with $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X}) = \Omega(\log_2(\frac{1}{\delta}))$. $\square$

We can replace $\mathbf{v}$ by a slightly modified version of itself, where each node is scaled by a factor $\frac{1}{N}$. Therefore, the quantum state associated to a path has an amplitude corresponding to $|i\rangle$: $\text{Tr}(|i\rangle\langle i| \omega(\epsilon))$ bounded by $\frac{1}{N}$. We will write $\mathbf{T}_h[i, T] = (\mathbf{x}, \mathbf{v})$ the resulting pair-valued tree and call it the Halving Tree of depth T . The index i indicates the Von Neumann measurement associated to $\mathbf{x}$.

B Proof of Theorem 3.3

Proof. Let $T \in [N]$. Define $\mathbf{v}$ as the complete binary tree of depth T such that $\forall t \in [T], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{v}_t(\epsilon) = \frac{1}{2^T}. \quad (11)$$

Furthermore, denote $\mathbf{x}$ (Figure 2) the complete binary tree of depth T such that $\forall t \in [T], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{x}_t(\epsilon) = |t-1\rangle\langle t-1|. \quad (12)$$

We refer to the pair $(\mathbf{x}, \mathbf{v})$ as the Von-Neumann tree $\mathbf{T}_{vn}$.

Given the Von Neumann tree, we now associate each path ϵ to a pure quantum state:

$$|\psi(\epsilon)\rangle = \frac{1}{\sqrt{K+1}} \sum_{i=0}^{T-2} \mathbf{1}_{\epsilon_{i+1}=1} |i\rangle + \frac{1}{\sqrt{K+1}} |N-1\rangle, \quad (13)$$

where $K = \sum_{i=0}^{T-2} \mathbf{1}_{\epsilon_{i+1}=1}$, and $\mathbf{1}_{\epsilon=1}$ is an indicator function which takes value 1 if $\epsilon = 1$ and 0 otherwise. Then, for $\delta \leq \frac{1}{2T}$, $\forall \epsilon \in \{\pm 1\}^{T-1}, \forall t \in [T]$:

$$\begin{aligned} \epsilon_t [\text{Tr}_{\omega(\epsilon)}(\mathbf{x}_t(\epsilon)) - \mathbf{v}_t(\epsilon)] &= \epsilon_t \left[\frac{\mathbf{1}_{\epsilon_t=1}}{K+1} - \frac{1}{2T} \right] \\ &\geq \delta. \end{aligned} \quad (14)$$

Thus, we show that the set $\mathcal{X}$ is δ -shattered by the hypothesis class $\mathcal{H}$ with $\text{sfat}_\delta(\mathcal{H}, \mathcal{X}) = \Omega(\min(\frac{1}{\delta}, 2^n))$. □

C Proof of Theorem 3.4

Proof. Without any loss of generality, we can chose the sample space to be a set of Von Neumann measurements $\mathcal{X} = \{|i\rangle\langle i| : i \in [0, N-1]\}$. Denote $t' = \lfloor \frac{t-1}{T} \rfloor$ and $\tilde{t} = t-1-Tt'$. We then define $\mathbf{x}$ as the complete binary tree of depth $T(N-1)$ such that $\forall t \in [T(N-1)], \forall \epsilon \in \{\pm 1\}^{T(N-1)-1}$,

$$\mathbf{x}_t(\epsilon) = |t'\rangle\langle t'|. \quad (15)$$

Furthermore, denote $\mathbf{v}$ the complete binary tree of depth $T(N-1)$ such that $\forall t \in [T(N-1)], \forall \epsilon \in \{\pm 1\}^{T(N-1)-1}$,

$$\mathbf{v}_t(\epsilon) = \frac{1}{2^{n+1}} (1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k}) \quad (16)$$

We refer to the pair $(\mathbf{x}, \mathbf{v})$ as the Von Neumann halving tree $\mathbf{T}_{vnh}$. The name follows from the fact that the trees $\mathbf{x}$ and $\mathbf{v}$ as defined in Equations (15) and (16) can be constructed by replacing each node on the t -th layer of $\mathbf{T}_{vn}$, for both $\mathcal{X}$ valued and real valued parts, by the corresponding parts of the halving tree $\mathbf{T}_h[t, T]$.

Given the Von Neumann halving tree we can now associate each path ϵ to a pure quantum state:

$$|\psi(\epsilon)\rangle = \sum_{i=0}^{N-2} \sqrt{a_i} |i\rangle + \sqrt{1 - \sum_{i=0}^{N-2} a_i} |N-1\rangle \quad (17)$$

where $\forall i \in [0, N-2]$,

$$a_i = \mathbf{v}_{(i+1)T}(\epsilon) + \frac{1}{2^{T+n+1}} \epsilon_{(i+1)T}. \quad (18)$$

