On Estimating Operator Norm Distance, with Optimal Trace Distance Estimation When One State Is Pure

Yupan Liu
School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne, Switzerland

Qisheng Wang
School of Computer Science, Shanghai Jiao Tong University, China

Zhan Yu
Centre for Quantum Technologies, National University of Singapore, Singapore

Abstract

We investigate the computational complexity of estimating the operator norm distance $\mathrm{T}_\infty(\rho_0, \rho_1)$, defined via the operator norm $\|A\|_\infty \coloneqq \sigma_{\max}(A)$, where $\sigma_{\max}(A)$ is the largest singular value of A, given $\mathrm{poly}(n)$-size state-preparation circuits of n-qubit quantum states ρ_0 and ρ_1. We provide efficient quantum estimators for the operator norm distance whose complexity is *independent* of the rank (and thus the dimension) of the states:

- When one state is pure, we establish an optimal quantum estimator using $\Theta(1/\epsilon)$ queries to the state-preparation circuits. Consequently, for constant additive error, say $\epsilon = 1/5$, our estimator runs in $\mathrm{poly}(n)$ time. Since the operator norm distance $\mathrm{T}_\infty(|\psi\rangle\langle\psi|, \rho)$ is *exactly half* of the trace distance $\mathrm{T}(|\psi\rangle\langle\psi|, \rho)$, our result gives a rank-independent query complexity for estimating $\mathrm{T}_\infty(|\psi\rangle\langle\psi|, \rho)$ and $\mathrm{T}(|\psi\rangle\langle\psi|, \rho)$, whereas the approaches due to van Apeldoorn, Cornelissen, Gilyén, and Nannicini (SODA 2023) and Wang and Zhang (TIT 2024) have query complexity scaling at least linearly with $\mathrm{rank}(\rho)$, which can be $\exp(n)$ in general. In addition, our query complexity matches the optimal bound when *both* states are pure by Wang (TIT 2024).
- For general quantum states, we also provide a quantum estimator using $\widetilde{O}(1/\epsilon^{3/2})$ queries to the state-preparation circuits, which shows that the corresponding promise problem is BQP-complete and improves the QMA upper bound sketched by Liu and Wang (ESA 2025). Together with an $\Omega(1/\epsilon)$ quantum query complexity lower bound, this leaves only square-root room for improvement.

The key intuition behind our estimators is that, when one state is pure, the pure state $|\psi\rangle$ has overlap at least $1/2$ with the top unit eigenvector of $|\psi\rangle\langle\psi| - \rho$, reflecting a structural feature specific to the operator norm distance.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Quantum query complexity

Keywords and phrases Quantum state testing, Operator norm distance, Trace distance

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.71

Related Version *Full Version*: `https://arxiv.org/abs/2607.03905v1` [39]

Funding *Yupan Liu*: Supported by funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).
Qisheng Wang: Supported by startup funding from Shanghai Jiao Tong University.
Zhan Yu: Supported by the CQT Young Researcher Career Development Grant.

Acknowledgements ChatGPT was used interactively to explore possible approaches to the research problems addressed in main technical sections, identify relevant references, and proofread the manuscript, while all writing, including mathematical statements and reasoning, was completed by the authors.

34th Annual European Symposium on Algorithms (ESA 2026).
Editors: Philip Bille, Seth Pettie, and Sabine Storandt; Article No. 71; pp. 71:1–71:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Testing the closeness between quantum states is a central topic in quantum property testing [42], which studies efficient quantum algorithms for testing properties of quantum objects and extends classical distribution testing (see e.g., [11]) to the non-commutative setting. This problem is also closely connected to the verification of quantum devices, such as Q_0 and Q_1, which prepare the corresponding quantum states ρ_0 and ρ_1. In (tolerant) quantum state testing, one seeks to design efficient quantum algorithms that determine whether ρ_0 is $2/5$-far from or $1/5$-close to ρ_1 with respect to a chosen closeness measure, typically the trace distance, which is induced by the ℓ_1 norm.

Importantly, the trace distance is one of the most widely used measures of closeness because it plays a central role in several research areas: quantum hypothesis testing and state discrimination [30, 31], information-theoretically secure quantum cryptography [5, 48, 47], classical and statistical zero-knowledge [50, 62, 63], and certification and benchmarking [21, 4, 17]. Beyond the ℓ_1 norm distance, namely the total variation distance and the trace distance, (tolerant) closeness testing between probability distributions or quantum states with respect to the ℓ_α norm distance, for instance,

$$\mathrm{T}_\alpha(\rho_0,\rho_1) \coloneqq \frac{1}{2}\|\rho_0-\rho_1\|_\alpha, \quad \text{where } \|A\|_\alpha \coloneqq \mathrm{tr}(|A|^\alpha)^{1/\alpha},$$

has also been studied in [54, 36]. Interestingly, for all real $\alpha \geq 1$ for which $\mathrm{T}_\alpha(\rho_0,\rho_1)$ is a distance metric, a dichotomy emerges:

- In the ℓ_1 norm case, the sample complexity, namely the number of samples drawn from the distributions, scales polynomially with the support size of those distributions [53, 12, 32]. A similar phenomenon holds for quantum states: the number of state copies required scales polynomially with the dimension or rank of the states [4].
- For all real $\alpha > 1$, the sample complexity becomes *independent* of the support size [54]. Analogously, rank-independent query and sample complexity upper bounds for estimating the quantum ℓ_α distance, defined via the Schatten norm, were also obtained in [36].

However, the case $\alpha = \infty$ remains mysterious. On the classical side, the support-size-independent (and thus dimension-independent) sample complexity upper bounds in [54] already include this endpoint case. On the quantum side, however, the operator-norm-distance case is still listed as an open problem in [36]:

> What are the query and sample complexity upper bounds for estimating the operator norm distance $\mathrm{T}_\infty(\rho_0,\rho_1)$? Can one obtain such bounds in a *rank-independent* manner, similar to the classical setting, or is rank dependence *unavoidable*?

Why is it important to study the operator norm distance? While the trace distance ($\alpha = 1$) serves as the most important case in the quantum ℓ_α distance hierarchy, defined via the Schatten norm, the operator norm distance ($\alpha = \infty$), namely

$$\mathrm{T}_\infty(\rho_0,\rho_1) \coloneqq \frac{1}{2}\|\rho_0-\rho_1\|_\infty,$$

is the most relaxed notion of closeness precisely because it is the endpoint of the hierarchy, and thus may be viewed as another particularly important case. The operator norm distance captures the largest spectral deviation between two states, and therefore arises as a closeness measure in quantum state tomography [10, 64, 1], as well as an intermediate analytic tool that can sometimes be converted into trace-distance guarantees in [26]. This spectral viewpoint is also closely connected to compressed-sensing and low-rank recovery approaches, which allow high-dimensional states to be reconstructed from far fewer measurements [35, 25]. Beyond measuring the closeness between states, related operator-norm notions also play an important role in quantum process and unitary tomography [29] and in quantifying the precision of Hamiltonian simulation [7, 28].

In this work, we provide an affirmative answer to these questions on estimating $\mathrm{T}_\infty(\rho_0, \rho_1)$. We next present our main results and discuss their significance.

1.1 Main results

We start by stating our first main result, where $\mathrm{T}(\rho_0, \rho_1) \coloneqq \mathrm{T}_{\alpha=1}(\rho_0, \rho_1) = \frac{1}{2}\operatorname{tr}|\rho_0 - \rho_1|$:

▶ **Theorem 1** (Optimal quantum estimator for the trace and operator norm distances when one state is pure, informal)**.** *Given quantum query access to the state-preparation circuits of two n-qubit quantum states, denoted by ρ and $|\psi\rangle\langle\psi|$, without prior knowledge of which input circuit prepares the pure state, the quantum query complexity for estimating $\mathrm{T}(\rho, |\psi\rangle\langle\psi|)$ or $\mathrm{T}_\infty(\rho, |\psi\rangle\langle\psi|)$ to within additive error ϵ is $\Theta(1/\epsilon)$.*

Here, the lower bounds in Theorem 1 follow directly from the corresponding quantum query lower bounds for estimating $\mathrm{T}(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|)$ and $\mathrm{T}_\infty(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|)$ to within additive error ϵ in the case where both states are pure, established in [55, Theorem V.2] and [36, Theorem 22(1)], respectively.

