On estimating operator norm distance, with optimal trace distance estimation when one state is pure

Yupan Liu[1] Qisheng Wang[2] Zhan Yu[3]

[1]École Polytechnique Fédérale de Lausanne

[2]Shanghai Jiao Tong University

[3]National University of Singapore

Available on arXiv:2607.03905.

European Symposium on Algorithms (ESA 2026), L'Aquila

What is quantum state testing

Basic ingredients in quantum computation:

- **Quantum states.** An n-qubit quantum state $\rho \in \mathbb{C}^{N \times N}$, where $N = 2^n$, is an N-dimensional positive semi-definite (PSD) matrix such that $\mathrm{Tr}(\rho) = 1$.
  - **Pure states.** An n-qubit state is *pure* if $\rho = |\psi\rangle\langle\psi|$, where $|\psi\rangle \in \mathbb{C}^N$ and $\langle\psi|\psi\rangle = 1$.
  - **Purification**. For any n-qubit quantum state ρ on $\mathcal{H}_\mathsf{A}$, there exists a $2n$-qubit pure state $|\psi\rangle$ on $\mathcal{H}_A \otimes \mathcal{H}_B$ such that $\mathrm{Tr}_\mathsf{B}(|\psi\rangle\langle\psi|) = \rho$.
- **Quantum gate**. Elementary quantum gates G_i (from some universal gateset) are unitary matrices acting on one or two qubits, e.g., $G_i \in \{\mathsf{CNOT}, \mathsf{Had}, \mathsf{T}\}$:

$$|0\rangle^{\otimes n} \xrightarrow{G_1} G_1 |0\rangle^{\otimes n} \xrightarrow{G_2} G_2 G_1 |0\rangle^{\otimes n} \rightarrow \cdots$$

- **Measurement**. Projective measurement in the computational basis $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$:

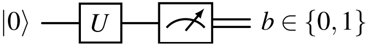

Task: Closeness testing of quantum states

Given two state-preparation circuits Q_0 and Q_1 ("quantum devices") that prepare (purifications of) n-qubit quantum states $\rho_0 \in \mathbb{C}^{N \times N}$ and $\rho_1 \in \mathbb{C}^{N \times N}$, respectively. Decide whether $\mathrm{dist}(\rho_0, \rho_1) \geq a(n)$ or $\mathrm{dist}(\rho_0, \rho_1) \leq b(n)$.

What is quantum state testing (Cont.)

Task: Closeness testing of quantum states

Given two state-preparation circuits Q_0 and Q_1 ("quantum devices") that prepare (the purification of) n-qubit quantum states $\rho_0 \in \mathbb{C}^{N\times N}$ and $\rho_1 \in \mathbb{C}^{N\times N}$, respectively. Decide whether $\mathrm{dist}(\rho_0,\rho_1) \geq a(n)$ or $\mathrm{dist}(\rho_0,\rho_1) \leq b(n)$.

- Quantum devices Q_b for $b \in \{0,1\}$ can be given either as a query oracle (*black-box model*) or a sequence of $\mathrm{poly}(n)$ elementary quantum gates (*white-box model*).
- The most canonical choices of closeness measures are:
  - **Trace distance** $\mathrm{T}(\rho_0,\rho_1) = \frac{1}{2}\mathrm{Tr}(|\rho_0-\rho_1|)$.
  - **Total variation distance** $\mathrm{TV}(D_0,D_1) = \frac{1}{2}\sum_x |D_0(x)-D_1(x)|$.

Typical goal. Minimize the "complexity" of ρ_b (or its corresponding Q_b) for $b \in \{0,1\}$:

Type of query access	Complexity measure
Black-box model	Query complexity (the number of queries)
White-box model	Complexity class

Generalizing the closeness measures via the Schatten α-norm

Generalization. Define the quantum ℓ_α distance via the Schatten norm:

$$\mathrm{T}_\alpha(\rho_0,\rho_1) := \frac{1}{2}\|\rho_0-\rho_1\|_\alpha = \frac{1}{2}\mathrm{Tr}(|\rho_0-\rho_1|^\alpha)^{1/\alpha}.$$

