Tight Bounds for Quantum State Certification with Incoherent Measurements

Sitan Chen
Department of EECS
UC Berkeley
Berkeley, USA
sitanc@berkeley.edu

Brice Huang
Department of EECS
Massachusetts Institute of Technology
Cambridge, USA
bmhuang@mit.edu

Jerry Li
Machine Learning Foundations Group
Microsoft Research
Redmond, USA
jerrl@microsoft.com

Allen Liu
Department of EECS
Massachusetts Institute of Technology
Cambridge, USA
cliu568@mit.edu

***Abstract*—We consider the problem of quantum state certification, where we are given the description of a mixed state $\sigma \in \mathbb{C}^{d\times d}$, n copies of a mixed state $\rho \in \mathbb{C}^{d\times d}$, and $\varepsilon > 0$, and we are asked to determine whether $\rho = \sigma$ or whether $\|\rho - \sigma\|_1 > \varepsilon$. When σ is the maximally mixed state $\frac{1}{d} I_d$, this is known as mixedness testing. We focus on algorithms which use incoherent measurements, i.e. which only measure one copy of ρ at a time. Unlike those that use entangled, multi-copy measurements, these can be implemented without persistent quantum memory and thus represent a large class of protocols that can be run on current or near-term devices.**

For mixedness testing, there is a folklore algorithm which uses incoherent measurements and only needs $O(d^{3/2}/\varepsilon^2)$ copies. The algorithm is non-adaptive, that is, its measurements are fixed ahead of time, and is known to be optimal for non-adaptive algorithms. However, when the algorithm can make arbitrary incoherent measurements, the best known lower bound is only $\Omega(d^{4/3}/\varepsilon^2)$ [5], and it has been an outstanding open problem to close this polynomial gap. In this work:

- **We settle the copy complexity of mixedness testing with incoherent measurements and show that $\Omega(d^{3/2}/\varepsilon^2)$ copies are necessary. This fully resolves open questions of [15] and [5].**
- **We show that the instance-optimal bounds for state certification to general σ first derived in [7] for non-adaptive measurements also hold for arbitrary incoherent measurements.**

Qualitatively, our results say that adaptivity does not help at all for these problems. Our results are based on new techniques that allow us to reduce the problem to understanding the concentration of certain matrix martingales, which we believe may be of independent interest.

***Index Terms*—Quantum learning, state certification, property testing, identity testing, instance optimality**

I. Introduction

Quantum mixedness testing, and more generally quantum state certification, are two of the most basic and fundamental tasks in quantum property testing. In quantum state certification, the learner is given n copies of a mixed state $\rho \in \mathbb{C}^{d\times d}$, and an explicit description of a mixed state $\sigma \in \mathbb{C}^{d\times d}$, and the objective is to distinguish with probability at least 0.99 between the case where $\rho = \sigma$ or if it is ε-far from σ in trace distance.[1] Mixedness testing is the special case of state certification where $\sigma = \frac{1}{d} I_d$, i.e., when the target state is the maximally mixed state.

Mixedness testing and state certification are the natural quantum analogues of uniformity testing and identity testing, respectively, two of the most well-studied problems in distribution testing. From a more practical point of view, state certification is also a key subroutine which allows experimentalists to verify the outcomes of their quantum experiments. For instance, if an algorithmist wishes to check that a quantum algorithm with quantum output is correctly outputting the right state, then this is exactly the problem of state certification.

Despite the fundamental nature of the problems, it was not until relatively recently that the copy complexity of state certification and mixedness testing were first understood. The seminal paper of [12] first demonstrated that $n = \Theta(d/\varepsilon^2)$ copies were necessary and sufficient to solve mixedness testing. Follow-up work of [3] later demonstrated that $n = O(d/\varepsilon^2)$ is also sufficient for the more general problem of state certification. Combined with the lower bound for mixedness testing, this resolved the copy complexity of state certification, in the worst case over σ.

However, a major downside of the estimators which achieve these copy complexities is that they require heavily entangled measurements over the joint state $\rho^{\otimes n}$. This poses a number of challenges to porting these algorithms into practical settings. First, the descriptions of the measurements are quite large (as the overall joint state is of size $d^n \times d^n$), and cannot

SC was supported in part by NSF Award 2103300. BH was supported by an NSF graduate research fellowship, a Siebel scholarship, NSF awards DMS-2022448 and CCF-1940205, and NSF TRIPODS award 1740751. AL was supported in part by an NSF Graduate Research Fellowship and a Fannie and John Hertz Foundation Fellowship.

[1] Note that by standard bootstrapping arguments the choice of constant here is arbitrary, and can be any constant larger than $1/2$. This only changes the sample complexity by constant factors.

2575-8454/22/$31.00 ©2022 IEEE
DOI 10.1109/FOCS54457.2022.00118

be implemented on current (or near-term) quantum devices. Second, the measurements require that all n copies of ρ are simultaneously present. In many realistic settings, such as streaming settings where one copy of ρ is given to the algorithm at a time, this would require that the quantum device be able to store all of these copies in persistent quantum memory. Such a task is also out of reach for current or near-term quantum devices, in essentially any non-trivial regime of the parameters, especially when one considers that d is exponential in the number of qubits in the system!

An appealing class of algorithms which avoids both these issues, and which can be implemented on real world noisy intermediate-scale quantum (NISQ) devices, are algorithms which only rely on *incoherent (a.k.a. unentangled) measurements*. In contrast to general protocols which perform arbitrary measurements on the joint state over all n copies, these algorithms only apply measurements to one copy of ρ at a time, although these measurements can possibly be adaptively chosen based on the (classical) outcomes of the previous measurements. Consequently, these measurements are performed on much smaller states, and moreover, can be performed without any quantum memory.

A. Optimal lower bounds for mixedness testing

For these reasons, there has been a considerable amount of attention in recent years devoted to understanding the statistical power of algorithms that only use incoherent measurements, which was also posed as an open problem in Wright's thesis [15]. A recent work of [5] demonstrated that if the measurements are additionally chosen non-adaptively, then $n = \Theta(d^{3/2}/\varepsilon^2)$ copies are necessary and sufficient to solve mixedness testing. They also demonstrated that *any* algorithm using incoherent measurements—even those chosen adaptively—must use at least $n = \Omega(d^{4/3}/\varepsilon^2)$ copies. In other words, there is a polynomial separation between the power of algorithms with and without quantum memory for this problem. Still, this left a gap between the best known upper and lower bounds for mixedness testing with incoherent measurements. This begs the question:

Can we fully characterize the copy complexity of mixedness testing with incoherent measurements?

Closing this gap was posed as an open question in the work of [5].

Underlying this question is another, more qualitative one, regarding the power of adaptivity. Indeed, a recurring theme in a number of different quantum learning settings is that while proving tight lower bounds against adaptive algorithms is quite challenging, the state-of-the-art algorithms almost always tend to be the "obvious" non-adaptive strategies. A very interesting meta-question is understanding for which natural quantum learning problems (if any) adaptivity helps at all for algorithms that use incoherent measurements.

Our first main contribution is fully resolve this question for mixedness testing: we prove that adaptivity does not improve the sample complexity at all, except possibly up to constant factors.

