# Quantum Composite Hypothesis Testing with Small Error

**Chenghua Liu**
Institute of Software, Chinese Academy of Sciences
University of Chinese Academy of Sciences
Beijing, China
liuch.russell@gmail.com

**Qisheng Wang**
School of Computer Science
Shanghai Jiao Tong University
Shanghai, China
QishengWang1994@gmail.com

## Abstract

We present a method of quantum composite hypothesis testing with small error, which enables us to establish quantum lower bounds in nonparametric statistics and high-dimensional functional estimation. We extend the phase estimation lower-bound argument of Mande and de Wolf (ESA 2023) to composite hypotheses for one-parameter unitary families. As an application, we settle the quantum complexities for a series of problems in property testing and functional estimation with small error:

- For $\ell_2$-closeness testing, we show that the approach due to Luo et al. (*IEEE Trans. Inf. Theory* 2024) is optimal.
- For Tsallis entropy estimation, where $q = 2$ refers to the Gini impurity, we show that the approaches due to Buhrman et al. (*Phys. Rev. Lett.* 2001) and Ekert et al. (*Phys. Rev. Lett.* 2002) are optimal for integer $q \geq 2$ and the approach due to Chen et al. (ICALP 2026) is near-optimal for real $q \geq 1.5$.
- For pure-state Uhlmann fidelity and trace distance estimation, we show that the approach due to Wang (*IEEE Trans. Inf. Theory* 2024) is optimal.

## 1 Introduction

Hypothesis testing is a fundamental problem in statistics. A testing problem is specified by two mutually exclusive hypotheses: a null hypothesis $H_0$ and an alternative hypothesis $H_1$. Given observations generated under one of the two hypotheses, the goal is to decide which hypothesis is true with a prescribed error probability. Characterizing the complexity of such testing problems is central to information theory, statistics, and learning theory. One of the classical lower-bound techniques is Le Cam's two-point method [20, 29], where both hypotheses are point hypotheses. This idea has been extended to composite hypotheses, for example through the method of two fuzzy hypotheses [25]. Such techniques have played an important role in proving minimax lower bounds for nonparametric statistics and functional estimation; see, for example, [25, 19].

Quantum computing has brought significant speedups to property testing problems through its ability to concentrate probability mass more efficiently (e.g., via quantum amplitude estimation [6]). A growing line of work has developed quantum algorithms for distributional property testing and related functional-estimation tasks [7, 10, 4, 17, 21, 28, 22]. However, the quantum complexity of many fundamental testing problems remains open, especially in the small error regime. A notable feature of the quantum setting is that the dependence on the error probability $\delta$ varies substantially across problems. For example, (i) logarithmic dependence: quantum phase estimation to within additive error $\varepsilon$ with small error $\delta$ requires query complexity $\Theta(\log(1/\delta)/\varepsilon)$ [23]; (ii) sub-logarithmic dependence: quantum search over $n$ elements with small error $\delta$ has query complexity $\Theta(\min\{n, \sqrt{n\log(1/\delta)}\})$ [8]; and (iii) no dependence: there are promised exact cases in which the error parameter disappears:

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

in amplitude amplification, if the initial marked fraction is exactly $1/4$, then a single Grover iterate succeeds with probability 1 [18, 6]. This behavior is quite different from the classical case, and should be compared with the classical i.i.d. setting. For testing two fixed distributions $P$ and $Q$, the optimal error exponent is governed by the Chernoff information [12]; in the local regime, this leads to the familiar $\Theta(\log(1/\delta)/d_{\mathrm{H}}^2(P,Q))$ scaling, as also reflected in the Hellinger distance lower bound framework of Bar-Yossef et al. [2, 1]. At the same time, even in classical distribution testing, small error bounds are not always obtained by simply repeating a constant error tester. For example, identity testing over an $n$-element domain has constant error sample complexity $\Theta(\sqrt{n}/\varepsilon^2)$ [24, 26], whereas its optimal high-confidence complexity is $\Theta((\sqrt{n\log(1/\delta)}+\log(1/\delta))/\varepsilon^2)$ [13]. The latter improves over black-box amplification of a constant error tester, which would require $O(\sqrt{n}\log(1/\delta)/\varepsilon^2)$ samples [13]. More generally, recent work on high-probability distribution testing treats the dependence on $\delta$ as part of the main complexity characterization [14]. Thus, in the quantum setting, where logarithmic, sub-logarithmic, and even absent $\delta$-dependence can all occur, it is natural to study small error quantum complexities with $\delta$ treated as an intrinsic complexity parameter.

In this paper, we present a framework for quantum composite hypothesis testing with small error, which allows us to systematically establish *tight* quantum lower bounds for a series of property testing and functional estimation tasks, including $\ell_2$-closeness testing, Tsallis entropy estimation, and pure-state fidelity and trace distance estimation.

### 1.1 Main results

The central problem we study in this paper is the composite hypothesis testing for trigonometrically linear unitaries. A matrix $M(\theta)$ with parameter $\theta$ is said to be trigonometrically linear in $\theta$, if $M(\theta) = Ae^{-i\theta} + Be^{i\theta} + C$ where $A$, $B$, and $C$ are matrices independent of $\theta$. This includes the oracle families used below, which consist of a phase gate composed with fixed unitaries. Each query increases the trigonometric degree of the amplitudes by at most one. The Remez inequality bounds the degree needed for an acceptance probability to be at least $1-\delta$ on $\Theta_0$ and at most $\delta$ on $\Theta_1$.

**Definition 1.1** (Composite hypothesis testing for trigonometrically linear unitaries)**.** *Let $\Theta := [-\pi,\pi)$, and let $\Theta_0,\Theta_1 \subseteq \Theta$ be two nonempty disjoint sets. Given a quantum unitary oracle $U(\theta)$ that is trigonometrically linear in an unknown parameter $\theta \in \Theta$, the task $\mathrm{HT}(\Theta_0,\Theta_1)$ is, under the promise that $\theta \in \Theta_0 \cup \Theta_1$, to determine which of the following two hypotheses holds:*

- *Null hypothesis $H_0$: the parameter $\theta$ belongs to the null hypothesis set $\Theta_0$.*
- *Alternative hypothesis $H_1$: the parameter $\theta$ belongs to the alternative hypothesis set $\Theta_1$.*

Our main result is a lower bound for this composite hypothesis testing problem.

**Theorem 1.1** (Quantum composite hypothesis lower bound)**.** *Let $\Theta_0,\Theta_1 \subseteq \Theta = [-\pi,\pi)$ be nonempty disjoint measurable hypothesis sets with $|\Theta_1| \geq 3\pi/2$. For any $\delta \in (0,1/2)$, any quantum algorithm that solves $\mathrm{HT}(\Theta_0,\Theta_1)$ with both Type-I and Type-II errors at most $\delta$ requires at least*

$$\frac{1}{8(|\Theta| - |\Theta_1|)} \log\left(\frac{1-\delta}{\delta}\right)$$

*queries to $U(\theta)$, where $|\cdot|$ denotes the Lebesgue measure on $[-\pi,\pi)$.*

Theorem 1.1 is the main technical tool behind our applications, yielding a series of tight lower bounds for fundamental problems in quantum property testing and functional estimation; see the applications below and Section 4.

