

Unconditional Quantum PAC Learning Advantage in the Large Error Regime

Igor C. Oliveira*

Department of Computer Science
University of Warwick

October 16, 2026

Abstract

We establish an unconditional super-polynomial separation between quantum and classical PAC learning of Boolean functions under the uniform distribution in the large error regime. For every advantage $\lambda(n) = n^{-\omega(1)}$, we construct a concept class that a polynomial-time quantum learner learns to agreement $1/2 + \lambda(n)$ with success probability at least $1 - 1/n^k$. The learner uses only quantum random examples and outputs a polynomial-size deterministic classical circuit. In contrast, no polynomial-time classical learner achieves agreement $1/2 + \lambda(n)/n^k$ with success probability greater than $1/n^k$, even with membership queries.

To achieve this, we develop a nearly-optimal quantum learner that uses quantum random examples to produce a hypothesis achieving a prescribed correlation with an arbitrary Boolean function. The learner generalizes Fourier sampling from parity hypotheses to quadratic hypotheses. While standard Fourier sampling is insensitive to the sign of the correlation, we use interference to bias sampling toward positively correlated hypotheses. This is combined with a careful analysis of the resulting distribution to boost both the learner's success probability and its agreement with the target function. We believe that the resulting quantum learning algorithm might be of independent interest.

Finally, although a robust quantum PAC learning advantage holds unconditionally in the large error regime, we note that an analogous separation between quantum and classical learning in the small error regime would imply that $\text{NP} \not\subseteq \text{BPP}$.

*E-mail: igor.oliveira@warwick.ac.uk

Contents

1	Introduction	3
1.1	Results	5
1.2	Techniques	8
1.3	Related Work	12
2	Preliminaries	14
2.1	Basic Notation	14
2.2	Classical Learning Theory	15
2.3	Quantum Learning Algorithms	15
2.4	Probability	16
2.5	4-Wise Independence	17
2.6	Cryptography	18
3	Quantum Learning Algorithms: Proof of Theorem 1.2	19
3.1	Technical Background	19
3.1.1	Quadratic Functions: A Suitable 4-Wise Independent Family	19
3.1.2	Derandomized Classical Learner with $p^* \leq 1/2$	21
3.2	The Correlation-Biased Quantum Sampler: Proof of Item 1	22
3.2.1	Hypothesis Family $\mathcal{H}$ and Isometry W	22
3.2.2	Quantum Algorithm $\mathcal{Q}$	23
3.2.3	Correctness	25
3.3	Boosting Correlation and Success Probability: Proof of Item 2	27
4	Limitations of Classical Learners: Proof of Theorem 1.4	30
4.1	Unconditional Information Theoretic Lower Bound for $\mathcal{C} = \text{ALL}$: Proof of Item 1	30
4.2	Conditional Computational Lower Bound for $\mathcal{C} = \text{SIZE}[\text{poly}]$: Proof of Item 2	31
4.2.1	Strong Pseudorandom Functions	31
4.2.2	Lower Bound	32
5	Quantum PAC Learning Advantage: Proof of Theorem 1.5	33
5.1	Refined Query Lower Bound for Padded Functions	34
5.2	Unconditional Separation	36
5.3	Conditional Separation for $\mathcal{C} \subseteq \text{SIZE}[\text{poly}]$	37
A	Fourier-Based Quantum Learner for $\mathcal{C} = \text{ALL}$ with $p^* \leq 1/2$	40
B	Orienting a Weak Hypothesis is Hard	42
C	Limits of Parity Functions and of Pairwise Independence	42

1 Introduction

We study quantum learning of Boolean functions under the uniform distribution in the *large error regime*, where the required agreement with the target function is only slightly greater than $1/2$. As explained below, even an exponentially small advantage over random guessing can be significant in this setting, both as a source of quantum learning separations and as a potential tool for proving complexity lower bounds.