Theorem I.1 (Informal, see Theorem III.1)**.** *The copy complexity of mixedness testing using incoherent measurements is* $n = \Theta(d^{3/2}/\varepsilon^2)$.

By completely pinning down the copy complexity of mixedness testing with incoherent measurements, this answers open questions of [15] and [5]. Qualitatively, our theorem states that adaptivity does not help the copy complexity of this problem whatsoever.

B. Instance-optimal lower bounds for state certification.

We next turn to state certification. Because mixedness testing is a special case of state certification, Theorem I.1 immediately implies that $n = \Omega(d^{3/2}/\varepsilon^2)$ copies are necessary for state certification, in the worst case over all choices of the reference state σ. This, coupled with a matching upper bound from [7, Lemma 6.2], resolves the copy complexity of state certification with incoherent measurements for worst-case σ.

However, it should be clear that this bound is not the correct bound for all possible σ. For instance, when σ is pure, it is not hard to see that $\Theta(1/\varepsilon^2)$ copies are sufficient and necessary. This raises the natural question: what is the copy complexity of state certification with incoherent measurements, as a function of the reference state σ? This is the quantum analogue of the (classical) distribution testing problem of obtaining instance optimal bounds for identity testing against a known distribution over d elements [1], [2], [4], [8], [10], [14]. In the classical version of the problem, there is a known distribution p over $\{1, \dots, d\}$, and we are given samples from a distribution q. We are asked to distinguish between the case wher $p = q$, and the case when $\|p - q\|_1 > \varepsilon$. A landmark result of [14] states that the sample complexity of this question is (essentially) characterized by the $\ell_{2/3}$-quasinorm of p.

In this work, we ask whether or not a similar characterization can be obtained for the quantum version of the question. Prior work of [7] demonstrated such a characterization, but under the caveat that the measurements are chosen non-adaptively. At a high level, they showed that the copy complexity of the problem is governed by the *fidelity* between σ and the maximally mixed state. More precisely, they showed that if $\overline{\sigma}$ and $\underline{\sigma}$ are states given by zeroing out eigenvalues of σ that have total mass at most $\Theta(\varepsilon^2)$ and $\Theta(\varepsilon)$ respectively and normalizing, then the copy complexity with non-adaptive measurements, denoted n, satisfies

$$\widetilde{\Omega}\left(\frac{d \cdot \underline{d}_{\mathsf{eff}}^{1/2}}{\varepsilon^2} \cdot F(\underline{\sigma}, \frac{1}{d} I_d)\right) \leq n \leq \widetilde{O}\left(\frac{d \cdot \overline{d}_{\mathsf{eff}}^{1/2}}{\varepsilon^2} \cdot F(\overline{\sigma}, \frac{1}{d} I_d)\right), \tag{1}$$

where $\underline{d}_{\mathsf{eff}}$ (resp. $\overline{d}_{\mathsf{eff}}$) is the "effective dimension" of the problem, namely, the rank of $\underline{\sigma}$ (resp. $\overline{\sigma}$). In the same work, they also gave lower bounds for arbitrary (possibly adaptive) incoherent measurements, but, like with mixedness testing, these lower bounds were looser and did not match the corresponding upper bound. In light of this, we ask:

Can we give an instance-optimal characterization of the copy complexity of state certification with incoherent measurements?

Our second main contribution is to give such a characterization:

Theorem I.2 (Informal, see Theorem 8.1 of full version)**.** *For any σ, and ε sufficiently small, the copy complexity of state certification w.r.t. σ using incoherent measurements is upper and lower bounded by* (1)*.*

We regard this as strong evidence that, as with mixedness testing, adaptivity does not help for state certification. It is not always a tight bound, as there are states for which the upper and lower bounds in (1) can differ by polynomial factors for some choices of ε, and so this bound can be loose, even in the non-adaptive setting. Still, we conjecture that for all σ, the copy complexity of state certification to σ with incoherent and non-adaptive measurements is the same as that with arbitrary incoherent measurements. Indeed, when ε is sufficiently small compared to the smallest nonzero eigenvalue of σ, our bounds are tight up to constant factors.

C. Our techniques.

We achieve our new lower bounds via a new proof technique which we believe may be of independent interest. As with other lower bounds in this area, we reduce to a "one-versus-many" distinguishing problem. To construct this instance, prior work leveraged the natural quantum analogue of Paninski's famous construction in the lower bound for (classical) uniformity testing [13] – namely, an additive perturbation by a multiple of $UZU^\dagger$, where U is a Haar random matrix and $Z = \mathrm{diag}(1, \ldots, -1, \ldots)$ has equally many $+1$s and -1s.

We instead use a different hard instance based on Gaussian perturbations. While this introduces a number of additional technical challenges, the key advantage of this instance is that the likelihood ratio for this instance has a very clean, self-similar form (see Lemma III.5). This allows us to essentially reduce the problem into one of understanding the concentration of a certain matrix martingale defined by the learning process, as well as an auxiliary matrix balancing question. We can then use classical tools from scalar and matrix concentration to demonstrate that the likelihood ratio is close to 1 with high probability over all possible outcomes of the learning algorithm, which yields our desired lower bound.

Not only does this framework dramatically simplify many of the difficult concentration calculations in prior work such as [5], it also has the conceptual advantage that it never requires a *pointwise* bound on the likelihood ratio. To our knowledge, all prior lower bounds against adaptive algorithms in this literature required some worst-case pointwise bound on the likelihood ratio. For some problems, e.g. shadow tomography [6], this was already sufficient to prove tight lower bounds. However, for mixedness testing, a worst-case bound cannot be sufficient, and from a technical perspective, the fact that [5] had to balance between their (much tighter) average case bound on the likelihood ratio and this (fairly large) worst-case bound to control the contribution of certain tail events was why their overall lower bound was loose. Consequently, we believe that this martingale-based technique may also yield tight lower bounds for a number of other problems in the literature.

Roadmap. After providing some technical preliminaries in Section II, including basics on showing lower bounds for algorithms that use incoherent measurements, we give a complete proof of Theorem I.1 in Section III. We defer the proof of Theorem I.2 to the full version, available at https://arxiv.org/abs/2204.07155.

II. Preliminaries

Throughout, let ρ denote the unknown state, and let $\rho_{\mathsf{mm}} = \frac{1}{d} I_d$ denote the maximally mixed state.

We now define the standard measurement formalism, which is the way algorithms are allowed to interact with the unknown quantum state ρ.