▶ **Theorem 2** (Quantum estimator for the operator norm distance, informal)**.** *Given quantum query access to the state-preparation circuits of n-qubit quantum states ρ_0 and ρ_1, there exists an explicit quantum algorithm that estimates $\mathrm{T}_\infty(\rho_0, \rho_1)$ to within additive error ϵ with query complexity $\widetilde{O}(1/\epsilon^{3/2})$, where $\widetilde{O}(\cdot)$ suppresses poly-logarithmic factors.*

It is worth noting that the quantum query complexity lower bound $\Omega(1/\epsilon)$ in Theorem 1 also applies to Theorem 2, and thus our upper bound leaves only a *square-root* gap. Moreover, the quantum query upper bound in Theorem 2 is strictly better than the corresponding upper bounds known for every real $\alpha > 1$, and is in fact *super-quadratically* better except for the case of $\alpha = 2$. Indeed, for even integers $\alpha \geq 2$, the standard upper bound is $O(1/\epsilon^\alpha)$, following directly from [9, 18]; whereas for general real $\alpha > 1$ (including odd integers $\alpha \geq 3$), the best known upper bound is $O\big(1/\epsilon^{\alpha+1+\frac{1}{\alpha-1}}\big)$ as shown in [36]. These comparisons further support the view that the operator norm distance, as the endpoint of the quantum ℓ_α distance hierarchy, is its *easiest* case.

Prior to this work, there are two classes of known approaches that can be used to estimate the operator norm distance given state-preparation circuits:

- **General case:** The operator norm distance between two quantum states can be estimated directly by quantum state tomography in the operator norm distance, whose current best approach is due to [1, Theorem 45] with query complexity $\widetilde{O}(d/\epsilon)$, where $d = 2^n$ is the dimension of the quantum states and ϵ is the desired precision. In comparison, our estimator in Theorem 2 with query complexity $\widetilde{O}(1/\epsilon^{3/2})$ removes the dependence on d. In particular, when the state-preparation circuits are of size $\mathrm{poly}(n)$ (which is usually assumed when considering the computational complexity of quantum state testing, e.g., [62, 49]) and $\epsilon \geq 1/\mathrm{poly}(n)$ (e.g., $\epsilon = 0.1$ is a constant), our estimator in Theorem 2 runs in time $\mathrm{poly}(n)$, whereas the approach based on [1] runs in time $\exp(n)$, therefore achieving an *exponential* speedup.
- **When one state is pure:** In this case, the operator norm distance is equivalent to the trace distance (up to a constant factor). There are quantum estimators for trace distance given in [56, 57] and the current best query complexity is $\widetilde{O}(r/\epsilon^2)$ due to [57], where r

is the larger rank of *both* quantum states. When one state, say ρ_0, is pure, the trace distance estimator has query complexity $\widetilde{O}(\mathrm{rank}(\rho_1)/\epsilon^2)$. In comparison, our estimator in Theorem 1 with query complexity $\widetilde{O}(1/\epsilon)$ removes the dependence on the larger rank. In particular, when the state-preparation circuits are of size $\mathrm{poly}(n)$ and $\epsilon \geq 1/\mathrm{poly}(n)$, our estimator in Theorem 1 runs in time $\mathrm{poly}(n)$, whereas the approach based on [57] runs in time $\exp(n)$ (note that r can be as large as $d = 2^n$), therefore achieving an *exponential* speedup. When both states are pure, there are optimal quantum estimators for trace distance [55, 19] with query complexity $\Theta(1/\epsilon)$ and our estimator in Theorem 1 reproduces their results.

It should be noted that our results in Theorems 1 and 2 further lead to new quantum sample complexities of estimating the operator norm distance by the quantum sample-to-query lifting method recently developed in [58, 59, 52, 13]. Specifically, Theorem 1 implies a sample complexity of $\Theta(1/\epsilon^2)$ for estimating the trace distance when one state is pure (see Theorem 17), which (i) reproduces the result in [60] when both states are pure and (ii) matches the sample complexity lower bound in [55, Theorem B.2]; Theorem 2 implies a sample complexity of $\widetilde{O}(1/\epsilon^3)$ for estimating $\mathrm{T}_\infty(\rho_0, \rho_1)$ (see Theorem 20), which *exponentially* improves the prior tomography-based approach due to [14, Theorem V.4] with sample complexity $O(d/\epsilon^2)$.

Let QSD_∞ denote the promise problem of testing the closeness between quantum states with respect to the operator norm distance, namely, deciding whether $\mathrm{T}_\infty(\rho_0, \rho_1)$ is at least $2/5$ or at most $1/5$.[1] We also consider the following restricted variants $\textsc{OnePure}\mathrm{QSD}_\infty$ and, more generally, $\textsc{OnePure}\mathrm{QSD}_\alpha$ for the quantum ℓ_α distance, in which at least one of ρ_0 and ρ_1 is pure; and $\textsc{Pure}\mathrm{QSD}$, in which both states are pure.

When the state-preparation circuits have polynomial-size descriptions, the quantum algorithms in Theorems 1 and 2 run in time polynomial in n. Therefore, in the white-box input model, these query complexity upper bounds imply polynomial-time quantum algorithms, i.e., BQP upper bounds, for the corresponding promise problems, improving the prior QMA upper bound sketched in [36, Footnote 8]. Combining these upper bounds with the BQP-hardness results established for $\textsc{Pure}\mathrm{QSD}_\alpha$ when $1 \leq \alpha \leq \infty$ in [36, Theorem 21], and using the fact that the trace distance and the operator norm distance correspond to the two endpoints of the quantum ℓ_α distance hierarchy, we obtain the following corollaries:

▶ **Corollary 3.** QSD_∞ *is* BQP*-complete.*

▶ **Corollary 4.** *For every* $\alpha \in [1, \infty]$*,* $\textsc{OnePure}\mathrm{QSD}_\alpha$ *is* BQP*-complete.*

1.2 Proof techniques

We begin by explaining the intuition behind our quantum estimators for the operator norm distance. One natural approach is based on phase estimation and proceeds as follows:

(1) Apply an appropriate version of phase estimation,[2] constructed based on a unitary dilation of $\Delta := (\rho_0 - \rho_1)/2$,[3] to a state supported on the *top eigenspace* of Δ, which corresponds to the largest eigenvalue $\lambda_{\max}(\Delta)$.

---

[1] See Definition 8 for a formal definition.

[2] For example, an approach with query complexity $\widetilde{O}(\sqrt{d}/\epsilon)$ is known in [2, Lemma 50] based on phase estimation, equipped with appropriate versions of Hamiltonian simulation and quantum minimum finding, where d is the dimension of Δ.

[3] Given the state-preparation circuits of ρ and $|\psi\rangle\langle\psi|$, or oracle access to these circuits, one can efficiently construct a unitary dilation of Δ via the techniques of Linear Combination of Unitaries (LCU) [15, 6] and purified density matrix [40]. See Lemma 10 for an explicit construction.

(2) Repeat the same construction for $-\Delta$, and then take the maximum of the two resulting estimates, thereby obtaining an estimate of $\sigma_{\max}(\Delta) = \max\{\lambda_{\max}(\Delta), \lambda_{\max}(-\Delta)\}$.

However, in general, preparing such a state in the top eigenspace appears to require additional witness information. Consequently, this approach seems to yield only a QMA upper bound, as sketched in [36, Footnote 8]. To turn this idea into an actual efficient algorithm, we must therefore resolve two issues: (i) find a quantum state with *sufficiently large* overlap with the top eigenspace of Δ; and (ii) identify the appropriate version of phase estimation.

The input states provide a warm start when one state is pure. We first address the former issue in the setting where one state is pure, and our construction even extends naturally to estimating the trace distance. Previous optimal quantum query and sample algorithms for estimating the trace distance between two pure states $|\psi_0\rangle$ and $|\psi_1\rangle$, developed in [60, 55] (and implicitly in [19]), rely crucially on the following identity, which is implicit in [20] and can be viewed as capturing the *two-dimensional geometry of pure states*:

$$\mathrm{T}(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|)^2 + \mathrm{F}(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|)^2 = 1. \tag{1}$$

Since the square of the fidelity, $\mathrm{F}(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|)^2$, between pure states is simply the overlap $|\langle\psi_0|\psi_1\rangle|^2$ and the largest singular value of $|\psi_0\rangle\langle\psi_0| - |\psi_1\rangle\langle\psi_1|$ is

$$\sigma_{\max}(|\psi_0\rangle\langle\psi_0| - |\psi_1\rangle\langle\psi_1|) = \sqrt{1 - |\langle\psi_0|\psi_1\rangle|^2},$$

we obtain the following identity relating the trace distance and the operator norm distance:

$$\mathrm{T}(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|) = \sigma_{\max}(|\psi_0\rangle\langle\psi_0| - |\psi_1\rangle\langle\psi_1|) = 2 \cdot \mathrm{T}_\infty(|\psi_0\rangle\langle\psi_0|, |\psi_1\rangle\langle\psi_1|). \tag{2}$$