**Relation to Mande and de Wolf.** Mande and de Wolf [23, Section 5] prove their small error phase estimation lower bound using trigonometric polynomials and the Remez inequality. They consider the phase oracle $U_\theta = I - (1-e^{i\theta})|0\rangle\langle 0|$ and distinguish $\theta = 0$ from phases outside an interval around zero. Theorem 1.1 applies the same proof to trigonometrically linear unitary families and disjoint measurable hypothesis sets. The bound depends on the total measure of $S = \Theta \setminus \Theta_1$, which may be a union of intervals. In particular, the $\ell_2$ and Tsallis examples below can also be treated by a symmetry reduction or a direct adaptation of their argument. Our theorem gives a common statement for these reductions and the other applications.

### 1.2 Applications

We next present the main applications of our method; these are summarized in Table 1.

Table 1: Summary of quantum query complexities. The middle column gives constant error bounds from the cited algorithms; the Hellinger upper bound follows by applying square-root amplitude estimation [27]. The last column gives the small error bounds proved here.

| Problem | | Prior results (constant error) | This paper |
|---|---|---|---|
| $\ell_2$-closeness testing | | $\Theta(1/\varepsilon)$ [22] | $\Theta(\log(1/\delta)/\varepsilon)$ |
| Tsallis entropy estimation | integer $q \geq 2$ | $\Theta(1/\varepsilon)$ [9, 15, 6] | $\Theta(\log(1/\delta)/\varepsilon)$ |
| | real $q \geq 1.5$ | $\widetilde{O}(1/\varepsilon)$ [11] | $\widetilde{\Theta}(\log(1/\delta)/\varepsilon)$ |
| Pure-state Uhlmann fidelity and trace distance estimation | | $\Theta(1/\varepsilon)$ [27] | $\Theta(\log(1/\delta)/\varepsilon)$ |
| Identity testing in Hellinger distance (point-mass reference) | | $O(1/\varepsilon)$ via [27] | $\Theta(\log(1/\delta)/\varepsilon)$ |
| Fidelity to a pure state and $\sqrt{\mathrm{tr}(\rho\sigma^2)}$ estimation | | $\Theta(1/\varepsilon)$ [16] | $\Theta(\log(1/\delta)/\varepsilon)$ |

**$\ell_2$-closeness testing.** The first application is $\ell_2$-closeness testing: given unknown distributions $p$ and $q$, the task is to distinguish, under a promise, whether $p = q$ or $\|p - q\|_2 \geq \varepsilon$. Here, for a distribution $p = (p_1, p_2, \ldots, p_n)$, we assume query access to the quantum oracle $\mathcal{O}_p$ such that

$$\mathcal{O}_p|0\rangle = \sum_{i=1}^{n} \sqrt{p_i}|i\rangle|\varphi_i\rangle,$$

where the $|\varphi_i\rangle$'s are normalized auxiliary states. For $\ell_2$-closeness testing, we are given such access to both $\mathcal{O}_p$ and $\mathcal{O}_q$. This model is the standard quantum input model for discrete distributions, which is known as the purified quantum query access model [17].

**Corollary 1.2.** *The quantum query complexity of $\ell_2$-closeness testing with small error $\delta$ is $\Theta(\log(1/\delta)/\varepsilon)$. In other words, there is a quantum algorithm for $\ell_2$-closeness testing with small error $\delta$ with query complexity $O(\log(1/\delta)/\varepsilon)$ and any quantum algorithm for $\ell_2$-closeness testing with small error $\delta$ requires query complexity $\Omega(\log(1/\delta)/\varepsilon)$.*

The constant error version of this problem was recently settled by Luo et al. [22], who gave a succinct quantum tester with query complexity $O(1/\varepsilon)$ and proved that its dependence on $\varepsilon$ is optimal. Our result strengthens this optimality to the small error regime: it shows that the logarithmic dependence on $1/\delta$ is also necessary. Thus the natural high-confidence version of the tester of [22] is optimal in its joint dependence on $\varepsilon$ and $\delta$.

**Tsallis entropy estimation.** The second application is $q$-Tsallis entropy estimation. Under the same purified quantum query-access model, given query access to $\mathcal{O}_p$ for an unknown distribution $p$ over $[n]$, the task is to estimate the power sum $F_q(p) = \sum_{i=1}^{n} p_i^q$ to within additive error $\varepsilon$. Estimating $F_q(p)$ is equivalent, up to an affine transformation, to estimating the $q$-Tsallis entropy $\mathrm{S}_q(p) = \frac{1-F_q(p)}{q-1}$; in particular, the case $q = 2$ corresponds to the Gini impurity. We obtain the following small error query complexity for this problem.

**Corollary 1.3.** *The quantum query complexity of $q$-Tsallis entropy estimation with small error $\delta$ is $\widetilde{\Theta}(\log(1/\delta)/\varepsilon)$ for any $q \geq 1.5$.*[1] *Moreover, if $q \geq 2$ is integer, then the complexity is $\Theta(\log(1/\delta)/\varepsilon)$.*

For integer $q \geq 2$, the functional $F_q(p)$ is a power sum and can be estimated using swap-test and cyclic-permutation primitives, as in the approaches of Buhrman et al. [9] and Ekert et al. [15]; combined with amplitude estimation [6], these give the matching upper bound. Our corollary proves that these integer-order algorithms are optimal, including their dependence on the small error

[1] $\widetilde{\Theta}(\cdot)$ suppresses polylogarithmic factors in $\varepsilon$.

probability $\delta$. For real orders $q \geq 1.5$, the recent quantum multi-level estimation framework of Chen et al. [11] gives near-optimal algorithms for such distribution functionals. Our lower bound shows that, in this range of $q$, their dependence on $\varepsilon$ and $\delta$ is optimal up to polylogarithmic factors.

**Pure-state fidelity and trace distance estimation.** Our framework also applies to estimation problems for quantum objects. The third application concerns two pure-state closeness measures: Uhlmann fidelity and trace distance. Given state-preparation query access to two unknown pure states $|\varphi\rangle$ and $|\phi\rangle$, namely quantum unitary oracles $\mathcal{O}_\varphi$ and $\mathcal{O}_\phi$ such that $\mathcal{O}_\varphi|0\rangle = |\varphi\rangle$ and $\mathcal{O}_\phi|0\rangle = |\phi\rangle$, the tasks are to estimate, to additive error $\varepsilon$, the Uhlmann fidelity $\mathrm{F}(|\phi\rangle, |\varphi\rangle) = |\langle\varphi|\phi\rangle|$ and the trace distance $\mathrm{T}(|\phi\rangle, |\varphi\rangle) = \sqrt{1 - |\langle\varphi|\phi\rangle|^2}$. We obtain the following tight small error complexity for both estimation tasks.