Definition II.1 (Positive operator valued measurement (POVM), see e.g. [11])**.** *A positive operator valued measurement $\mathcal{M}$ is a finite collection of psd matrices $\mathcal{M} = \{M_z\}_{z \in \mathcal{Z}}$ satisfying $\sum_z M_z = I_d$. When a state ρ is measured using $\mathcal{M}$, we get a draw from a classical distribution over $\mathcal{Z}$, where we observe z with probability $\mathrm{Tr}(\rho M_z)$. Afterwards, the quantum state is destroyed.*

Next, we formally define what we mean by an algorithm that uses incoherent measurements. Intuitively, such an algorithm operates as follows: given n copies of ρ, it iteratively measures the i-th copy using a POVM (which could depend on the results of previous measurements), records the outcome, and then repeats this process on the $(i+1)$-th copy. After having performed all n measurements, it must output a decision based on the (classical) sequence of outcomes it has received. More formally, such an algorithm can be represented as a tree:

Definition II.2 (Tree representation, see e.g. [6])**.** *Fix an unknown d-dimensional mixed state ρ. A learning algorithm that only uses n incoherent, possibly adaptive, measurements of ρ can be expressed as a rooted tree $\mathcal{T}$ of depth n satisfying the following properties:*

- *Each node is labeled by a string of vectors $\boldsymbol{x} = (x_1, \ldots, x_t)$, where each x_i corresponds to measurement outcome observed in the i-th step.*
- *Each node $\boldsymbol{x}$ is associated with a probability $p^\rho(\boldsymbol{x})$ corresponding to the probability of observing $\boldsymbol{x}$ over the course of the algorithm. The probability for the root is 1.*
- *At each non-leaf node, we measure ρ using a rank-1 POVM $\{\omega_x d \cdot xx^\dagger\}_x$ to obtain classical outcome $x \in \mathbb{S}^{d-1}$. The children of $\boldsymbol{x}$ consist of all strings $\boldsymbol{x}' = (x_1, \ldots, x_t, x)$ for which x is a possible POVM outcome.*
- *If $\boldsymbol{x}' = (x_1, \ldots, x_t, x)$ is a child of $\boldsymbol{x}$, then*

$$p^\rho(\boldsymbol{x}') = p^\rho(\boldsymbol{x}) \cdot \omega_x d \cdot x^\dagger \rho x. \tag{2}$$

- *Every root-to-leaf path is length-n. Note that $\mathcal{T}$ and ρ induce a distribution over the leaves of $\mathcal{T}$.*

We briefly note that in this definition, we assume that the POVMs are always rank-1. It is a standard fact that this is without loss of generality (see e.g. [6, Lemma 4.8]).

A. Notation.

Given $z \in \mathbb{R}$, we use z_- to denote $-\min(z, 0)$. We use $\wedge$ to denote minimum. We use $f \lesssim g$ to denote $f = O(g)$ and $f \ll g$ to denote $f = o(g)$. We will always implicitly assume a sufficiently large system; for example, if $f \gg g$ we will assume where necessary that $f \geq 100g$. We use $f = \widetilde{O}(g)$ (resp. $f = \widetilde{\Omega}(g)$) to denote that there exists some absolute constant c for which $f = O(g \cdot \log^c g)$ (resp. $f = \Omega(g/\log^c g)$).

Given a vector v, we use $\|v\|_p$ to denote its ℓ^p norm; when $p = 2$, we sometimes drop the subscript. Given a matrix M, we use $\|M\|_{\mathrm{op}}$ or $\|M\|$ to denote its operator norm, $\|M\|_1$ to denote its trace norm, and $\|M\|_F$ to denote its Frobenius norm.

For a string $\boldsymbol{x} = (x_1, \ldots, x_n)$, we let $\boldsymbol{x}_{\sim i}$ and $\boldsymbol{x}_{\sim i,j}$ denote the string with the i-th index removed and the string with the i-th and j-th indices removed. For any set $S \subseteq [n]$, we let $\boldsymbol{x}_S$ denote the string restricted to the entries in S.

We will work with the following random matrix ensemble:

Definition II.3 (Trace-centered Gaussian orthogonal ensemble (GOE))**.** *For $d \in \mathbb{N}$, let $G \sim \mathrm{GOE}(d)$, that is, $G \in \mathbb{R}^{d\times d}$ is symmetric with upper diagonal entries sampled independently from $\mathcal{N}(0, 1/d)$ and diagonal entries sampled independently from $\mathcal{N}(0, 2/d)$.*

Define $M = G - \frac{\mathrm{Tr}(G)}{d} I_d$. We say that M is a trace-centered GOE matrix *and denote its distribution* $\mathrm{GOE}^*(d)$*. For $U \subseteq \mathbb{R}^{d\times d}$, $\overline{M}$ is a U-*truncated trace-centered GOE matrix *if it is drawn from* $\mathrm{GOE}^*(d)$ *conditioned on $\overline{M} \in U$. We denote the distribution of $\overline{M}$ by* $\mathrm{GOE}^*_U(d)$.

Our result for state certification uses the following notion of fidelity.

Definition II.4 (Fidelity between two quantum states)**.** *The fidelity of quantum states $\rho, \sigma \in \mathbb{C}^{d\times d}$ is $F(\rho, \sigma) = (\mathrm{Tr}\sqrt{\rho^{1/2}\sigma\rho^{1/2}})^2$.*

Our lower bounds are based on Le Cam's two-point method which we briefly review here. The following is an elementary result in binary hypothesis testing:

Fact II.5 (See e.g. Theorem 4.3 from [16])**.** *Given distributions p_0, p_1 over a domain $\mathcal{S}$, if $d_{\mathrm{TV}}(p_0, p_1) < 1/3$, there is no $\mathcal{A} : \mathcal{S} \to \{0, 1\}$ for which $\Pr_{x\sim p_i}[\mathcal{A}(x) = i] \geq 2/3$ for both $i = 0, 1$.*

Now consider a state distinguishing task of the form

$$H_0 : \rho = \sigma \qquad \text{and} \qquad H_1 : \rho = \sigma_M,$$

where σ_M is a random state sampled from some distribution $\mathcal{D}$ over the set of states satisfying $\|\sigma - \sigma_M\|_1 > \epsilon$. Recall from Definition II.2 that a learning algorithm that uses n incoherent measurements corresponds to a tree $\mathcal{T}$ of depth n, and $\rho = \sigma$ and $\rho = \sigma_M$ induce distributions p_0 and p_M on the leaves of this tree. We can use Fact II.5 to reduce proving a copy complexity lower bound for state certification with respect to σ, which is a worst-case guarantee over all possible input states ρ, to bounding $d_{\mathrm{TV}}(p_0, \mathbb{E}_M[p_M])$, which is an average-case bound.

Lemma II.6 (Le Cam's two-point method, see e.g. Lemma 1 in [17])**.** *If there is a distribution $\mathcal{D}$ over states satisfying $\|\sigma - \sigma_M\|_1 > \epsilon$ for which $d_{\mathrm{TV}}(p_0, \mathbb{E}_M[p_M]) \leq 1/3$ for any tree $\mathcal{T}$ of depth n, then any algorithm $\mathcal{A}$ using incoherent measurements for state certification with respect to σ must make more than n incoherent measurements to achieve success probability at least $2/3$.*

Proof. Suppose to the contrary there existed such an algorithm $\mathcal{A}$ using at most n incoherent measurements, and let p_0 and p_M denote the distributions over the leaves of the tree corresponding to $\mathcal{A}$ when $\rho = \sigma$ and $\rho = \sigma_M$ respectively. Suppose when it succeeds, $\mathcal{A}$ outputs 0 when $\rho = \sigma$ and 1 when $\|\rho - \sigma\|_1 > \epsilon$. Let $p_1 \triangleq \mathbb{E}_{M\sim\mathcal{D}}[p_M]$. Because $\mathcal{A}$ successfully outputs 1 with probability 2/3 when given as input the state σ_M for any M, $2/3 \leq \mathbb{E}_M[\Pr_{\boldsymbol{x}\sim p_M}[\mathcal{A}(\boldsymbol{x}) = 1]] = \mathbb{E}_{\boldsymbol{x}\sim p_1}[\mathcal{A}(\boldsymbol{x}) = 1]$. Similarly, $2/3 \leq \mathbb{E}_{\boldsymbol{x}\sim p_0}[\mathcal{A}(\boldsymbol{x}) = 0]$. By Fact II.5, this would contradict the bound on $d_{\mathrm{TV}}(p_0, p_1)$. □

III. Lower Bound for Mixedness Testing

In this section we prove the following theorem, which is the formal version of Theorem I.1.