Trace distance ($\alpha = 1$). The closeness testing problem in this case is *hard*, with complexity *linearly* depending on the rank r of the quantum states:

- The quantum query complexity for estimating $\mathrm{T}(\rho_0,\rho_1)$ to within additive error ε is $\widetilde{O}(r/\varepsilon^2)$ [Wang–Zhang'23] and $\widetilde{\Omega}(r/\varepsilon)$ [Wang'26, Tang–Wright–Zhandry'25].
- The promise problem Quantum State Distinguishability (QSD$[a,b]$) is QSZK-complete* [Watrous'02, Watrous'05], and it is widely believed that BQP $\subsetneq$ QSZK.
  - The QSZK containment holds only in the *polarizing regime* $a(n)^2 - b(n) > 1/O(\log n)$.
- When *both states are pure*, there exist explicit optimal quantum estimators for $\mathrm{T}(|\psi_0\rangle\langle\psi_0|,|\psi_1\rangle\langle\psi_1|)$ with query complexity $\Theta(1/\varepsilon)$ [Wang'24, Fang–Wang'25].

Generalizing the closeness measures via the Schatten α-norm (Cont.)

Even $\alpha \in \{2,4,\cdots\}$. The closeness testing problem in this case is *easy*, with complexity *independent* of the rank r of the quantum states:

- The query complexity for estimating $\mathrm{Tr}(\rho_0\rho_1)$ to within additive error ε is $O(1/\varepsilon)$ via the SWAP test [Buhrman–Cleve–Watrous–de Wolf'01], and thus the promise problem is in BQP.
  - This directly applies to the case $\alpha = 2$, since $\mathrm{Tr}\big((\rho_0-\rho_1)^2\big) = \mathrm{Tr}(\rho_0^2) + \mathrm{Tr}(\rho_1^2) - 2\mathrm{Tr}(\rho_0\rho_1)$.
- Similar techniques [Ekert–Alves–Oi–Horodecki–Horodecki–Kwek'02] can estimate $\mathrm{Tr}(\rho_1\rho_2\cdots\rho_k)$ for integers $k > 1$, and solve the case of even integers α.

Real $\alpha > 1$. Recent advances establish the *rank-independent* bounds:

- For *real* orders $\alpha > 1$, the query complexity for estimating $\mathrm{T}_\alpha(\rho_0,\rho_1)$ to within additive error ε is $O(1/\varepsilon^{\alpha+1+\frac{1}{\alpha-1}}) = \mathrm{poly}(1/\varepsilon)$ [**L.**–Wang'25].
  - Consequently, when $a(n) - b(n) \geq 1/\mathrm{poly}(n)$, the promise problem $\mathrm{QSD}_\alpha[a,b]$ is in BQP.
- For *fixed* orders $1 \leq \alpha \leq \infty$, PureQSD$_\alpha$ is BQP-hard, even with *constant* precision. As a corollary, for *real* orders $\alpha > 1$, QSD$_\alpha$ is BQP-complete.

Generalizing the closeness measures via the Schatten α-norm (Cont.[2])

What about the complexity of estimating the classical ℓ_α distance?

Similarly, define the classical ℓ_α distance $\mathrm{TV}_\alpha(D_0, D_1) := \frac{1}{2}\left(\sum_x |D_0(x) - D_1(x)|^\alpha\right)^{1/\alpha}$.

For *fixed* orders $\alpha \in (1, \infty]$, classical ℓ_α distance admits estimation to additive error ε with $\mathrm{poly}(1/\varepsilon)$ dependence that is *independent* of the support size of distributions D_0 and D_1, and improves as α increases [Waggoner'14].

Open problems in (**L.** & Wang, ESA 2025)

While $\mathrm{T}_\alpha(\rho_0, \rho_1)$ and its powered version $\Lambda_\alpha(\rho_0, \rho_1) := \frac{1}{2}\mathrm{Tr}(|\rho_0 - \rho_1|^\alpha)$ are *almost interchangeable* for real orders $\alpha > 1$, their behavior differs significantly when $\alpha = \infty$:

- $\mathrm{T}_\infty(\rho_0, \rho_1)$ corresponds to the largest singular value $\sigma_{\max}\left(\frac{\rho_0 - \rho_1}{2}\right)$.
- $\Lambda_\infty(\rho_0, \rho_1) \in \left\{0, \frac{1}{2}, 1\right\}$ for any quantum states ρ_0, ρ_1, and it is nonzero if and only if the states are orthogonal with at least one of them being pure.