**Corollary 1.4.** *The quantum query complexity of both pure-state Uhlmann fidelity estimation and trace distance estimation with small error $\delta$ is $\Theta(\log(1/\delta)/\varepsilon)$.*

For constant error, Wang [27] developed separate optimal quantum algorithms for pure-state fidelity estimation and trace distance estimation, each using $\Theta(1/\varepsilon)$ queries to the state-preparation circuits and quadratically improving the folklore $O(1/\varepsilon^2)$ approaches based on SWAP tests. Our corollary extends these constant error optimality results to the small error regime, showing that each of the two estimation tasks has query complexity $\Theta(\log(1/\delta)/\varepsilon)$.

**Further applications.** We also prove a tight small error bound for identity testing in Hellinger distance when the reference distribution is a point mass, with an equivalent statement in terms of the Bhattacharyya coefficient (Section 4.4). Our pure-state fidelity lower bound also gives matching small error bounds for estimating fidelity to a pure state and $\sqrt{\mathrm{tr}(\rho\sigma^2)}$, extending the constant error results of Fang and Wang [16] (Section 4.5).

### 1.3 Techniques

**Composite testing through trigonometric query families.** Our lower bound framework is inspired by the classical theory of composite hypothesis testing, which has long served as a systematic source of information-theoretic lower bounds in statistics. We bring this perspective into the quantum query setting by formulating a composite testing problem for trigonometrically linear unitary families. Concretely, rather than reducing a problem to two isolated hard instances, we embed it into a one-parameter family of quantum oracles $\{U(\theta)\}_{\theta\in\Theta}$ and ask an algorithm to distinguish a null parameter set $\Theta_0$ from a large alternative parameter set $\Theta_1$. This formulation is particularly well suited to the small error regime: the lower bound is governed not merely by a pairwise separation between hard instances, but by how sharply a quantum algorithm can separate $\Theta_0$ from $\Theta_1$ when $\Theta_1$ occupies almost all of the parameter space. Classical lower bounds based on Chernoff information concern i.i.d. observations. In the quantum query model, an algorithm may interleave coherent oracle calls with arbitrary unitaries; we control such algorithms by bounding the degree of their acceptance probability.

The reason this composite viewpoint becomes powerful in the quantum setting is that many standard quantum oracles depend on an underlying phase parameter through trigonometric polynomials. For the oracle families used in our applications, every matrix entry of $U(\theta)$ is trigonometrically linear in $\theta$. It follows from the quantum polynomial method [3] that the acceptance probability of any $T$-query quantum algorithm is a real trigonometric polynomial of degree at most $2T$ (cf. [23]). Thus, a successful small error tester would give a low-degree trigonometric polynomial that is close to one on the null side and close to zero on almost all of the parameter space. The exponential-form Remez inequality (see Theorem 3.1) shows that such a separation is impossible unless the degree grows logarithmically in $1/\delta$ and inversely with the measure of the parameter region not covered by the alternative hypothesis. In the notation of our main theorem, this yields a lower bound of order $\Omega(\log(1/\delta)/(|\Theta| - |\Theta_1|))$.

**Instantiating the framework.** To apply the framework, the main task is to embed a concrete testing or estimation problem into a one-parameter trigonometrically linear oracle family and to choose the hypothesis sets so that the alternative region occupies almost all of $\Theta$. Let us illustrate this with $\ell_2$-closeness testing. We fix $p^\star = (1/2, 1/2, 0, \ldots, 0)$, and let the second distribution vary as $q^{(\theta)} = (\cos^2(\theta/2), \sin^2(\theta/2), 0, \ldots, 0)$, which is generated by the Hadamard–phase–Hadamard unitary $W_\theta = H\,\mathrm{diag}(1, e^{i\theta})H$. This gives a trigonometrically linear oracle family, and the $\ell_2$

distance satisfies $\|q^{(\theta)} - p^\star\|_2 = |\cos(\theta)|/\sqrt{2}$. Thus the null case occurs at $\theta^\star = \pi/2$. We choose $\Theta_0 = \{\pi/2\}$ and $\Theta_1 = \{\theta : \|q^{(\theta)} - p^\star\|_2 \geq \varepsilon\}$. The exceptional region outside $\Theta_1$ is exactly the set where $|\cos(\theta)| < \sqrt{2}\,\varepsilon$, whose measure is $\Theta(\varepsilon)$. Therefore, Theorem 1.1 gives the lower bound $\Omega(\log(1/\delta)/\varepsilon)$. The other applications follow the same pattern: for Tsallis entropy estimation, we choose a two-point family whose power sum varies linearly near a reference phase; for pure-state fidelity and trace distance estimation, the Hadamard–phase–Hadamard family makes the relevant quantities simple trigonometric functions of $\theta$. In each case, the exceptional region has measure $O(\varepsilon)$, and the composite-testing lower bound yields the desired small error dependence.

### 1.4 Discussion

We developed a quantum composite hypothesis testing framework for trigonometrically linear unitary families, turning small exceptional sets in phase space into tight small error quantum lower bounds. Together with known algorithms, this framework settles the small error complexities of several quantum testing and estimation problems considered in this paper.

Several directions remain open. First, it would be interesting to develop a quantum analogue of the method of two fuzzy hypotheses [25]. Our present framework separates a null parameter set from a large alternative set along a deterministic one-parameter oracle family. In contrast, classical minimax lower-bound theory contains several complementary tools, including Le Cam's method, Assouad's lemma, Fano's method, and the method of two fuzzy hypotheses, which use carefully chosen finite sets, hypercubes, packings, or priors over composite hypotheses to prove lower bounds [20, 29, 25, 19]. A quantum fuzzy-hypothesis principle would aim to compare the behavior of quantum algorithms under two distributions over oracle instances or states, rather than under a single localized exceptional set. Such a principle could handle hard instances that arise from averaging over structured families of oracles or states, and may bring quantum minimax lower bounds closer to their classical counterparts.

Second, our framework is intrinsically one-dimensional: the hard family is parametrized by a single phase $\theta$, and the lower bound is controlled by the measure of an exceptional subset of $[-\pi, \pi)$. Many high-dimensional statistical problems have hard instances with several independent perturbation directions. A natural next step is to develop multi-parameter extensions of the present approach. For phase-parametrized oracle families, this would lead to multivariate trigonometric polynomials; for other parametrizations, ordinary or structured multivariate polynomials may be more appropriate. The main challenge is to find higher-dimensional analogues of the Remez principle that convert the geometry of exceptional sets into quantum query lower bounds.

Third, our results reinforce the view that the dependence on the error probability $\delta$ should be treated as an intrinsic part of the complexity, rather than merely as a black-box amplification factor. This phenomenon is already familiar in classical distribution testing: the optimal high-confidence complexity of identity testing improves over naive amplification of the constant error tester [13]. In the quantum setting, the landscape appears even richer. The problems studied in this paper have logarithmic dependence on $1/\delta$, as in quantum phase estimation [23], while quantum search exhibits a different high-confidence dependence on $\delta$ [8]. For a fixed oracle family and exceptional set $S$, Theorem 1.1 gives a lower bound proportional to $\log((1-\delta)/\delta)/|S|$. This dependence does not explain the $\sqrt{\log(1/\delta)}$ factor in the optimal complexity of quantum search [8].