Then, for $\delta \leq \frac{1}{2^{n+T+1}}$, we can show that $\forall \epsilon \in \{\pm 1\}^{T(N-1)-1}, \forall t \in [T(N-1)]$:

$$\begin{aligned} \epsilon_t [\text{Tr}_{\omega(\epsilon)}(\mathbf{x}_t(\epsilon)) - \mathbf{v}_t(\epsilon)] &= \epsilon_t [\mathbf{v}_{(t'+1)T}(\epsilon) + \frac{1}{2^{T+n+1}} \epsilon_{(t'+1)T} - \frac{1}{2^{n+1}} (1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k})] \\ &= \epsilon_t [\frac{1}{2^{n+1}} (1 + \sum_{k=1}^T \epsilon_{k+Tt'} 2^{-k}) - \frac{1}{2^{n+1}} (1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k})] \\ &= \frac{\epsilon_t}{2^{n+1}} (\sum_{k=\tilde{t}+1}^T \epsilon_{k+Tt'} 2^{-k}) \\ &= \frac{1}{2^{n+1}} (\frac{1}{2^{\tilde{t}+1}} + \sum_{k=\tilde{t}+2}^T \epsilon_{k+Tt'} 2^{-k}) \\ &\geq \delta \end{aligned} \quad (19)$$

In particular, let $k \in \mathbb{N}^*$. Taking $\delta = \frac{1}{N^{1+\frac{1}{k}}}$ and $T = \lfloor \log_2(\frac{1}{4\delta N}) \rfloor$, we get $TN = \Omega(n\delta^{-\frac{1}{1+\frac{1}{k}}})$. Therefore, we have shown that the set $\mathcal{X}$ is δ -shattered by the Hypothesis class $\mathcal{H}_n$ with $\text{sfat}_\delta(\mathcal{H}_n) = \Omega(\frac{n}{\delta^\eta}), \forall \eta < 1$. $\square$

D Proof of Lemma 3.5

To prove Lemma 3.5, we rely on Theorem 7 in Grone et al. [1984]. We begin by introducing the necessary definitions. Let $\mathcal{G} = (V, E)$ be a finite undirected graph. A *cycle* in $\mathcal{G}$ is a sequence of distinct vertices $v_1, v_2, \dots, v_s \in V$ such that $\{v_i, v_{i+1}\} \in E$ for all $i \in [s-1]$, and $\{v_s, v_1\} \in E$. A cycle is said to be *minimal* if and only if it has no chord, where a chord is an edge $\{v_i, v_j\} \in E$ with $|i-j| > 1$ and $\{i, j\} \neq \{1, s\}$.

A matrix ω is said to be $\mathcal{G}$ -partial when its entries ω_{ij} are determined if and only if $\{i, j\} \in E$, while other elements are undetermined. A $\mathcal{G}$ -partial matrix ω is said to be non-negative if and only if (a) $\omega_{ij} = \overline{\omega_{ji}}, \forall \{i, j\} \in E$ and (b) for any clique $\mathcal{C}$ of $\mathcal{G}$, the principal submatrix of ω corresponding to $\mathcal{C}$ (which has entries corresponding to $\mathcal{C}$) is positive semidefinite. Recall that a clique $\mathcal{C}$ of $\mathcal{G}$ is a complete subgraph of $\mathcal{G}$. The corresponding principal submatrix is obtained by keeping only the indices in $\mathcal{C}$.

A *completion* of a $\mathcal{G}$ -partial matrix ω is a full Hermitian matrix M such that $M_{ij} = \omega_{ij}$ for all $\{i, j\} \in E$. We say that M is a non-negative completion if and only if M is also positive semidefinite. A graph $\mathcal{G}$ is said to be *completable* if and only if any $\mathcal{G}$ -partial non-negative matrix has a non-negative completion. With these definitions in place, we proceed to state the relevant results in Grone et al. [1984].

Lemma D.1 (Grone et al. [1984]). *A graph $\mathcal{G}$ is completable if and only if every minimal cycle in the graph is of length < 4 .*

Now that we have covered the necessary background, we can proceed to prove Lemma 3.5.

Proof. Let ω be a partial matrix as stated in Lemma 3.5. Consider the graph $\mathcal{G} = (V, E)$, where $V = [N]$ and $E = \{\{i, j\}, i \in [N], j \in \{1, i\}\} \cup \{\{i, i\}, i \in [N]\}$.




Figure 4: Graph representation of $\mathcal{G} = (V, E)$

One can easily check that $\mathcal{G}$ is completable using Lemma D.1. Therefore, if we prove that ω is non-negative $\mathcal{G}$ -partial, it will have a positive semidefinite Hermitian completion. Combining with the first two conditions in Lemma 3.5, we can show that the completion of ω is Hermitian, positive semidefinite, with trace equal to 1, and therefore is a density matrix. Now to show that ω is non-negative $\mathcal{G}$ -partial, let $\mathcal{C} = \{1, i\}$ be a clique of $\mathcal{G}$. Since $|w_{1i}| \leq \frac{1}{2\sqrt{N-1}}$ by assumption, the principal submatrix of ω corresponding to $\mathcal{C}$, i.e.,

$$\begin{pmatrix} \frac{1}{2} & \omega_{1i} \\ \overline{\omega_{1i}} & \frac{1}{2(N-1)} \end{pmatrix} \quad (20)$$

can be shown to be non-negative. $\square$

$$\begin{bmatrix} \frac{1}{2} & w_{12} & \cdots & \cdots & w_{1N} \\ w_{12} & \frac{1}{2(N-1)} & ? & \cdots & ? \\ \vdots & ? & \frac{1}{2(N-1)} & \cdots & ? \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ w_{1N} & ? & ? & \cdots & \frac{1}{2(N-1)} \end{bmatrix}$$

Figure 5: General form of the partial matrix described in Lemma 3.5. The interrogation marks denote the unspecified entries in the matrix.