Synergies Between Learning and Complexity. The latter motivation comes from a general connection established by [OS17, AGG⁺21]. They showed that, for any typical Boolean circuit class $\mathcal{C}$ and choice of the parameter γ , learning every polynomial size $\mathcal{C}$ -circuit under the uniform distribution with agreement $1/2 + \gamma$ in quantum time

$$T(n) \leq \frac{\gamma(n)^2 2^n}{n^{\omega(1)}} \quad (1)$$

implies $\text{BQE} \not\subseteq \mathcal{C}[n^k]$ for every fixed k , where $\text{BQE} \triangleq \text{BQTIME}[2^{O(n)}]$, and $\mathcal{C}[n^k]$ denotes the class of functions computable by $\mathcal{C}$ -circuits of size n^k . This provides an approach to complexity lower bounds via quantum algorithm design.¹ We note that a precedent for obtaining circuit lower bounds from algorithms is Williams' celebrated proof [Wil14] that $\text{NE} \not\subseteq \text{ACC}^0$ (poly-size bounded-depth circuits extended with mod m gates), which uses faster classical algorithms for circuit satisfiability. In contrast, it remains open to separate BQE from ACC^0 , or to obtain any explicit super-polynomial circuit size lower bound against TC_2^0 (poly-size depth-2 circuits extended with linear threshold gates).

In the setting of *polynomial time*, the aforementioned learning connection singles out the *exponentially small* correlation regime: a learning agreement of just $1/2 + n^{\omega(1)} \cdot 2^{-n/2}$ would suffice for such circuit lower bounds. This suggests the following question.

Can classical or quantum algorithms running in polynomial time learn expressive concept classes under the uniform distribution with advantage $n^{\omega(1)} \cdot 2^{-n/2}$?

Achieving even this extremely weak learning guarantee remains an intriguing open problem for prominent circuit classes such as ACC^0 and TC_2^0 .

The discussion above motivates the investigation of the power of classical and quantum learners in the large error regime. Our main contribution is an unconditional separation between quantum and classical PAC learning at exponentially small advantage $\Omega(2^{-n/2})$ and, through additional ideas, a robust extension of this separation at inverse superpolynomial advantages $n^{-\omega(1)}$. To achieve this, we develop a quantum learning algorithm that might be of independent interest. Finally, in contrast to the large error regime, we observe that obtaining a similar quantum vs classical learning separation in the small error regime requires complexity-theoretic assumptions.

Quantum Learning Background. We consider the PAC learning model under the uniform distribution. We work with uniform quantum random examples of an unknown Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, namely copies of

$$|\text{QEX}_f\rangle = 2^{-n/2} \sum_{x \in \{0, 1\}^n} |x\rangle |f(x)\rangle.$$

Our quantum algorithms use only these examples and make no quantum membership queries. The quantum learner outputs a deterministic classical circuit h .

¹From this perspective, quantum learning is viewed merely as a tool for separating uniform quantum computation from non-uniform Boolean circuit classes. Its usefulness for complexity theory need not depend on practical considerations of the quantum learning model, and the resulting lower bound statement makes no reference to learning algorithms.

For Boolean functions $f, h: \{0, 1\}^n \rightarrow \{0, 1\}$, define their *agreement under the uniform distribution* by

$$\text{agreement}(f, h) \triangleq \Pr_{x \sim \mathcal{U}_n} [f(x) = h(x)].$$

We call γ an *advantage* when the agreement is at least $1/2 + \gamma$.

It is important to distinguish two probabilities: the *agreement probability* $\text{agreement}(f, h)$ defined above, which captures the error of the hypothesis h with respect to the target function f , and the *success probability* of the learning algorithm, i.e., the probability that the learner produces a hypothesis with the required agreement.

For additional context and background on quantum learning theory, we refer to [AdW17].

Remark 1.1 (Quantum Fourier Sampling). Fourier sampling is a central tool for quantum learning under the uniform distribution. It allows a learner to use uniform quantum random examples to sample efficiently from the squared Fourier spectrum of the target’s sign representation $F(x) \triangleq (-1)^{f(x)}$. As highlighted by [AGS21], Fourier sampling is a common ingredient in several quantum learning algorithms. For instance, this technique and its extensions underlie quantum algorithms for learning parities [BV97], DNF formulas under the uniform and bounded product distributions [BJ98, KRS19], juntas [AS07], and functions with sparse Fourier spectra [ACL⁺19]. These applications make Fourier sampling a natural starting point for investigating the power of quantum examples in the large error regime.

