Learning Distributions over Quantum Measurement Outcomes

Weiyuan Gong [1] **Scott Aaronson** [2]

Abstract

Shadow tomography for quantum states provides a sample efficient approach for predicting the measurement outcomes of quantum systems. However, these shadow tomography procedures yield poor bounds if there are more than two outcomes per measurement. In this paper, we consider a general problem of learning properties from quantum states: given an unknown d-dimensional quantum state ρ and M unknown quantum measurements $\mathcal{M}_1, ..., \mathcal{M}_M$ with $K \geq 2$ outcomes, estimating the probability distribution for applying $\mathcal{M}_i$ on ρ to within total variation distance ϵ. Compared to the special case when $K = 2$, we have to learn unknown distributions instead of values. Here, we propose an online shadow tomography procedure that solves this problem with high success probability requiring $\tilde{O}(K \log^2 M \log d/\epsilon^4)$ copies of ρ. We further prove an information-theoretic lower bound showing that at least $\Omega(\min\{d^2, K + \log M\}/\epsilon^2)$ copies of ρ are required to solve this problem with high success probability. Our shadow tomography procedure requires sample complexity with only logarithmic dependence on M and d and is sample-optimal concerning the dependence on K.

1. Introduction

The statistical learning theory problem of extracting information based on empirical observations is of fundamental importance in a number of fields. In quantum physics, a fundamental problem is to obtain the properties of a quantum system based on statistical results from quantum measurements. A general method to obtain the full information of an unknown d-dimensional quantum state ρ, called *quantum state tomography*, completely recovers the density matrix to within a small error. This task is proved to require $\Omega(d^2)$ copies of ρ from the information-theoretical perspective (Haah et al., 2017; O'Donnell & Wright, 2016; Chen et al., 2022b). These state tomography procedures have been pushed to the limit of their capabilities after the recent advances in experimental quantum platforms (Preskill, 2018), recalling that the dimension of the quantum state $d = 2^n$ increases exponentially in the number of qubits.

However, demanding full descriptions of quantum states may be excessive for concrete quantum problems. Following this conceptual different line of research, the *quantum shadow tomography* problem developed by (Aaronson, 2019; 2007) considers the case when we are given an unknown d-dimensional quantum state and M known quantum events regarded as a two-outcome quantum measurement that outputs 1 (or "accept") with probability $\mathrm{Tr}(E_i\rho)$ and outputs 0 (or "reject") otherwise. The goal is to estimate each expectation $\mathbb{E}_\rho[E_i] = \mathrm{Tr}(E_i\rho)$ to within additive error $\pm\epsilon$. This shadow tomography problem for two-outcome POVMs is a quantum analog of the classical *adaptive data analysis* (Dwork et al., 2015; 2010), which can be solved with poly$(\log M, \log d, 1/\epsilon)$ samples (Bassily et al., 2021). Recently, (Bădescu & O'Donnell, 2021) proved the best known upper bound on sample complexity for this problem as $N = \tilde{O}(\log^2 M \log d/\epsilon^4)$.

In quantum mechanics, the prediction of some intriguing properties requires quantum measurements with $K > 2$ measurement outcomes. In addition, in practical tasks, many quantum properties are represented by the expectation values of quantum observables, which are obtained by computing the probability distribution of each measurement outcome and averaging the eigenvalue corresponding to the outcome. In this case, the measurement $\mathcal{M}$ outputs results $j = 1, ..., K$ with probability $\mathbb{E}_\rho[E_j] = \mathrm{Tr}(E_j\rho)$, which are expectations of quantum events $E_1, ..., E_K$ that satisfy $\sum_{j=1}^K E_j = \mathbb{I}$. Our goal is to approximate the probability distribution over the outcomes of $\mathcal{M}$ within total variation distance ϵ. Recalling that K can be as large as $\Theta(d)$ in the extreme case, it is an important factor to concern in practical shadow tomography. In this paper, we study the shadow tomography problem of K-outcome quantum measurements, which can be formulated as follows

Problem 1.1 (Shadow Tomography of K-outcome Measurements)**.** We consider an unknown d-dimensional quantum state, as well as M quantum measurements $\mathcal{M}_1, ..., \mathcal{M}_M$,

[1] IIIS, Tsinghua University [2] Department of Computer Science, University of Texas at Austin. Correspondence to: Weiyuan Gong <wygong8@gmail.com>.

Proceedings of the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

each of which has K results and outputs the j-th result with probability $\mathrm{Tr}(E_{i,j}\rho)$ for $i \in [M]$ and $j \in [K]$. We denote $\boldsymbol{p}_i$ the probability distribution $(\mathrm{Tr}(E_{i,1}\rho), ..., \mathrm{Tr}(E_{i,K}\rho))$ after measurement $\mathcal{M}_i$. Our goal is to output M probability distributions $\boldsymbol{b}_1, ..., \boldsymbol{b}_M$ defined on the K-outcomes such that the total variation distance $d_{TV}(\boldsymbol{p_i}, \boldsymbol{b}_i) \leq \epsilon$ with success probability at least $1 - \delta$.

The quantum events $E_{i,j}$ for $j \in [K]$ corresponding to the quantum measurement $\mathcal{M}_i$ is defined to satisfy the constraint $0 \preceq E_{i,j} \preceq \mathbb{I}$ and $\sum_{j=1}^{K} E_{i,j} = \mathbb{I}$. We remark that Problem 1.1 can also be extended to other metrics such as Euclidean norm or infinity norm and obtain a sample complexity without K dependence. We choose the total variation distance because it has a direct connection with predicting expectation values of observables: by bounding the total variation distance, the error for the expectation value of any observable O can be bounded with poly($\|O\|_\infty$) overhead.

1.1. Main Results

The first main result of this paper is to propose an algorithm to solve this shadow tomography problem of K-outcome measurements. We prove the following sample-complexity upper bound for our algorithm.

Theorem 1.2 (Shadow Tomography of K-outcome Measurements)**.** *Problem 1.1 is solvable using*

$$N = \tilde{O}\left(\frac{\log(1/\delta)}{\epsilon^4} \cdot K \cdot \log^2 M \cdot \log d\right)$$

copies of ρ. Here, the $\tilde{O}$ hides a poly$(\log\log M, \log\log D, \log(1/\epsilon), \log K)$ factor. The procedure is fully explicit and online.

We provide an overview of proof for this theorem in Section 1.2. The detailed proof is technically involved and provided in Section 3 and Section 4. Theorem 1.2 indicates that we can learn the probability distribution of M quantum measurements of K outcomes using sample complexity that depends logarithmically on M and d but linearly on K. Considering the parameters M, d, and ϵ, our algorithm has the same dependence as the best known upper bound for 2-outcome case (Bădescu & O'Donnell, 2021). The dependence on K is the most important result in this work. Compared to directly regarding each quantum event $E_{i,j}$ as a two-outcome quantum measurement and approximating the expectation $\mathrm{Tr}(E_{i,j}\rho)$ to within additive error $2\epsilon/K$, our algorithm reduces the dependence on K from $\tilde{O}(K^4)$ to $\tilde{O}(K)$. Notice that in some extreme cases, K can be as large as $\Theta(d)$, which is exponential in system size n, our algorithm reduces the number of copies required to perform the shadow tomography task significantly. Although the complexity of our algorithm still has an $\tilde{O}(K)$ dependence on K, we emphasize that this dependence is necessary by the following information-theoretic lower bound on Problem 1.1:

Theorem 1.3. *Any strategy for Problem 1.1—i.e., for estimating all $\boldsymbol{p}_i = (\mathrm{Tr}(E_{i,1}\rho), ..., \mathrm{Tr}(E_{i,K}\rho))$ of $\mathcal{M}_i$ to within total variation distance ϵ for all $i \in [M]$, with success probability at least (say) $2/3$—requires at least*

$$N \geq \Omega\left(\frac{\min\{d^2, K + \log M\}}{\epsilon^2}\right)$$

copies of unknown d-dimensional quantum state ρ.

We provide the sketch of proof for this lower bound in Section 1.2 and leave the detailed proof in Section 5. The lower bound is obtained by an information-theoretic argument developed by (Flammia et al., 2012) and refined by further works (Aaronson, 2019; Huang et al., 2020; Haah et al., 2017). The proof exploits Holevo's theorem (Holevo, 1973) and Fano's inequality (Fano, 1949). Even in the special case where there is only $M = 1$ entirely classical measurement and the unknown quantum state is also classical, learning a distribution on $[K]$ to within total variation distance ϵ still requires $O(K/\epsilon^2)$ samples (Canonne, 2020). This result can be understood as the information required to approximate the probability distribution scales linearly with the dimension of the distribution. By comparing the lower bound in Theorem 1.3 and the upper bound in Theorem 1.2, we can conclude that our algorithm for shadow tomography of K-outcome quantum measurement is optimal concerning the dependence on K.

1.2. Technical Overview

Our shadow tomography procedure involves the combination of two ideas: solving a quantum distribution threshold search problem using $O(K \log^2 M/\epsilon^2)$ samples in each iteration and performing an online learning procedure that has at most $O(\log d/\epsilon^2)$ iterations.

The first step in this work concerns a problem which we call the *quantum distribution threshold search* problem. We formulate this problem as below:

Problem 1.4 (Quantum Distribution Threshold Search)**.** Suppose we are given

- Parameters $0 < \epsilon, \delta < \frac{1}{2}$;
- Unentangled copies of an unknown d-dimensional quantum state ρ.
- A list of M d-dimensional POVMs $\mathcal{M}_1, ..., \mathcal{M}_M$ each of K outcomes corresponding to quantum events $E_{i,j}$, where $i \in [M]$, $j \in [K]$, and $\sum_{j=1}^{K} E_{i,j} = \mathbb{I}$. We denote $\boldsymbol{p}_i = (\mathrm{Tr}(E_{i,1}\rho), ...,$ $\mathrm{Tr}(E_{i,K}\rho))$ to be the actual distribution over the measurement outcomes of $\mathcal{M}_i$.