While Equation (1) is specific to the case in which both states are *pure*, the coincidence between the trace distance and the operator norm distance in Equation (2) extends to the more general setting in which only one state is pure, as proven in [37, Theorem 18(2)]:

$$\mathrm{T}(\rho, |\psi\rangle\langle\psi|) = 2 \cdot \mathrm{T}_\infty(\rho, |\psi\rangle\langle\psi|). \tag{3}$$

Therefore, in the single-pure-state setting, estimating the trace distance is *equivalent* to estimating the operator norm distance. We now show that there exists a quantum state with sufficiently large overlap with the top eigenspace of Δ, as established by our *single-pure-state warm-start lemma* (Lemma 12), which shows that $|\psi\rangle$ itself already provides a *warm start*. Indeed, assuming without loss of generality that $\rho_0 = |\psi\rangle\langle\psi|$ and $\rho_1 = \rho$, the state $|\psi\rangle$ has squared overlap at least $1/2$ with the eigenvector $|v_+\rangle$ corresponding to the *unique* positive eigenvalue $\lambda_{\max}(\Delta)$, and therefore the largest eigenvalue, of Δ:

$$|\langle v_+|\psi\rangle|^2 \geq \frac{1}{2} + \|\Delta\|_\infty > \frac{1}{2}. \tag{4}$$

Combining this warm start with an appropriate version of phase estimation, we obtain the $O(1/\epsilon)$ quantum query complexity upper bound, matching the pure-vs.-pure case in [55]. The underlying reason is essentially a *rank-one perturbation* phenomenon: since ρ is positive semidefinite, the restriction of Δ to the subspace orthogonal to $|\psi\rangle$ is negative semidefinite. Consequently, any eigenvector corresponding to a positive eigenvalue must have substantial overlap with $|\psi\rangle$, and in particular Δ has *exactly one* positive eigenvalue when $\Delta \neq 0$. This fact also explains the intuition behind Equations (2) and (3).

Furthermore, the proof of this warm-start lemma relies crucially on the fact that $\rho^2 \preceq \rho$, which becomes an equality when ρ is pure. Therefore, the bound in Equation (4) is tight, and is already saturated when both the states are pure. In this sense, this rank-one perturbation phenomenon can be viewed as a natural counterpart to Equation (1) when the assumption that both states are pure is relaxed to the case where only one state is pure.

From naïve phase estimation to maximum phase estimation. To resolve the latter issue, one might first consider the naïve approach of combining Kitaev's standard phase estimation [33] with Hamiltonian simulation, such as that in [7]: namely, one simulates $e^{it\Delta}$ and recovers an extremal eigenvalue of Δ from the corresponding eigenphase. While this approach is conceptually natural, it is not well suited to our setting. Indeed, estimating an eigenvalue to within additive error ϵ in this way requires resolving phases to precision $O(\epsilon)$, and thus using simulation time $t = \Theta(1/\epsilon)$. Implementing such a long-time evolution then incurs an additional factor proportional to t, yielding only a $\widetilde{O}(1/\epsilon^2)$-type quantum query upper bound, which is quadratically worse than the optimal bound in Theorem 1.

The right choice, instead, is to work with a unitary dilation of Δ and then apply maximum phase estimation [41] (see also Lemma 14) after qubitization [40] (see Lemma 13 for the specific version we need). Under the resulting monotone correspondence between eigenvalues and eigenphases, the largest eigenvalue of Δ is converted into the largest eigenphase of the associated unitary. Hence, the problem becomes that of estimating a maximum phase from a state having nontrivial overlap with the top eigenspace.

Consequently, in the single-pure-state setting, the warm-start lemma (Lemma 12) provides exactly such a state with constant overlap, and this is precisely what enables the maximum-phase-estimation primitive to recover $\lambda_{\max}(\Delta)$, and hence both $\mathrm{T}_\infty(\rho, |\psi\rangle\langle\psi|)$ and $\mathrm{T}(\rho, |\psi\rangle\langle\psi|)$, within additive error ϵ using $O(1/\epsilon)$ queries, as stated in Theorem 1.

The input states provide an eigenvalue-scaled overlap. While the same version of phase estimation can also be used to estimate the operator norm distance in the general setting, the key remaining issue is to identify a quantum state with sufficiently large overlap with the top eigenspace of $\Delta = (\rho_0 - \rho_1)/2$.

To this end, rather than using the input states ρ_0 and ρ_1 directly, we work with their purifications $|\psi_0\rangle$ and $|\psi_1\rangle$. Let $\Pi_{\max}^{(\Delta)}$ denote the projector onto the top eigenspace of Δ, which may have rank greater than 1 since the top eigenvalue of Δ need not be *unique*. Then, the overlap between $|\psi_0\rangle$ and this eigenspace, defined via the Euclidean vector norm $\|\cdot\|_2$, is

$$\big\|\big(I \otimes \Pi_{\max}^{(\Delta)}\big)|\psi_0\rangle\big\|_2^2 = \mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\rho_0\big). \tag{5}$$

It is also worth noting that Equation (5) already captures the warm-start phenomenon in the single-pure-state setting (cf. Lemma 12). Indeed, when $\rho_0 = |\psi\rangle\langle\psi|$ is pure, Equation (4) immediately yields the following *constant*-overlap bound:

$$\big\|\big(I \otimes \Pi_{\max}^{(\Delta)}\big)|\psi_0\rangle\big\|_2^2 = \mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\rho_0\big) \geq \frac{1}{2}. \tag{6}$$

When neither ρ_0 nor ρ_1 is pure, however, this constant-overlap argument from Equation (6) no longer applies. Instead, using the projection identity $\Pi_{\max}^{(\Delta)}\Delta\Pi_{\max}^{(\Delta)} = \lambda_{\max}(\Delta)I$, together with the overlap formula in Equation (5), we obtain the weaker but fully general lower bound established in our *top eigenspace overlap lemma* (Lemma 18):

$$\big\|\big(I \otimes \Pi_{\max}^{(\Delta)}\big)|\psi_0\rangle\big\|_2^2 \geq 2\,\mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\Delta\big) = 2\lambda_{\max}(\Delta)\,\mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\big) \geq 2\lambda_{\max}(\Delta). \tag{7}$$

Consequently, unlike in the single-pure-state setting, there is in general no constant-overlap warm start. Since the overlap guarantee in Equation (7) scales only with $2\lambda_{\max}(\Delta)$, the success probability of a single run can be as small as $O(\epsilon)$ at the target precision. To boost the success probability to a constant, say 2/3, we therefore combine this procedure with amplitude amplification and estimation [8] (as in maximum phase estimation [41]). This subtlety incurs an additional multiplicative factor of $\widetilde{O}(1/\sqrt{\epsilon})$, leading to the desired quantum query upper bound $\widetilde{O}(1/\epsilon^{3/2})$ in Theorem 2.

1.3 Discussion and open problems

A natural question is whether our warm-start lemma (see Lemma 12) can be extended to a more general setting. Here, we provide a simple example showing that such a warm start already fails for a pair of rank-two states, thereby revealing an immediate obstruction:

$$\rho_0^{(\eta)} \coloneqq (1-\eta)|0\rangle\langle 0| + \eta|1\rangle\langle 1| \quad \text{and} \quad \rho_1^{(\eta)} \coloneqq (1-\eta)|0\rangle\langle 0| + \eta|2\rangle\langle 2|.$$

A direct calculation shows that $\mathrm{T}_\infty\big(\rho_0^{(\eta)}, \rho_1^{(\eta)}\big) = \eta/2$ and $\mathrm{T}\big(\rho_0^{(\eta)}, \rho_1^{(\eta)}\big) = \eta$. Moreover, the extremal positive and negative eigenvalues of $\rho_0^{(\eta)} - \rho_1^{(\eta)}$ are η and $-\eta$, respectively, with the corresponding eigenvectors $|1\rangle$ and $|2\rangle$. Consequently, the overlaps between these eigenvectors and the input states are $\langle 1|\rho_0^{(\eta)}|1\rangle = \eta$ and $\langle 2|\rho_1^{(\eta)}|2\rangle = \eta$. Since η can be chosen arbitrarily small, there is *no constant* $c > 0$ such that every pair of rank-two states admits an extremal eigenvector of $\rho_0 - \rho_1$ with overlap at least c with one of the input states, or even with a linear combination of them. This motivates our first open problem:

(i) Can one characterize the largest family of mixed states for which one can still prepare a state that has constant overlap with the top eigenspace of $\rho_0 - \rho_1$?

A second direction concerns optimality. Our current bounds in Theorem 2 for the general setting still leave a square-root gap, leading to the following question:

(ii) Can the quantum query and sample complexity of estimating the operator norm distance be made optimal? More generally, can analogous optimal results be achieved for estimating the quantum ℓ_α distance for real $\alpha > 1$ [36]?