Understanding which quantum testing and estimation problems admit small error behavior beyond black-box amplification remains an interesting direction.

Finally, the present applications all rely on relatively simple phase-parametrized hard families. It would be useful to identify other settings in which such families exist and in which the relevant quantity has controlled local variation near a reference phase. This may lead to further small error lower bounds for related quantum statistical estimation problems.

## 2 Preliminaries

We briefly recall the basic quantum notation and the query model used throughout the paper. A qubit is a unit vector in $\mathbb{C}^2$, and an $n$-qubit pure state is a unit vector in $(\mathbb{C}^2)^{\otimes n}$. We use Dirac notation and

write the computational basis as $\{|x\rangle : x \in \{0,1\}^n\}$. If a state $|\psi\rangle = \sum_x \alpha_x |x\rangle$ is measured in the computational basis, the outcome is $x$ with probability $|\alpha_x|^2 = |\langle x|\psi\rangle|^2$.

A quantum oracle is a unitary operator $\mathcal{O}$ on a finite-dimensional Hilbert space. A $T$-query quantum algorithm $\mathcal{A}$ with oracle access to $\mathcal{O}$ consists of oracle-independent unitaries $U_0, U_1, \ldots, U_T$, together with $T$ oracle calls $\mathcal{O}_1, \ldots, \mathcal{O}_T$, where each $\mathcal{O}_j$ may be $\mathcal{O}$, $\mathcal{O}^\dagger$, or a controlled version thereof. Starting from a fixed initial state $|0\rangle$, the algorithm produces the final state $|\psi\rangle = U_T \mathcal{O}_T U_{T-1} \cdots U_1 \mathcal{O}_1 U_0 |0\rangle$. It then performs a final measurement in the computational basis and outputs $x$ with probability $\Pr[\mathcal{A}^{\mathcal{O}} = x] = |\langle x|\psi\rangle|^2$. More generally, for an event $\Gamma$ consisting of several possible outcomes, we write $\Pr[\mathcal{A}^{\mathcal{O}} \in \Gamma] = \sum_{x \in \Gamma} |\langle x|\psi\rangle|^2$. The query complexity of $\mathcal{A}$ is the number $T$ of oracle calls.

## 3 Quantum Composite Hypothesis Testing

In this section, we prove the main composite-testing lower bound, Theorem 1.1. Let $\Theta = [-\pi, \pi)$. For a measurable set $E \subseteq \Theta$, we write $|E|$ for its Lebesgue measure. We say that a function $p : \Theta \to \mathbb{C}$ is a trigonometric polynomial of degree at most $d$ if it can be written as $p(\theta) = \sum_{k=-d}^{d} c_k e^{ik\theta}$.

The proof can be viewed as a trigonometric analogue of the quantum polynomial method. Along a trigonometrically linear unitary family $U(\theta)$, any $T$-query algorithm has acceptance probability given by a real trigonometric polynomial of degree at most $2T$. Thus, a small error algorithm for the composite testing problem would induce a low-degree trigonometric polynomial separating the null hypothesis set from a large alternative hypothesis set. The following Remez-type inequality quantifies the obstruction to such a separation.

**Theorem 3.1** (Exponential-form Remez inequality, [5, Theorem 5.1.2])**.** *Let $p : \Theta \to \mathbb{R}$ be a real trigonometric polynomial of degree at most $d$, and let $E \subseteq \Theta$ be measurable with $|E| \leq \pi/2$. Then*

$$\sup_{\theta \in \Theta} |p(\theta)| \leq \exp(4d|E|) \sup_{\theta \in \Theta \setminus E} |p(\theta)|.$$

We now turn to the proof of Theorem 1.1.

*Proof of Theorem 1.1.* Let $\mathcal{A}$ be a $T$-query quantum algorithm that solves $\mathrm{HT}(\Theta_0, \Theta_1)$ with both Type-I and Type-II errors at most $\delta$. Let $\Gamma$ be the event that $\mathcal{A}$ outputs the null hypothesis $H_0$. For each $\theta \in \Theta$, define $f(\theta) := \Pr[\mathcal{A}^{U(\theta)} \in \Gamma]$.

We first record the trigonometric-polynomial structure of $f$, following the degree-counting argument of [23, Section 5]. By deferred measurement, we may assume that $\mathcal{A}$ performs only a final measurement. After $t$ oracle calls, every basis-state amplitude is a trigonometric polynomial of degree at most $t$ in $\theta$. Indeed, this is trivial for $t = 0$, and each query to $U(\theta)$, $U(\theta)^\dagger$, or a controlled version thereof increases the degree by at most one, since the entries of $U(\theta)$ are trigonometrically linear in $\theta$; the oracle-independent unitaries do not change the degree. Thus every final amplitude $\alpha_x(\theta)$ has degree at most $T$, and hence

$$f(\theta) = \sum_{x \in \Gamma} |\alpha_x(\theta)|^2$$

is a real trigonometric polynomial of degree at most $2T$. Indeed, multiplying frequencies $e^{ik\theta}$ and $e^{-i\ell\theta}$ from an amplitude and its conjugate produces only frequencies $k - \ell \in [-2T, 2T]$.

Now choose any reference point $\theta^\star \in \Theta_0$, and set $S = \Theta \setminus \Theta_1$. Since $\Theta_0 \cap \Theta_1 = \emptyset$, we have $\theta^\star \in S$. Moreover, the assumption $|\Theta_1| \geq 3\pi/2$ gives $|S| = |\Theta| - |\Theta_1| \leq \pi/2$. The Type-I error guarantee gives $f(\theta^\star) \geq 1 - \delta$, while the Type-II error guarantee gives $f(\theta) \leq \delta$ for every $\theta \in \Theta_1 = \Theta \setminus S$.

Apply Theorem 3.1 to the real trigonometric polynomial $g(\theta) = f(\theta)/\delta$, whose degree is at most $2T$. Since $|g(\theta)| \leq 1$ for every $\theta \in \Theta \setminus S$, we obtain

$$\frac{1-\delta}{\delta} \leq g(\theta^\star) \leq \sup_{\theta \in \Theta} |g(\theta)| \leq \exp(8T|S|).$$

Rearranging yields

$$T \geq \frac{1}{8|S|} \log\left(\frac{1-\delta}{\delta}\right) = \frac{1}{8(|\Theta| - |\Theta_1|)} \log\left(\frac{1-\delta}{\delta}\right),$$

as claimed. □

# 4 Applications

We now instantiate Theorem 1.1 for several testing and estimation problems. In each application, the key step is a carefully chosen one-parameter trigonometrically linear oracle family that converts the desired testing or estimation task into a composite hypothesis testing problem. We then choose hypothesis sets $\Theta_0, \Theta_1 \subseteq \Theta$ so that the alternative set covers all but a small exceptional region, and Theorem 1.1 yields the desired lower bound.

## 4.1 $\ell_2$-closeness testing

We first consider $\ell_2$-closeness testing. Our hard family fixes one distribution and varies the other along a two-dimensional phase family for which the $\ell_2$ distance is linear around the null phase.