For arbitrary Boolean targets, Fourier sampling already gives a polynomial time learner achieving advantage $\Omega(2^{-n/2})$ with constant success probability [AGG⁺21, Appendix A] (see Section A for an exposition). With constant probability, the sampled parity has a Fourier coefficient of sufficiently large magnitude. However, the sampling probability depends on the square of that coefficient and therefore does not reveal its sign. Even when the magnitude is sufficiently large, randomly choosing between the parity and its negation succeeds only half the time. Thus the resulting learner’s success probability is at most $1/2$. Recovering the sign of an arbitrary hypothesis with exponentially small correlation is itself costly, even with quantum membership queries (see Section B).

Our Approach in a Nutshell. At this scale, the *success probability* is essential to establishing a quantum learning advantage. As we explain below, a classical polynomial-time algorithm making no queries can also achieve advantage $\Omega(2^{-n/2})$ with constant success probability below $1/2$. In contrast, we show that a quantum learner can cross the success threshold $1/2$ in polynomial time, while any classical learner meeting the same guarantee requires $\Omega(2^{n/2})$ membership queries, regardless of its running time. This establishes an *exponential separation* between quantum and classical learning at exponentially small advantage.

To obtain this guarantee, we develop a learner that generalizes Fourier sampling from parity hypotheses to quadratic hypotheses and uses interference to bias the sampling distribution toward positively correlated hypotheses. This allows us to show that a single quantum random example suffices to achieve the required advantage with constant success probability strictly greater than $1/2$.

Through a careful analysis of the distribution of quadratic functions sampled by the polynomial-time quantum learner, it is possible to boost the learner’s success probability and its agreement with the target function using additional random examples and larger running time. In turn, this enables us to employ a standard padding argument to translate the exponential separation between quantum and classical learners to inverse superpolynomial advantages $n^{-\omega(1)}$ in the setting of polynomial-time learning algorithms.

The rest of this introduction is organized as follows. We formally state our results in Section 1.1 below. A detailed overview of our techniques appears in Section 1.2. Finally, we discuss related work in Section 1.3.

1.1 Results

In this section, we explain our results in detail. Let ALL_n denote the class of all Boolean functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$.

Quantum Upper Bounds. In our first result, we show how to modify the quantum Fourier sampling procedure in order to obtain a polynomial-time quantum learner that succeeds with probability $p^* > 1/2$. In addition, we show that the agreement and success probabilities can be boosted using additional quantum random examples and larger running time.

Theorem 1.2 (Quantum Learning Algorithms). *There are universal constants $a_0, a_1, C > 0$ and $p^* > 1/2$ such that the following statements hold.*

- (i) *There is a quantum algorithm $\mathcal{Q}$ that, given $n \in \mathbb{N}$ and a single quantum random example for $f \in \text{ALL}_n$, runs in time $\text{poly}(n)$ and outputs a deterministic classical Boolean circuit h of size $O(n^2)$ such that*

$$\Pr_{\mathcal{Q}} \left[\text{agreement}(f, h) \geq \frac{1}{2} + a_0 \cdot 2^{-n/2} \right] \geq p^*.$$

Concretely,² one may take $a_0 = 10^{-3}$ and $p^* = 0.53$.

- (ii) *There is a quantum algorithm $\mathcal{Q}^+$ that, given $n \in \mathbb{N}$ and $\eta \in (0, 1/2)$ satisfying $\eta \geq C \cdot 2^{-n}$, and using $O(1/\eta)$ independent quantum random examples for $f \in \text{ALL}_n$, runs in time*

$$T(n, \eta) = \text{poly}(n) \cdot \frac{1}{\eta}$$

and outputs a deterministic classical Boolean circuit h of size at most $T(n, \eta)$ such that³

$$\Pr_{\mathcal{Q}^+} \left[\text{agreement}(f, h) \geq \frac{1}{2} + \frac{a_1}{\sqrt{\eta}} \cdot 2^{-n/2} \right] \geq 1 - \eta.$$

Theorem 1.2 Item (ii) gives a tradeoff between accuracy, confidence, and running time. In terms of the number m of quantum random examples, we obtain advantage $\Omega(\sqrt{m}/2^n)$, failure probability $O(1/m)$, and running time $m \cdot \text{poly}(n)$. Thus an advantage γ can be obtained in time $\text{poly}(n) \cdot \gamma^2 \cdot 2^n$ throughout the range from a constant multiple of $2^{-n/2}$ to a sufficiently small positive constant.⁴