- A list of M threshold vectors $\boldsymbol{\theta}_i = (\theta_{i,1}, ..., \theta_{i,K})$, where $\theta_{i,j} \in [0,1]$ and $\sum_{j=1}^{K} \theta_{i,j} = 1$.

the algorithm outputs either:

- $d_{TV}(\boldsymbol{p}_{i^*}, \boldsymbol{\theta}_{i^*}) > 3\epsilon/4$ for some particular i^*; or
- $d_{TV}(\boldsymbol{p}_i, \boldsymbol{\theta}_i) \leq \epsilon$ for any i.

Our goal is to minimize the number of copies required to ensure we output correctly with success probability at least $1-\delta$.

A similar problem for the case of $K = 2$ was originally called a gentle search procedure in (Aaronson, 2019). Later, it was renamed as the quantum threshold search problem in (Bădescu & O'Donnell, 2021) since the *gentle* measurement assumption (Aaronson & Rothblum, 2019) is not necessary. It is proven that the quantum threshold problem can be solved using $\tilde{O}(\log^2 M/\epsilon^2)$ copies of ρ with probability at least (say) $3/4$ (Bădescu & O'Donnell, 2021). Yet, it is worthwhile to mention that Problem 1.4 is not a direct extension of the quantum threshold search problem. Even at $K = 2$, the requirement of Problem 1.4 is a two-side bound instead of a one-side bound in the quantum threshold search problem. In this paper, we provide an algorithm that can solve Problem 1.4 for any $K \geq 2$:

Theorem 1.5. *Problem 1.4 (Quantum Distribution Threshold Search) is solvable using*

$$N = \tilde{O}\left(\frac{\log(1/\delta)}{\epsilon^2} \cdot K \cdot \log^2 M\right)$$

copies of ρ.

We provide proof of this theorem in Section 3. When $K = 2$, our upper bound for the quantum distribution threshold search problem reduces to the same bound for the quantum threshold search problem. We remark that the K dependence in the sample complexity bound we provide in Theorem 1.2 directly comes from the K dependence in solving the quantum distribution threshold search problem.

Given our quantum distribution threshold search algorithm, the second step is to employ a black-box reduction to an online quantum state learning algorithm. In the special case when $K = 2$, the bound is obtained by (Aaronson et al., 2018). The formal version of our result in online learning distributions is provided as follows:

Theorem 1.6. *Let ρ be an unknown d-dimensional quantum state, as well as $\mathcal{M}_1, \mathcal{M}_2, ..., \mathcal{M}_t, ...$ be a sequence of K-outcome POVMs each consisting quantum events $E_{t,j}$ for $j \in [K]$. We denote $\boldsymbol{p}_t = (\mathrm{Tr}(E_{t,1}\rho), ..., \mathrm{Tr}(E_{t,K}\rho))$ to be the actual probability distribution when we apply $\mathcal{M}_t$ on ρ. We are provided with a probability distribution $\boldsymbol{b}_t$ after each measurement $\mathcal{M}_t$ with $d_{TV}(\boldsymbol{p}_t, \boldsymbol{b}_t) \leq \epsilon/4$. There exists a strategy for outputting hypothesis states $\omega_1, \omega_2, ...$ such that the probability distribution $\boldsymbol{\mu}_t$, which is obtained by applying $\mathcal{M}_t$ on ω_t, deviates more than $3\epsilon/4$ from $\boldsymbol{p}_t$ for at most $T = O(\log d/\epsilon^2)$ iterations t (also called "bad iterations").*

We provide the proof for Theorem 1.6 in Section 4.1 adapting the template of the *Regularized Follow-the-Leader algorithm* (RFTL; see, for example, in (Hazan et al., 2016)). However, our online learning procedure is not a direct extension of the standard template as we have to modify the loss function to measure the total variation distance between two distributions instead of the ℓ_1 loss between two values in (Aaronson et al., 2018; Chen et al., 2022c). We exploit the property of the POVM to provide an upper bound for the total regret of the learning procedure. We can then combine our quantum distribution threshold search algorithm with this online setting to prove the sample complexity for our shadow tomography procedure of K-outcome quantum measurements. We start with the maximally mixed state $\mathbb{I}/d$. In each iteration, we first perform the quantum distribution threshold search algorithm to find an i^* such that the total variation distance between $\boldsymbol{\mu}_t$ and $\boldsymbol{p}_t$ is larger than $3\epsilon/4$. We can then use $\tilde{O}(K/\epsilon^2)$ samples to estimate $\boldsymbol{b}_t$ with high success probability and update the hypothesis. As there are at most $O(\log d/\epsilon^2)$ "bad iterations", we finally reach the complexity bound in Theorem 1.2. The overall computational complexity of our shadow tomography protocol is estimated to be $O(KM \cdot \mathrm{poly}(K, 1/\epsilon, \log d)) + d^{O(1)}$. We further show the extension of Theorem 1.2 to the case of predicting quantum observables in Appendix D.

To prove the lower bound, we first fix M different quantum measurements. We then find a set (known as packing net (Haah et al., 2017)) of size $2^{K/2}M$ consisting of mixed states $\{\rho_1, ..., \rho_N\}$ such that we can use our shadow tomography procedure to distinguish between any pair of states chosen from this packing net, which requires $\log\big(2^{K/2}M\big) = \Theta(K + \log M)$ bits of information. We further show using Holevo's theorem (Nielsen & Chuang, 2002) that we can at most obtain $O(\epsilon^2)$ bits of information from any quantum states chosen from this set. Therefore, the sample complexity is bounded below by $\Omega((K+\log M)/\epsilon^2)$ to make it possible to obtain the information. We emphasize that, different from the previous lower bounds (Aaronson, 2019; Aaronson & Rothblum, 2019; Huang et al., 2020) which either adapt a classical bound (Ullman et al., 2018; Bun et al., 2018; Vadhan, 2017), or just consider a set containing M quantum states corresponding to the M measurements (observables), our construction of the set exploit a coding parameter $\boldsymbol{z}$ for the states to introduce a K dependence.

1.3. Related Works

Shadow tomography of two-outcome quantum measurements. A first related topic is the shadow tomography of two-outcome quantum measurements (Aaronson, 2019; Bădescu & O'Donnell, 2021). It is proven that only $O(\log^2 M \log d/\epsilon^4)$ copies of ρ can estimate the expectation value of M two-outcome quantum measurement. When $K = 2$, the sample complexity in Theorem 1.2 reduces to this bound. To extend this result to the case of $K > 2$, a straightforward approach is to regard each quantum event as a two-outcome quantum measurement and estimate each expectation value $\mathrm{Tr}(E_{i,j}\rho)$ to within additive error $2\epsilon/K$. However, this approach requires an additional cost of $\tilde{O}(K^4)$ using the state-of-art shadow tomography algorithms. Compared to this direct extension, our approach only requires sample complexity that increases linear with K and is proven to be optimal concerning the dependence on K.

Quantum observable estimation using classical shadow. (Huang et al., 2020) and (Chen et al., 2022a) considered the task of estimating quantum functions (or the expectations of quantum operators). Given ρ an unknown d-dimensional state, as well as M quantum operators $O_1, ..., O_M$, they provided a strategy that can approximate the expectation value for each operator $\mathrm{Tr}(O_i\rho)$ to within additive error ϵ with high success probability $O(\log M 2^k/\epsilon^2)$ copies of ρ, where k is the locality of the observable. Their protocol requires neither quantum memory nor joint measurements that simultaneously measure states of the form $\rho^{\otimes k}$. However, the sample complexity may increase exponentially with the system size $n = \log d$. Our algorithm, however, can provide sample-efficient shadow tomography when the number of measurement outcome scales polynomially with system size n, regardless of whether the measurement is global.

Quantum state tomography. It is proven that there exists a sample-optimal algorithm that can perform state tomography for an unknown quantum state ρ of rank $r \leq d$ using $O(dr/\epsilon^2)$ copies of ρ (Haah et al., 2017). Although the shadow tomography procedure in this paper does not require full information of ρ, the information obtained in this procedure increases linearly with K. In the extreme case, when we perform a quantum measurement on the computational basis—i.e., there are d possible outcomes corresponding to all possible n-bit classical strings $\boldsymbol{x}$ chosen from $\{0,1\}^n$. To perform shadow tomography on this measurement, we require $\Omega(d/\epsilon^2)$ copies of ρ. By performing this measurement, we can obtain a full description of any pure states. This bound is the same as the sample complexity required for state tomography for pure states.

1.4. Broader impact

This work focuses on the theory of quantum shadow tomography based on online learning, and as far as we see, we do not anticipate its potential negative societal impact. Nevertheless, it might have a positive impact on researchers who are interested in understanding the theoretical underpinnings of online learning applications and statistical learning.

2. Preliminaries

2.1. Classical Probability Theory

We consider two probability distributions $\mathcal{D} = (p_x)_x$ and $\mathcal{D}' = (q_x)_x$ on K-dimensional space, we will use the following two distance measures between them. The *total variation distance* between $\mathcal{D}$ and $\mathcal{D}'$ is defined by

$$d_{TV}(\mathcal{D}, \mathcal{D}') = \frac{1}{2}\sum_x |p_x - q_x|.$$

We also consider another distance measure that is commonly used for vectors. The *Euclidean norm* of the distance between the two distributions is defined by

$$\|\mathcal{D} - \mathcal{D}'\|_2 = \left(\sum_x (p_x - q_x)^2\right)^{1/2}.$$

The Euclidean norm is not commonly used in probability theory. We employ it as the intermediate tool when using the concentration inequalities on random vectors. To connect among these norms, we notice that for any probability distribution $\mathcal{D}$ and $\mathcal{D}'$, the following inequality holds

$$d_{TV}(\mathcal{D}, \mathcal{D}') \leq \frac{\sqrt{K}}{2}\|\mathcal{D} - \mathcal{D}'\|_2. \tag{1}$$

2.2. Quantum Preliminaries

Here, we briefly review some basic notations and concepts in quantum information. More details can be found, for example, in (Nielsen & Chuang, 2002).