1.4 Related works

Quantum state tomography has been studied with respect to trace distance, fidelity (Bures distance), Frobenius distance (Hilbert–Schmidt distance), and operator norm distance in the literature [27, 44, 14, 1, 46, 51, 45]. The closeness testing of quantum states, also known as the quantum state certification, was studied with respect to trace distance, fidelity, ℓ_2 distance, and ℓ_3 distance in [4, 22]. The tolerant closeness testing, which is equivalent to closeness estimation, was studied for trace distance, fidelity, and ℓ_α distance in [16, 61, 23, 56, 57, 34, 38].

2 Preliminaries

We assume that the reader is familiar with the basics of quantum computation and quantum information. For background on these subjects, we refer to the standard textbook [43].

Throughout the paper, we use the following notation. First, for any positive integer n, let $[n] \coloneqq \{1, 2, \ldots, n\}$. Second, $\widetilde{O}(f)$ stands for $O(f \operatorname{polylog}(f))$, whereas $\widetilde{\Omega}(f)$ stands for $\Omega(f/\operatorname{polylog}(f))$. Finally, for a matrix A, the Schatten α-norm is defined by

$$\|A\|_\alpha \coloneqq \big(\mathrm{tr}(|A|^\alpha)\big)^{1/\alpha} = \Big(\mathrm{tr}\Big(\big(A^\dagger A\big)^{\alpha/2}\Big)\Big)^{1/\alpha}.$$

2.1 Closeness measures for quantum states

We begin by defining the trace distance:

▶ **Definition 5** (Trace distance). *Let ρ_0 and ρ_1 be two quantum states on the same Hilbert space. The trace distance between ρ_0 and ρ_1 is defined by*

$$\mathrm{T}(\rho_0, \rho_1) \coloneqq \frac{1}{2}\|\rho_0 - \rho_1\|_1 = \frac{1}{2}\mathrm{tr}(|\rho_0 - \rho_1|).$$

Next, we define the operator norm distance, a quantity that has also been used in [26, 1], where the corresponding definition is only implicit, while an explicit formal definition is given as a special case of [36, Definition 2.3]:

▶ **Definition 6** (Operator norm distance). *Let ρ_0 and ρ_1 be two quantum states on the same Hilbert space. The operator norm distance between ρ_0 and ρ_1 is defined by*

$$\mathrm{T}_\infty(\rho_0, \rho_1) \coloneqq \frac{1}{2}\|\rho_0 - \rho_1\|_\infty = \frac{1}{2}\max\{\lambda_{\max}(\rho_0 - \rho_1), \lambda_{\max}(\rho_1 - \rho_0)\}.$$

Here, $\lambda_{\max}(A)$ denotes the largest eigenvalue of the matrix A.

Notably, the operator norm distance arises as the endpoint case ($\alpha = \infty$) of the quantum ℓ_α distance [36], which generalizes both the trace distance ($\alpha = 1$) and the Hilbert-Schmidt distance ($\alpha = 2$) via the Schatten norm. In addition, both the trace distance and the operator norm distance are distance metrics, and hence satisfy the triangle inequality, as stated in [36, Lemma 2.4] (see also [3, Proposition 1.16]).

Moreover, the inequalities that relate the operator norm distance to the trace distance were established in [36], as stated in Lemma 7. Interestingly, when one of the states is pure, the trace distance is *exactly* twice the operator norm distance.

▶ **Lemma 7** (T_∞ vs. T, adapted from [36, Theorem 4.2(2)]). *Let ρ_0 and ρ_1 be quantum states on the same Hilbert space. Then, the following bounds hold:*

$$2 \cdot \mathrm{T}_\infty(\rho_0, \rho_1) \leq \mathrm{T}(\rho_0, \rho_1) \leq 2\min\{\mathrm{rank}(\rho_0), \mathrm{rank}(\rho_1)\} \cdot \mathrm{T}_\infty(\rho_0, \rho_1).$$

Furthermore, when $\rho_0 = |\psi\rangle\langle\psi|$ and $\rho_1 = \rho$, the inequality becomes an identity:

$$2 \cdot \mathrm{T}_\infty(|\psi\rangle\langle\psi|, \rho) = \mathrm{T}(|\psi\rangle\langle\psi|, \rho).$$

2.2 Closeness testing of quantum states via state-preparation circuits

We first define closeness testing for quantum states with respect to the quantum ℓ_α distance, which generalizes the Quantum State Distinguishability Problem (QSD) introduced in [62, Section 3.3], as stated in Definition 8. We denote this promise problem by $\mathrm{QSD}_\alpha[a, b]$ and also consider the following two restricted variants:

▶ **Definition 8** (Quantum State Distinguishability Problem with Schatten α-norm, QSD_α, adapted from [36, Definition 17]). *Let Q_0 and Q_1 be quantum circuits acting on m qubits ("input length") and having n specified output qubits ("output length"), where $m(n)$ is polynomial in n. Let ρ_i denote the quantum state obtained by running Q_i on the state $|0\rangle^{\otimes m}$ and tracing out the non-output qubits. Let $a(n)$ and $b(n)$ be efficiently computable functions. The task is to decide whether:*

- Yes: *The pair of quantum circuits (Q_0, Q_1) satisfies $\mathrm{T}_\alpha(\rho_0, \rho_1) \geq a(n)$;*
- No: *The pair of quantum circuits (Q_0, Q_1) satisfies $\mathrm{T}_\alpha(\rho_0, \rho_1) \leq b(n)$.*

Furthermore, we consider two restricted versions:

(1) PureQSD: *Both ρ_0 and ρ_1 are pure states, and the input length may be larger than the output length.*

(2) OnePureQSD: *At least one of ρ_0 and ρ_1 is a pure state, without prior knowledge of which input circuit prepares the pure state.*

In this work, we consider the *purified quantum access input model*, as originally introduced in [62], in both white-box and black-box settings:

- **White-box model**: The input of the problem QSD_α consists of descriptions of quantum circuits Q_0 and Q_1 that are *polynomial-size*. In other words, each circuit description consists of a polynomial number of one- and two-qubit gates.
- **Black-box model**: Instead of having explicit circuit descriptions, one is given only query access to Q_b, $Q_b^\dagger$, controlled-Q_b, and controlled-$Q_b^\dagger$, for $b \in \{0,1\}$. For convenience, we use "a query to Q_b" to refer to a query to Q_b, $Q_b^\dagger$, controlled-Q_b, and controlled-$Q_b^\dagger$.

In addition to query complexity, which is defined in the black-box model, *sample complexity* refers to the number of state copies of ρ_0 and ρ_1 required to test the closeness between these states with respect to the chosen closeness measure.

2.3 Quantum algorithmic toolkit

We will use the notion of unitary dilation, which may be viewed as an *error-free* version of block-encoding with normalization factor 1 [40, 24].

▶ **Definition 9** (Unitary dilation). *Let A be an n-qubit operator on subsystem A with $\|A\|_\infty \leq 1$. An $(n+m)$-qubit unitary operator U on subsystem AB is said to be a unitary dilation of A with m ancillary qubits, if $\Pi U \Pi = A_{\mathsf{A}} \otimes |0\rangle\langle 0|_{\mathsf{B}}$, where $\Pi = I_{\mathsf{A}} \otimes |0\rangle\langle 0|_{\mathsf{B}}$.*

▶ **Lemma 10** (Unitary dilation from state-preparation circuits). *Given quantum unitary oracles Q_0 and Q_1 that prepare the purifications of two unknown quantum states ρ_0 and ρ_1, respectively, we can implement a unitary dilation U_Δ of $\Delta := (\rho_0 - \rho_1)/2$ using 1 query to each of controlled-Q_0, controlled-Q_1, and their inverses. Moreover, if each of Q_0 and Q_1 uses a ancillary qubits, then U_Δ uses $(2a+1)$ ancillary qubits.*

The proof of Lemma 10 can be found in the full version [39, Section 2.3].

Quantum sample-to-query lifting. We mention the quantum sample-to-query lifting method, which enables us to convert quantum query complexity to quantum sample complexity. This lifting method, used implicitly before in [23, 57], was pointed out in [58] and later refined in [59, 52]. Here, we use the version in [52].

▶ **Lemma 11** (Quantum sample-to-query lifting, adapted from [52, Theorem 1.5] and [13, Corollary 1.4]). *For any quantum algorithm $\mathcal{A}$ that uses q queries to the state-preparation circuits of some unknown quantum states, there is another quantum algorithm $\mathcal{A}'$ using $O(q^2)$ samples of those unknown quantum states such that the output distribution of $\mathcal{A}'$ is 0.99-close to that of $\mathcal{A}$.*

3 Efficient quantum algorithms for estimating trace and operator norm distances when one state is pure

This section focuses on the setting in which one of the two input states is pure. We first prove a structural warm-start property specific to this case in Section 3.1, then combine it with qubitization (Lemma 13) and maximum phase estimation (Lemma 14) to derive a quantum query algorithm in Section 3.2. Finally, the sample upper bound follows immediately in Section 3.3.