Remark 1.3 (On the Optimality of the Parameters in **Theorem 1.2** Item (ii)). Relative to the criterion of [AGG⁺21] for complexity lower bounds corresponding to Equation (1), the amplified quantum learner achieves running time $\text{poly}(n) \cdot \gamma^2 \cdot 2^n$ as opposed to $n^{-\omega(1)} \cdot \gamma^2 \cdot 2^n$. This dependence on γ is essentially optimal for ALL_n when the learner must output a classical circuit. Indeed, by a standard Chernoff bound, a fixed hypothesis h agrees with a uniformly random Boolean function on a $1/2 + \gamma$ fraction of inputs with probability at most $\exp(-2\gamma^2 2^n)$, and a simple counting argument shows that covering every target in ALL_n requires $\Omega(\gamma^2 2^n)$ bits of hypothesis description.

²We do not optimize the constant p^* in our presentation in order to keep the quantum algorithm and its analysis simpler.

³The hypothesis h is a depth-3 Boolean circuit of the form $\text{MAJ}_{O(1/\eta)} \circ \oplus_{O(n^2)} \circ \text{AND}_2$, where subscripts denote the fan-in of each gate type.

⁴For larger advantages, learning the whole truth table from $\text{poly}(n) \cdot 2^n$ measured examples gives the same time bound up to polynomial factors.

Classical Lower Bounds. We show next that classical learners cannot match the guarantee in [Theorem 1.2](#) Item (i) in subexponential time. We prove an unconditional query lower bound when the target is an arbitrary Boolean function, and a conditional computational lower bound for targets computed by poly-size circuits.

For a size bound $s: \mathbb{N} \rightarrow \mathbb{N}$, we let $\text{SIZE}[s(n)]$ denote the set of Boolean functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$ computed by a Boolean circuit of size at most $s(n)$.

The computational lower bound uses the standard cryptographic hypothesis that one-way functions remain secure against subexponential-time classical algorithms. This assumption and its relation to strong pseudorandom functions appear, for example, in the natural proofs framework [\[RR97\]](#) and in the study of quantum advantages relative to efficiently computable oracles [\[AC17, Section 7\]](#).

Formally, we say that a uniform family of functions $f_n: \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ computable in deterministic polynomial time is a *subexponentially secure one-way function* if there is a constant $\varepsilon > 0$ such that, for every (classical) randomized algorithm $\mathcal{I}$ running in time $\leq 2^{n^\varepsilon}$ and for every large enough n ,

$$\Pr_{x \sim \mathcal{U}_n, y=f_n(x), \mathcal{I}}[\mathcal{I}(1^n, y) \in f_n^{-1}(y)] \leq n^{-\omega(1)}.$$

Theorem 1.4 (Limitations of Classical Learners). *The following statements hold.*

- (i) *Let $\mathcal{A}$ be a randomized algorithm that makes at most q membership queries to $f \in \text{ALL}_n$ and outputs a deterministic classical Boolean circuit h .⁵ For every $\gamma > q/2^{n+1}$, there is a function $f \in \text{ALL}_n$ such that*

$$\Pr_{\mathcal{A}} \left[\text{agreement}(f, h) \geq \frac{1}{2} + \gamma \right] \leq \frac{1}{2}.$$

In particular, when learning ALL_n it is not possible to achieve correlation $\Omega(2^{-n/2})$ with success probability $> 1/2$ using $o(2^{n/2})$ classical membership queries.

- (ii) *Assume the existence of subexponentially secure one-way functions. There is $k \geq 1$ such that, for every pair of constants $a_0 > 0$ and $p^* > 1/2$, no randomized algorithm $\mathcal{A}$ has the following property:*

For every large enough $n \in \mathbb{N}$ and every $f \in \text{SIZE}[n^k]$, the algorithm runs in time $o(2^{n/2})$, makes classical membership queries, and outputs a deterministic classical Boolean circuit h satisfying

$$\Pr_{\mathcal{A}} \left[\text{agreement}(f, h) \geq \frac{1}{2} + a_0 \cdot 2^{-n/2} \right] \geq p^*.$$

In other words, under the existence of subexponentially secure OWFs, it is not possible to achieve correlation $\Omega(2^{-n/2})$ and constant success probability $p^* > 1/2$ for every $f \in \text{SIZE}[n^{O(1)}]$ in time $o(2^{n/2})$.