A matrix $A \in \mathbb{C}^{d\times d}$ is said to be a *Hermitian* matrix if $A^\dagger = A$, where $A^\dagger$ denotes the conjugate transpose of A. We write $A \succeq 0$ when the Hermitian operator A is positive semidefinite. We write $A \succeq B$ when $A - B \succeq 0$. We use $\mathbb{I}$ for the identity matrix and the dimension can be understood from the context.

In quantum mechanics, a d-dimensional quantum state vector is described by a unit vector $\psi = (\psi_1, \ldots, \psi_d)^\top$ denoted by $|\psi\rangle$ with the Dirac symbol $|\cdot\rangle$, in a complex Hilbert space $\mathbb{C}^d$. The computational basis of $\mathbb{C}^d$ are defined as $\{|i\rangle\}_{i=1}^d$, where $|i\rangle = (0, \ldots, 0, 1, 0, \ldots, 0)^\top$ is the vector with the i-th entry being 1 and other entries being 0. A d-dimensional general *quantum state* can be written as a matrix $\rho \in \mathbb{C}^{d\times d}$ with $\rho \succeq 0$ and $\mathrm{Tr}(\rho) = 1$. If ρ has rank 1, it is a pure state and can be written as an outer product $|\psi\rangle\langle\psi|$ of a complex vector $|\psi\rangle$. Equivalently, we can write ρ as a convex combination for outer products of different pure states (without loss of generality, there can be at most d orthogonal pure states):

$$\rho = \sum_{i=1}^{d} p_i \left|\psi_i\right\rangle \left\langle\psi_i\right|,$$

where $\sum_{i=1}^{d} p_i = 1$ and $p_i \geq 0$ for arbitrary $i \in [d]$. This representation can be interpreted as a probability distribution over each pure state $|\psi_i\rangle$. In the special case when ρ is diagonal, it represents a classical probability distribution over orthogonal computational basis $|1\rangle, ..., |d\rangle$. The *maximally mixed state* $\mathbb{I}/d$ corresponds to the uniform distribution over $|1\rangle\langle 1|, ..., |d\rangle\langle d|$.

A *quantum observable*, or a *quantum operator*, is a d-dimensional Hermitian matrix $O \in \mathbb{C}^{d\times d}$. A quantum observable is a real-valued property of the physical systems. Given a quantum state ρ, the *expectation* of O with respect to ρ is defined by

$$\mathbb{E}_\rho[O] = \mathrm{Tr}(O\rho).$$

A *quantum event* is a quantum operator that satisfies $0 \preceq E \preceq \mathbb{I}$, i.e., a Hermitian operator with eigenvalues chosen from $[0, 1]$. The expectation value of a quantum event $\mathbb{E}_\rho[E]$ can be interpreted as a probability assigned by quantum state ρ to E. We further call E a *projector* for the special case when $E^2 = E$ and all eigenvalues for E are Boolean values 0 and 1.

A *quantum measurement* $\mathcal{M}$, or a *positive operator-valued measure (POVM)*, is a sequence $\mathcal{M} = (E_1, ..., E_K)$ of quantum events with $\sum_{j=1}^{K} E_j = \mathbb{I}$. According to the linearity of trace, we can obtain that

$$\sum_{j=1}^{K} \mathbb{E}_\rho[E_j] = \mathbb{E}_\rho\left[\sum_{j=1}^{K} E_j\right] = \mathbb{E}_\rho[\mathbb{I}] = 1.$$

Given a quantum state ρ, a POVM determines a probability distribution $\mathcal{D} = \{p_j\}_j$ on $[K]$ defined by $p_j = \mathbb{E}_\rho[E_j] = \mathrm{Tr}(E_j\rho)$.

2.3. Online Learning Settings and Regrets

In the context of online learning quantum states considered in Theorem 1.6, we are given a sequence of quantum measurements $\mathcal{M}_1, \mathcal{M}_2, ...$ in each iteration t. The learner constructs a hypothesis state $\omega_t \in \mathbb{C}^{d\times d}$ in each iteration. Given the quantum measurement $\mathcal{M}_t$, the learner calculates the distribution after applying $\mathcal{M}_t$ on the hypothesis state ω_t as $\boldsymbol{\mu}_t = (\mathrm{Tr}(E_{t,1}\rho), ..., \mathrm{Tr}(E_{t,K}\rho))$, which is known as a ”prediction”.

The learner then obtains feedback from the measurement $\mathcal{M}$. The simplest feedback can be a random variable Y_t chosen from value $[K] = \{1, ..., K\}$ for different outcomes. In this paper, the learner obtains feedback by performing a quantum distribution threshold search to find whether $d_{TV}(\boldsymbol{\mu}_t, \boldsymbol{p}_t)$ is larger than some tolerance threshold, where $\boldsymbol{p}_t = (\mathrm{Tr}(E_{t,1}\rho), ..., \mathrm{Tr}(E_{t,K}\rho))$ is the actual probability distribution for the unknown state ρ.

If the quantum distribution threshold search procedure does not output a t such that $d_{TV}(\boldsymbol{\mu}_t, \boldsymbol{p}_t) > 3\epsilon/4$, the learner accepts the prediction and set it as the final result. If the quantum distribution threshold search procedure outputs such a t, the learner starts an update procedure. The learner first estimates a probability distribution $\boldsymbol{b}_t$. According to Eq. (7), the learner can guarantee that $d_{TV}(\boldsymbol{b}_t, \boldsymbol{p}_t) \leq \epsilon/4$ with high probability by using $O(K/\epsilon^2)$ copies of ρ. Then, the learner defines a loss function that measures the total variation distance between the “bad prediction” $\boldsymbol{\mu}_t$ and $\boldsymbol{b}_t$ as:

$$\ell_t(\boldsymbol{\mu}_t) := \frac{1}{2}\sum_{j=1}^{K} |\mathrm{Tr}(E_{t,j}\omega_t) - b_{t,j}|, \tag{2}$$

where $b_{t,j}$ denotes the j-th entry of $\boldsymbol{b}_t$. The learner updates the hypothesis $\omega_t \to \omega_{t+1}$ based on the loss functions, measurements, and feedback before the current iteration.

Our goal is to design a strategy such that the learner’s total loss is minimized. Suppose there are in total T iterations, we want to find a strategy such that the learner’s total loss is not much more than that of the strategy which outputs the same quantum hypothesis φ in each iteration, where φ is chosen as the minimization of the total loss *with perfect hindsight.* Formally, we define the *regret* R_T to be the difference between values of total loss for these two strategies as

$$R_T := \sum_{t=1}^{T} \ell_t(\boldsymbol{\mu}_t) - \min_{\varphi\in\mathcal{H}} \sum_{t=1}^{T} \ell_t(\boldsymbol{\mu}_\varphi), \tag{3}$$

where $\boldsymbol{\mu}_\varphi = (\mathrm{Tr}(E_{t,1}\varphi), ..., \mathrm{Tr}(E_{t,K}\varphi))$ is the probability distribution after applying $\mathcal{M}_t$ on φ and $\mathcal{H}$ denotes the set of all d-dimensional quantum states. We remark that the sequence of measurements $\mathcal{M}_t$ can be arbitrary, even adversarial, based on the learner’s prior actions.

3. Quantum Distribution Threshold Search

In this section, we prove Theorem 1.5. In Section 4, we will use this procedure as feedback in the online learning procedure of our shadow tomography algorithm for K-outcome POVMs. Our starting point is the following expectation estimation lemma.

Lemma 3.1. *Let ρ be an unknown d-dimensional state and $\mathcal{M}$ be a K-outcome POVM. The probability distribution over the outcomes for applying $\mathcal{M}$ to ρ is $\boldsymbol{p} = (\mathrm{Tr}(E_1\rho), ..., \mathrm{Tr}(E_K\rho))$. We choose parameters $0 < \epsilon, \delta < \frac{1}{2}$. Then there exists $N = K\log(1/\delta)/\epsilon^2$ such that, for any d-dimensional quantum states ρ,*

$$\Pr\left(d_{TV}(\boldsymbol{p}, \boldsymbol{p}') \geq \frac{\epsilon}{8}\right) \leq \delta,$$

where $\boldsymbol{p}' = (p'_1, ..., p'_K)$ is the empirical distribution by applying $\mathcal{M}$ to the joint state $\rho^{\otimes N}$

Moreover, there exists a quantum event B such that for any K-dimensional distribution $\boldsymbol{\tau}$

$$d_{TV}(\boldsymbol{p}, \boldsymbol{\tau}) > \epsilon \Rightarrow \mathbb{E}_{\rho^{\otimes N}}[B] > 1 - \delta,$$
$$d_{TV}(\boldsymbol{p}, \boldsymbol{\tau}) \leq \frac{3\epsilon}{4} \Rightarrow \mathbb{E}_{\rho^{\otimes N}}[B] \leq \delta.$$

We provide the proof for Lemma 3.1 in Appendix B. Moreover, we can observe that if E_i's are projectors, then $A_{\boldsymbol{k}}$'s are also projectors. We can prove that B is a summation of $A_{\boldsymbol{k}}$'s. thus is also a projector. By using this lemma, we reduce the shadow tomography procedure of a K-outcome POVM to two-outcome ones.