3.1 The input states provide a warm start when one state is pure

We begin with the structural statement that makes the single-pure-state setting easier:

▶ **Lemma 12** (Single-pure-state warm-start lemma). *Let ρ be a quantum state, and let $|\psi\rangle\langle\psi|$ be a pure state on the same Hilbert space. Define the Hermitian operator $\Delta \coloneqq (|\psi\rangle\langle\psi| - \rho)/2$. Then either $\Delta = 0$, or Δ has the unique positive eigenvalue equal to $\|\Delta\|_\infty$. If $|v_+\rangle$ is the unique unit eigenvector corresponding to this eigenvalue, then*

$$|\langle v_+|\psi\rangle|^2 \geq \frac{1}{2} + \|\Delta\|_\infty > \frac{1}{2}.$$

Moreover, equality holds when ρ is also pure and $\Delta \neq 0$.

Proof. For every pure state $|\psi^\perp\rangle$ orthogonal to $|\psi\rangle$, we have

$$\langle\psi^\perp|\Delta|\psi^\perp\rangle = -\langle\psi^\perp|\rho|\psi^\perp\rangle/2 \leq 0, \tag{8}$$

since ρ is positive semidefinite. Hence, the restriction of Δ to $\Pi_\perp$, namely

$$\Pi_\perp \Delta \Pi_\perp, \quad \text{where } \Pi_\perp \coloneqq I - |\psi\rangle\langle\psi|,$$

is negative semidefinite. It follows that the positive eigenspace of Δ has dimension at most 1: indeed, if $|\phi_1\rangle$ and $|\phi_2\rangle$ were two linearly independent positive-eigenvalue eigenvectors, then some nonzero linear combination of them would be orthogonal to $|\psi\rangle$, contradicting Equation (8).

If Δ had no positive eigenvalue, then Δ is negative semidefinite, namely $\Delta \preceq 0$. Since $\mathrm{tr}(\Delta) = 0$, this would force $\Delta = 0$. Therefore, if $\Delta \neq 0$, then Δ has the unique positive eigenvalue, say $\tau > 0$. Since Δ is traceless, its negative eigenvalues sum to $-\tau$, and hence each of them is at least $-\tau$. Because Δ is Hermitian, we obtain:

$$\tau = \|\Delta\|_\infty. \tag{9}$$

Now let $|v_+\rangle$ be the unique unit eigenvector corresponding to the eigenvalue τ. We have:

$$\Delta|v_+\rangle = \tau|v_+\rangle \quad \text{and} \quad 2\cdot\Delta|v_+\rangle = \langle\psi|v_+\rangle|\psi\rangle - \rho|v_+\rangle. \tag{10}$$

Since ρ is positive semidefinite and $\mathrm{tr}(\rho) = 1$, we have $0 \preceq \rho \preceq I$, and hence $\rho^2 \preceq \rho$. Therefore, Equation (10) implies the following:

$$\begin{aligned} \|\rho|v_+\rangle\|_2^2 &= \langle v_+|\rho^2|v_+\rangle && (11a)\\ &\leq \langle v_+|\rho|v_+\rangle = |\langle\psi|v_+\rangle|^2 - 2\tau. && (11b)\end{aligned}$$

On the other hand, a direct calculation shows that

$$\begin{aligned} \|\rho|v_+\rangle\|_2^2 &= \|\langle\psi|v_+\rangle|\psi\rangle - 2\tau|v_+\rangle\|_2^2 && (12a)\\ &= |\langle\psi|v_+\rangle|^2 + 4\tau^2 - 4\tau|\langle\psi|v_+\rangle|^2. && (12b)\end{aligned}$$

Consequently, combining Equations (9), (11), and (12) gives

$$2\|\Delta\|_\infty + 1 = 2\tau + 1 \leq 2|\langle\psi|v_+\rangle|^2, \tag{13}$$

which establishes the desired bound after rearranging the terms.

To see the equality condition, note that if $\rho = |\phi\rangle\langle\phi|$ is pure, then the identity $\rho^2 = |\phi\rangle\langle\phi| = \rho$, and it follows that both Equations (11) and (13) become identities, as desired. ◀

3.2 Quantum query algorithm when one state is pure

Beyond presenting the main result of this section, we need the following lemmas. We begin with the following version of qubitization that converts eigenvalues of the Hermitian operator into eigenphases of an associated unitary dilation in a form suitable for phase estimation:

▶ **Lemma 13** (Qubitization for unitary dilation)**.** *Let U be an $(n+a)$-qubit unitary dilation of an n-qubit Hermitian operator H. Denote $\Pi = |0\rangle\langle 0|^{\otimes a} \otimes I_n$, $R = 2\Pi - I_{n+a}$, and $W = URU^\dagger R$, where I_n denotes the n-qubit identity operator. Let $|v_j\rangle$ be a unit eigenvector of H with eigenvalue $\lambda_j \in [-1,1]$, i.e., $H|v_j\rangle = \lambda_j|v_j\rangle$ for each $1 \le j \le 2^n$. Denote $|\psi_j\rangle = |0\rangle^{\otimes a} \otimes |v_j\rangle$. If $|\lambda_j| < 1$, then define*

$$|\psi_j^\perp\rangle \coloneqq \frac{U|\psi_j\rangle - \lambda|\psi_j\rangle}{\sqrt{1-\lambda_j^2}}, \quad |\phi_j^\pm\rangle \coloneqq \frac{1}{\sqrt{2}}\big(|\psi_j\rangle \mp i|\psi_j^\perp\rangle\big),$$

and the following holds:

1. $\Pi|\psi_j\rangle = |\psi_j\rangle$, $\Pi|\psi_j^\perp\rangle = 0$, $U|\psi_j\rangle = \lambda_j|\psi_j\rangle + \sqrt{1-\lambda_j^2}|\psi_j^\perp\rangle$,
2. *Let $\theta_j = \arccos(\lambda_j) \in (0,\pi)$, then*

$$\begin{aligned} W|\psi_j\rangle &= \cos(2\theta_j)|\psi_j\rangle + \sin(2\theta_j)|\psi_j^\perp\rangle, \\ W|\psi_j^\perp\rangle &= -\sin(2\theta_j)|\psi_j\rangle + \cos(2\theta_j)|\psi_j^\perp\rangle, \\ W|\phi_j^\pm\rangle &= e^{\pm i2\theta}|\phi_j^\pm\rangle. \end{aligned}$$

If $|\lambda_j| = 1$, then $U|\psi_j\rangle = \lambda|\psi_j\rangle$ and $W|\psi_j\rangle = |\psi_j\rangle$.

Moreover, let $\mathcal{P} \coloneqq \mathrm{ran}(\Pi) = \mathrm{span}\{\, |\psi_j\rangle : 1 \le j \le 2^n \,\}$, $\mathcal{Q} \coloneqq U\mathcal{P} = \mathrm{ran}(U\Pi U^\dagger)$, and $\mathcal{M} = (\mathcal{P}+\mathcal{Q})^\perp$, where $\mathrm{ran}(A) \coloneqq \{\, A|x\rangle : |x\rangle \in \mathrm{dom}(A) \,\}$, then every nonzero vector $|\varphi\rangle \in \mathcal{M}$ is an eigenvector of W with eigenvalue 1, i.e., $W|\varphi\rangle = |\varphi\rangle$.

It is noteworthy that, compared to the qubitization in [40], our version of qubitization requires *no* ancillary qubits and produces a different rotation angle. The detailed proof of Lemma 13 can be found in the full version [39, Section 3.2].