E Proof of Theorem 3.6

Proof. Set $\mathcal{X} = \{E_{0,i}, i \in [N-1]\}$ where $E_{0,i} = \frac{1}{2}(|0\rangle\langle 0| + |i\rangle\langle i| + |0\rangle\langle i| + |i\rangle\langle 0|)$. Define $t' = \lfloor \frac{t-1}{T} \rfloor$ and $\tilde{t} = t-1 - Tt'$ for $T > 0$. Write $\mathbf{x}$ the complete binary tree of depth $T(N-1)$ such that $\forall t \in [T(N-1)], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{x}_t(\epsilon) = E_{0,t'+1}, \quad (21)$$

and $\mathbf{v}$ the complete binary tree of depth $T(N-1)$ such that $\forall t \in [T(N-1)], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{v}_t(\epsilon) = \frac{1}{4\sqrt{N-1}} \left(1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k}\right). \quad (22)$$

The pair $(\mathbf{x}, \mathbf{v})$ resembles the Von Neumann halving tree constructed in the previous section, with the key difference being that the nodes in the $\mathcal{X}$ valued part of the new tree are now labelled by a different class of measurement operators. We will replace $\mathbf{v}$ with a slightly modified version $\tilde{\mathbf{v}}$:

$$\tilde{\mathbf{v}}_t(\epsilon) = \mathbf{v}_t(\epsilon) + \frac{1}{4} \left(1 + \frac{1}{N-1}\right). \quad (23)$$

We now associate every path ϵ to any nonnegative completion $\omega(\epsilon)$ of $\text{part}(a_1, \dots, a_{N-1})$, where $\forall i \in [1, N-1]$,

$$a_i = \mathbf{v}_{iT}(\epsilon) + \frac{1}{2^{T+2}\sqrt{N-1}} \epsilon_{(i+1)T}. \quad (24)$$

Recall that the partial matrix $\text{part}(a_1, \dots, a_{N-1})$ satisfies all the conditions in Lemma 3.5 and hence can be completed to a valid density matrix $\omega(\epsilon)$.

Then, for $\delta \leq \frac{1}{2^{T+2}\sqrt{N-1}}$, we have that $\forall \epsilon \in \{\pm 1\}^{T(N-1)-1}, \forall t \in [T(N-1)]$,

$$\begin{aligned} & \epsilon_t [\text{Tr}_{\omega(\epsilon)}(\mathbf{x}_t(\epsilon)) - \tilde{\mathbf{v}}_t(\epsilon)] \\ &= \epsilon_t [\mathbf{v}_{(t'+1)T}(\epsilon) + \frac{1}{2^{T+2}\sqrt{N-1}} \epsilon_{(t'+2)T} - \frac{1}{4\sqrt{N-1}} (1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k})] \\ &= \epsilon_t \left[\frac{1}{4\sqrt{N-1}} (1 + \sum_{k=1}^T \epsilon_{k+Tt'} 2^{-k}) - \frac{1}{4\sqrt{N-1}} (1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k}) \right] \\ &= \frac{\epsilon_t}{4\sqrt{N-1}} \left(\sum_{k=\tilde{t}+1}^T \epsilon_{k+Tt'} 2^{-k} \right) \\ &= \frac{1}{4\sqrt{N-1}} \left(\frac{1}{2^{\tilde{t}+1}} + \sum_{k=\tilde{t}+2}^T \epsilon_{k+Tt'} 2^{-k} \right) \\ &\geq \delta. \end{aligned} \quad (25)$$

By taking $\delta = \frac{1}{N^{\frac{1}{2} + \frac{1}{k}}}$ and $T = \lfloor \log_2(\frac{1}{4\delta\sqrt{N}}) \rfloor$, we get $TN = \Omega(n\delta^{-\frac{2}{1+k}})$. Therefore, we have shown that the set $\mathcal{X}$ is δ -shattered by the Hypothesis class $\mathcal{H}_n$ with $\text{sfat}_\delta(\mathcal{H}_n) = \Omega(\frac{n}{\delta^\eta}), \forall \eta < 2$. $\square$