Now, we begin to prove Theorem 1.5. Notice that the assumptions on $E_{i,j}$ is a quantum event in Problem 1.4 while the assumptions for $E_{i,j}$ is a projector in Lemma A.1, we have to first reduce the theorem to the case of projectors. Let $\rho \in \mathbb{C}^{d\times d}$ be the unknown quantum state, and $E_{i,j}$ be the quantum events for $i \in [M]$ and $j \in [K]$. We can achieve this through Naimark's theorem(see, for example, (Riesz & Nagy, 2012; Akhiezer & Glazman, 2013)). This theorem demonstrates that a quantum event $E \in \mathbb{C}^{d\times d}$ can be reduced to a projector Π on the space $\mathbb{C}^{2d\times 2d}$, such that for arbitrary ρ,

$$\mathbb{E}_{\rho\otimes|0\rangle\langle 0|}[\Pi] = \mathbb{E}_{\rho}[E].$$

Therefore, we assume that $E_{i,j}$ are projectors in the following proofs. Suppose we are given M K-outcome POVMs $\mathcal{M}_1, ..., \mathcal{M}_M$ and M threshold vectors $\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_M$. We first apply Lemma 3.1 with parameters $\delta = 1/4$ and $\tau = \boldsymbol{\theta}_i$ for each measurement $\mathcal{M}_i$. Therefore, we can find some $N_0 = O(K/\epsilon^2)$ such that each measurement $\mathcal{M}_i$ can be replaced by a quantum event $B_i \in (\mathbb{C}^{d\times d})^{\otimes N_0}$ satisfying

- if $d_{TV}(\boldsymbol{p}_i, \boldsymbol{\theta}_i) > \epsilon$, $\mathbb{E}_{\rho^{\otimes N_0}}[B_i] > 3/4$;
- if $d_{TV}(\boldsymbol{p}_i, \boldsymbol{\theta}_i) \leq 3\epsilon/4$, $\mathbb{E}_{\rho^{\otimes N_0}}[B_i] \leq 1/4$;

Here $\boldsymbol{p}_i$ is the actual distribution after applying $\mathcal{M}_i$ on ρ. Since $E_{i,j}$'s are projectors, quantum events B_i are also projectors.

Algorithm 1 RFTL for Quantum Tomography of K-outcome POVMs

```
Input: $T$, $\eta < \frac{1}{2}$.
Set $\omega_1 := \mathbb{I}/d$.
for $t = 1, ..., T$ do
    Predict $\omega_t$. Consider the loss function $\ell_t : \mathbb{R}^{K-1} \to \mathbb{R}$ given by measurement $\mathcal{M}_t$ : $\ell_t(\mathrm{Tr}(E_{t,1}\varphi), ..., \mathrm{Tr}(E_{t,K-1}\varphi))$. It has the same value with the loss function defined in Eq. (2). Let $\partial\ell_t/\partial x_j$ be a sub-derivative of $\ell_t$ with respect to $x_j$ for $j \in [K-1]$. Define
        $\nabla_t := \sum_{j=1}^{K-1} \frac{\partial \ell_t}{\partial(\mathrm{Tr}(E_{t,j})\omega_t)} E_{t,j}.$    (4)
    Update decision according to the RFTL rule with von Neumann entropy by $\omega_{t+1} :=$:
        $\arg\min_{\varphi\in\mathcal{H}} \left\{\eta \sum_{s=1}^{t} \mathrm{Tr}(\nabla_s \varphi) + \sum_{i=1}^{d} \lambda_i(\varphi)\log\lambda_i(\varphi)\right\},$    (5)
    where $\lambda_i(A)$ denotes the $i$-th eigenvalue of Hermitian matrix $A \in \mathbb{C}^{d\times d}$, and $\mathcal{H} \subset \mathbb{C}^{d\times d}$ denotes all $d$-dimensional quantum states.
end for
```

We then apply Lemma A.1 by setting each B_i to be the projectors we have just constructed and unknown state to be $\rho' = \rho^{\otimes N_0}$. If the algorithm outputs i^* such that $\mathbb{E}_{\rho'}[B_{i^*}] > 1/4$, we have $d_{TV}(\boldsymbol{p}_{i^*}, \boldsymbol{\theta}_{i^*}) > 3\epsilon/4$. Otherwise, we can guarantee that $d_{TV}(\boldsymbol{p}_i, \boldsymbol{\theta}_i) \leq \epsilon$ for all $i \in [M]$ with high probability.

4. Shadow Tomography of K-outcome POVMs

In this section, we first prove Theorem 1.6 in Section 4.1. We then prove the upper bound in Theorem 1.2.

4.1. Online Learning of Quantum States

We suppose there are in total T iterations where the learner performs an update procedure. In the update procedure, the learner follows the template of the Regularized Follow-the-Leader algorithm (RFTL) as the following Algorithm 1.

Algorithm 1 employs von Neumann entropy, which relates to the matrix exponentiated gradient algorithm (Tsuda et al., 2005). We remark that the loss function defined in Eq. (4) of the RFTL algorithm is slightly different from the definition in Eq. (2) in that it takes a vector of $(K-1)$ entries instead of K entries. This is because the input vectors in Eq. (2) are supposed to be probability distributions such that the summation of all entries is 1. Therefore, there are only

$(K-1)$ free parameters. We rewrite the loss function with an input vector containing only free entries as Eq. (4).

According to the definition of regret in Eq. (3), we now provide the following regret bound on this RFTL algorithm.

Lemma 4.1. *Setting $\eta = \sqrt{\log d/8T}$, the regret R_T of Algorithm 1 is bounded by $4\sqrt{(2\log 2)T\log d}$.*

The proof for Lemma 4.1 is provided in Appendix C. We then prove Theorem 1.6. We consider the case that the RFTL is triggered when the prediction $\boldsymbol{\mu}_t = (\mathrm{Tr}(E_{t,1}\omega_t), ..., \mathrm{Tr}(E_{t,K}\omega_t))$ deviates from the actual probability distribution $\boldsymbol{p}_t = (\mathrm{Tr}(E_{t,1}\rho), ..., \mathrm{Tr}(E_{t,K}\rho))$ for more than $3\epsilon/4$—i.e.,$d_{TV}(\boldsymbol{\mu}_t, \boldsymbol{p}_t) > 3\epsilon/4$. As the provided distribution $\boldsymbol{b}_t$ satisfies $d_{TV}(\boldsymbol{b}_t, \boldsymbol{p}_t) \leq \epsilon/4$, the loss function ℓ_t is at least $\epsilon/2$ by triangle inequality.

We then consider using the real distribution in each iteration, the loss function is at most $\epsilon/4$ in each iteration. By the regret bound, we have

$$\frac{\epsilon}{2}T \leq \frac{\epsilon}{4}T + 4\sqrt{2T\log d}.$$

Therefore, we can obtain the upper bound on T as $T \leq O(\log d/\epsilon^2)$.

4.2. Online Shadow Tomography of K-outcome POVMs

We now prove Theorem 1.2 using Theorem 1.5 and Theorem 1.6. We describe our online shadow tomography procedure for K-outcome POVMs below.

Given the requirement parameters ϵ, δ and the number of measurements M, we first define the following ancillary parameters

$$T_0 = \left\lceil \frac{C_0 \log d}{\epsilon^2} \right\rceil + 1, \quad \delta_0 = \frac{\delta}{2T_0},$$
$$N_0 = \frac{C_1 K \log(1/\delta_0)}{\epsilon^2} \log^2 M,$$
$$N_b = \frac{C_2 K \log(1/\delta_0)}{\epsilon^2} \log^2 M,$$

where C_0, C_1, and C_2 are three parameters that scale at most $\mathrm{poly}(\log\log M, \log\log D, \log(1/\epsilon), \log K)$. The number of copies of ρ will be $N = T_0(N_0 + N_b)$, which is indeed

$$N = \tilde{O}\left(\frac{\log(1/\delta)}{\epsilon^4} \cdot K \cdot \log^2 M \log d\right),$$

where $\tilde{O}$ hides a $\mathrm{poly}(\log\log M, \log\log D, \log(1/\epsilon), \log K)$ factor.

After receiving N copies of ρ, our algorithm first divides these states equally into T_0 batches, each consisting N_0 states. We prepare two joint states $\rho^{\otimes N_0}$ and $\rho^{\otimes N_b}$ using each batch. Each batch is used for the update procedure in a "bad iteration" in our online learning procedure.

To begin with, the learner initializes the hypothesis state $\omega_0 = \mathbb{I}/d$. In each iteration t, it chooses a fresh batch of states and runs the quantum distribution threshold search algorithm using joint state $\rho^{\otimes N_0}$. The threshold is chosen to be the probability distribution $\boldsymbol{\mu}_i$ after applying $\mathcal{M}_i$ for $i \in [M]$ on the hypothesis ω_t. According to Theorem 1.5, we can always find such C_1 to solve this quantum distribution search problem with success probability at least $1-\delta_0$.

If the quantum distribution threshold search declares that for all $i \in [M]$, $d_{TV}(\boldsymbol{\mu}_i, \boldsymbol{p}_i) \leq \epsilon$. Then we have successfully found a hypothesis such that the probability distributions after applying all K-outcome POVMs on this hypothesis are at most ϵ from that of the unknown state ρ.

If the quantum distribution threshold search outputs i^* where $d_{TV}(\boldsymbol{p}_{i^*}, \boldsymbol{\mu}_{i^*}) > 3\epsilon/4$. We use $\rho^{\otimes N_b}$ for an estimation $\boldsymbol{b}_{i^*}$ of the probability distribution after applying $\mathcal{M}_{i^*}$ on ρ. According to Eq. (7), we can always find C_2 such that with probability at least $1-\delta_0$, one can bound the total variation distance $d_{TV}(\boldsymbol{p}_{i^*}, \boldsymbol{b}_{i^*}) \leq \epsilon/4$. We supply this $\boldsymbol{b}_{i^*}$ to the learner and the learner employs the Algorithm 1 to update the hypothesis state into ω_{t+1}. Furthermore, the remaining copies in the current batch will be abandoned. The learner will move into the next iteration with a new batch.

According to Theorem 1.6, the number of "bad iterations" is bounded by $O(\log d/\epsilon^2)$. If there is no failure in any of the rounds, we can always find C_0 such that we can guarantee that for all $i \in [M]$, $d_{TV}(\boldsymbol{\mu}_i, \boldsymbol{p}_i) \leq \epsilon$ after the online procedure, where $\boldsymbol{\mu}_i$ is obtained by applying $\mathcal{M}_i$ for $i \in [M]$ on the hypothesis ω_{T_0}. Now we calculate the failure probability in this procedure. In each iteration, the success probability for the quantum distribution threshold search and the calculation of $\boldsymbol{b}_{i^*}$ are both at least $1-\delta_0$. By the union bound, the probability for failure after T_0 iterations is bounded by $2T_0\delta_0 = \delta$.