Question: What is the computational complexity of estimating $\mathrm{T}_\infty(\rho_0, \rho_1)$?

1. Quantum state testing with respect to the operator norm distance
2. Main results
3. Proof techniques
4. Open problems

Main results

Theorem 1 (Optimal quantum estimators when one state is pure).
Given quantum query access to the state-preparation circuits of two states ρ and $|\psi\rangle\langle\psi|$, *without prior knowledge of which 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 $\Theta(1/\varepsilon)$.

- Key observation [L.–Wang'25]: T and T_∞ coincide *when one state is pure*:

$$\mathrm{T}(\rho, |\psi\rangle\langle\psi|) = \frac{1}{2}\mathrm{Tr}(|\rho - |\psi\rangle\langle\psi||) = \sigma_{\max}(\rho - |\psi\rangle\langle\psi|) = 2\,\mathrm{T}_\infty(\rho, |\psi\rangle\langle\psi|).$$

- Remove *the rank dependence* in the approaches due to (i) pure-state tomography [van Apeldoorn–Cornelissen–Gilyén–Nannicini'22], and (ii) general estimator for T [Wang–Zhang'23].

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

- Consequently, the promise problem QSD_∞ is BQP-complete.

A natural approach for estimating T_∞ and two challenges

Recall that $\mathrm{T}_\infty(\rho_0,\rho_1) = \sigma_{\max}(\Delta) = \max\{\lambda_{\max}(\Delta), \lambda_{\max}(-\Delta)\}$, where $\Delta := (\rho_0 - \rho_1)/2$.

Natural approach for estimating T_∞ via phase estimation

1. Apply an appropriate version of *phase estimation*, constructed using a unitary dilation of Δ, to a state supported on the *top eigenspace* of Δ (corresponding to $\lambda_{\max}(\Delta)$).
2. Repeat Step 1 for $-\Delta$, then take the maximum of the two resulting estimates of $\lambda_{\max}(\Delta)$ and $\lambda_{\max}(-\Delta)$, thereby obtaining an estimate of $\sigma_{\max}(\Delta)$.

As noted in [L.–Wang'25], preparing a top-eigenspace state appears to *require a quantum witness* in general, so this direct approach suggests only a QMA upper bound. ☹

To turn this approach into an efficient algorithm, there are two challenges:

1. Find a quantum state with *sufficiently large* overlap with the top eigenspace of Δ.
2. Identify the *appropriate* version of phase estimation.

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

We resolve Challenge ❶ by the following lemma:

Lemma 3 (Single-pure-state warm-start lemma). Assume $\rho_0 = |\psi\rangle\langle\psi|$ and $\rho_1 = \rho$. Define $\Delta := (|\psi\rangle\langle\psi| - \rho)/2$. If $\Delta \neq 0$, then Δ has a unique positive eigenvalue $\lambda_{\max}(\Delta) = \|\Delta\|_\infty$. Let $|v_+\rangle$ be the unique unit eigenvector corresponding to $\lambda_{\max}$. Then,

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

Proof intuition underlying Lemma 3:

- Uniqueness from a *rank-one perturbation* viewpoint. On the subspace orthogonal to $|\psi\rangle$, Δ is negative semidefinite; hence Δ has at most *one* positive eigenvalue:

$$\Delta\,|v_+\rangle = \|\Delta\|_\infty\,|v_+\rangle \quad \text{and} \quad 2\,\Delta\,|v_+\rangle = |\psi\rangle\,\langle\psi|v_+\rangle - \rho\,|v_+\rangle.$$