Quantum PAC Learning Advantage. Together, the two theorems give an *unconditional* exponential separation between one quantum random example and $\Omega(2^{n/2})$ classical membership queries, and a conditional separation between quantum polynomial time and classical time $o(2^{n/2})$ for a fixed polynomial-size circuit class. Note that the computational separation requires no specific number-theoretic assumption, such as the hardness of factoring or computing discrete logarithms. In particular, our results show that an exponential quantum vs classical PAC learning separation for a concept class $\mathcal{C} \subseteq \text{SIZE}[\text{poly}]$ can be obtained from a *generic cryptographic assumption*.

Next, we build on the results described above to translate the exponential quantum advantage at the exponentially small correlation regime into a robust polynomial advantage in the agreement regime $1/2 + \lambda(n)$, for any “constructive” function $\lambda(n) = 1/n^{\omega(1)}$. We say that a function $\lambda: \mathbb{N} \rightarrow \mathbb{R}^+$ is constructive

⁵This statement places no restriction on the running time of $\mathcal{A}$ or on the size of h .

if $\lambda(n)$ can be computable from 1^n in time $\text{poly}(n)$.⁶ For example, [Theorem 1.5](#) below applies to $\lambda(n) = n^{-\log n}$, $\lambda(n) = n^{-\log \log n}$, $\lambda(n) = n^{-\log^* n}$, etc.

Theorem 1.5 (Robust Quantum PAC Learning Advantage at Inverse Superpolynomial Correlation). *Let $\lambda(n)$ be any constructive function satisfying $n^{\omega(1)} \cdot 2^{-n/2} \leq \lambda(n) \leq n^{-\omega(1)}$, and let $k \geq 1$ be an arbitrary constant. Then there is a concept class $\mathfrak{C} = \{\mathcal{C}_n\}_{n \geq 1}$ for which the following holds.*

- *There is a quantum algorithm $\mathcal{A}$ that, for every sufficiently large n and every $f \in \mathcal{C}_n$, uses quantum random examples, runs in $\text{poly}(n)$ time, and outputs a deterministic classical circuit h satisfying*

$$\Pr_{\mathcal{A}} \left[\text{agreement}(f, h) \geq \frac{1}{2} + \lambda(n) \right] \geq 1 - \frac{1}{n^k}.$$

- *For every randomized membership-query polynomial-time learner $\mathcal{B}$, for every sufficiently large n there is $f \in \mathcal{C}_n$ such that*

$$\Pr_{\mathcal{B}} \left[\text{agreement}(f, h) \geq \frac{1}{2} + \frac{\lambda(n)}{n^k} \right] \leq \frac{1}{n^k}.$$

Moreover, if subexponentially secure one-way functions exist, the concept class $\mathfrak{C}$ can be chosen to satisfy $\mathfrak{C} \subseteq \text{SIZE}[\text{poly}]$.

To our knowledge, [Theorem 1.5](#) establishes the first unconditional super-polynomial separation in the time complexity of PAC learning Boolean functions under the uniform distribution. The quantum learner runs in polynomial time, uses only quantum random examples, and outputs a classical hypothesis, while no polynomial-time classical learner achieves the same guarantee even with membership queries. Moreover, the separation is robust, in the sense that a polynomial-time classical learner is not able to achieve a weaker agreement guarantee under a significantly lower success probability.

Remark 1.6 (The Small Error Regime). We remark that obtaining a similar robust separation in the agreement regime $1/2 + \lambda(n)$, where $\lambda(n) \geq 1/\text{poly}(n)$, with classical circuit hypotheses would imply that $\text{NP} \not\subseteq \text{BPP}$. In this sense, [Theorem 1.5](#) is essentially optimal for unconditional separations of this form, given our current knowledge of complexity theory.