▶ **Lemma 14** (Quantum maximum phase estimation [41, Lemma 4.3])**.** *Let $\gamma \in (0,1)$. Suppose we have a d-dimensional unitary U with unknown eigenvectors $|\phi_1\rangle, \dots, |\phi_d\rangle$ and associated eigenphases $\theta_1, \dots, \theta_d \in [0, 2\pi - 2\epsilon)$. Suppose we also have a state preparation unitary S such that*

$$S|0\rangle = \sum_{j=1}^{d} \alpha_j |\phi_j\rangle|\varphi_j\rangle,$$

where $\sum_{j:\theta_j=\theta_{\max}} |\alpha_j|^2 \ge \gamma^2$ and $|\varphi_j\rangle$ are arbitrary (normalized) states. Then there exists a quantum algorithm $\mathsf{MaxPhaseEst}(U,S,\epsilon,\gamma)$ *that uses at most $O(1/\gamma)$ queries to S, and $O(\log(1/\gamma)/\gamma\epsilon)$ queries to U, and with probability at least 0.99 outputs $\tilde{\theta}$ such that $|\tilde{\theta} - \theta_{\max}| \le \epsilon$. Moreover, if $\sum_{j:\theta_j=\theta_{\max}} |\alpha_j|^2 < \gamma^2$, then with probability at least 0.99, $|\tilde{\theta} - \theta_j| \le \epsilon$ for some j.*

▶ **Lemma 15.** *For a pure state $|\psi\rangle\langle\psi|$ and a state ρ, it holds that*

$$\mathrm{T}_\infty(|\psi\rangle\langle\psi|, \rho) = \lambda_{\max}\left(\frac{|\psi\rangle\langle\psi| - \rho}{2}\right).$$

More generally, we have

$$\lambda_{\max}\left(\frac{|\psi\rangle\langle\psi| - \rho}{2}\right) \ge \lambda_{\max}\left(\frac{\rho - |\psi\rangle\langle\psi|}{2}\right).$$

Proof. Let $\Delta := (|\psi\rangle\langle\psi| - \rho)/2$. If $\Delta = 0$, the lemma holds trivially. If $\Delta \neq 0$, by Lemma 12, Δ has a unique positive eigenvalue equal to $\|\Delta\|_\infty$. The sum of all negative eigenvalues of Δ (with multiplicities) is equal to $-\|\Delta\|_\infty$ due to $\mathrm{tr}(\Delta) = 0$. Therefore, no negative eigenvalue is less than $-\|\Delta\|_\infty$, which implies that $\lambda_{\max}(\Delta) \geq \lambda_{\max}(-\Delta)$. ◀

3.2.1 The quantum query algorithm

Now we establish the main result of this section:

▶ **Theorem 16.** *Given quantum query access to the state-preparation circuits Q_0 and Q_1 of quantum states ρ_0 and ρ_1 with the promise that one of the two states is pure, there exists an explicit quantum algorithm that estimates $\mathrm{T}(\rho_0, \rho_1)$ and $\mathrm{T}_\infty(\rho_0, \rho_1)$ to within additive error ϵ, using $O(1/\epsilon)$ queries to Q_0 and Q_1.*

We present the formal algorithm corresponding to Theorem 16 in Algorithm 1. The full proof of Theorem 16, including the correctness and performance analysis of Algorithm 1, is provided in the full version [39, Section 3.3.1].

Algorithm 1 Operator norm distance estimation when one state is pure.

Input: $(n+a)$-qubit unitary oracles Q_0 and Q_1 that prepare purifications of n-qubit quantum states ρ_0 and ρ_1, precision $\epsilon \in (0, 0.1)$.
Output: An estimate of $\mathrm{T}_\infty(\rho_0, \rho_1)$.
1: Denote $\Delta := (\rho_0 - \rho_1)/2$ and $R := 2\Pi - I_{n+2a+2}$, where $\Pi := |0\rangle\langle 0|^{\otimes(2a+2)} \otimes I_n$.
2: **for** $k \in \{0, 1\}$ **do**
3: Construct an $(n + 2a + 2)$-qubit unitary dilation U_{A_k} of $A_k := (I - (-1)^k\Delta)/2$.
4: Construct a qubitization unitary $W_{A_k} := U_{A_k} R (U_{A_k})^\dagger R$.
5: $\tilde{\theta}^{(k)}_{\max} \leftarrow \mathsf{MaxPhaseEst}(-W_{A_k}, I_{2a+2} \otimes Q_k, \epsilon/2, \gamma)$, where $\gamma = 1/2$.
6: **end for**
7: $\tilde{\theta}_{\max} \leftarrow \max\{\tilde{\theta}^{(0)}_{\max}, \tilde{\theta}^{(1)}_{\max}\}$.
8: **return** $\tilde{\lambda}_{\max} \leftarrow 1 - 2\cos((\pi - \tilde{\theta}_{\max})/2)$.

3.3 Quantum sample algorithm when one state is pure

▶ **Theorem 17.** *When one of the quantum states ρ_0 and ρ_1 is pure, there is a quantum algorithm that estimates $\mathrm{T}(\rho_0, \rho_1)$ and $\mathrm{T}_\infty(\rho_0, \rho_1)$ to within additive error ϵ, using $O(1/\epsilon^2)$ samples of ρ_0 and ρ_1.*

Proof. This is obtained immediately by Lemma 11 using the query complexity $O(1/\epsilon)$ given in Theorem 16, which gives the sample complexity $O(1/\epsilon^2)$. ◀

4 Efficient quantum algorithms for estimating operator norm distance

We now turn to the general setting, in which neither input state is assumed to be pure. The same overall strategy still applies, but the constant-overlap warm start from Section 3 is no longer available and therefore must be replaced by a weaker overlap guarantee. In particular, by combining the overlap lemma in Section 4.1 with the qubitization framework (Lemma 13), we obtain general quantum query and sample algorithms for estimating the operator norm distance, which are presented respectively in Sections 4.2 and 4.3.

4.1 The input states provide an eigenvalue-scaled overlap

To handle arbitrary mixed states, we next prove a substitute for the warm-start lemma by working with purifications of the input states. The resulting overlap bound is no longer constant, but it scales with the top eigenvalue and is sufficient for the general algorithm.

▶ **Lemma 18** (Top eigenspace overlap lemma). *Let $|\psi_0\rangle$ and $|\psi_1\rangle$ be $(n+a)$-qubit purifications of n-qubit quantum states ρ_0 and ρ_1, respectively. Define $\Delta := (\rho_0 - \rho_1)/2$ and let $\Pi_{\max}^{(\Delta)}$ be the projector onto the top eigenspace of Δ, which corresponds to its maximum eigenvalue $\lambda_{\max}(\Delta)$. Then, it holds that:*

$$\big\|\big(\Pi_{\max}^{(\Delta)} \otimes I_a\big)|\psi_0\rangle\big\|_2^2 \geq 2\lambda_{\max}(\Delta).$$

Moreover, if $\lambda_{\max}(\Delta) > 0$, then equality holds if and only if

$$\operatorname{rank}\big(\Pi_{\max}^{(\Delta)}\big) = 1 \quad \textit{and} \quad \Pi_{\max}^{(\Delta)}\rho_1\Pi_{\max}^{(\Delta)} = 0.$$

In particular, if $\rho_0 = |\phi\rangle\langle\phi|$ is pure, then equality holds if and only if ρ_1 and $|\phi\rangle$ are orthogonal.

Proof. Since $|\psi_0\rangle$ and $|\psi_1\rangle$ are purifications of ρ_0 and ρ_1, respectively, for any n-qubit operator M acting on the system register, it follows that:

$$\langle\psi_i|(M \otimes I_a)|\psi_i\rangle = \operatorname{tr}(M\rho_i), \qquad i \in \{0, 1\}.$$

Therefore, we have

$$\begin{aligned}\big\|\big(\Pi_{\max}^{(\Delta)} \otimes I_a\big)|\psi_0\rangle\big\|_2^2 &= \langle\psi_0|\big(\Pi_{\max}^{(\Delta)} \otimes I_a\big)|\psi_0\rangle &(14a)\\ &= \operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\rho_0\big). &(14b)\end{aligned}$$

On the other hand, since $\Pi_{\max}^{(\Delta)}$ projects onto the top eigenspace of Δ, which corresponds to the eigenvalue $\lambda_{\max}$, we have

$$\Pi_{\max}^{(\Delta)}\Delta = \Delta\Pi_{\max}^{(\Delta)} = \lambda_{\max}(\Delta)\Pi_{\max}^{(\Delta)}.$$

Taking the trace gives

$$\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\Delta\big) = \lambda_{\max}(\Delta)\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\big). \tag{15}$$

Since $\rho_1 \succeq 0$ and $\Pi_{\max}^{(\Delta)}$ is a projector, we have $\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\rho_1\big) \geq 0$, and hence

$$\begin{aligned}\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\rho_0\big) &\geq \operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\rho_0\big) - \operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\rho_1\big) &(16a)\\ &= \operatorname{tr}\big(\Pi_{\max}^{(\Delta)}(\rho_0 - \rho_1)\big) &(16b)\\ &= 2\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\Delta\big) &(16c)\\ &= 2\lambda_{\max}(\Delta)\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\big) &(16d)\\ &\geq 2\lambda_{\max}(\Delta). &(16e)\end{aligned}$$

Here, the third line follows by substituting $\Delta = (\rho_0 - \rho_1)/2$ into Equation (15), and the last line uses the fact that $\operatorname{tr}\big(\Pi_{\max}^{(\Delta)}\big) = \operatorname{rank}\big(\Pi_{\max}^{(\Delta)}\big) \geq 1$, because $\Pi_{\max}^{(\Delta)} \neq 0$.