Theorem III.1. *Let $d \gg 1$ and $0 < \varepsilon \leq 1/12$. Any algorithm using incoherent measurements that can distinguish between $\rho = \rho_{\mathsf{mm}}$ and $\|\rho - \rho_{\mathsf{mm}}\|_1 > \varepsilon$ with probability at least $2/3$ must use at least $n = \Omega(d^{3/2}/\varepsilon^2)$ copies of ρ.*

By the upper bound in [5], this is tight up to constant factors. Also, by standard amplification arguments, the choice of constant in the success probability is arbitrary, and can be taken to be any constant which is strictly larger than $1/2$.

Formally, we consider the task of distinguishing between the following two alternatives:

$$H_0 : \rho = \frac{1}{d} I_d \qquad \text{and} \qquad H_1 : \rho = \frac{1}{d}(I_d + \varepsilon\overline{M}).$$

Here, $\overline{M} \sim \mathrm{GOE}^*_U(d)$ for the U given by Lemma III.2 below.

Lemma III.2. *There exists $U \subseteq \mathbb{R}^{d\times d}$ such that if $M \sim \mathrm{GOE}^*(d)$, then $\Pr[M \notin U] \leq \exp(-\Omega(d))$ and on the event $M \in U$, we have $\|M\|_{\mathrm{op}} \leq 3$ and $\|M\|_1 \geq d/12$.*

Note that this lemma ensures that under H_1, ρ is psd (and thus a valid quantum state) and has trace distance $\Omega(\varepsilon)$ to $\frac{1}{d} I_d$. Its proof is by standard spectral bounds on the GOE ensemble and can be found in Appendix A of the full version of this paper. Theorem III.1 follows from the following theorem.

Theorem III.3. *Let $d \gg 1$ and $0 < \varepsilon \leq 1/12$. Any algorithm using incoherent measurements that distinguishes between H_0*

and H_1 with success probability at least $2/3$ must use at least $n = \Omega(d^{3/2}/\varepsilon^2)$ copies of ρ.

Take any learning tree $\mathcal{T}$ corresponding to an algorithm for this task that uses n incoherent measurements. Recalling the terminology from Definition II.2, we let p_0 and p_1 denote the distributions over leaves of $\mathcal{T}$ induced by ρ under H_0 and H_1 respectively. In the rest of this section, we assume $n \ll d^{3/2}/\varepsilon^2$ and will prove $d_{\mathrm{TV}}(p_0, p_1) = o(1)$. It is clear that this implies Theorem III.3.

For a sequence of unit vectors $\boldsymbol{x} = (x_1, \ldots, x_n)$, we define the likelihood ratio $L^*(\boldsymbol{x}) \triangleq p_1(\boldsymbol{x})/p_0(\boldsymbol{x})$. Note that

$$L^*(\boldsymbol{x}) = \mathop{\mathbb{E}}_{\overline{M} \sim \mathrm{GOE}^*_U(d)} \left[\prod_{i=1}^n \left(1 + \varepsilon x_i^\dagger \overline{M} x_i\right) \right].$$

Similarly define

$$L(\boldsymbol{x}) \triangleq \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[\prod_{i=1}^n \left(1 + \varepsilon x_i^\dagger \overline{M} x_i\right) \right]. \tag{3}$$

This is an estimate for the likelihood ratio $L^*(\boldsymbol{x})$ where the conditioned Gaussian integral is replaced by a true Gaussian integral. Most of the computations in this section will be done in terms of $L(\boldsymbol{x})$; the proof of Theorem III.3 below quantifies that $L(\boldsymbol{x})$ is a close approximation of $L^*(\boldsymbol{x})$.

We will somewhat abuse notation and write $L(\boldsymbol{z})$ for any sequence of unit vectors $\boldsymbol{z} = (z_1, \ldots, z_t)$ of length not necessarily n. This is defined identically to (3). We also write $L(\boldsymbol{x}, \boldsymbol{x})$ to denote the value of L on input $(x_1, x_1, x_2, x_2, \ldots, x_n, x_n)$.

The main ingredient in the proof of Theorem III.3 is the following high-probability bound on L evaluated at the leaves of $\mathcal{T}$.

Proposition III.4. *There exists a subset S of the leaves of $\mathcal{T}$ such that $\Pr_{p_0}[S] = 1 - o(1)$ and for all $\boldsymbol{x} \in S$, $|L(\boldsymbol{x}) - 1| = o(1)$ and $L(\boldsymbol{x}, \boldsymbol{x}) \ll e^{\sqrt{d}}$.*

Let us first prove Theorem III.3 assuming Proposition III.4.

Proof of Theorem III.3. Let U be as in Lemma III.2. Define

$$\overline{L}(\boldsymbol{x}) = \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[\mathbb{1}\{M \in U\} \prod_{i=1}^n \left(1 + \varepsilon x_i^\dagger M x_i\right) \right].$$

It is clear that $L^*(\boldsymbol{x}) = \Pr[U]^{-1} \overline{L}(\boldsymbol{x})$. For all $\boldsymbol{x} \in S$, by Cauchy-Schwarz

$$\begin{aligned}
&|L(\boldsymbol{x}) - \overline{L}(\boldsymbol{x})| \\
&= \left| \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[\mathbb{1}\{M \notin U\} \prod_{i=1}^n \left(1 + \varepsilon x_i^\dagger M x_i\right) \right] \right| \\
&\le \Pr[U^c]^{1/2} \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[\prod_{i=1}^n \left(1 + \varepsilon x_i^\dagger M x_i\right)^2 \right]^{1/2} \\
&= \sqrt{\Pr[U^c] L(\boldsymbol{x}, \boldsymbol{x})} = o(1).
\end{aligned}$$

Here we use that $\Pr[U^c] \le \exp(-\Omega(d))$ and $L(\boldsymbol{x}, \boldsymbol{x}) \ll e^{\sqrt{d}}$. Since $|L(\boldsymbol{x}) - 1| = o(1)$, we have $\overline{L}(\boldsymbol{x}) = 1 + o(1)$ and

$$\begin{aligned}
|L^*(\boldsymbol{x}) - 1| &\le |L^*(\boldsymbol{x}) - \overline{L}(\boldsymbol{x})| + |\overline{L}(\boldsymbol{x}) - 1| \\
&= \frac{\Pr[U^c]}{\Pr[U]} \overline{L}(\boldsymbol{x}) + o(1) = o(1).
\end{aligned}$$