The reason is that polynomial-time agnostic learning of polynomial-size circuits is possible under the average-case easiness assumption $\text{DistNP} \subseteq \text{AvgBPP}$ [[GKLO22](#)]. We use two features of our separation: the quantum learner outputs a deterministic classical circuit, and classical hardness persists at the lower agreement $1/2 + \lambda(n)/2$.

In more detail, let $\mathfrak{C} = \{\mathcal{C}_n\}_{n \geq 1}$ be an arbitrary class of Boolean functions. Suppose that $\mathfrak{C}$ admits a polynomial-time quantum learner that, for every sufficiently large n and every $f \in \mathcal{C}_n$, outputs a deterministic classical circuit h with $\text{agreement}(f, h) \geq 1/2 + \lambda(n)$ with probability at least $2/3$, where $\lambda(n) \geq 1/n^a$ for some fixed $a > 0$. The running time bound implies that, for some fixed $\ell > 0$, every such f admits a circuit D_f of size at most n^ℓ achieving this agreement. Under $\text{DistNP} \subseteq \text{AvgBPP}$, [[GKLO22](#), Theorem 56] gives a classical agnostic learner for $\text{SIZE}[n^\ell]$ under the uniform distribution, running in time $\text{poly}(n, 1/\varepsilon)$. Using membership queries to obtain uniform labeled examples, and choosing $\varepsilon = 1/(2n^a)$, this learner outputs, with probability at least $2/3$, a hypothesis whose agreement with f is at least the maximum agreement achievable by a circuit in $\text{SIZE}[n^\ell]$ minus ε , and therefore at least $1/2 + \lambda(n)/2$. This contradicts the robust classical lower bound. Thus such a separation implies $\text{DistNP} \not\subseteq \text{AvgBPP}$, and hence $\text{NP} \not\subseteq \text{BPP}$.

⁶It is enough that there is an efficient algorithm that approximates $\lambda(n)$. We refer to [Section 5](#) for the details.

1.2 Techniques

We now describe our main ideas and techniques in more detail, including how they can be combined to derive [Theorem 1.5](#). First, we review quantum Fourier sampling and its limitations, then explain how to achieve a similar guarantee using a classical algorithm. While neither approach achieves success probability higher than $1/2$, these considerations reveal useful insights that are explored in our results.

Review of Quantum Fourier Sampling and Its Limitations. As mentioned above, a standard technique in quantum learning theory is *Fourier sampling*. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be an unknown function, and suppose we have access to a quantum random example

$$|\text{QEX}_f\rangle = \frac{1}{\sqrt{N}} \sum_{x \in \{0, 1\}^n} |x\rangle |f(x)\rangle,$$

where $N \triangleq 2^n$. Let $F(x) \triangleq (-1)^{f(x)}$ be the sign representation of f . Then, an efficient computation allows us to produce the state

$$\sum_{S \in \{0, 1\}^n} \hat{F}(S) |S\rangle,$$

where $\hat{F}(S) \in [-1, 1]$ is the *correlation* between F and the parity function χ_S . This means that a measurement in the computational basis samples a parity $\chi_S(x) = (-1)^{S \cdot x}$ with probability $\hat{F}(S)^2$. Using that the $N = 2^n$ squared Fourier coefficients sum to 1, it is not hard to show that the sampled parity has correlation of magnitude $\Omega(N^{-1/2})$ with constant probability. We refer to [Section A](#) for the details.

The difficulty is its *orientation*: the parity may be positively or negatively correlated with the target. Squaring the Fourier coefficient erases its sign, and the phase states for F and $-F$ differ only by a global phase. Consequently, the sampling outcome cannot distinguish these two targets. Choosing uniformly between the sampled parity and its negation gives the required positive correlation with constant success probability, but this probability is always at most $1/2$. Moreover, at exponentially small correlation, recovering the correct orientation by a separate test is expensive even with quantum membership queries (see [Section B](#)).

An Efficient Classical Learner with Success Probability at Most $1/2$. The same correlation scale can be reached classically *without any queries*: simply output a uniformly random truth table. For any fixed target, its correlation is the average of N independent unbiased signs. The sum fluctuates on the scale $\sqrt{N}$, so positive correlation $\Omega(N^{-1/2})$ occurs with constant probability. However, symmetry again bounds this success probability by $1/2$. Moreover, a random truth table requires $N = 2^n$ bits to describe, so this observation alone does not give an *efficient learner* with a succinct hypothesis.