**Theorem 4.1** (Quantum lower bound for $\ell_2$-closeness testing)**.** *For every $\varepsilon \in (0, 1/4]$, $\delta \in (0, 1/3]$, and $n \geq 2$, consider the following promise problem: given purified quantum query access to two unknown distributions $p, q$ on $[n]$, decide whether*

$$p = q \quad \text{or} \quad \|p - q\|_2 \geq \varepsilon$$

*with success probability at least $1 - \delta$. Then every quantum algorithm for this problem requires $\Omega(\log(1/\delta)/\varepsilon)$ queries.*

*Proof.* Fix $p^\star = (1/2, 1/2, 0, \ldots, 0)$. Let

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

be the Hadamard gate, and set $W_\theta = H \operatorname{diag}(1, e^{i\theta}) H$ and $\widetilde{W}_\theta = W_\theta \oplus I_{n-2}$.

Let $q^{(\theta)}$ be the distribution prepared by the purified oracle $\widetilde{W}_\theta$. Then $q^{(\theta)} = (\cos^2(\theta/2), \sin^2(\theta/2), 0, \ldots, 0)$, and every entry of $\widetilde{W}_\theta$ is trigonometrically linear in $\theta$.

Define $g(\theta) = \left\| q^{(\theta)} - p^\star \right\|_2$. A direct calculation gives $g(\theta) = |\cos\theta|/\sqrt{2}$. We choose $\Theta_0 = \{\pi/2\}$ and $\Theta_1 = \{\theta \in \Theta : g(\theta) \geq \varepsilon\}$. Then $\Theta_0 \cap \Theta_1 = \emptyset$, and the exceptional region is $\Theta \setminus \Theta_1 = \{\theta \in \Theta : |\cos\theta| < \sqrt{2}\,\varepsilon\}$. Hence $|\Theta| - |\Theta_1| = 4 \arcsin(\sqrt{2}\,\varepsilon) = O(\varepsilon)$. For $\varepsilon \leq 1/4$, this measure is at most $4 \arcsin(\sqrt{2}/4) < \pi/2$, so $|\Theta_1| \geq 3\pi/2$.

Now suppose there is a $T$-query tester for the original $\ell_2$-closeness problem. Restricting it to the family $(p^\star, q^{(\theta)})$ gives an algorithm for $\mathrm{HT}(\Theta_0, \Theta_1)$: it outputs $H_0$ when the tester declares $p = q$, and $H_1$ otherwise. The oracle for $p^\star$ is fixed and can be absorbed into the oracle-independent part of the circuit, so the only parameter-dependent oracle is $\widetilde{W}_\theta$. By Theorem 1.1,

$$T \geq \Omega\left( \frac{1}{|\Theta| - |\Theta_1|} \log\left(\frac{1}{\delta}\right) \right) = \Omega\left( \frac{1}{\varepsilon} \log\left(\frac{1}{\delta}\right) \right).$$

□

The distance vanishes at $\theta = \pm\pi/2$ and grows linearly near both points, so the excluded set consists of two intervals of total length $O(\varepsilon)$.

## 4.2 Tsallis entropy estimation

We next consider $q$-Tsallis entropy estimation. For fixed $q > 0$ with $q \neq 1$, estimating the $q$-Tsallis entropy $S_q(p) = (1 - F_q(p))/(q-1)$ is equivalent, up to a known affine transformation, to estimating the power sum $F_q(p) = \sum_i p_i^q$. The case $q = 2$, corresponding to the Gini impurity, already illustrates the construction: on a two-point family $p^{(\theta)} = (x(\theta), 1 - x(\theta))$ with $x(\theta) = u_0 + v_0 \sin\theta$ and $u_0 \neq 1/2$, the quantity $F_2(p^{(\theta)}) = x(\theta)^2 + (1 - x(\theta))^2$ has a nonzero derivative at the reference phase. The same idea works for every fixed $q > 0$, $q \neq 1$, by keeping $x(\theta)$ inside an interval contained in $(1/2, 1)$, where the derivative of $x^q + (1 - x)^q$ is bounded away from zero.

**Theorem 4.2** (Quantum lower bound for estimating $F_q$)**.** *Fix $q > 0$ with $q \neq 1$. There exists a constant $c_q > 0$ such that the following holds. For every $\varepsilon \in (0, c_q]$, $\delta \in (0, 1/3]$, and $n \geq 2$, consider the problem of estimating $F_q(p) = \sum_{i=1}^{n} p_i^q$ to additive error $\varepsilon$ with success probability at least $1 - \delta$, given purified quantum query access to an unknown distribution $p$ on $[n]$. Then every quantum algorithm for this problem requires $\Omega(\log(1/\delta)/\varepsilon)$ queries.*

*Proof.* We construct a two-point family $p^{(\theta)} = (x(\theta), 1 - x(\theta), 0, \ldots, 0)$ whose first coordinate stays inside a fixed interval contained in $(1/2, 1)$. Let $\phi = \pi/10$, $\psi = \pi/20$, $D_\theta = \mathrm{diag}(1, e^{i\theta})$, and

$$A = \begin{pmatrix} \cos\phi & -i\sin\phi \\ -i\sin\phi & \cos\phi \end{pmatrix}, \qquad B = \begin{pmatrix} \cos\psi & -\sin\psi \\ \sin\psi & \cos\psi \end{pmatrix}.$$

Set $W_\theta = AD_\theta B$ and $\widetilde{W}_\theta = W_\theta \oplus I_{n-2}$, and let $p^{(\theta)}$ be the distribution prepared by $\widetilde{W}_\theta$. The family $\widetilde{W}_\theta$ is trigonometrically linear. Writing $x(\theta) = p_1^{(\theta)}$, a direct calculation gives $x(\theta) = u_0 + v_0 \sin\theta$, where

$$u_0 = \cos^2\phi\cos^2\psi + \sin^2\phi\sin^2\psi,$$
$$v_0 = 2\cos\phi\cos\psi\sin\phi\sin\psi > 0.$$

Moreover,

$$x(\theta) \in I \coloneqq [\cos^2(3\pi/20), \cos^2(\pi/20)] \subset (1/2, 1)$$

for every $\theta \in \Theta$.

Let $h_q(x) = x^q + (1-x)^q$; then $F_q(p^{(\theta)}) = h_q(x(\theta))$. For $x \in I$, we have $x > 1 - x$. Consider the following two cases.

- If $q > 1$, then $t \mapsto t^{q-1}$ is strictly increasing, so $h_q'(x) = q(x^{q-1} - (1-x)^{q-1}) > 0$ on $I$.
- If $0 < q < 1$, then $t \mapsto t^{q-1}$ is strictly decreasing, so $h_q'(x) < 0$ on $I$.