Finally, we consider the computational complexity of our shadow tomography procedure. In each iteration, we have to implement a series of $O(M)$ measurements on a batch of $O(K/\epsilon^2)$ joint samples to perform a quantum distribution threshold search and compute K terms for the gradient ∇_t to update the hypothesis in the RFTL protocol. As the iteration number is bounded by $O(\log d/\epsilon^2)$, the overall computational complexity is bounded by $O(KM \cdot \mathrm{poly}(K, 1/\epsilon, \log d)) + d^{O(1)}$.

5. The Lower Bound

We now show that any shadow tomography procedure for K-outcome POVMs requires at least $\Omega(\min\{D^2, K + \log(M)\}/\epsilon^2)$ copies of ρ. We set $D := \lfloor \min\{d, \sqrt{\log_2 M + K}\} \rfloor$ and suppose the unknown state is D-dimensional mixed state. We choose some constant $c \in (0,1)$ and set $L = \lfloor c^{D^2-K} \rfloor$. We will have

L quantum measurements of K outcomes for $L \leq M$. Notice that the probability for each outcome of the quantum measurement can be regarded as the expectation for a quantum event, there are in total $L \cdot K$ quantum events.

We choose $L \cdot K$ subspaces $\{S_{1,1}, ..., S_{1,K}\}$,..., $\{S_{L,1}, ..., S_{L,K}\}$ from $\mathbb{C}^{D\times D}$ for L POVMs independently and Haar-randomly such that $\dim(S_{i,j}) = D/K$ for any $i \in [L]$ and $j \in [K]$. The K subspaces in each set $\{S_{i,1}, ..., S_{i,K}\}$ are orthogonal. We denote $\mathbb{P}_{i,j}$ to be the projection to $S_{i,j}$ and $\rho_{i,j} = K\mathbb{P}_{i,j}/D$ to be the maximally mixed state projected onto $S_{i,j}$. As long as we choose a c that is close enough to 1, we can always find a choice over $S_{i,j}$'s with success probability $1 - o(1)$ such that

$$\left|\mathrm{Tr}(P_{i,j}\rho_{i',j'}) - \frac{1}{K}\right| \leq \frac{1}{2K} \tag{6}$$

for $i \neq i'$ according to Lemma A.2. We fix such a choice over $S_{i,j}$.

Without loss of generality, we assume that K is even. Now, we consider constructing the following states using a classical bit string $\boldsymbol{z} = (z_1, ..., z_{K/2})$ of $K/2$ bits as

$$\rho_i(\boldsymbol{z}) := \sum_{j=1}^{K/2}\left[\frac{1-50\epsilon z_j}{K}\rho_{i,2j-1} + \frac{1+50\epsilon z_j}{K}\rho_{i,2j}\right].$$

We consider applying measurement $\mathcal{M}_i$ on state $\rho_i(\boldsymbol{z})$. The $(2j-1)$-th and the $2j$-th entry for the probability distribution are

$$\mathrm{Tr}(E_{i,2j-1}\rho_i(\boldsymbol{z})) = \frac{1-50\epsilon z_i}{K},$$
$$\mathrm{Tr}(E_{i,2j}\rho_i(\boldsymbol{z})) = \frac{1+50\epsilon z_i}{K}.$$

We consider applying $\mathcal{M}_{i'}$ on state $\rho_i(\boldsymbol{z})$ for $i \neq i'$. According to Eq. (6), the j-th entry of the probability distribution is

$$\mathrm{Tr}(E_{i',j}\rho_i(\boldsymbol{z})) \leq \frac{3}{4}\cdot\frac{1+50\epsilon}{K} + \frac{1}{4}\cdot\frac{1-50\epsilon}{K} = \frac{1+25\epsilon}{K},$$
$$\mathrm{Tr}(E_{i',j}\rho_i(\boldsymbol{z})) \geq \frac{3}{4}\cdot\frac{1-50\epsilon}{K} + \frac{1}{4}\cdot\frac{1+50\epsilon}{K} = \frac{1-25\epsilon}{K}.$$

Now, we fix a measurement $\mathcal{M}_i$. If we apply this measurement to two quantum states $\rho_i(\boldsymbol{z}_1)$ and $\rho_{i'}(\boldsymbol{z}_2)$ for $i \neq i'$. The total variation distance for the two probability distributions is at least $25\epsilon/2$. It follows that, if we can estimate the probability distribution after a $\mathcal{M}_i$ to within total variation distance ϵ, we can immediately estimate $i \in [L]$ for the unknown state $\rho_i(\boldsymbol{z})$.

We then consider applying this measurement to two quantum states $\rho_i(\boldsymbol{z}_1)$ and $\rho_i(\boldsymbol{z}_2)$ for $\boldsymbol{z}_1 \neq \boldsymbol{z}_2$. Since each single difference on one entry in $\boldsymbol{z}_1$ and $\boldsymbol{z}_2$ will contribute $\frac{100\epsilon}{n}$ to the total variation distance between the two probability distributions, the distance is at least ϵ if more than 1% of the entries are different. Therefore, we can distinguish between such $\boldsymbol{z}_1$ and $\boldsymbol{z}_2$ with more than 1% different entries if we can estimate the probability distribution after a $\mathcal{M}_i$ to within total variation distance ϵ.

Suppose we choose i and $\boldsymbol{z}$ uniformly and randomly, then such choice contains $\log_2(2^{K/2}M) = \Omega(K + \log_2(M))$ bits of classical information. Suppose we require N copies of ρ to perform a shadow tomography procedure of K-outcome POVMs. Let

$$\zeta := \mathbb{E}_{i\in[L],\boldsymbol{z}\in\{0,1\}^{K/2}}\left[\rho_i(\boldsymbol{z})^{\otimes N}\right].$$

In order to make learning i and 99% of the entries for $\boldsymbol{z}$ from ζ information-theoretically possible, the mutual information $I(\zeta : i, \boldsymbol{z})$ must be at least $\Omega(K + \log_2(M))$. As both i and $\boldsymbol{z}$ are classical, we have

$$\begin{aligned} I(\zeta : i, \boldsymbol{z}) &= S(\zeta) - S(\zeta|i) \\ &= S(\zeta) - S(\rho_i(\boldsymbol{z})^{\otimes N}) \\ &\leq N(\log_2 D - S(\rho_i(\boldsymbol{z})), \end{aligned}$$

where $S(\cdot)$ is the von Neumann entropy. Now, we calculate the term $S(\rho_i(\boldsymbol{z}))$. Let $\lambda_{i,\boldsymbol{z},1}, ..., \lambda_{i,\boldsymbol{z},D}$ be the eigenvalues for $\rho_i(\boldsymbol{z})$. By applying a unitary transformation that diagonalizes $\rho_i(\boldsymbol{z})$ rotating to a basis that contains half of the projectors, we can observe that half the $\lambda_{i,\boldsymbol{z},j}$'s are $(1+50\epsilon)/D$ and the other half of the $\lambda_{i,\boldsymbol{z},j}$'s are $(1-50\epsilon)/D$. Hence,

$$\begin{aligned} S(\rho_i(\boldsymbol{z})) &= \sum_{j=1}^{D}\lambda_{i,\boldsymbol{z},j}\log_2\left(\frac{1}{\lambda_{i,\boldsymbol{z},j}}\right) \\ &\geq \log_2 D - O(\epsilon^2). \end{aligned}$$

Therefore, the mutual information can be bounded by

$$I(\zeta : i, \boldsymbol{z}) = O(N\epsilon^2).$$

As learning i and 99% of the entries for $\boldsymbol{z}$ requires $\Omega(K + \log_2(M))$ bits of classical information, we conclude that

$$N = \Omega\left(\frac{D^2}{\epsilon^2}\right) = \Omega\left(\frac{\min\{d^2, K + \log(M)\}}{\epsilon^2}\right).$$

6. Conclusion

This work theoretically established the exact dependence on K for shadow tomography of K-outcome quantum measurements and proposed the explicit algorithm that learns these distributions with sample complexity optimal in K. To the best of our knowledge, K is the only parameter we can obtain exact dependence for sample complexity in the context of shadow tomography. We conclude by discussing a few possible future directions.

- Can we develop an algorithm or provide a tight sample complexity that has smaller dependence on $\log M$, $\log d$, and $1/\epsilon$. In addition, it is interesting to explore whether the lower bound in our setting is multiplicative ($\Omega(K \cdot \log M)$) or additive ($\Omega(K + \log M)$). Also, can we find a trade-off relation between the sample and the time complexity?
- Can we achieve a better query complexity for shadow tomography with access to a unitary oracle that prepares the state? It has been shown that we can achieve quantum speedups in similar tasks (Huggins et al., 2021; van Apeldoorn et al., 2022) using quantum mean estimation (Hamoudi, 2021; Cornelissen et al., 2022).
- Can we find a shadow tomography procedure that has polynomial dependence on $\log d$, $\log M$, and $\log K$ for some family of $\mathcal{M}_i$'s and ρ's that are commonly considered in practical experiments?

Acknowledgements

S.A. was supported by a Vannevar Bush Fellowship from the US Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons ”It from Qubit” collaboration. The authors thank William Kretschmer, Sitan Chen, Tongyang Li, Dong-Ling Deng, and Weikang Li for helpful discussion, and anonymous reviewers for constructive comments.

References

Aaronson, S. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, 2007.

Aaronson, S. Shadow tomography of quantum states. *SIAM Journal on Computing*, 49(5):STOC18–368, 2019.

Aaronson, S. and Rothblum, G. N. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 322–333, 2019.

Aaronson, S., Chen, X., Hazan, E., Kale, S., and Nayak, A. Online learning of quantum states. *Advances in neural information processing systems*, 31, 2018.