To obtain a polynomial-time learner, we (partially) derandomize the truth table construction using limited independence, a standard technique from the theory of computational pseudorandomness (see, e.g., [\[Vad12\]](#)). We say that a family $\mathcal{H} = \{h_a\}_{a \in \{0, 1\}^r}$ of functions $h_a: \{0, 1\}^m \rightarrow \{0, 1\}$ is *k-wise independent* if, for a uniformly random seed $a \sim \{0, 1\}^r$, its values at any set of at most k distinct inputs are independent uniform bits. While the second moment identifies the natural scale of correlation, it could be carried by a few exceptionally large values while almost all hypotheses have correlation zero. In contrast, one can show that a fourth moment bound rules out this scenario, ensuring that useful magnitudes occur with constant probability. The important point is that full independence is unnecessary for these calculations: a 4-wise independent family $\mathcal{H}$ suffices to reproduce both moments.

In more detail, for $H_a = (-1)^{h_a}$ and a fixed target F , it is possible to show using the 4-wise independence of $\mathcal{H}$ that the normalized correlation Z_a satisfies

$$\mathbb{E}_a[Z_a^2] = 1 \quad \text{and} \quad \mathbb{E}_a[Z_a^4] \leq 3, \quad \text{where} \quad Z_a \triangleq \sqrt{N} \operatorname{corr}(H_a, F) = \frac{1}{\sqrt{N}} \sum_x F(x) H_a(x).$$

Then, applying the Paley-Zygmund Inequality (Lemma 2.1) to Z_a^2 gives a constant lower bound on $\Pr_a[|Z_a| \geq 1/2]$. If the family also contains each hypothesis and its negation with equal weight, half of this probability comes from *positive* correlations.

An explicit 4-wise independent family with polynomial length seeds and polynomial size circuits therefore gives an efficient classical learner making zero queries, whose success probability is bounded below by a positive constant and above by $1/2$ (Proposition 3.2). While the learner still chooses a polynomial size random seed, this partial derandomization reduces the randomness and hypothesis size, as desired.

Before we describe how to go beyond success probability $1/2$ for advantage $\gamma = \Omega(2^{-n/2})$ using a quantum learner, we briefly explain why a classical randomized learner that makes subexponentially many membership queries cannot attain this guarantee.

Why $o(2^{n/2})$ Classical Queries Cannot Cross the $1/2$ Barrier. The proof of Theorem 1.4 Item (i) follows a standard query lower bound argument. As above, we use F and H to denote the sign representations of Boolean functions f and h , respectively. Choose the target f uniformly from ALL_n , fix the learner's random coins, and condition on its complete query transcript. Even when the queries are adaptive, the labels at unqueried inputs remain independent uniform bits, and the output hypothesis h is fixed. The queried inputs contribute at most q to the sum defining $N \cdot \operatorname{corr}(H, F)$, and the remaining contribution is a sum of independent unbiased signs, hence symmetric about zero.

Achieving agreement $1/2 + \gamma$ requires this total sum to be at least $2N\gamma$. If $2N\gamma > q$, success therefore requires the unseen contribution to be strictly positive, an event of probability at most $1/2$. Averaging over transcripts, coins, and targets yields a target on which the learner's success probability is at most $1/2$. In particular, $q = o(\sqrt{N})$ cannot support a guarantee above $1/2$ at advantage $\Omega(N^{-1/2})$. Intuitively, the observed labels can shift the correlation by only q/N , which is too small to overcome the uncertainty on the relevant scale.

Correlation-Biased Quantum Sampling and Quadratic Hypotheses. To prove Theorem 1.2 Item (i), we combine the succinct hypotheses of the classical learner with an extension of the ideas behind Fourier sampling.