Finally,

$$\begin{aligned}
d_{\mathrm{TV}}(p_0, p_1) &= 2 \mathop{\mathbb{E}}_{\boldsymbol{x} \sim p_0} [(L^*(\boldsymbol{x}) - 1)_-] \\
&= 2 \mathop{\mathbb{E}}_{\boldsymbol{x} \sim p_0} [\mathbb{1}\{\boldsymbol{x} \in S\}(L^*(\boldsymbol{x}) - 1)_-] \\
&\quad + 2 \mathop{\mathbb{E}}_{\boldsymbol{x} \sim p_0} [\mathbb{1}\{\boldsymbol{x} \notin S\}(L^*(\boldsymbol{x}) - 1)_-] \\
&\le 2 \sup_{\boldsymbol{x} \in S} (L^*(\boldsymbol{x}) - 1)_- + 2 \Pr_{p_0}[S^c] = o(1). \qquad \square
\end{aligned}$$

A. Recursive evaluation of likelihood ratio

Let $\boldsymbol{z} = (z_1, \ldots, z_t)$ be a sequence of unit vectors. Recall that $\boldsymbol{z}_{\sim i}$ and $\boldsymbol{z}_{\sim i,j}$ denote $\boldsymbol{z}$ with z_i and z_i, z_j omitted. The following lemma gives a recursive formula for $L(\boldsymbol{z})$.

Lemma III.5. *The following identity holds.*

$$L(\boldsymbol{z}) = L(\boldsymbol{z}_{\sim t}) + \frac{2\varepsilon^2}{d^2} \sum_{i=1}^{t-1} \left[\left(d \langle z_i, z_t \rangle^2 - 1 \right) L(\boldsymbol{z}_{\sim i,t}) \right].$$

The proof is based on Isserlis' theorem, which we record below. For k even, let $\mathrm{PMat}(k)$ denote the set of perfect matchings of $\{1, \ldots, k\}$.

Theorem III.6 (Isserlis' theorem, [9])**.** *Let $g = (g_1, \ldots, g_k)$ be a jointly Gaussian vector. If k is odd, then $\mathbb{E}[\prod_{i=1}^k g_i] = 0$. If k is even, then*

$$\mathbb{E}\left[\prod_{i=1}^k g_i \right] = \sum_{\{\{a_1, b_1\}, \ldots, \{a_{k/2}, b_{k/2}\}\} \in \mathrm{PMat}(k)} \prod_{i=1}^{k/2} \mathbb{E}[g_{a_i} g_{b_i}].$$

Proof of Lemma III.5. For a set $S \subseteq [t]$ with $|S|$ even, let $\mathrm{PMat}(S)$ denote the set of perfect matchings of S. For even $k \le t$, let $\mathrm{Mat}(t, k)$ denote the set of matchings of $[t]$ consisting of $k/2$ pairs. We compute that

$$\begin{aligned}
L(\boldsymbol{z}) &= \sum_{S \subseteq [t]} \varepsilon^{|S|} \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[\prod_{i \in S} (z_i^\dagger M z_i) \right] \\
&= \sum_{k=0}^{\lfloor t/2 \rfloor} \varepsilon^{2k} \sum_{\{\{a_1, b_1\}, \ldots, \{a_k, b_k\}\} \in \mathrm{Mat}(t, 2k)} \\
&\qquad \prod_{i=1}^k \mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[(z_{a_i}^\dagger M z_{a_i})(z_{b_i}^\dagger M z_{b_i}) \right] \\
&= \sum_{k=0}^{\lfloor t/2 \rfloor} \left(\frac{2\varepsilon^2}{d^2} \right)^k \sum_{\{\{a_1, b_1\}, \ldots, \{a_k, b_k\}\} \in \mathrm{Mat}(t, 2k)} \\
&\qquad \prod_{i=1}^k \left(d \langle z_{a_i}, z_{b_i} \rangle^2 - 1 \right). \qquad (4)
\end{aligned}$$

In the final step we use that for unit vectors $x, y \in \mathbb{C}^d$,

$$\mathop{\mathbb{E}}_{M \sim \mathrm{GOE}^*(d)} \left[(x^\dagger M x)(y^\dagger M y) \right] = \frac{2}{d^2} (d \langle x, y \rangle^2 - 1),$$

which can be verified by direct computation. The lemma follows by partitioning the summands in (4) based on whether t appears in the matching, and if so which $i \in \{1, \dots, t-1\}$ it is paired with. □

B. High probability bound on likelihood ratio at leaves

This subsection gives the main part of the proof of Proposition III.4. For any sequence of unit vectors $\boldsymbol{z} = (z_1, \dots, z_t)$, define

$$H(\boldsymbol{z}) = \sum_{i=1}^{t} \left(d z_i z_i^\dagger - I_d \right) \frac{L(\boldsymbol{z}_{\sim i})}{L(\boldsymbol{z})},$$
$$K(\boldsymbol{z}) = \sum_{i=1}^{t} \left(d z_i z_i^\dagger - I_d \right).$$

The function H enters our calculations by the following rewriting of Lemma III.5:

$$\frac{L(\boldsymbol{z})}{L(\boldsymbol{z}_{\sim t})} = 1 + \frac{2\varepsilon^2}{d^2} \cdot z_t^\dagger H(\boldsymbol{z}_{\sim t}) z_t. \tag{5}$$

If $\boldsymbol{z} = \boldsymbol{x}_{\le t} \triangleq (x_1, \dots, x_t)$ is a prefix of $\boldsymbol{x} \sim p_0$, then $\frac{L(\boldsymbol{z})}{L(\boldsymbol{z}_{\sim t})} = \frac{L(\boldsymbol{x}_{\le t})}{L(\boldsymbol{x}_{\le t-1})}$ is one step in the likelihood ratio martingale. As we will see (proof of Claim III.10) below, the multiplicative fluctuation of this step is

$$\mathop{\mathbb{E}}_{x_t}\left[\left(\frac{L(\boldsymbol{x}_{\le t})}{L(\boldsymbol{x}_{\le t-1})} \right)^2 \right] = 1 + O\left(\frac{\varepsilon^4}{d^5} \right) \left\| H(\boldsymbol{x}_{\le t-1}) \right\|_F^2.$$

Thus, an upper bound on $\|H(\boldsymbol{z})\|_F$ over all prefixes $\boldsymbol{z}$ of $\boldsymbol{x}$ controls the fluctuations of the likelihood ratio martingale. Because the matrices $H(\boldsymbol{z})$ are hard to control directly, we will use $K(\boldsymbol{z})$ as a proxy for $H(\boldsymbol{z})$. The following lemma quantifies this relationship, showing that if $K(\boldsymbol{z})$ is bounded in Frobenius norm, $H(\boldsymbol{z})$ is bounded at the same scale.

Lemma III.7. *Suppose $1 \ll \gamma \ll d/(\varepsilon^2 n^{1/2})$. If $\boldsymbol{z} = (z_1, \dots, z_t)$ is a sequence of unit vectors satisfying $t \le n$ and $\|K(\boldsymbol{z})\|_F \le n^{1/2} d\gamma$, then $\|H(\boldsymbol{z})\|_F \le 3n^{1/2} d\gamma$.*

Note that this lemma is a "deterministic" statement about a sequence of vectors. We will prove this in Subsection III-C. The following lemma bounds $K(\boldsymbol{z})$ in Frobenius norm uniformly over all prefixes $\boldsymbol{z}$ of $\boldsymbol{x}$. We will prove this lemma in Subsection III-D by mimicking the proof of Doob's L^2 maximal inequality for the matrix valued martingale $K(\boldsymbol{x}_{\le t})$.