Akhiezer, N. I. and Glazman, I. M. *Theory of linear operators in Hilbert space*. Courier Corporation, 2013.

Audenaert, K. M. and Eisert, J. Continuity bounds on the quantum relative entropy. *Journal of mathematical physics*, 46(10):102104, 2005.

Bădescu, C. and O'Donnell, R. Improved quantum data analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1398–1411, 2021.

Bassily, R., Nissim, K., Smith, A., Steinke, T., Stemmer, U., and Ullman, J. Algorithmic stability for adaptive data analysis. *SIAM Journal on Computing*, 50(3):STOC16–377, 2021.

Bhatia, R. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.

Bun, M., Ullman, J., and Vadhan, S. Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938, 2018.

Canonne, C. L. A short note on learning discrete distributions. *arXiv preprint arXiv:2002.11457*, 2020.

Carlen, E. A. and Lieb, E. H. Remainder terms for some quantum entropy inequalities. *Journal of Mathematical Physics*, 55(4):042201, 2014.

Chen, S., Cotler, J., Huang, H.-Y., and Li, J. Exponential separations between learning with and without quantum memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 574–585. IEEE, 2022a.

Chen, S., Huang, B., Li, J., Liu, A., and Sellke, M. Tight bounds for state tomography with incoherent measurements. *arXiv:2206.05265*, 2022b.

Chen, X., Hazan, E., Li, T., Lu, Z., Wang, X., and Yang, R. Adaptive online learning of quantum states. *arXiv:2206.00220*, 2022c.

Cornelissen, A., Hamoudi, Y., and Jerbi, S. Near-optimal quantum algorithms for multivariate mean estimation. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 33–43, 2022.

Dwork, C., Rothblum, G. N., and Vadhan, S. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 51–60. IEEE, 2010.

Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. L. Preserving statistical validity in adaptive data analysis. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 117–126, 2015.

Fano, R. M. *The transmission of information*. Massachusetts Institute of Technology, Research Laboratory of Electronics . . . , 1949.

Flammia, S. T., Gross, D., Liu, Y.-K., and Eisert, J. Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New Journal of Physics*, 14(9):095022, 2012.

Gross, D. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.

Haah, J., Harrow, A. W., Ji, Z., Wu, X., and Yu, N. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017.

Hamoudi, Y. Quantum sub-gaussian mean estimator. *arXiv preprint arXiv:2108.12172*, 2021.

Hayden, P., Leung, D. W., and Winter, A. Aspects of generic entanglement. *Communications in mathematical physics*, 265(1):95–117, 2006.

Hazan, E. et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

Holevo, A. S. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973.

Huang, H.-Y., Kueng, R., and Preskill, J. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

Huggins, W. J., Wan, K., McClean, J., O'Brien, T. E., Wiebe, N., and Babbush, R. Nearly optimal quantum algorithm for estimating multiple expectation values. *arXiv preprint arXiv:2111.09283*, 2021.

Kohler, J. M. and Lucchi, A. Sub-sampled cubic regularization for non-convex optimization. In *International Conference on Machine Learning*, pp. 1895–1904. PMLR, 2017.

Nielsen, M. A. and Chuang, I. Quantum computation and quantum information, 2002.

O'Donnell, R. and Wright, J. Efficient quantum tomography. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 899–912, 2016.

Preskill, J. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

Riesz, F. and Nagy, B. S. *Functional analysis*. Courier Corporation, 2012.

Smith, A. Lecture notes for the algorithmic foundations of adaptive data analysis. *Stability and adaptive analysis I*, Lecture 7-10, 2017.

Tsuda, K., Rätsch, G., and Warmuth, M. K. Matrix exponentiated gradient updates for on-line learning and bregman projection. *Journal of Machine Learning Research*, 6 (Jun):995–1018, 2005.

Ullman, J., Smith, A., Nissim, K., Stemmer, U., and Steinke, T. The limits of post-selection generalization. *Advances in Neural Information Processing Systems*, 31, 2018.

Vadhan, S. The complexity of differential privacy. In *Tutorials on the Foundations of Cryptography*, pp. 347–450. Springer, 2017.

van Apeldoorn, J., Cornelissen, A., Gilyén, A., and Nannicini, G. Quantum tomography using state-preparation unitaries. *arXiv preprint arXiv:2207.08800*, 2022.

A. Auxiliary lemmas

Lemma A.1 (Lemma 4.2, (Bădescu & O'Donnell, 2021))**.** *Suppose we are given an unknown d-dimensional quantum state ρ, and M quantum projectors $B_1, ..., B_M \in \mathbb{C}^{d\times d}$.There exists an algorithm using $O(\log^2 M \log(1/\delta))$ copies of ρ outputting either*

- $\mathbb{E}_\rho[B_{i^*}] = \mathrm{Tr}(B_{i^*}\rho) > 1/4$ *for some particular i^*; or*
- $\mathbb{E}_\rho[B_i] \leq 3/4$ *for any i,*

with success probability at least $1 - \delta$.

The proof of this theorem employs the χ^2-stable threshold reporting technique, which is a quantum version of classical statistical results fitting into the adaptive data analysis framework. We omit the details here and refer to (Smith, 2017), for example, for the related background.

Lemma A.2 (Hayden, Leung, and Winter (Hayden et al., 2006))**.** *Let S and T be two subspaces of $\mathbb{C}^{d\times d}$ with dimension d_1 and d_2. We denote $\mathbb{P}_S$ and $\mathbb{P}_T$ to be projectors on subspaces S and T. Consider $\rho_S = \frac{1}{d_1}\mathbb{P}_S$ to be the maximally mixed state projected onto S. If we fix T and randomly choose S, then*

$$\Pr\left[\left|\mathrm{Tr}(\mathbb{P}_T\rho_S) - \frac{d_2}{d}\right| \geq \frac{c_0 d_2}{d}\right] \leq \exp\left(-\cdot\frac{c_0^2 d_1 d_2}{6\ln 2}\right).$$

B. Expectation Estimation

We will need a concentration lemma for random vectors, which is an extension of the vector Bernstein inequality (Theorem 6) in (Gross, 2011; Kohler & Lucchi, 2017).

Lemma B.1. *Let $\mathbf{x}_1, ..., \mathbf{x}_m$ be independent K-dimensional vector-valued random variables. We assume that each random vector is zero-mean, uniformly bounded, and has bounded variance, i.e.,*

$$\mathbb{E}[\mathbf{x}_i] = 0 \text{ and } \|\mathbf{x}_i\|_\infty \leq \mu \text{ as well as } \mathbb{E}\left[\|\mathbf{x}_i\|_2^2\right] \leq \sigma^2$$

for some constants $\mu, \sigma > 0$. Suppose that parameters satisfies $0 < \epsilon < \sigma^2/\mu$, then we have

$$\Pr\left\{\left\|\frac{1}{m}\sum_{i=1}^m \mathbf{x}_i\right\|_2 \geq \epsilon\right\} \leq \exp\left(-m\cdot\frac{\epsilon^2}{8\sigma^2} + C\right),$$

for some positive constant C.

Proof. Theorem 6 in (Gross, 2011) indicates that for independent, zero-mean random vectors

$$\Pr\left\{\left\|\sum_{i=1}^m \mathbf{x}_i\right\| \geq t + \sqrt{V}\right\} \leq \exp\left(-\frac{t^2}{4V}\right),$$

where $V = \sum_{i=1}^m \mathbb{E}\left[\|\mathbf{x}_i\|_2^2\right]$ is the sum of variances for random vectors. We define $\epsilon = t + \sqrt{V}$ and rewrite the above inequality as

$$\Pr\left\{\left\|\sum_{i=1}^m \mathbf{x}_i\right\| \geq \epsilon\right\} \leq \exp\left(-\frac{1}{4}\left(\frac{\epsilon}{\sqrt{V}} - 1\right)^2\right) \leq \exp\left(-\frac{\epsilon^2}{8V} + \frac{1}{4}\right).$$

Since the sum of variance V can be bounded by $m\sigma^2$ according to our assumption, we can finally obtain the following inequality

$$\Pr\left\{\left\|\frac{1}{m}\sum_{i=1}^m \mathbf{x}_i\right\| \geq \epsilon\right\} \leq \exp\left(-\frac{1}{4}\left(\frac{\epsilon}{\sqrt{V}} - 1\right)^2\right) \leq \exp\left(-m\cdot\frac{\epsilon^2}{8\sigma^2} + \frac{1}{4}\right).$$

By choosing the constant $C = \frac{1}{4}$, we finish the proof for this lemma. □

Now we consider sampling from a probability distribution $\boldsymbol{p} = (p_1, ..., p_K)$ for m times. For the i-th sample where $i \in [m]$, we obtain one sample $\hat{\boldsymbol{p}}_i = (\hat{p}_i^1, ..., \hat{p}_i^K)$ with only one entry 1 and the other entries 0. We set $\mathbf{x}_i = \boldsymbol{p} - \hat{\boldsymbol{p}}_i$. Then $\mathbf{x}_i$ is centered because $\mathbb{E}(\mathbf{x}_i) = \mathbb{E}[\boldsymbol{p} - \hat{\boldsymbol{p}}_i] = 0$. Each entry of $\mathbf{x}_i$ is bounded below by 1 and

$$\mathbb{E}[\|\mathbf{x}_i\|^2] = \sigma^2 = 1 - \sum_{j=1}^{K} p_j^2 < 1.$$

Therefore, by Lemma B.1, we can guarantee that

$$\Pr\left(\left\|\frac{1}{m}\sum_{i=1}^{m}\hat{\boldsymbol{p}}_i - \boldsymbol{p}\right\|_2 \geq \epsilon\right) \leq \delta,$$

as long as we choose $m \geq O(\log(1/\delta)/\epsilon^2)$. To bound the total variation distance $d_{TV}(\frac{1}{m}\sum_{i=1}^{m}\hat{\boldsymbol{p}}_i, \boldsymbol{p})$ between the empirical distribution and the actual distribution, we combine the bound in Eq. (1) with Lemma B.1 and obtain:

$$\Pr\left(d_{TV}\left(\frac{1}{m}\sum_{i=1}^{m}\hat{\boldsymbol{p}}_i, \boldsymbol{p}\right) \geq \epsilon\right) \leq \Pr\left(\left\|\frac{1}{m}\sum_{i=1}^{m}\hat{\boldsymbol{p}}_i - \boldsymbol{p}\right\|_2 \geq \frac{2\epsilon}{\sqrt{K}}\right) \leq \exp\left(-m \cdot \frac{4\epsilon^2}{K} + \frac{1}{4}\right).$$