F Proof of Theorem 4.1

Proof. Set $\mathcal{X} = \{E_{0,i}, i \in [N-1]\}$ where $E_{0,i} = \frac{1}{2}(|0\rangle\langle 0| + |i\rangle\langle i| + |0\rangle\langle i| + |i\rangle\langle 0|)$. Define $t' = \lfloor \frac{t-1}{T} \rfloor$ and $\tilde{t} = t-1-Tt'$ for $T > 0$. Write $\mathbf{x}$ the complete binary tree of depth $T(N-1)$ such that $\forall t \in [T(N-1)], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{x}_t(\epsilon) = E_{0,t'+1}. \quad (26)$$

and $\mathbf{v}$ the complete binary tree of depth T such that $\forall t \in [T], \forall \epsilon \in \{\pm 1\}^{T-1}$,

$$\mathbf{v}_t(\epsilon) = \frac{1}{4\sqrt{N-1}}(1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k}). \quad (27)$$

We can associate every path ϵ to a pure state

$$|\psi(\epsilon)\rangle = \frac{1}{\sqrt{2}}|0\rangle + \sum_{i=1}^{N-2} a_i|i\rangle + \left(\frac{1}{2} - \sum_{i=1}^{N-2} a_i^2\right)|N-1\rangle \quad (28)$$

where,

$$a_i = \mathbf{v}_{iT}(\epsilon) + \frac{1}{2^{T+2}\sqrt{N-1}}\epsilon_{(i+1)T}. \quad (29)$$

Here $i \in [1, N-2]$. Let $\mathbf{w}_t = \frac{1}{\sqrt{2}}\mathbf{v}_t + \frac{1}{2}(\frac{1}{2} + a_{t'+1}^2)$. Then, for $\delta \leq \frac{1}{2^{T+2}\sqrt{2(N-1)}}$, we have that $\forall \epsilon \in \{\pm 1\}^{T(N-2)-1}, \forall t \in [T(N-2)]$,

$$\begin{aligned} \epsilon_t[\text{Tr}_{|\psi(\epsilon)\rangle}(\mathbf{x}_t(\epsilon)) - \mathbf{w}_t(\epsilon)] &= \frac{\epsilon_t}{\sqrt{2}}[\mathbf{v}_{(t'+1)T}(\epsilon) + \frac{1}{2^{T+2}\sqrt{N-1}}\epsilon_{(t'+2)T} - \frac{1}{4\sqrt{N-1}}(1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k})] \\ &= \frac{\epsilon_t}{\sqrt{2}}[\frac{1}{4\sqrt{N-1}}(1 + \sum_{k=1}^T \epsilon_{k+Tt'} 2^{-k}) - \frac{1}{4\sqrt{N-1}}(1 + \sum_{k=1}^{\tilde{t}} \epsilon_{k+Tt'} 2^{-k})] \\ &= \frac{\epsilon_t}{4\sqrt{2(N-1)}}(\sum_{k=\tilde{t}+1}^T \epsilon_{k+Tt'} 2^{-k}) \\ &= \frac{1}{4\sqrt{2(N-1)}}(\frac{1}{2^{\tilde{t}+1}} + \sum_{k=\tilde{t}+2}^T \epsilon_{k+Tt'} 2^{-k}) \\ &\geq \delta \end{aligned} \quad (30)$$

In particular, let $k \in \mathbb{N}^*$. By taking $\delta = \frac{1}{N^{\frac{1}{2}+k}}$ and $T = \lfloor \log_2(\frac{1}{4\delta\sqrt{N}}) \rfloor$, we get $TN = \Omega(n\delta^{-\frac{2}{1+\frac{2}{k}}})$. Therefore, we have shown that the set $\mathcal{X}$ is δ -shattered by the Hypothesis class $\mathcal{H}$ with $\text{sfat}_\delta(\mathcal{H}, \mathcal{X}) = \Omega(\frac{n}{\delta^\eta}) \forall \eta < 2$. $\square$

G Proof of Corollary 4.2

Proof for the L_1 loss. From Rakhlin et al. [2015a,b], the minimax regret is lower bound by the sequential fat-shattering dimension as:

$$\mathcal{V}_T \geq \frac{1}{4\sqrt{2}} \sup_{\delta > 0} \left\{ \sqrt{\delta^2 T \min\{\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X}), T\}} \right\}$$

provided that the loss function under consideration is the L_1 -loss. Now combining this result with the lower bound on $\text{sfat}_\delta(\mathcal{H}_n, \mathcal{X})$ established in Theorem 3.6 we get $\mathcal{V}_T = \Omega(\sqrt{nT})$. $\square$

Proof for the L_2 loss. From Equation (1), and taking the y_t to be Rademacher random variables, we get:

$$\begin{aligned}
\mathcal{V}_T &\geq \left\langle \sup_{x_t} \inf_{h_t \in \mathcal{H}} \mathbb{E}_{y_t} \right\rangle_{t=1}^T \left[\sum_{t=1}^T (h_t(x_t) - y_t)^2 - \inf_{h \in \mathcal{H}} \sum_{t=1}^T (h(x_t) - y_t)^2 \right] \\
&\geq \left\langle \sup_{x_t} \right\rangle_{t=1}^T \left[\inf_{h_t \in \mathcal{H}} \sum_{t=1}^T (1 + h_t^2) + \mathbb{E}_{y_t} \sup_{h \in \mathcal{H}} \sum_{t=1}^T -(h(x_t) - y_t)^2 \right] \\
&\geq \left\langle \sup_{x_t} \right\rangle_{t=1}^T \left[\sum_{t=1}^T 1 + \mathbb{E}_{y_t} \sup_{h \in \mathcal{H}} \sum_{t=1}^T -(h(x_t) - y_t)^2 \right] \\
&= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \left[\mathbb{E}_{y_t} \sup_{h \in \mathcal{H}} \sum_{t=1}^T 1 - (h(x_t) - y_t)^2 \right] \\
&= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \left[\mathbb{E}_{y_t} \sup_{h \in \mathcal{H}} \sum_{t=1}^T 2y_t h(x_t) - h^2(x_t) \right] \\
&= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \frac{1}{2} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T 2h(x_t) - h^2(x_t) + \sup_{h \in \mathcal{H}} \sum_{t=1}^T -2h(x_t) - h^2(x_t) \right] \\
&\geq \left\langle \sup_{x_t} \right\rangle_{t=1}^T \frac{1}{2} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T h(x_t)(2 - h(x_t)) + \left(\sum_{t=1}^T -2h(x_t) - h^2(x_t) \right)_{h=\text{Tr} \frac{1}{N} I_N} \right] \\
&\geq \left\langle \sup_{x_t} \right\rangle_{t=1}^T \frac{1}{2} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T h(x_t) + \sum_{t=1}^T -\frac{2}{N} - \frac{1}{N^2} \right] \\
&\geq \mathcal{R}_T(\mathcal{H}) - \frac{3T}{2N}
\end{aligned} \tag{31}$$

Where we used the fact that

$$\begin{aligned}
\mathcal{R}_T(\mathcal{H}) &= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \mathcal{R}_T(x, \mathcal{H}) \\
&= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \mathbb{E}_y \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T y_t h_t(x_t) \right] \\
&= \left\langle \sup_{x_t} \right\rangle_{t=1}^T \frac{1}{2} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T h_t(x_t) - \sup_{h \in \mathcal{H}} \sum_{t=1}^T h_t(x_t) \right] \\
&\leq \left\langle \sup_{x_t} \right\rangle_{t=1}^T \frac{1}{2} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T h_t(x_t) \right]
\end{aligned} \tag{32}$$

□

H Proof of Theorem 5.1

Proof. All we need to do in order to lower bound $\bar{\mathcal{V}}_T$ is to exhibit a particular strategy $(\mathcal{D}_t)_{t \in [T]}$ the adversary can adopt. We will then show that such a strategy guarantees a minimax regret greater than $\Omega(\sqrt{nT})$, independently of the strategy picked by the learner.

It was shown in Aaronson [2007] that n -qubit quantum states, considered as a hypothesis class, have $(\delta, \frac{\delta^2}{n})$ -fine-shattering dimension greater than $D = \lfloor \frac{n}{5\delta^2} \rfloor$. Let $T \in \mathbb{N}^*$ and write $\delta = \sqrt{\frac{n}{5T}}$. This means that there exists $(E_t)_{t \in [D]} \in \mathcal{X}^D$ such that $\forall (\epsilon_t) \in \{\pm 1\}^D, \exists \rho \in \mathcal{C}_n, \forall t \in [D]$:

$$\mathrm{Tr}(E_t \rho) \in \begin{cases} \left[\frac{1}{2} - \delta - \frac{\delta^2}{n}, \frac{1}{2} - \delta \right] & \text{if } \epsilon_t = -1, \\ \left[\frac{1}{2} + \delta, \frac{1}{2} + \delta + \frac{\delta^2}{n} \right] & \text{if } \epsilon_t = 1. \end{cases} \quad (33)$$

We will denote by $\rho : 2^D \rightarrow \mathcal{C}_n$ the function that associates every path to the corresponding quantum state in Equation (33). We can now show that if the adversary picks E_t at every round t , no matter the strategy chosen by the learner, there exists a quantum state ρ that will guarantee a loss greater than $\sqrt{nT}$.

For every $t \in [D]$, write $\epsilon_t = \mathrm{sgn} \left[\mathbb{E}_{\omega \sim \mathcal{Q}_t} \left(\frac{1}{2} - \mathrm{Tr}(E_t \omega) \right) \right]$. Then, we clearly have that

$$\begin{aligned} \left\langle \mathbb{E}_{\omega_t \sim \mathcal{Q}} \right\rangle_{t=1}^T \left[\sum_{t=1}^T |\mathrm{Tr}(E_t \omega_t) - \mathrm{Tr}(E_t \rho)| \right] &\geq D\delta \\ &\geq \delta(T-1) \\ &\geq \sqrt{\frac{nT}{5}} - 1 \end{aligned} \quad (34)$$

Remark H.1. The last thing we need to check is that the feedback received by the learner doesn't give *much more* information than the interval of size $\frac{\delta^2}{n}$ in which lies $\mathrm{Tr}(E_t \rho)$, as defined in Equation (33). If $y_t \sim \mathcal{U}([\mathrm{Tr}(E_t \rho) - \epsilon, \mathrm{Tr}(E_t \rho) + \epsilon])$ for instance, this amounts to $\epsilon \geq \frac{\delta^2}{n}$, *id est* $T \geq \frac{1}{5\epsilon}$. For a general distribution, picking T big enough such that the probability of $y_t \in [\mathrm{Tr}(E_t \rho) - \epsilon, \mathrm{Tr}(E_t \rho) + \epsilon]$ is greater than $\frac{1}{2}$ will only divide the lower bound by 2.