Lemma III.8. *If $\boldsymbol{x} \sim p_0$, then $\mathbb{E}\left[\sup_{1 \le t \le n} \|K(\boldsymbol{x}_{\le t})\|_F^2 \right] \lesssim nd^2$.*

We will now prove Proposition III.4 assuming Lemmas III.7 and III.8. We set α, β to be slowly-growing functions such that $1 \ll \alpha \ll \beta \ll d^{3/2}/(\varepsilon^2 n) \wedge d/(\varepsilon^2 n^{1/2})$, and furthermore $\alpha^2 \ll d^{3/2}/(\varepsilon^2 n)$. This is possible because $n \ll d^{3/2}/\varepsilon^2$.

Let $\boldsymbol{x} \sim p_0$. For $1 \le t \le n$, define the filtration $\mathcal{F}_t = \sigma(\boldsymbol{x}_{\le t})$ and the sequences

$$H_t = H(\boldsymbol{x}_{\le t}), \quad K_t = K(\boldsymbol{x}_{\le t}), \quad \Phi_t = L(\boldsymbol{x}_{\le t}).$$

Consider the time

$$\tau = \inf\left\{ t : \|K_t\|_F > n^{1/2} d\alpha \text{ or } |\Phi_t - 1| > \frac{\varepsilon^2 n}{d^{3/2}} \beta \right\} \cup \{\infty\},$$

which is clearly a stopping time with respect to $\mathcal{F}_t$. Also define the stopped sequence $\Psi_t = \Phi_{t \wedge \tau}$.

Claim III.9. *With probability $1 - o(1)$, $\|K_t\|_F \le n^{1/2} d\alpha$ for all $1 \le t \le n$.*

Proof. By Lemma III.8,

$$\Pr\left[\sup_{1 \le t \le n} \|K_t\|_F > n^{1/2} d\alpha \right] \le \frac{\mathbb{E}\left[\sup_{1 \le t \le n} \|K_t\|_F^2 \right]}{nd^2\alpha^2} \lesssim \alpha^{-2} = o(1).$$ □

Claim III.10. *With probability $1 - o(1)$, $|\Psi_n - 1| \le \frac{\varepsilon^2 n}{d^{3/2}} \beta$.*

Proof. Note that Ψ_t is a multiplicative martingale: if $\tau \le t-1$ then certainly $\mathbb{E}[\frac{\Psi_t}{\Psi_{t-1}} | \mathcal{F}_{t-1}] = 1$, and if $\tau > t-1$, (5) implies

$$\mathbb{E}\left[\frac{\Psi_t}{\Psi_{t-1}} | \mathcal{F}_{t-1} \right] = 1 + \frac{2\varepsilon^2}{d^2} \mathbb{E}\left[x_t^\dagger H_{t-1} x_t | \mathcal{F}_{t-1} \right] = 1,$$

using that

$$\begin{aligned} \mathbb{E}\left[x_t^\dagger H_{t-1} x_t | \mathcal{F}_{t-1} \right] &= \sum_{x_t} \omega_{x_t} (x_t^\dagger H_{t-1} x_t) \\ &= \left\langle H_{t-1}, \sum_{x_t} \omega_{x_t} x_t x_t^\dagger \right\rangle \\ &= \langle H_{t-1}, I_d/d \rangle = 0. \end{aligned} \tag{6}$$

We next bound the quadratic increment $\mathbb{E}[(\frac{\Psi_t}{\Psi_{t-1}})^2 | \mathcal{F}_{t-1}]$. If $\tau \le t-1$ this is 1, and otherwise

$$\mathbb{E}\left[\left(\frac{\Psi_t}{\Psi_{t-1}} \right)^2 | \mathcal{F}_{t-1} \right] = 1 + \frac{4\varepsilon^4}{d^4} \mathbb{E}\left[(x_t^\dagger H_{t-1} x_t)^2 | \mathcal{F}_{t-1} \right]$$

since the linear term vanishes by (6). The remaining expectation can be bounded by

$$\begin{aligned} \mathbb{E}\left[(x_t^\dagger H_{t-1} x_t)^2 | \mathcal{F}_{t-1} \right] &= \sum_{x_t} \omega_{x_t} x_t^\dagger H_{t-1} (x_t x_t^\dagger) H_{t-1} x_t \\ &\le \sum_{x_t} \omega_{x_t} x_t^\dagger H_{t-1}^2 x_t \\ &= \left\langle H_{t-1}^2, \sum_{x_t} \omega_{x_t} x_t x_t^\dagger \right\rangle \\ &= \frac{1}{d} \|H_{t-1}\|_F^2. \end{aligned}$$

Moreover, since $\tau > t-1$, $\|K_{t-1}\|_F \le n^{1/2} d\alpha$ and Lemma III.7 implies $\|H_{t-1}\|_F \le 3n^{1/2} d\alpha$. Thus,

$$\mathbb{E}\left[\left(\frac{\Psi_t}{\Psi_{t-1}} \right)^2 | \mathcal{F}_{t-1} \right] \le 1 + \frac{36\varepsilon^4 n}{d^3} \alpha^2. \tag{7}$$

So, for all $1 \le t \le n$,

$$\mathbb{E}[\Psi_t^2] = \mathbb{E}\left[\mathbb{E}\left[\left(\frac{\Psi_t}{\Psi_{t-1}}\right)^2 |\mathcal{F}_{t-1}\right]\Psi_{t-1}^2\right] \le \left(1 + \frac{36\varepsilon^4 n}{d^3}\alpha^2\right)\mathbb{E}[\Psi_{t-1}^2],$$

and therefore

$$\mathbb{E}[\Psi_t^2] \le \exp\left(\frac{36\varepsilon^4 n^2}{d^3}\alpha^2\right) \le 2$$

since $\frac{\varepsilon^4 n^2}{d^3}\alpha^2 \ll 1$. Moreover, (recalling $\mathbb{E}[\Psi_t] = 1$ and (7))

$$\begin{aligned}\mathbb{E}[(\Psi_t - 1)^2] &= \mathbb{E}\left[\mathbb{E}\left[\left(\frac{\Psi_t}{\Psi_{t-1}}\right)^2 |\mathcal{F}_{t-1}\right]\Psi_{t-1}^2 - 1\right] \\ &\le \frac{36\varepsilon^4 n}{d^3}\alpha^2\,\mathbb{E}[\Psi_{t-1}^2] + \mathbb{E}[(\Psi_{t-1}-1)^2] \\ &\le \frac{72\varepsilon^4 n}{d^3}\alpha^2 + \mathbb{E}[(\Psi_{t-1}-1)^2],\end{aligned}$$

so by induction

$$\mathbb{E}[(\Psi_n - 1)^2] \le \frac{72\varepsilon^4 n^2}{d^3}\alpha^2.$$

Thus

$$\Pr\left[|\Psi_n - 1| > \frac{\varepsilon^2 n}{d^{3/2}}\beta\right] \le \frac{\mathbb{E}\left[|\Psi_n - 1|^2\right]}{\frac{\varepsilon^4 n^2}{d^3}\beta^2} \le \frac{72\alpha^2}{\beta^2} = o(1).$$