Hence, we can bound $d_{TV}(\frac{1}{m}\sum_{i=1}^{m}\hat{\boldsymbol{p}}_i, \boldsymbol{p})$ below ϵ with probability at least $1 - \delta$ if

$$m \geq O\left(\frac{\log(1/\delta)}{\epsilon^2} \cdot K\right). \tag{7}$$

We assign an index for every single copy ρ in the joint state $\rho^{\otimes N}$ and assume each single ρ occupies a "register". For all N-bit classical strings $\boldsymbol{x} = (x_1, ..., x_N) \in [K]^N$, we define quantum events $E_{\boldsymbol{x}} = E_{x_1} \otimes .. \times E_{x_N}$ to be the tensor product of quantum event E_{x_i} in the i-th register. It is easy to verify that $\sum_{\boldsymbol{x} \in \{0,1\}^N} E_x = \mathbb{I}$. For all K-dimensional positive integer arrays $\boldsymbol{k} = (k_1, ..., k_K)$ with $\sum_{j=1}^{K} k_j = N$, we define quantum event $A_{\boldsymbol{k}}$ to be

$$A_{\boldsymbol{k}} = \sum_{\substack{\boldsymbol{x} \in [K]^{\otimes N} \\ [\text{num of } x_i = j] = k_j}} E_x.$$

Then the empirical approximation $\boldsymbol{p}'$ is chosen as $\boldsymbol{p}' = \boldsymbol{k}/N$. Since each entry k_i of $\boldsymbol{k}$ is distributed as $\text{Binomial}(N, \text{Tr}(E_i\rho))$, we can bound the following probability using Eq. (7):

$$\Pr\left(d_{TV}(\boldsymbol{p}, \boldsymbol{p}') \geq \frac{\epsilon}{8}\right) \leq \delta \tag{8}$$

as $N = O(K\log(1/\delta)/\epsilon^2)$.

We define a function $f : [0,1]^{\otimes K} \to \{0,1\}$ by

$$f(\boldsymbol{t}) = \begin{cases} 1, \ d_{TV}(\boldsymbol{t}, \boldsymbol{\tau}) \geq \frac{7\epsilon}{8}, \\ 0, \ \text{otherwise.} \end{cases}$$

Based on this function, we define quantum event B by

$$B = \sum_{\substack{\boldsymbol{k} \\ k_1 + ... + k_K = N}} f\left(\frac{\boldsymbol{k}}{N}\right) A_{\boldsymbol{k}}.$$

As each entry k_i of $\boldsymbol{k}$ is distributed as $\text{Binomial}(N, \text{Tr}(E_i\rho))$, we can observe that

$$\mathbb{E}_{\rho^{\otimes N}}[B] = \Pr\left(d_{TV}(\boldsymbol{p}', \boldsymbol{\tau}) \geq \frac{7\epsilon}{8}\right).$$

Recall the guarantee in Eq. (8). The condition $d_{TV}(\boldsymbol{p}, \boldsymbol{\tau}) > \epsilon$ implies that $d_{TV}(\boldsymbol{p}', \boldsymbol{\tau}) \geq 7\epsilon/8$ by triangle inequality. Hence,

$$\mathbb{E}_{\rho^{\otimes N}}[B] = \Pr\left(d_{TV}(\boldsymbol{p}', \boldsymbol{\tau}) \geq \frac{7\epsilon}{8}\right) \geq \Pr\left(d_{TV}(\boldsymbol{p}, \boldsymbol{p}') \leq \frac{\epsilon}{8}\right) \geq 1 - \delta.$$

Similarly, the condition $d_{TV}(\boldsymbol{p}, \boldsymbol{\tau}) \leq 3\epsilon/4$ implies that $d_{TV}(\boldsymbol{p}', \boldsymbol{\tau}) \leq 7\epsilon/8$. Hence,

$$\mathbb{E}_{\rho^{\otimes N}}[\overline{B}] = \Pr\left(d_{TV}(\boldsymbol{p}', \boldsymbol{\tau}) > \frac{7\epsilon}{8}\right) > \Pr\left(d_{TV}(\boldsymbol{p}, \boldsymbol{p}') \leq \frac{\epsilon}{8}\right) \geq 1 - \delta.$$

C. Proof of Lemma 4.1

We mainly follow the template of the proof for Theorem 3 in (Aaronson et al., 2018), but there are some differences since the loss function is different. We first observe that the loss function $\ell_t(\mathrm{Tr}(E_{t,1}\varphi), ...,$
$\mathrm{Tr}(E_{t,K-1}\varphi))$ is convex. There are at most two terms that contain each $\mathrm{Tr}(E_{t,j}\varphi)$ in the loss function when calculating the sub-derivative over each value $\mathrm{Tr}(E_{t,j}\varphi)$:

- The variance in the j-th entry: $1/2|\mathrm{Tr}(E_{t,j}\varphi) - b_{t,j}|$;
- The variance in the last entry: $1/2|\mathrm{Tr}(E_{t,K}\varphi) - b_{t,K}|$ as $\mathrm{Tr}(E_{t,K}\varphi) = 1 - \sum_{j=1}^{K-1} \mathrm{Tr}(E_{t,j}\varphi)$.

Therefore, the value of sub-derivative $\partial \ell_t / \partial(\mathrm{Tr}(E_{t,j}))$ is either ± 1 or 0. We can divide all indexes j of $E_{t,j}$ into three subsets $S_{t,1}, S_{t,-1}$, and $S_{t,0}$ such that the value of $\partial \ell_t / \partial(\mathrm{Tr}(E_{t,j}))$ is $1, -1$, and 0 for j chosen from $S_{t,1}, S_{t,-1}$, and $S_{t,0}$. We thus rewrite ∇_t as:

$$\nabla_t = \sum_{j \in S_{t,1}} E_{t,j} - \sum_{j \in S_{t,-1}} E_{t,j}.$$

Notice that $E_{t,j}$ are projectors corresponding to different measurement outcomes and $\sum_{j=1}^{K} E_{t,j} = \mathbb{I}$, each $E_{t,j}$ are orthogonal and the spectral norm of any summation $\left\|\sum_{j\in[K]} E_{t,j}\right\| \leq 1$. We can thus bound the spectral norm of ∇_t below by

$$\|\nabla_t\| \leq \left\|\sum_{j \in S_{t,1}} E_{t,j}\right\| + \left\|\sum_{j \in S_{t,-1}} E_{t,j}\right\| \leq 2.$$

In the following, we denote $\boldsymbol{\mu}_t = \mathrm{Tr}(E_{t,1}\omega_t), ..., \mathrm{Tr}(E_{t,K-1}\omega_t)$ and $\boldsymbol{\tau}_t = \mathrm{Tr}(E_{t,1}\varphi), ..., \mathrm{Tr}(E_{t,K-1}\varphi)$ for simplicity. Since ℓ_t is convex,

$$\ell_t(\boldsymbol{\mu}_t) - \ell_t(\boldsymbol{\tau}_t) \leq \nabla_t \cdot (\omega_t - \varphi)$$

holds for all $\varphi \in \mathcal{H}$, where $\cdot$ denotes the trace inner-product between complex matrices. Summing over t, we obtain

$$\sum_{t=1}^{T} [\ell_t(\boldsymbol{\mu}_t) - \ell_t(\boldsymbol{\tau}_t)] \leq \sum_{t=1}^{T} [\mathrm{Tr}(\nabla_t \omega_t) - \mathrm{Tr}(\nabla_t \varphi)].$$

We define $g_t(X) = \nabla_t \cdot X$ for $X \in \mathcal{H}$ and $H(X)$ to be the negative von Neumann Entropy of X. By Lemma 5.2 in (Hazan et al., 2016), we have

$$\sum_{t=1}^{T} [g_t(\omega_t) - g_t(\varphi)] \leq \sum_{t=1}^{T} \nabla_t \cdot (\omega_t - \omega_{t+1}) + \frac{1}{\eta} D_R^2 \tag{9}$$

for any $\varphi \in \mathcal{H}$, where $D_R^2 \coloneqq \max_{\varphi, \varphi' \in \mathcal{H}} \{R(\varphi) - R(\varphi')\}$. We define $\Phi_t(X) = \eta \sum_{s=1}^{t} \nabla_s \cdot X + R(X)$, then line 5 of Algorithm 1 finds the minimal value of $\Phi_t(X)$ in $\mathcal{H}$. To prove the theorem, we need the following two claims.

Claim C.1. For all $t \in \{1.,,,.T\}$, we have $\omega_t \succ 0$.