Finally, the fact that $\mathbb{E}(y_t) = \mathrm{Tr}(E_t \rho)$ allows us to conclude.

Therefore, $\bar{\mathcal{V}}_T = \Omega(\sqrt{nT})$. It is easy to see that $\bar{\mathcal{V}}_T \leq \mathcal{V}_T = O(\sqrt{nT})$. $\square$

I Regret bounds with sequential complexities in classical online learning

Notions of complexity for a given hypothesis class have traditionally been studied within the batch learning framework and are often characterized by the Rademacher complexity [Bartlett and Mendelson, 2003].

Definition I.1 (Rademacher complexity). Let $\mathcal{X}$ be a sample space with an associated distribution $\mathcal{D}$ and $\mathcal{H}$ be the hypothesis class. Let $(x_j)_{j \in [m]} \sim \mathcal{D}^m$ be a sequence of samples, sampled i.i.d. from $\mathcal{X}$. The Rademacher complexity can then be defined as:

$$\mathcal{R}_m(\mathcal{H}) = \mathbb{E}_{(x_j) \sim \mathcal{D}^m} \left[\frac{1}{m} \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{j=1}^m \epsilon_j h(x_j) \right] \right],$$

where $\epsilon = (\epsilon_1, \dots, \epsilon_m)$ are called Rademacher variables, that satisfy $P(\epsilon = +1) = P(\epsilon = -1) = 1/2$.

Perhaps not surprisingly, in addition to being an indicator for expressivity of a given hypothesis class, Rademacher complexity also upper bounds generalization error in the setting of batch learning [Bartlett and Mendelson, 2003].

Rademacher complexity generalises to sequential Rademacher complexity in the online setting [Rakhlin et al., 2015a]. In order to define sequential Rademacher complexity let us first define a $\mathcal{X}$ -valued complete binary tree.

Definition I.2 (Sequential Rademacher complexity). Let $\mathbf{x}$ be a $\mathcal{X}$ -valued complete binary tree of depth T . The sequential Rademacher complexity of a hypothesis class $\mathcal{H}$ on the tree $\mathbf{x}$ is then given as:

$$\mathfrak{R}_T(\mathcal{H}, \mathbf{x}) = \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T \epsilon_t h(\mathbf{x}_t(\epsilon)) \right] \right].$$

The $\mathbf{x}$ dependence of the sequential Rademacher complexity can be subsequently removed by considering the supremum over all $\mathcal{X}$ -valued trees of depth T : $\mathfrak{R}_T(\mathcal{H}) = \sup_{\mathbf{x}} \mathfrak{R}_T(\mathcal{H}, \mathbf{x})$. Similar to how Rademacher complexity upper bounds the generalization error, the sequential Rademacher complexity was shown to upper bound the minimax regret [Rakhlin et al., 2015a]. For the case of supervised learning, the following relation holds:

$$\mathcal{V}_T \leq 2LT\mathfrak{R}_T(\mathcal{H}). \quad (35)$$

Here, L comes from the fact that the loss function considered is L -Lipschitz.

The growth of sequential Rademacher complexities has been shown to be influenced by other related notions of sequential complexities. One prominent example is the sequential fat-shattering dimension [Rakhlin et al., 2015a]. It was shown in Rakhlin et al. [2015a] that this dimension serves as an upper bound to the sequential Rademacher complexity, which subsequently provides a bound on the minimax regret as per Equation (35). Similarly, recall that the minimax regret can be lower bounded by the sequential fat-shattering dimension (Equation (3)), provided that $\ell_t(h_t(x_t), y_t) = |h_t(x_t) - y_t|$ and that $\mathcal{P}$ is taken to be the whole set of all distributions on $\mathcal{X}$. In fact, it is also lower bounded by the Rademacher complexity [Rakhlin et al., 2015a,b].

$$\begin{aligned} \mathcal{V}_T &\geq \left\langle \sup_{\mathcal{D}_t \in \mathcal{P}} \mathbb{E}_{x_t \sim \mathcal{D}_t} \right\rangle_{t=1}^T \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{t=1}^T \epsilon_t h(\mathbf{x}_t(\epsilon)) \right] \\ &\geq \frac{1}{4\sqrt{2}} \sup_{\delta > 0} \left\{ \sqrt{\delta^2 T \min\{\text{sfat}_{\delta}(\mathcal{H}, \mathcal{X}), T\}} \right\}. \end{aligned} \quad (36)$$

where $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_T)$ are Rademacher variables.