In either case, $h_q'$ has no zero on $I$. Since $h_q'$ is continuous on the compact interval $I$, we have $m_q \coloneqq \min_{x \in I} |h_q'(x)| > 0$. Set $\kappa_q = m_q v_0$. By the mean value theorem,

$$|F_q(p^{(\theta)}) - F_q(p^{(0)})| = |h_q(x(\theta)) - h_q(x(0))| \geq m_q |x(\theta) - x(0)| = \kappa_q |\sin\theta|.$$

Set $c_q = \kappa_q/8$, and assume $\varepsilon \leq c_q$. Define

$$S = \{\theta \in \Theta : |F_q(p^{(\theta)}) - F_q(p^{(0)})| \leq 2\varepsilon\}.$$

The previous inequality implies $S \subseteq \{\theta : |\sin\theta| \leq 2\varepsilon/\kappa_q\}$, and hence $|S| \leq 4\arcsin(2\varepsilon/\kappa_q) = O_q(\varepsilon)$. Also, since $2\varepsilon/\kappa_q \leq 1/4$, we have $|S| \leq 4\arcsin(1/4) < \pi/2$.

Take $\Theta_0 = \{0\}$ and $\Theta_1 = \Theta \setminus S$. Suppose a $T$-query estimator $\widetilde{F}$ is $\varepsilon$-accurate for $F_q$ with success probability at least $1 - \delta$. We turn it into a tester for $\mathrm{HT}(\Theta_0, \Theta_1)$ by outputting $H_0$ if and only if $|\widetilde{F} - F_q(p^{(0)})| \leq \varepsilon$. At $\theta = 0$, this event occurs with probability at least $1 - \delta$. For every $\theta \in \Theta_1$, we have $|F_q(p^{(\theta)}) - F_q(p^{(0)})| > 2\varepsilon$, so by the triangle inequality an $\varepsilon$-accurate estimate cannot lead to output $H_0$. Thus the Type-II error is at most $\delta$. By Theorem 1.1,

$$T \geq \Omega_q\left(\frac{1}{|S|}\log\left(\frac{1}{\delta}\right)\right) = \Omega_q\left(\frac{1}{\varepsilon}\log\left(\frac{1}{\delta}\right)\right).$$

□

The choice $I \subset (1/2, 1)$ ensures that $|h_q'|$ is bounded below by a positive constant. Hence the phases for which $|F_q(p^{(\theta)}) - F_q(p^{(0)})| \leq 2\varepsilon$ lie in intervals of total length $O_q(\varepsilon)$ around 0 and $\pi$ modulo $2\pi$.

### 4.3 Pure-state fidelity and trace distance estimation

We next apply the framework to pure-state functional estimation. A single two-dimensional phase family suffices for both Uhlmann fidelity $\mathrm{F}(|\phi\rangle, |\varphi\rangle) = |\langle\varphi|\phi\rangle|$ and trace distance $\mathrm{T}(|\phi\rangle, |\varphi\rangle) = \sqrt{1 - |\langle\varphi|\phi\rangle|^2}$, since both quantities are simple functions of the overlap between the two states.

**Theorem 4.3** (Quantum lower bound for pure-state Uhlmann fidelity estimation)**.** *For every $\varepsilon \in (0, 1/32]$ and $\delta \in (0, 1/3]$, consider the following problem: given state-preparation query access to two unknown pure states $|\phi\rangle$ and $|\varphi\rangle$, output an estimate $\widetilde{F}$ satisfying $|\widetilde{F} - \mathrm{F}(|\phi\rangle, |\varphi\rangle)| \le \varepsilon$ with success probability at least $1 - \delta$. Then every quantum algorithm for this problem requires $\Omega(\log(1/\delta)/\varepsilon)$ queries.*

*Proof.* Fix $|\phi\rangle = |0\rangle$, let $W_\theta = H \,\mathrm{diag}(1, e^{i\theta})H$, and set $|\varphi^{(\theta)}\rangle = W_\theta|0\rangle$. The oracle preparing $|\phi\rangle$ is fixed, while $W_\theta$ is trigonometrically linear. For $\theta \in [-\pi, \pi)$, the Uhlmann fidelity along this family is

$$F(\theta) = \left|\left\langle \phi \middle| \varphi^{(\theta)} \right\rangle\right| = |\langle 0|W_\theta|0\rangle| = \cos(\theta/2).$$

Take

$$\theta^\star = \pi/2, \quad a = F(\theta^\star) = 1/\sqrt{2}, \quad S = \{\theta \in \Theta : |F(\theta) - a| \le 2\varepsilon\}.$$

For $u = 2\varepsilon \le 1/16$, the set $S$ has measure $4(\arccos(a - u) - \arccos(a + u))$. Since the derivative of $\arccos$ is bounded on $[a - 1/16, a + 1/16]$, we get $|S| = O(\varepsilon)$; moreover, for $\varepsilon \le 1/32$, this measure is at most $\pi/2$. Let $\Theta_0 = \{\theta^\star\}$ and $\Theta_1 = \Theta \setminus S$.

Suppose a $T$-query estimator $\widetilde{F}$ is $\varepsilon$-accurate with probability at least $1 - \delta$. Define a tester by outputting $H_0$ if and only if $|\widetilde{F} - a| \le \varepsilon$. At $\theta = \theta^\star$, this happens with probability at least $1 - \delta$. For every $\theta \in \Theta_1$, we have $|F(\theta) - a| > 2\varepsilon$, so an $\varepsilon$-accurate estimate cannot lead to output $H_0$. Therefore the tester solves $\mathrm{HT}(\Theta_0, \Theta_1)$ with error at most $\delta$. By Theorem 1.1,

$$T \ge \Omega\left(\frac{1}{|S|}\log\left(\frac{1}{\delta}\right)\right) = \Omega\left(\frac{1}{\varepsilon}\log\left(\frac{1}{\delta}\right)\right).$$

□

We choose the reference value $1/\sqrt{2}$ because $F'(\theta) \neq 0$ at both phases $\theta = \pm\pi/2$ where it is attained. This gives $|S| = O(\varepsilon)$.

**Theorem 4.4** (Quantum lower bound for pure-state trace distance estimation)**.** *For every $\varepsilon \in (0, 1/8]$ and $\delta \in (0, 1/3]$, consider the following problem: given state-preparation query access to two unknown pure states $|\phi\rangle$ and $|\varphi\rangle$, output an estimate $\widetilde{T}$ satisfying $|\widetilde{T} - \mathrm{T}(|\phi\rangle, |\varphi\rangle)| \le \varepsilon$ with success probability at least $1 - \delta$. Then every quantum algorithm for this problem requires $\Omega(\log(1/\delta)/\varepsilon)$ queries.*

*Proof.* Use the same family as above: $|\phi\rangle = |0\rangle$ and $|\varphi^{(\theta)}\rangle = W_\theta|0\rangle$, where $W_\theta = H \,\mathrm{diag}(1, e^{i\theta})H$. The trace distance along this family is $D(\theta) = \mathrm{T}(|\phi\rangle, |\varphi^{(\theta)}\rangle) = |\sin(\theta/2)|$. Take

$$\theta^\star = 0, \quad S = \{\theta \in \Theta : D(\theta) \le 2\varepsilon\}, \quad \Theta_0 = \{0\}, \quad \Theta_1 = \Theta \setminus S.$$

For $\varepsilon \le 1/8$, we have

$$|S| = |\{\theta \in \Theta : |\sin(\theta/2)| \le 2\varepsilon\}| = 4\arcsin(2\varepsilon) = O(\varepsilon),$$

and $|S| \le 4\arcsin(1/4) < \pi/2$.