Therefore, $|\Psi_n - 1| \le \frac{\varepsilon^2 n}{d^{3/2}}\beta$ with probability $1 - o(1)$. □

Claim III.11. *If $\|K_n\|_F \le n^{1/2}d\alpha$, then $L(\boldsymbol{x},\boldsymbol{x}) \ll e^{\sqrt{d}}$.*

Proof. We refer the reader to Claim 6.11 of the full version of our paper. □

Proof of Proposition III.4. Define the event

$$S = \left\{\sup_{1\le t\le n}\|K_t\|_F \le n^{1/2}d\alpha \text{ and } |\Psi_n - 1| \le \frac{\varepsilon^2 n}{d^{3/2}}\beta\right\}.$$

By Claims III.9 and III.10, $\Pr_{p_0}[S] = 1 - o(1)$. We will show that if S holds, then $\tau = \infty$. Indeed, if $\tau = t < \infty$, then either $\|K_t\|_F > n^{1/2}d\alpha$ or $|\Phi_t - 1| > \frac{\varepsilon^2 n}{d^{3/2}}\beta$ holds. Since $\Psi_n = \Phi_t$, this contradicts S.

So, $\tau = \infty$ on S. This implies $|L(\boldsymbol{x}) - 1| = |\Phi_n - 1| \le \frac{\varepsilon^2 n}{d^{3/2}}\beta = o(1)$. Moreover $\|K_n\|_F \le n^{1/2}d\alpha$, so by Claim III.11 we have $L(\boldsymbol{x},\boldsymbol{x}) \ll e^{\sqrt{d}}$. □

C. Bounding $H(\boldsymbol{x}_{\le t})$ in Frobenius norm by bootstrapping

In this subsection, we prove Lemma III.7. The main idea of the proof is to begin with a preliminary bound on $\|H(\boldsymbol{x}_{\le t})\|_F$ and its sub-sums (Lemma III.13) and, using the self-similar nature of the likelihood ratio, bootstrap this bound into the desired result. Throughout this subsection, let $\boldsymbol{z} = (z_1, \ldots, z_t)$ be a sequence of unit vectors satisfying $t \le n$ and

$$\|K(\boldsymbol{z})\|_F \le n^{1/2}d\gamma \tag{8}$$

for some $1 \ll \gamma \ll d/(\varepsilon^2 n^{1/2})$.

The following lemma bounds a variant of $K(\boldsymbol{z})$ where we multiply each summand by an adversarial $b_i \in [-1,1]$. This will be used to control the discrepancy $H(\boldsymbol{z}) - K(\boldsymbol{z})$ in the bootstrapping argument.

Lemma III.12. *Uniformly over $b_1, \ldots, b_t \in [-1,1]$, we have*

$$\left\|\sum_{i=1}^t b_i\left(dz_i z_i^\dagger - I_d\right)\right\|_F \le n^{1/2}d\gamma + 2nd^{1/2}.$$

Proof. For any choice of $b_1, \ldots, b_t$,

$$\begin{aligned}\left\|\sum_{i=1}^t b_i\left(dz_i z_i^\dagger - I_d\right)\right\|_F &\le \left\|\sum_{i=1}^t b_i \cdot dz_i z_i^\dagger\right\|_F + \left\|\sum_{i=1}^t b_i \cdot I_d\right\|_F \\ &\le \left\|\sum_{i=1}^t dz_i z_i^\dagger\right\|_F + t\sqrt{d} \\ &\le \|K(\boldsymbol{z})\|_F + 2t\sqrt{d}.\end{aligned}$$

The result follows by $t \le n$ and (8). □

For $S \subseteq [t]$, let $\boldsymbol{z}_S = (z_i)_{i\in S}$. Further, let

$$H_S = \sum_{i\in S}\left(dz_i z_i^\dagger - I_d\right)\cdot\frac{L(\boldsymbol{z}_{S\setminus\{i\}})}{L(\boldsymbol{z}_S)},$$
$$K_S = \sum_{i\in S}\left(dz_i z_i^\dagger - I_d\right).$$

The following lemma gives a preliminary bound on $\|H_S\|_F$. In the proof of Lemma III.7, we will use this bound to control $\|H_S\|_F$ for $|S| = t - O(\log n)$, followed by the bootstrap argument over $O(\log n)$ recursive rounds to contract the bound to $O(n^{1/2}d)$.

Lemma III.13. *For all $S \subseteq [t]$, $\|H_S\|_F \le 2n^{1/2}d\gamma + 4nd^{1/2}$.*

Proof. For any fixed $\overline{M} \in U$ and unit vector z,

$$\varepsilon\left|z^\dagger\overline{M}z\right| \le \frac{1}{12}\cdot 3 \le \frac{1}{2},$$

so $1 + \varepsilon z^\dagger\overline{M}z \in [1/2, 3/2]$. Thus, for all i, $L(\boldsymbol{z}_S)/L(\boldsymbol{z}_{S\setminus\{i\}}) \in [1/2, 3/2]$, which implies

$$\frac{L(\boldsymbol{z}_{S\setminus\{i\}})}{L(\boldsymbol{z}_S)} \in [2/3, 2]. \tag{9}$$

Lemma III.12 gives

$$\frac{1}{2}\|H_S\|_F \le n^{1/2}d\gamma + 2nd^{1/2},$$

as desired. □

Proof of Lemma III.7. Let $D = \log\sqrt{n/d}$. If $t < D$, by (9),

$$\|H(\boldsymbol{z})\|_F \le \sum_{i=1}^t\left\|dz_i z_i^\dagger - I_d\right\|_F \cdot \frac{L(\boldsymbol{z}_{\sim i})}{L(\boldsymbol{z})} \le 2dD \ll n^{1/2}d\gamma$$

as desired. Otherwise $t \ge D$. We will prove by induction on $a \ge 0$ that if $S \subseteq [t]$ satisfies $|S| = t - D + a$, then

$$\|H_S\|_F \le \xi_a \triangleq 2n^{1/2}d\gamma + 4e^{-a}nd^{1/2}.$$

The base case $a = 0$ holds by Lemma III.13. Assume $a \geq 1$; by the inductive hypothesis and (5), for all $i \in S$

$$\left| \frac{L(\boldsymbol{z}_S)}{L(\boldsymbol{z}_{S\setminus\{i\}})} - 1 \right| \leq \frac{2\varepsilon^2}{d^2} \left\| H_{S\setminus i} \right\|_{\mathrm{op}} \leq \frac{2\varepsilon^2}{d^2} \xi_{a-1}.$$