- One can see the overlap bound $2\|\Delta\|_\infty + 1 \leq 2|\langle\psi|v_+\rangle|^2$ by

$$\|\rho\,|v_+\rangle\|_2^2 = \langle v_+|\,\rho^2\,|v_+\rangle \leq \langle v_+|\,\rho\,|v_+\rangle = |\langle\psi|v_+\rangle|^2 - 2\|\Delta\|_\infty,$$

$$\|\rho\,|v_+\rangle\|_2^2 = \|\langle\psi|v_+\rangle\,|\psi\rangle - 2\|\Delta\|_\infty\,|v_+\rangle\|_2^2 = |\langle\psi|v_+\rangle|^2 + 4\|\Delta\|_\infty^2 - 4\|\Delta\|_\infty|\langle\psi|v_+\rangle|^2.$$

Achieving efficient quantum estimators for T_∞

We next move to resolve Challenge ❷:

- The naïve approach is to combine *phase estimation* [Kitaev'95] and *Hamiltonian simulation* $e^{it\Delta}$ [Kitaev'95, Berry–Childs–Kothari'15], leading to a *quadratically worse* bound.
  - ◇ More specifically, recovering $\lambda_{\max}(\Delta)$ to within precision ε requires simulation time $t = \Theta(1/\varepsilon)$, yielding only an $\widetilde{O}(1/\varepsilon^2)$ quantum query upper bound.

Achieving efficient quantum estimators for T_∞

We next move to resolve Challenge ❷:

- The naïve approach is to combine *phase estimation* [Kitaev'95] and *Hamiltonian simulation* $e^{it\Delta}$ [Kitaev'95, Berry–Childs–Kothari'15], leading to a *quadratically worse* bound.
- The actual approach instead combines the *maximum phase estimation* [Mande–de Wolf'23] with *qubitization on a unitary dilation* W_Δ *of* Δ [Low–Chuang'19]; $\lambda_{\max}(\Delta)$ is thus converted into the largest eigenphase $\theta_{\max}(W_\Delta)$, leading to the desired $O(1/\varepsilon)$ bound.

To achieve the quantum estimator for T_∞ with bound $\widetilde{O}(1/\varepsilon^{3/2})$, we generalize Lemma 3:

Lemma 4 (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. Let $\Pi_{\max}^{(\Delta)}$ be the projector onto the top eigenspace of Δ that corresponds to $\lambda_{\max}(\Delta)$. Then,

$$\left\|\left(\Pi_{\max}^{(\Delta)} \otimes I_a\right)|\psi_0\rangle\right\|_2^2 = \mathrm{Tr}\left(\Pi_{\max}^{(\Delta)}\rho_0\right) \geq 2\lambda_{\max}(\Delta).$$

As the *vanishing* overlap can be *as small as* $O(\varepsilon)$, combining with *amplitude amplification & estimation* [Brassard–Høyer–Mosca–Tapp'00] incurs an extra $\widetilde{O}(1/\sqrt{\varepsilon})$ multiplicative factor.

Conclusions and open problems

Take-home messages

1. When *one input state is pure*, both T and T_∞ admit an *optimal* $\Theta(1/\varepsilon)$-query quantum estimator.
2. For general states, T_∞ admits a *rank-independent* $\widetilde{O}(1/\varepsilon^{3/2})$-query quantum estimator.
   - Consequently, QSD_∞ is BQP-complete.
   - Meanwhile, the quantum query complexity lower bound is $\Omega(1/\varepsilon)$ [**L.**–Wang'25].

Discussion and open problems

- Characterize the largest family of mixed states for which one can prepare a state with constant overlap with the top eigenspace of $\rho_0 - \rho_1$.

  A constant warm start already fails for the *rank-two* states:

  $$\rho_0^{(\eta)} = (1-\eta)\,|0\rangle\langle 0| + \eta\,|1\rangle\langle 1|, \qquad \rho_1^{(\eta)} = (1-\eta)\,|0\rangle\langle 0| + \eta\,|2\rangle\langle 2|.$$

  Their common component cancels, leaving the extremal directions $|1\rangle$ and $|2\rangle$, but $\langle 1|\,\rho_0^{(\eta)}\,|1\rangle = \langle 2|\,\rho_1^{(\eta)}\,|2\rangle = \eta$. Thus, as $\eta \to 0$, both warm-start overlaps vanish.

- Can the quantum query complexity for estimating T_∞ be made optimal?

Thanks!