Given an $\varepsilon$-accurate trace distance estimator $\widetilde{T}$, define a tester that outputs $H_0$ if and only if $\widetilde{T} \le \varepsilon$. At $\theta = 0$, the true trace distance is 0, so the tester outputs $H_0$ with probability at least $1 - \delta$. For every $\theta \in \Theta_1$, the true trace distance is larger than $2\varepsilon$, so an $\varepsilon$-accurate estimate cannot lead to output $H_0$. Therefore the tester solves $\mathrm{HT}(\Theta_0, \Theta_1)$ with error at most $\delta$. By Theorem 1.1,

$$T \ge \Omega\left(\frac{1}{|S|}\log\left(\frac{1}{\delta}\right)\right) = \Omega\left(\frac{1}{\varepsilon}\log\left(\frac{1}{\delta}\right)\right).$$

□

Here $D(\theta) = |\sin(\theta/2)|$ grows linearly in $|\theta|$ near zero, which gives $|S| = O(\varepsilon)$.

### 4.4 Identity testing in Hellinger distance

For distributions $p, r$, let $\mathrm{BC}(p,r) = \sum_i \sqrt{p_i r_i}$ be the Bhattacharyya coefficient. We use the normalized Hellinger distance $d_{\mathrm{H}}(p,r) = \left(\frac{1}{2}\sum_i (\sqrt{p_i} - \sqrt{r_i})^2\right)^{1/2} = \sqrt{1 - \mathrm{BC}(p,r)}$. We consider identity testing with the known reference distribution $p^\star = (1, 0, \ldots, 0)$ and purified quantum query access to $r$.

**Corollary 4.5.** *For $n \geq 2$, $0 < \delta \leq 1/3$, and $0 < \varepsilon \leq 1/8$, distinguishing $r = p^\star$ from $d_{\mathrm{H}}(r, p^\star) \geq \varepsilon$ with error probability at most $\delta$ has query complexity $\Theta(\log(1/\delta)/\varepsilon)$. Equivalently, for $0 < \eta \leq 1/64$, distinguishing $\mathrm{BC}(r, p^\star) = 1$ from $\mathrm{BC}(r, p^\star) \leq 1 - \eta$ has complexity $\Theta(\log(1/\delta)/\sqrt{\eta})$.*

*Proof.* For the lower bound, use $W_\theta = H\,\mathrm{diag}(1, e^{i\theta})H$, embedded in $n$ dimensions as in the proof of Theorem 4.1, to prepare $r^{(\theta)} = (\cos^2(\theta/2), \sin^2(\theta/2), 0, \ldots, 0)$. On $\Theta = [-\pi, \pi)$,

$$\mathrm{BC}(r^{(\theta)}, p^\star) = \cos(\theta/2), \qquad d_{\mathrm{H}}(r^{(\theta)}, p^\star) = \sqrt{2}\,|\sin(\theta/4)|.$$

Take $\Theta_0 = \{0\}$ and $\Theta_1 = \{\theta : d_{\mathrm{H}}(r^{(\theta)}, p^\star) \geq \varepsilon\}$. Then $|\Theta \setminus \Theta_1| = 8\arcsin(\varepsilon/\sqrt{2}) = O(\varepsilon)$, and this measure is at most $\pi/2$ for $\varepsilon \leq 1/8$. The lower bound follows from Theorem 1.1.

For the upper bound, let $h = d_{\mathrm{H}}(r, p^\star)$ and let $a = 1 - r_1$ be the probability of observing an outcome other than the first. Since $a = h^2(2 - h^2)$, the promise implies either $a = 0$ or $\sqrt{a} \geq \varepsilon$. Square-root amplitude estimation [27] estimates $\sqrt{a}$ to accuracy $\varepsilon/3$ using $O(1/\varepsilon)$ queries at constant error. Comparing the estimate with $\varepsilon/2$ decides the two cases; $O(\log(1/\delta))$ independent repetitions and a majority vote reduce the error to $\delta$. The equivalent statement follows from $d_{\mathrm{H}}^2 = 1 - \mathrm{BC}$ by setting $\varepsilon = \sqrt{\eta}$. □

Since a point mass is an admissible reference distribution, this also proves a worst-case lower bound for identity testing with a general known reference distribution. The matching upper bound above assumes that the reference is a point mass.

### 4.5 Fidelity to a pure state and $\sqrt{\mathrm{tr}(\rho\sigma^2)}$ estimation

Fang and Wang [16, Theorems 6 and 8] give constant error estimators for $\mathrm{F}(\rho, |\psi\rangle) = \sqrt{\langle\psi|\rho|\psi\rangle}$ and $\sqrt{\mathrm{tr}(\rho\sigma^2)}$, each using $O(1/\varepsilon)$ queries. Here $\rho$ and $\sigma$ are density operators, and purified quantum query access allows queries to a unitary preparing a purification of each input state, its inverse, and their controlled versions.

**Corollary 4.6.** *For $0 < \varepsilon \leq 1/32$ and $0 < \delta \leq 1/3$, estimating either $\mathrm{F}(\rho, |\psi\rangle)$ or $\sqrt{\mathrm{tr}(\rho\sigma^2)}$ to additive error $\varepsilon$ with failure probability at most $\delta$ has purified quantum query complexity $\Theta(\log(1/\delta)/\varepsilon)$.*

*Proof.* Run the respective estimator of [16] independently $O(\log(1/\delta))$ times and take the median, giving the upper bound. For the lower bound, restrict to $\rho = |\varphi\rangle\langle\varphi|$ and $\sigma = |\psi\rangle\langle\psi|$ with trivial purifying registers. Both quantities then equal $|\,\langle\varphi|\psi\rangle\,|$, so Theorem 4.3 applies. □

## Acknowledgments and Disclosure of Funding

This research was partially supported by the Quantum Science and Technology-National Science and Technology Major Project under Grant No. 2024ZD0300500. Qisheng Wang was supported by the Startup Funding from Shanghai Jiao Tong University.

## References


[1] Ziv Bar-Yossef. *The Complexity of Massive Data Set Computations*. PhD thesis, University of California, Berkeley, 2002. URL https://www.proquest.com/docview/304791145.

[2] Ziv Bar-Yossef, Ravi Kumar, and D. Sivakumar. Sampling algorithms: lower bounds and applications. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 266–275, New York, NY, USA, 2001. Association for Computing Machinery. ISBN 1581133499. doi: 10.1145/380752.380810. URL https://doi.org/10.1145/380752.380810.

[3] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, July 2001. ISSN 0004-5411. doi: 10.1145/502090.502097. URL https://doi.org/10.1145/502090.502097.

[4] Aleksandrs Belovs. Quantum Algorithms for Classical Probability Distributions. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms (ESA 2019)*, volume 144 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:11, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-124-5. doi: 10.4230/LIPIcs.ESA.2019.16. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2019.16.

[5] Peter Borwein and Tamas Erdelyi. *Polynomials and Polynomial Inequalities*, volume 161. Springer Science & Business Media, 1995.

[6] 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.