J Coupling

To establish regret bounds in smoothed online learning, Haghtalab et al. [2024], Block et al. [2022] introduced the concept of coupling. The key idea here is that if the distributions $(\mathcal{D}_t)_{t=1}^T$ are σ -smooth with respect to $\mathcal{D}$, we may pretend that in expectation the data is sampled i.i.d from $\mathcal{D}$ instead of $(\mathcal{D}_t)_{t=1}^T$. For a more formal description, define $\mathcal{B}_T(\sigma, \mathcal{D})$ to be the set of joint distributions $\mathcal{D}_{\wedge}$ on $\mathcal{X}^T$, where each marginal distribution $\mathcal{D}_t(\cdot | x_1, \dots, x_{t-1})$ is conditioned on the previous draws.

Definition J.1 (Coupling). A distribution $\mathcal{D}_{\wedge} \in \mathcal{B}_T(\sigma, \mathcal{D})$ is said to be coupled to independent random variables drawn according to $\mathcal{D}$ if there exists a probability measure Π with random variables $(x_t, Z_t^j)_{t \in [T], j \in [k]} \sim \Pi$ satisfying the following conditions:

1. $x_t \sim \mathcal{D}_t(\cdot | x_1, x_2, \dots, x_{t-1})$,
2. $\{Z_t^j\}_{t \in [T], j \in [k]} \sim \mathcal{D}^{\otimes kT}$,
3. With probability at least $1 - Te^{-\sigma k}$, we have $x_t \in \{Z_t^j\}_{j \in [k]} \forall t \in [T]$.

The last relation is particularly interesting and is used to derive the regret bounds for smoothed online quantum state learning.

K Proof of Theorem 5.3

For this proof, we will use the notion of Rademacher complexity and a few related results that can be found in Appendix I. We also use the concept of coupling described in Appendix J.

Consider the sequential Rademacher complexity $\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathbf{x})$ in Definition I.2 defined on the function class $\ell \circ \mathcal{H}$. For the purpose of this proof let us consider a slightly modified version of the sequential Rademacher complexity defined as:

$$\begin{aligned} \mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_{\wedge}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\wedge}} \mathfrak{R}_T(\ell \circ \mathcal{H}, \mathbf{x}) \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\wedge}} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_{\rho}(\mathbf{x}_t(\epsilon))) \right] \right], \end{aligned} \quad (37)$$

where $\mathcal{D}_\wedge \in \mathcal{B}_T(\sigma, \mathcal{D})$. Moreover we will call $\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{B}_T) = \sup_{\mathcal{D}_\wedge \in \mathcal{B}_T(\sigma, \mathcal{D})} \mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_\wedge)$. The key idea now is to relate $\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_\wedge)$ to $\mathcal{R}(\mathcal{H})$ (Rademacher complexity assuming i.i.d. data inputs; see Definition I.1). This can be achieved using the idea of coupling discussed in the previous section.

Lemma K.1. *Let $\mathcal{D}_\wedge \in \mathcal{B}_T(\sigma, \mathcal{D})$ be a distribution that is coupled to independent random variables drawn according to the distribution $\mathcal{D}$, as per Definition J.1. Let $\mathcal{H}$ be a Hypothesis class and ℓ be the loss function. Then we have $\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_\wedge) \leq T^2 e^{-\sigma k} + \mathcal{R}_{kT}(\ell \circ \mathcal{H})$*

Proof. Let A be an event that $\mathbf{x}_t(\epsilon) \in \{Z_t^j\}_{j=1}^k \forall t \in [T]$. Furthermore, let χ_A be the corresponding indicator function and χ_{A^c} be $1 - \chi_A$. Then we have:

$$\begin{aligned}
\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_\wedge) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_\wedge} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right] \right] \\
&= \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right] \right] \\
&= \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\chi_A \sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right] \right. \\
&\quad \left. + \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\chi_{A^c} \sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right] \right] \right] \\
&\leq \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\chi_A \sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right] \right] + T^2 e^{-\sigma k} \\
&\leq \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\chi_A \sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right. \right. \\
&\quad \left. \left. + \sum_{j: Z_t^j \neq \mathbf{x}_t(\epsilon)} \sum_{t=1}^T \mathbb{E}_{\epsilon_{jt}} \epsilon_{jt} \ell(h(Z_t^j), h_\rho(Z_t^j)) \right] \right] + T^2 e^{-\sigma k} \\
&\leq \mathbb{E}_{\mathbf{x} \sim \Pi} \left[\frac{1}{T} \mathbb{E}_{\epsilon} \left[\chi_A \sup_{h \in \mathcal{H}} \sum_{j=1}^T \epsilon_j \ell(h(\mathbf{x}_t(\epsilon)), h_\rho(\mathbf{x}_t(\epsilon))) \right. \right. \\
&\quad \left. \left. + \mathbb{E}_{\epsilon_{jt}} \sum_{j: Z_t^j \neq \mathbf{x}_t(\epsilon)} \sum_{t=1}^T \epsilon_{jt} \ell(h(Z_t^j), h_\rho(Z_t^j)) \right] \right] + T^2 e^{-\sigma k} \\
&\leq \mathbb{E}_{\mathcal{D}} \left[\frac{1}{T} \mathbb{E}_{\epsilon_{jt}} \left[\sup_{h \in \mathcal{H}} \sum_{j=1}^k \sum_{t=1}^T \epsilon_{jt} \ell(h(Z_t^j), h_\rho(Z_t^j)) \right] \right] + T^2 e^{-\sigma k} \\
&\leq T^2 e^{-\sigma k} + \mathcal{R}_{kT}(\ell \circ \mathcal{H}). \tag{38}
\end{aligned}$$