Since this upper bound is $o(1)$, we also have

$$\frac{L(\boldsymbol{z}_{S\setminus\{i\}})}{L(\boldsymbol{z}_S)} - 1 = \frac{3\varepsilon^2}{d^2} \xi_{a-1} b_i$$

for some $b_i \in [-1, 1]$. By Lemma III.12,

$$\begin{aligned}
&\left\| \sum_{i\in S} \left(dz_i z_i^\dagger - I_d \right) \left(\frac{L(\boldsymbol{z}_{S\setminus\{i\}})}{L(\boldsymbol{z}_S)} - 1 \right) \right\|_F \\
&= \frac{5\varepsilon^2}{d^2} \xi_{a-1} \left\| \sum_{i\in S} b_i \left(dz_i z_i^\dagger - I_d \right) \right\|_F \\
&\leq \left(\frac{3\varepsilon^2 n^{1/2}}{d} \gamma + \frac{6\varepsilon^2 n}{d^{3/2}} \right) \xi_{a-1} \leq e^{-1} \xi_{a-1},
\end{aligned}$$

using the hypotheses $\gamma \ll d/(\varepsilon^2 n^{1/2})$ and $n \ll d^{3/2}/\varepsilon^2$. Since

$$\begin{aligned}
\|K_S\|_F &\leq \|K(\boldsymbol{z})\|_F + \sum_{i\in[t]\setminus S} \left\| dz_i z_i^\dagger - I_d \right\|_F \\
&\leq n^{1/2} d\gamma + 2dD \leq \frac{101}{100} n^{1/2} d\gamma,
\end{aligned}$$

we have

$$\begin{aligned}
\|H_S\|_F &\leq \|K_S\|_F + \left\| \sum_{i\in S} \left(dz_i z_i^\dagger - I_d \right) \left(\frac{L(\boldsymbol{z}_{S\setminus\{i\}})}{L(\boldsymbol{z}_S)} - 1 \right) \right\|_F \\
&\leq \frac{101}{100} n^{1/2} d\gamma + e^{-1} \xi_{a-1} \leq \xi_a,
\end{aligned}$$

as $\frac{101}{100} + 2e^{-1} \leq 2$. This completes the induction. Finally,

$$\|H(\boldsymbol{z})\|_F = \left\| H_{[t]} \right\|_F \leq 2n^{1/2} d\gamma + 4e^{-D} n d^{1/2} \leq 3n^{1/2} d\gamma. \square$$

D. Uniform Frobenius bound on the $K(\boldsymbol{x}_{\leq t})$

The proof of Lemma III.8 mimics the proof of Doob's L^2 maximal inequality. Let $\boldsymbol{x} \sim p_0$, recall that $K_t = K(\boldsymbol{x}_{\leq t})$, and define $X = \sup_{1\leq t\leq n} \|K_t\|_F$.

Lemma III.14. *We have that* $\mathbb{E}[X^2] \leq 4\,\mathbb{E}[\|K_n\|_F^2]$

Proof. We will first upper bound $\Pr[X \geq x]$ for all $x > 0$. Consider the stopping time $\tau = \inf\{t : \|K_t\|_F \geq x\} \cup \{n\}$. Then,

$$\begin{aligned}
\Pr[X \geq x] &= \Pr[\|K_\tau\|_F \geq x] \\
&\leq x^{-1}\, \mathbb{E}[\|K_\tau\|_F \mathbb{1}\{\|K_\tau\|_F \geq x\}] \\
&\leq x^{-1}\, \mathbb{E}[\mathbb{E}[\|K_n\|_F |\mathcal{F}_\tau] \mathbb{1}\{\|K_\tau\|_F \geq x\}] \\
&= x^{-1}\, \mathbb{E}[\|K_n\|_F \mathbb{1}\{X \geq x\}].
\end{aligned}$$

The first estimate is by Markov's inequality, and the second is by convexity of the norm $\|\cdot\|_F$. Thus,

$$\begin{aligned}
\mathbb{E}[X^2] &= \int_0^\infty \Pr[X^2 \geq x]\ \mathrm{d}x = \int_0^\infty \Pr[X \geq x] 2x\ \mathrm{d}x \\
&\leq \int_0^\infty 2\,\mathbb{E}[\|K_n\|_F \mathbb{1}\{X \geq x\}]\ \mathrm{d}x = 2\,\mathbb{E}[\|K_n\|_F X] \\
&\leq 2\sqrt{\mathbb{E}\left[\|K_n\|_F^2\right] \mathbb{E}[X^2]}.
\end{aligned}$$

Rearranging yields the result. $\square$

Lemma III.15. *We have that* $\mathbb{E}\left[\|K_n\|_F^2\right] \leq nd^2$.

Proof. We can expand

$$\begin{aligned}
\mathbb{E}\left[\|K_n\|_F^2\right] &= \sum_{i=1}^n \mathbb{E}\left[\left\| dx_i x_i^\dagger - I_d \right\|_F^2 \right] \\
&\quad + 2 \sum_{1\leq i<j\leq n} \mathbb{E}\left[\left\langle dx_i x_i^\dagger - I_d, dx_j x_j^\dagger - I_d \right\rangle \right]. \qquad (10)
\end{aligned}$$

Since

$$\mathbb{E}\left[dx_j x_j^\dagger - I_d | \mathcal{F}_{j-1} \right] = \sum_{x_j} \omega_{x_j} (dx_j x_j^\dagger - I_d) = 0,$$

the second sum of (10) vanishes. For any unit vector x_i,

$$\left\| dx_i x_i^\dagger - I_d \right\|_F^2 = d^2 \|x_i\|^4 - 2d \|x_i\|^2 + d = d^2 - d.$$

Therefore $\mathbb{E}[\|K_n\|_F^2] \leq n(d^2 - d) \leq nd^2$. $\square$

Proof of Lemma III.8. Follows from Lemmas III.14 and III.15. $\square$

References

[1] Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, and Shengjun Pan. Competitive closeness testing. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 47–68. JMLR Workshop and Conference Proceedings, 2011.

[2] Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, Shengjun Pan, and Ananda Suresh. Competitive classification and closeness testing. In *Conference on Learning Theory*, pages 22–1. JMLR Workshop and Conference Proceedings, 2012.

[3] Costin Bădescu, Ryan O'Donnell, and John Wright. Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 503–514, 2019.

[4] Eric Blais, Clément L Canonne, and Tom Gur. Distribution testing lower bounds via reductions from communication complexity. *ACM Transactions on Computation Theory (TOCT)*, 11(2):1–37, 2019.

[5] Sebastien Bubeck, Sitan Chen, and Jerry Li. Entanglement is necessary for optimal quantum property testing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 692–703. IEEE, 2020.

[6] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential separations between learning with and without quantum memory. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 574–585. IEEE, 2022.

[7] Sitan Chen, Jerry Li, and Ryan O'Donnell. Toward instance-optimal state certification with incoherent measurements. *arXiv preprint arXiv:2102.13098*, 2021.

[8] Ilias Diakonikolas and Daniel M Kane. A new approach for testing properties of discrete distributions. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 685–694. IEEE, 2016.

[9] Leon Isserlis. On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. *Biometrika*, 12(1–2):134–139, 1918.

[10] Jiantao Jiao, Yanjun Han, and Tsachy Weissman. Minimax estimation of the l_1 distance. *IEEE Transactions on Information Theory*, 64(10):6672–6706, 2018.

[11] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[12] Ryan O'Donnell and John Wright. Quantum spectrum testing. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 529–538, 2015.

[13] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.

[14] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.

[15] John Wright. How to learn a quantum state. *Ph. D. dissertation*, 2016.

[16] Yihong Wu. Lecture notes on information-theoretic methods for high-dimensional statistics. http://www.stat.yale.edu/ yw562/teaching/it-stats.pdf, 2017.

[17] Bin Yu. Assouad, fano, and le cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.