[7] Sergey Bravyi, Aram W. Harrow, and Avinatan Hassidim. Quantum algorithms for testing properties of distributions. *IEEE Transactions on Information Theory*, 57(6):3971–3981, 2011. doi: 10.1109/TIT.2011.2134250.

[8] H. Buhrman, R. Cleve, R. De Wolf, and C. Zalka. Bounds for small-error and zero-error quantum algorithms. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*, pages 358–368, 1999. doi: 10.1109/SFFCS.1999.814607.

[9] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87:167902, Sep 2001. doi: 10.1103/PhysRevLett.87.167902. URL https://link.aps.org/doi/10.1103/PhysRevLett.87.167902.

[10] Sourav Chakraborty, Eldar Fischer, Arie Matsliah, and Ronald de Wolf. New Results on Quantum Property Testing. In Kamal Lodaya and Meena Mahajan, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, volume 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 145–156, Dagstuhl, Germany, 2010. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-939897-23-1. doi: 10.4230/LIPIcs.FSTTCS.2010.145. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.FSTTCS.2010.145.

[11] Kean Chen, Minbo Gao, Tongyang Li, Qisheng Wang, and Xinzhao Wang. Quantum multi-level estimation of functionals of discrete distributions. In *Proceedings of the 53rd International Colloquium on Automata, Languages, and Programming*, Leibniz International Proceedings in Informatics. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026. To appear.

[12] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952. URL http://www.jstor.org/stable/2236576.

[13] Ilias Diakonikolas, Themis Gouleakis, John Peebles, and Eric Price. Sample-Optimal Identity Testing with High Probability. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics*

*(LIPIcs)*, pages 41:1–41:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-076-7. doi: 10.4230/LIPIcs.ICALP.2018.41. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2018.41.

[14] Ilias Diakonikolas, Themis Gouleakis, Daniel M. Kane, John Peebles, and Eric Price. Optimal testing of discrete distributions with high probability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 542–555, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380539. doi: 10.1145/3406325.3450997. URL https://doi.org/10.1145/3406325.3450997.

[15] Artur K. Ekert, Carolina Moura Alves, Daniel K. L. Oi, Michał Horodecki, Paweł Horodecki, and L. C. Kwek. Direct estimations of linear and nonlinear functionals of a quantum state. *Physical Review Letters*, 88:217901, May 2002. doi: 10.1103/PhysRevLett.88.217901. URL https://link.aps.org/doi/10.1103/PhysRevLett.88.217901.

[16] Wang Fang and Qisheng Wang. Optimal Quantum Algorithm for Estimating Fidelity to a Pure State. In Anne Benoit, Haim Kaplan, Sebastian Wild, and Grzegorz Herman, editors, *33rd Annual European Symposium on Algorithms (ESA 2025)*, volume 351 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:12, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.ESA.2025.4. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2025.4.

[17] András Gilyén and Tongyang Li. Distributional Property Testing in a Quantum World. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:19, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-134-4. doi: 10.4230/LIPIcs.ITCS.2020.25. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2020.25.

[18] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996. doi: 10.1145/237814.237866.

[19] Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015. doi: 10.1109/TIT.2015.2412945.

[20] Lucien LeCam. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, pages 38–53, 1973. URL http://www.jstor.org/stable/2958155.

[21] Tongyang Li and Xiaodi Wu. Quantum query complexity of entropy estimation. *IEEE Transactions on Information Theory*, 65(5):2899–2921, 2019. doi: 10.1109/TIT.2018.2883306.

[22] Jingquan Luo, Qisheng Wang, and Lvzhou Li. Succinct quantum testers for closeness and k-wise uniformity of probability distributions. *IEEE Transactions on Information Theory*, 70 (7):5092–5103, 2024. doi: 10.1109/TIT.2024.3393756.

[23] Nikhil S. Mande and Ronald de Wolf. Tight Bounds for Quantum Phase Estimation and Related Problems. In Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman, editors, *31st Annual European Symposium on Algorithms (ESA 2023)*, volume 274 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 81:1–81:16, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-295-2. doi: 10.4230/LIPIcs.ESA.2023.81. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2023.81.

[24] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008. doi: 10.1109/TIT.2008.928987.

[25] Alexandre B. Tsybakov. *Nonparametric estimators*, pages 1–76. Springer New York, New York, NY, 2009. ISBN 978-0-387-79052-7. doi: 10.1007/978-0-387-79052-7_1.

[26] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017. doi: 10.1137/151002526. URL https://doi.org/10.1137/151002526.

[27] 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.

[28] Xinzhao Wang, Shengyu Zhang, and Tongyang Li. A quantum algorithm framework for discrete probability distributions with applications to rényi entropy estimation. *IEEE Transactions on Information Theory*, 70(5):3399–3426, 2024. doi: 10.1109/TIT.2024.3382037.

[29] Bin Yu. Assouad, Fano, and Le Cam. In David Pollard, Erik Torgersen, and Grace L. Yang, editors, *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, pages 423–435. Springer New York, NY, 1997. doi: 10.1007/978-1-4612-1880-7_29.

## NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: The abstract and introduction state the main contribution.

2. **Limitations**

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

   Answer: [Yes]

   Justification: The limitations and future directions are discussed in the Discussion subsection.

3. **Theory assumptions and proofs**

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

   Answer: [Yes]

   Justification: The assumptions are stated in the theorem statements, and complete proofs are provided in Sections 3 and 4.

4. **Experimental result reproducibility**

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

   Answer: [N/A]

   Justification: The paper does not include experiments; its contributions are theoretical lower bounds.

5. **Open access to data and code**

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

   Answer: [N/A]

   Justification: The paper does not include experiments, datasets, or code artifacts needed to support the main claims.

6. **Experimental setting/details**

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

   Answer: [N/A]

   Justification: The paper contains no experimental evaluation, training procedure, or test setting.

7. **Experiment statistical significance**

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

   Answer: [N/A]

   Justification: The paper contains no experiments, so statistical significance reporting is not applicable.

8. **Experiments compute resources**

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

   Answer: [N/A]

Justification: The paper contains no computational experiments.

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: The work is theoretical and does not involve human subjects, datasets, deployed systems, or safety-sensitive data. We have reviewed the NeurIPS Code of Ethics and are not aware of any violation.

10. **Broader impacts**

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

    Answer: [N/A]

    Justification: The paper is foundational theoretical work on quantum query lower bounds. It does not introduce a deployed system, dataset, model, or application with direct societal impact.

11. **Safeguards**

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

    Answer: [N/A]

    Justification: The paper does not release data, models, or systems with misuse risk.

12. **Licenses for existing assets**

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

    Answer: [N/A]

    Justification: The paper does not use existing datasets, code packages, or models as research assets. Existing scientific works are cited in the bibliography.

13. **New assets**

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

    Answer: [N/A]

    Justification: The paper does not introduce or release new datasets, code, models, or other assets.

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

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

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or human-subject research.

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

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

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or human-subject research.

16. **Declaration of LLM usage**

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

Answer: [N/A]

Justification: LLMs are not used as an important, original, or non-standard component of the research methodology. Any use for writing, editing, or formatting does not affect the core theoretical results.