Combining Equation (14) and Equation (16) yields the desired overlap bound:

$$\big\|\big(\Pi_{\max}^{(\Delta)} \otimes I_a\big)|\psi_0\rangle\big\|_2^2 \geq 2\lambda_{\max}(\Delta).$$

To see the equality condition in the case $\lambda_{\max}(\Delta) > 0$, we examine when both inequalities in Equation (16) are tight, which corresponds to the following conditions

$$\mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\rho_1\big) = 0 \quad \text{and} \quad \mathrm{tr}\big(\Pi_{\max}^{(\Delta)}\big) = 1. \tag{17}$$

Therefore, the first condition means that ρ_1 has no support on the top eigenspace of Δ, equivalently $\Pi_{\max}^{(\Delta)}\rho_1\Pi_{\max}^{(\Delta)} = 0$; while the second condition means that $\mathrm{rank}\big(\Pi_{\max}^{(\Delta)}\big) = 1$.

In particular, when ρ_0 is pure, Equation (17) simplifies to the condition that $\rho_0 = |\psi\rangle\langle\psi|$ is orthogonal to ρ_1, more precisely, $\langle\psi|\rho_1|\psi\rangle = 0$. ◀

4.2 Quantum query algorithm for estimating operator norm distance

We now present the main result of this section:

▶ **Theorem 19.** *Given quantum query access to the state-preparation circuits Q_0 and Q_1 of quantum states ρ_0 and ρ_1, respectively, there exists an explicit quantum algorithm that estimates $\mathrm{T}_\infty(\rho_0, \rho_1)$ to within additive error ϵ, using $O\big(\log(1/\epsilon)/\epsilon^{3/2}\big)$ queries to Q_0 and Q_1.*

We now present the formal algorithm corresponding to Theorem 19 in Algorithm 2. The detailed proof of Theorem 19, including the correctness and performance analysis of Algorithm 2, is provided in the full version [39, Section 4.2].

Algorithm 2 Operator norm distance estimation.

Input: $(n+a)$-qubit unitary oracles Q_0 and Q_1 that prepare purifications of n-qubit quantum states ρ_0 and ρ_1, precision $\epsilon \in (0, 0.1)$.
Output: An estimate of $\mathrm{T}_\infty(\rho_0, \rho_1)$.
1: Denote $\Delta := (\rho_0 - \rho_1)/2$ and $R := 2\Pi - I_{n+2a+2}$, where $\Pi := |0\rangle\langle 0|^{\otimes(2a+2)} \otimes I_n$.
2: **for** $k \in \{0, 1\}$ **do**
3: Construct an $(n + 2a + 2)$-qubit unitary dilation U_{A_k} of $A_k := (I - (-1)^k\Delta)/2$.
4: Construct a qubitization unitary $W_{A_k} := U_{A_k}R(U_{A_k})^\dagger R$.
5: $\tilde{\theta}_{\max}^{(k)} \leftarrow \mathsf{MaxPhaseEst}(-W_{A_k}, I_{2a+2} \otimes Q_k, \epsilon, \gamma)$, where $\gamma = \sqrt{\epsilon}$.
6: $\tilde{\lambda}_{\max}^{(k)} \leftarrow \max\{0, 1 - 2\cos((\pi - \tilde{\theta}_{\max}^{(k)})/2)\}$.
7: **end for**
8: **return** $\max\{\tilde{\lambda}_{\max}^{(0)}, \tilde{\lambda}_{\max}^{(1)}\}$.

4.3 Quantum sample algorithm for estimating operator norm distance

▶ **Theorem 20.** *There is a quantum algorithm that estimates $\mathrm{T}_\infty(\rho_0, \rho_1)$ to within additive error ϵ, using $O(\log^2(1/\epsilon)/\epsilon^3)$ samples of ρ_0 and ρ_1.*

Proof. This bound is obtained immediately by Lemma 11 using the query complexity $O(\log(1/\epsilon)/\epsilon^{3/2})$ given in Theorem 19, which gives the sample complexity $O(\log^2(1/\epsilon)/\epsilon^3)$. ◀

References

1 Joran van Apeldoorn, Arjan Cornelissen, András Gilyén, and Giacomo Nannicini. Quantum tomography using state-preparation unitaries. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1265–1318, 2023. doi:10.1137/1.9781611977554.ch47.

2 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: better upper and lower bounds. *Quantum*, 4:230, 2020. doi:10.22331/q-2020-02-14-230.

3 Guillaume Aubrun and Stanisław J. Szarek. *Alice and Bob Meet Banach: The Interface of Asymptotic Geometric Analysis and Quantum Information Theory*, volume 223 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2017. doi:10.1090/surv/223.

4 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. doi:10.1145/3313276.3316344.

5 Michael Ben-Or, Michał Horodecki, Debbie W. Leung, Dominic Mayers, and Jonathan Oppenheim. The universal composable security of quantum key distribution. In *Theory of Cryptography, Second Theory of Cryptography Conference, TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings*, pages 386–406. Springer, 2005. doi:10.1007/978-3-540-30576-7_21.

6 Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114(9):090502, 2015. doi:10.1103/PhysRevLett.114.090502.

7 Dominic W. Berry, Andrew M. Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *Proceedings of the 56th IEEE Annual Symposium on Foundations of Computer Science*, pages 792–809, 2015. doi:10.1109/FOCS.2015.54.

8 Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In Samuel J. Lomonaco, Jr. and Howard E. Brandt, editors, *Quantum Computation and Information*, volume 305 of *Contemporary Mathematics*, pages 53–74. AMS, 2002. doi:10.1090/conm/305/05215.

9 Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. doi:10.1103/PhysRevLett.87.167902.

10 Tony Cai, Donggyu Kim, Yazhen Wang, Ming Yuan, and Harrison H Zhou. Optimal large-scale quantum state tomography with pauli measurements. *The Annals of Statistics*, 44(2):682–712, 2016. doi:10.1214/15-AOS1382.

11 Clément L. Canonne. A survey on distribution testing: your data is big. but is it blue? In *Theory of Computing Library*, number 9 in Graduate Surveys, pages 1–100. University of Chicago, 2020. ECCC:TR15-063. doi:10.4086/toc.gs.2020.009.

12 Siu-On Chan, Ilias Diakonikolas, Paul Valiant, and Gregory Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1203. SIAM, 2014. doi:10.1137/1.9781611973402.88.

13 Kean Chen, Qisheng Wang, and Zhicheng Zhang. A list of complexity bounds for property testing by quantum sample-to-query lifting, 2025. ArXiv preprint. doi:10.48550/arXiv.2512.01971.

14 Sitan Chen, Brice Huang, Jerry Li, Allen Liu, and Mark Sellke. When does adaptivity help for quantum state learning? In *Proceedings of the 64th Annual Symposium on Foundations of Computer Science*, pages 391–404. IEEE, 2023. doi:10.1109/FOCS57990.2023.00029.

15 Andrew M. Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Information and Computation*, 12(11–12):901–924, 2012. doi:10.26421/QIC12.11-12-1.

16 Patrick J. Coles, Marco Cerezo, and Lukasz Cincio. Strong bound between trace distance and Hilbert-Schmidt distance for low-rank states. *Physical Review A*, 100(2):022103, 2019. doi:10.1103/physreva.100.022103.

17 Jens Eisert, Dominik Hangleiter, Nathan Walk, Ingo Roth, Damian Markham, Rhea Parekh, Ulysse Chabaud, and Elham Kashefi. Quantum certification and benchmarking. *Nature Reviews Physics*, 2(7):382–390, 2020. doi:10.1038/s42254-020-0186-4.

18 Artur K. Ekert, Carolina Moura Alves, Daniel K. L. Oi, Michał Horodecki, Paweł Horodecki, and Leong Chuan Kwek. Direct estimations of linear and nonlinear functionals of a quantum state. *Physical Review Letters*, 88(21):217901, 2002. doi:10.1103/physrevlett.88.217901.

19 Wang Fang and Qisheng Wang. Optimal quantum algorithm for estimating fidelity to a pure state. In *Proceedings of the 33rd Annual European Symposium on Algorithms (ESA)*, volume 351 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ESA.2025.4.

20 Christopher A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999. doi:10.1109/18.761271.

21 Alexei Gilchrist, Nathan K. Langford, and Michael A. Nielsen. Distance measures to compare real and ideal quantum processes. *Physics Review A*, 71:062310, 2005. doi:10.1103/PhysRevA.71.062310.

22 András Gilyén and Tongyang Li. Distributional property testing in a quantum world. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference*, pages 25:1–25:19, 2020. doi:10.4230/LIPIcs.ITCS.2020.25.