□

Here the first inequality comes from the fact that coupling constructed in the previous section bounds the probability of $\mathbf{x}_t(\epsilon) \notin \{Z_t^j\}_{j=1}^k$ at least for one value of t . The second inequality follows from the fact that σ_t has a zero mean while the third one follows Jensen's inequality.

Proof of Theorem 5.3. As per Equation (35) $\mathcal{V}_T$ is upper bounded by the sequential Rademacher complexity $\mathfrak{R}_T(\mathcal{H})$. Therefore it suffices to establish an upper bound on the latter in the smoothed setting considered here. Equation (38) bounds the distribution dependent sequential Rademacher complexity with the standard notion of Rademacher complexity which assumes independent samples. Assuming the loss function to be L -Lipschitz, the quantity $\mathcal{R}_{kT}(\ell \circ \mathcal{H})$ can be upper bounded as:

$$\mathcal{R}_{kT}(\ell \circ \mathcal{H}) \leq L \mathcal{R}_{kT}(\mathcal{H}). \tag{39}$$

The sequential Rademacher complexity can be bounded above by the sequential fat-shattering dimension. Likewise, one can establish upper bounds on $\mathcal{R}_{kT}(\mathcal{H})$:

$$\mathcal{R}_{kT}(\mathcal{H}) \leq \inf_{\alpha > 0} \left\{ 4\alpha kT + 12\sqrt{kT} \int_{\alpha}^1 \sqrt{K \text{fat}_{c\delta}(\mathcal{H}) \log \frac{2}{\delta}} d\delta \right\}, \quad (40)$$

where K and c are constants. Note here that $\text{fat}_{\delta}(\mathcal{H})$ unlike its sequential counterparts assume independent data. Therefore one can recover the definition of $\text{fat}_{\delta}(\mathcal{H})$ from their sequential version by replacing the $\mathcal{X}$ and $\mathbb{R}$ -valued tree by the sets $\mathcal{X}$ and $\mathbb{R}$. Combining Equations (38) to (40), and setting $k = \frac{2 \log T}{\sigma}$ we get:

$$\mathfrak{R}_T(\ell \circ \mathcal{H}, \mathcal{D}_{\wedge}) \leq 1 + L \inf_{\alpha > 0} \left\{ \frac{8\alpha T \log T}{\sigma} + 12\sqrt{\frac{2T \log T}{\sigma}} \int_{\alpha}^1 \sqrt{K \text{fat}_{c\delta}(\mathcal{H}) \log \frac{2}{\delta}} d\delta \right\}. \quad (41)$$

Now using the relation that $\text{fat}_{\delta}(\mathcal{H}_n) = O(n/\delta^2)$ when $\mathcal{H}_n = \{\text{Tr}_{\omega}, \omega \in \mathcal{C}_n\}$ and setting $K = c = 1$ we get:

$$\mathfrak{R}_T(\ell \circ \mathcal{H}_n, \mathcal{D}_{\wedge}) \leq 1 + L \inf_{\alpha > 0} \left\{ \frac{8\alpha T \log T}{\sigma} + 12\sqrt{\frac{2nT \log T}{\sigma}} \int_{\alpha}^1 \sqrt{\frac{1}{\delta^2} \log \frac{2}{\delta}} d\delta \right\}. \quad (42)$$

Finally to eliminate the infimum in Equation (42) we recall that $\alpha \in [0, 1]$ and therefore for any function f on α we get $\inf_{\alpha} f(\alpha) \leq f(\alpha = \alpha^*)$; $\alpha^* \in [0, 1]$. Thus setting $\alpha = \sqrt{\frac{n\sigma}{T \log T}}$, we get:

$$\mathfrak{R}_T(\ell \circ \mathcal{H}_n, \mathcal{D}_{\wedge}) = O\left(\sqrt{\frac{nT \log T}{\sigma}}\right). \quad (43)$$

□

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: In addition do the abstract, the paper's contributions and scope are detailed in Section 1.1.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: All the assumptions are explicitly stated in the Theorems. In addition, the limitations of the paper are discussed in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All the complete proofs can be found in the appendices.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: This paper does not include any experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: This paper does not include any experiment.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: This paper does not include any experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: This paper does not include any experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: This paper does not include any experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We affirm that we have read the NeurIPS Code of Ethics and that our research fully conforms to its guidelines.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Our paper is purely theoretical and does not involve deployed systems, real-world data, or direct societal applications. As such, we believe that it does not present immediate positive or negative societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work does not involve the release of data or models, and poses no foreseeable risk of misuse, as it is purely theoretical and focused on quantum state learning.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.