Proof. Consider a Hermitian matrix $P \in \mathbb{C}^{d\times d}$ with zero minimal eigenvalue—i.e.,$\lambda_{\text{min}} = 0$. Suppose $P = VQV^\dagger$, where Q is a diagonal matrix with real entries as the eigenvalues of P. Assume $Q_{1,1} = \lambda_{\text{max}}(P)$ and $Q_{d,d} = \lambda_{\min}(P) = 0$. We consider a different matrix $P' = VQ'V^\dagger$ such that $Q'_{1,1} = Q_{1,1} - \epsilon$, $Q'_{i,i} = Q_{i,i}$ for $i \in \{2, ..., d-1\}$, and $Q'_{d,d} = \epsilon$ for $\epsilon < \lambda_{\max}(P)$. We then prove that there exists $\epsilon > 0$ that satisfies $\Phi_t(P') \leq \Phi_t(P)$. By expanding both sides of the inequality, we need to prove an equivalent inequality

$$A \cdot (P' - P) \leq \alpha \log \alpha - (\alpha - \epsilon)\log(\alpha - \epsilon) - \epsilon \log \epsilon,$$

where $A = \eta \sum_{s=1}^t \nabla_s$ and $\alpha = \lambda_{\max}(P) = Q_{1,1}$. Notice that $\|A\| \leq \eta \sum_{s=1}^t \|\nabla_s\| \leq 2\eta t$. The left side of the inequality can be bounded using Generalized Cauchy-Schwartz inequality (Bhatia, 2013) as

$$A \cdot (P - P') \leq 2\eta t \|P - P'\|_{\text{Tr}} \leq 4\epsilon\eta t.$$

where $\|A\|_{\text{Tr}}$ is the trace norm for matrix A. As $\log \epsilon \to -\infty$ when $\epsilon \to 0$, there exists a small enough ϵ such that $4\eta t \leq \log \alpha - \log \epsilon$. Therefore, we have

$$4\eta t\epsilon \leq \epsilon \log \alpha - \epsilon \log \epsilon \leq \alpha \log \alpha - (\alpha - \epsilon)\log(\alpha - \epsilon) - \epsilon \log \epsilon.$$

This indicates that there exists ϵ that is small enough such that $\Phi_t(P') \leq \Phi_t(P)$. If P has more than one zero eigenvalues, we can repeat the proof and construct the matrix P'. As ω_t is a minimal point of Φ_{t-1} and $\omega_1 \succeq 0$, we have $\omega_t \succeq 0$ for all t. □

Now, we can focus on $X \succeq 0$ and write $R(X) = \text{Tr}(X \log X)$. We can further calculate the gradient of $\Phi_t(X)$ as

$$\nabla \Phi_t(X) = \eta \sum_{s=1}^t \nabla_s + \mathbb{I} + \log X.$$

Here, we assume that the function $\Phi_t(X)$ is defined over real symmetric matrices. We can further prove the following claim.

Claim C.2. For all $t \in \{1, ..., T-1\}$, $\nabla \Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) \geq 0$.

Proof. We inversely assume that $\nabla \Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) < 0$. We choose a parameter $a \in (0, 1)$ and construct $\overline{X} = (1-a)\omega_{t+1} + a\omega_t$. Then $\overline{X} \succeq 0$ is also a density matrix. We denote $\Delta = \overline{X} - \omega_{t+1} = a(\omega_t - \omega_{t+1})$. According to Theorem 2 in (Audenaert & Eisert, 2005), we have

$$\begin{aligned}\Phi_t(\overline{X}) - \Phi_t(\omega_{t+1}) &\leq a\nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) + \frac{\text{Tr}(\Delta^2)}{\lambda_{\min}(\omega_{t+1})} \\ &= a\nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) + \frac{a^2\,\text{Tr}\big((\omega_t - \omega_{t+1})^2\big)}{\lambda_{\min}(\omega_{t+1})}.\end{aligned}$$

Then we divide the above inequality by a on both sides and get

$$\frac{\Phi_t(\overline{X}) - \Phi_t(\omega_{t+1})}{a} \leq \nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) + \frac{a\,\text{Tr}\big((\omega_t - \omega_{t+1})^2\big)}{\lambda_{\min}(\omega_{t+1})}.$$

Since we assume that $\nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) < 0$, we can always choose some small enough a such that the right-hand side is negative while the left-hand side is always positive since $\Phi_t(\overline{X}) > \Phi_t(\omega_{t+1})$. This leads to a contradiction. Therefore, we have proved that $\nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}) \geq 0$. □

We define

$$B_{\Phi_t}(\omega_t || \omega_{t+1}) \coloneqq \Phi_t(\omega_t) - \Phi_t(\omega_{t+1}) - \nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_{t+1}).$$

By Pinsker inequality (Carlen & Lieb, 2014), we have

$$\frac{1}{2}\|\omega_t - \omega_{t+1}\|_{\mathrm{Tr}}^2 \leq \mathrm{Tr}(\omega_t \log \omega_t) - \mathrm{Tr}(\omega_t \log \omega_{t+1}) = B_{\Phi_t}(\omega_t||\omega_{t+1}).$$

Using Claim C.2 and $\Phi_{t-1}(\omega_t) \leq \Phi_{t-1}(\omega_{t+1})$, we have

$$\begin{aligned} B_{\Phi_t}(\omega_t||\omega_{t+1}) &= \Phi_t(\omega_t) - \Phi_t(\omega_{t+1}) - \nabla\Phi_t(\omega_{t+1}) \cdot (\omega_t - \omega_[t+1]) \\ &\leq \Phi_t(\omega_t) - \Phi_t(\omega_{t+1}) \\ &= \Phi_{t-1}(\omega_t) - \Phi_{t-1}(\omega_{t+1}) + \eta\nabla_t \cdot (\omega_t - \omega_{t+1}) \\ &\leq \eta\nabla_t \cdot (\omega_t - \omega_{t+1}). \end{aligned}$$

Therefore,

$$\frac{1}{2}\|\omega_t - \omega_{t+1}\|_{\mathrm{Tr}}^2 \leq \eta\nabla_t(\omega_t - \omega_{t+1}).$$

By Generalized Cauchy-Schwartz inequality, we have

$$\begin{aligned} \nabla_t \cdot (\omega_t - \omega_{t+1}) &\leq \|\nabla_t\|\|\omega_t - \omega_{t+1}\|_{\mathrm{Tr}} \\ &\leq \|\nabla_t\|\sqrt{2\eta\nabla \cdot (\omega_t - \omega_{t+1})} \\ &\leq 2\eta\|\nabla_t\|^2 \\ &\leq 8\eta. \end{aligned}$$

We combine this inequality with Eq. (9) and reach the following bound

$$\sum_{t=1}^{T} \nabla_t \cdot (\omega_t - \varphi) \leq 8\eta T + \frac{1}{\eta} D_R^2.$$

We take $\eta = \frac{D_R}{2\sqrt{2T}}$. Observe that $D_R^2 \leq \log d$ according to the definition of von Neumann entropy, the value for η is

$$\eta = \sqrt{\frac{\log d}{8T}}.$$

The corresponding regret bound is

$$\sum_{t=1}^{T} [\ell_t(\boldsymbol{\mu}_t) - \ell_t(\boldsymbol{\tau}_t)] \leq \sum_{t=1}^{T} \nabla_t \cdot (\omega_t - \varphi) \leq 4\sqrt{2T \log d}.$$

D. An Exemplary Application

Here, we provide some applications of our shadow tomography procedure of K-outcome POVMs. In quantum mechanics, we are sometimes interested in the expectation value of quantum operators $\{O_i\}_{i=1}^M$:

$$o_i = \langle O_i \rangle = \mathrm{Tr}(O_i \rho),$$

given an unknown quantum state ρ. Suppose we perform a quantum measurement $\mathcal{M}_i$ that has K outcomes to estimate the expectation value o_i. Then the following corollary holds by using our shadow tomography procedure

Corollary D.1. *We consider an unknown d-dimensional quantum state, as well as M quantum operators $O_1, ..., O_M$. Assume we can measure each operator O_i using a quantum measurement $\mathcal{M}$ of K results. Then there exists a strategy that can approximate the expectation of each operator $\mathrm{Tr}(O_i \rho)$ within additive error ϵ using*

$$N = \tilde{O}\left(\frac{\max_i \|O_i\|^4}{\epsilon^4} \cdot K \cdot \log^2 M \log d\right)$$

copies of ρ. Here, $\|\cdot\|$ is the spectral norm. The success probability is at least $1 - \delta$.

To prove this corollary, we can divide the procedure into two steps.

In the first step, we approximate the distribution after we apply each measurement $\mathcal{M}_i$ within total variation distance $\epsilon / \max_i \|O_i\|$, which requires N copies of ρ according to Theorem 1.2.

Next, we calculate the expectation value using the distribution we obtained. The additive error for the expectation of $O_{i'}$ is bounded above by

$$\|O_{i'}\| \cdot \frac{\epsilon}{\max_i \|O_i\|} \leq \epsilon.$$

As an example, we consider a n-qubit quantum states that is $d = 2^n$-dimensional. We want to measure the expectation value for the operators $\{S_{\hat{n}_i}\}_{i=1}^{M}$ which measures the spin along $\hat{n}_i$ directions as

$$S_{\hat{n}_i} = \sum_{k=1}^{n} \sigma_{\hat{n}_i}^{k} \bigotimes_{k' \neq k} \mathbb{I}^k$$

where $\sigma_{\hat{n}_i}^{k}$ denotes the spin operator along $\hat{n}_i$ on the k-th operator and $\mathbb{I}^k$ denotes the identity operator on the k-th qubit. Each measurement $\mathcal{M}_i$ has $K = n + 1$ outcomes. The quantum event corresponding to each outcome $n - 2k$ for $k = 0, 1, ..., n$ can be written as a projector

$$A_{n-2k} = \sum_{\substack{x \in \{0,1\}^n \\ |x| = x - 2k}} |x\rangle \langle x| \,,$$

where $|x|$ represents the Hamming weight for string x. We can calculate the spectral norm $\|\cdot\|$ and the Hilbert-Schmidt norm $\|\cdot\|_{\text{HS}}$ of $S_{\hat{n}_i}$ by

$$\begin{aligned} \|S_{\hat{n}_i}\|_{\text{HS}} &= n2^n, \\ \|S_{\hat{n}_i}\| &= n. \end{aligned}$$

Therefore, we can approximate the expectation value for $\{S_{\hat{n}_i}\}_{i=1}^{M}$ using

$$N = \tilde{O}\left(\frac{\log^7 d}{\epsilon^4} \cdot \log^2 M\right)$$

copies of ρ according to Corollary D.1, which scales only poly-logarithmic on d. However, directly using classical shadow exponential number of samples.