23 András Gilyén and Alexander Poremba. Improved quantum algorithms for fidelity estimation. ArXiv e-prints, 2022. arXiv:2203.15993.

24 András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019. doi:10.1145/3313276.3316366.

25 David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011. doi:10.1109/TIT.2011.2104999.

26 Mădălin Guţă, Jonas Kahn, Richard Kueng, and Joel A. Tropp. Fast state tomography with optimal error bounds. *Journal of Physics A: Mathematical and Theoretical*, 53(20):204001, 2020. doi:10.1088/1751-8121/ab8111.

27 Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017. Preliminary version in *STOC 2016*. doi:10.1109/TIT.2017.2719044.

28 Jeongwan Haah, Matthew B. Hastings, Robin Kothari, and Guang Hao Low. Quantum algorithm for simulating real time evolution of lattice Hamiltonians. *SIAM Journal on Computing*, 52(6):S18–250, 2023. Preliminary version in *FOCS 2018*. doi:10.1137/18M1231511.

29 Jeongwan Haah, Robin Kothari, Ryan O'Donnell, and Ewin Tang. Query-optimal estimation of unitary channels in diamond distance. In *Proceedings of the 64th IEEE Annual Symposium on Foundations of Computer Science (FOCS 2023)*, pages 363–390. IEEE, 2023. doi:10.1109/FOCS57990.2023.00028.

30 Carl W. Helstrom. Detection theory and quantum mechanics. *Information and Control*, 10(3):254–291, 1967. doi:10.1016/S0019-9958(67)90302-6.

31 Alexander S. Holevo. Statistical decision theory for quantum systems. *Journal of Multivariate Analysis*, 3(4):337–394, 1973. doi:10.1016/0047-259X(73)90028-6.

32 Jiantao Jiao, Yanjun Han, and Tsachy Weissman. Minimax estimation of the l_1 distance. *IEEE Transactions on Information Theory*, 64(10):6672–6706, 2018. doi:10.1109/TIT.2018.2846245.

33 A. Yu. Kitaev. Quantum measurements and the Abelian stabilizer problem. ArXiv e-prints, 1995. arXiv:quant-ph/9511026.

34 François Le Gall, Yupan Liu, and Qisheng Wang. Space-bounded quantum state testing via space-efficient quantum singular value transformation. To appear in *computational complexity*, 2026. doi:10.1007/s00037-025-00284-5.

35 Yi-Kai Liu. Universal low-rank matrix recovery from Pauli measurements. In *Proceedings of the 25th Annual Conference on Neural Information Processing Systems (NIPS 2011)*, pages 1638–1646, 2011. arXiv:1103.2816.

36 Yupan Liu and Qisheng Wang. On estimating the quantum ℓ_α distance. In *Proceedings of the 33rd Annual European Symposium on Algorithms (ESA 2025)*, volume 351 of *LIPIcs*, pages 105:1–105:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. `doi:10.4230/LIPIcs.ESA.2025.105`.

37 Yupan Liu and Qisheng Wang. On estimating the trace of quantum state powers. *IEEE Transactions on Information Theory*, pages 1–1, 2026. Preliminary version in *SODA 2025*. `doi:10.1109/TIT.2026.3683891`.

38 Yupan Liu and Qisheng Wang. Optimal fidelity estimation when one state is pure via algorithmic uhlmann transform. ArXiv preprint, 2026. `arXiv:2608.10674`.

39 Yupan Liu, Qisheng Wang, and Zhan Yu. On estimating operator norm distance, with optimal trace distance estimation when one state is pure. ArXiv e-prints, 2026. `arXiv:2607.03905v1`.

40 Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. `doi:10.22331/q-2019-07-12-163`.

41 Nikhil S. Mande and Ronald de Wolf. Tight Bounds for Quantum Phase Estimation and Related Problems. *Quantum*, 10:2140, 2026. Preliminary version in *ESA 2023*. `doi:10.22331/q-2026-06-15-2140`.

42 Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. In *Theory of Computing Library*, number 7 in Graduate Surveys, pages 1–81. University of Chicago, 2016. `doi:10.4086/toc.gs.2016.007`.

43 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010. `doi:10.1017/CBO9780511976667`.

44 Ryan O'Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2016)*, pages 899–912. ACM, 2016. `doi:10.1145/2897518.2897544`.

45 Angelos Pelecanos, Jack Spilecki, Ewin Tang, and John Wright. Mixed state tomography reduces to pure state tomography. To appear in *the Proceedings of the 67th Annual IEEE Symposium on Foundations of Computer Science*, 2026. `arXiv:2511.15806`.

46 Angelos Pelecanos, Jack Spilecki, and John Wright. The debiased Keyl's algorithm: A new unbiased estimator for full state tomography. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 1266–1277, 2026. `doi:10.1145/3798129.3800837`.

47 Christopher Portmann and Renato Renner. Security in quantum cryptography. *Reviews of Modern Physics*, 94:025008, 2022. `doi:10.1103/RevModPhys.94.025008`.

48 Renato Renner and Robert König. Universally composable privacy amplification against quantum adversaries. In *Theory of Cryptography, Second Theory of Cryptography Conference, TCC 2005, Cambridge, MA, USA, February 10-12, 2005, Proceedings*, pages 407–425. Springer, 2005. `doi:10.1007/978-3-540-30576-7_22`.

49 Soorya Rethinasamy, Rochisha Agarwal, Kunal Sharma, and Mark M. Wilde. Estimating distinguishability measures on quantum computers. *Physical Review A*, 108(1):012409, 2023. `doi:10.1103/PhysRevA.108.012409`.

50 Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):196–249, 2003. Preliminary version in *FOCS 1997*. `ECCC:TR00-084`. `doi:10.1145/636865.636868`.

51 Thilo Scharnhorst, Jack Spilecki, and John Wright. Optimal lower bounds for quantum state tomography. ArXiv preprints, 2025. `doi:10.48550/arXiv.2510.07699`.

52 Ewin Tang, John Wright, and Mark Zhandry. Quantum oracles, weak and strong. To appear in *the Proceedings of the 67th Annual IEEE Symposium on Foundations of Computer Science*, 2026. `arXiv:2510.07622`.

53 Gregory Valiant and Paul Valiant. Estimating the unseen: Improved estimators for entropy and other properties. *Journal of the ACM*, 64(6):37:1–37:41, 2017. Preliminary version in *NIPS 2013*. `doi:10.1145/3125643`.

54 Bo Waggoner. ℓ_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 347–356, 2015. doi:10.1145/2688073.2688095.

55 Qisheng Wang. Optimal trace distance and fidelity estimations for pure quantum states. *IEEE Transactions on Information Theory*, 70(12):8791–8805, 2024. doi:10.1109/TIT.2024.3447915.

56 Qisheng Wang, Ji Guan, Junyi Liu, Zhicheng Zhang, and Mingsheng Ying. New quantum algorithms for computing quantum entropies and distances. *IEEE Transactions on Information Theory*, 70(8):5653–5680, 2024. doi:10.1109/TIT.2024.3399014.

57 Qisheng Wang and Zhicheng Zhang. Fast quantum algorithms for trace distance estimation. *IEEE Transactions on Information Theory*, 70(4):2720–2733, 2024. doi:10.1109/TIT.2023.3321121.

58 Qisheng Wang and Zhicheng Zhang. Quantum lower bounds by sample-to-query lifting. *SIAM Journal on Computing*, 54(5):1294–1334, 2025. doi:10.1137/24M1638616.

59 Qisheng Wang and Zhicheng Zhang. Time-efficient quantum entropy estimator via samplizer. *IEEE Transactions on Information Theory*, 71(12):9569–9599, 2025. Preliminary version in *ESA 2024*. doi:10.1109/TIT.2025.3576137.

60 Qisheng Wang and Zhicheng Zhang. Sample-optimal quantum estimators for pure-state trace distance and fidelity via samplizer. In *Proceedings of the 53rd International Colloquium on Automata, Languages, and Programming*, pages 154:1–154:21, 2026. doi:10.4230/LIPIcs.ICALP.2026.154.

61 Qisheng Wang, Zhicheng Zhang, Kean Chen, Ji Guan, Wang Fang, Junyi Liu, and Mingsheng Ying. Quantum algorithm for fidelity estimation. *IEEE Transactions on Information Theory*, 69(1):273–282, 2023. doi:10.1109/TIT.2022.3203985.

62 John Watrous. Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468. IEEE, 2002. doi:10.1109/SFCS.2002.1181970.

63 John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009. Preliminary version in *STOC 2006*. doi:10.1137/060670997.

64 Dong Xia and Vladimir Koltchinskii. Estimation of low rank density matrices: Bounds in schatten norms and other distances. *Electronic Journal of Statistics*, 10:2717–2745, 2016. doi:10.1214/16-EJS1192.