The key is to retain the known branch that the usual Fourier procedure discards when preparing a phase state. To make this explicit, write

$$|u\rangle \triangleq \frac{1}{\sqrt{N}} \sum_x |x\rangle, \quad |F\rangle \triangleq \frac{1}{\sqrt{N}} \sum_x F(x) |x\rangle.$$

Applying a Hadamard gate to the label qubit of a quantum example gives

$$|\text{QEX}_f\rangle \mapsto \frac{1}{\sqrt{2}}(|u\rangle|0\rangle + |F\rangle|1\rangle).$$

The branch labelled 0 is the known uniform state, independent of the target, while the branch labelled 1 contains the target values as phases. The usual Fourier procedure measures the label qubit and keeps the

outcome 1, which occurs with probability $1/2$. It then uses the isolated phase state $|F\rangle$. Since $|-F\rangle = -|F\rangle$, replacing the target by its complement changes this state only by a global phase, which cannot affect any subsequent measurement statistics.

In the joint state, however, replacing F by $-F$ changes only the second summand: the reference branch $|u\rangle|0\rangle$ stays fixed. The sign change is therefore a *relative phase* between the two branches, rather than a global phase of the entire state. This preserves the distinction between the two orientations, but retaining the branches is only the first step. To make that distinction affect the output distribution, we must keep the label qubit coherent and later transform both branches into a common measurement basis, so that their amplitudes interfere. Measuring the label in the computational basis at this stage, or discarding it, would remove the coherence between the branches and lose this information.

The known branch is especially convenient as a coherent reference because it can be transformed into a single fixed basis state. Place the label qubit first and apply a Hadamard transform to the input register only when the label is 0. This sends $|0\rangle|u\rangle$ to $|0\rangle|0^n\rangle$, while leaving the target branch unchanged. Calling this fixed reference state $|\perp\rangle$ and writing $|x\rangle_{\text{lab}} = |1\rangle|x\rangle$, we obtain

$$\frac{1}{\sqrt{2}}|\perp\rangle + \frac{1}{\sqrt{2N}} \sum_x F(x)|x\rangle_{\text{lab}}.$$

We note that this preparation uses the original quantum example without any measurement or additional access to the target Boolean function.

Jumping ahead, a suitable quantum transform W described below maps the entire label space into a hypothesis register, with no separate branch label retained in an unobserved register. The reference and target contributions to each hypothesis amplitude can then add coherently, allowing their relative sign to favor the positively correlated orientation.

In more detail, our hypothesis family $\mathcal{H} = \{h_a\}_{a \in \{0,1\}^r}$ consists of all quadratic Boolean polynomials

$$h_a(x) = a_\emptyset \oplus \bigoplus_i a_i x_i \oplus \bigoplus_{i < j} a_{ij} x_i x_j.$$

A seed contains $r = 1 + n + \binom{n}{2}$ coefficients, and each hypothesis has a circuit of size $O(n^2)$. While affine functions provide pairwise independence, it is possible to show that the quadratic function family $\mathcal{H}$ is 4-wise independent. The details appear in [Section 3.1.1](#).

Write $M = 2^r$ for the number of seeds. The aforementioned quantum transform W converts the reference branch into a uniform superposition of seeds and the target branch into a superposition whose amplitude at seed a is proportional to the normalized correlation Z_a defined above. Up to a common normalization, these contributions are 1 and Z_a . Keeping the two branches coherent adds their amplitudes before measurement, giving amplitude $(1 + Z_a)/\sqrt{2M}$ and hence the exact output distribution

$$P_F(a) = \frac{(1 + Z_a)^2}{2M}.$$

We note that the reference branch alone would sample hypotheses uniformly, just as the classical learner does. The target branch alone would assign weights proportional to Z_a^2 , generalizing Fourier sampling to quadratics but still erasing the sign. Together they produce $(1 + Z_a)^2 = 1 + 2Z_a + Z_a^2$. Crucially, the linear term $2Z_a$ favors positive correlations: for opposite hypotheses with normalized correlations z and $-z$, where $z > 0$, the weights differ by $(1 + z)^2 - (1 - z)^2 = 4z$. The correct orientation is therefore more likely to be sampled.

To complete the proof sketch of [Theorem 1.2](#) Item (i), we briefly discuss the implementation of W and the achieved success probability. The quadratic family $\mathcal{H}$ described above supports this transform efficiently using a Hadamard transform over its coefficient register. We refer to the full construction in

Section 3.2. The remaining issue is to show that the preference for positive correlations is strong enough to raise success above $1/2$. Here the fourth moment bound $\mathbb{E}_a[Z_a^4] \leq 3$ stated above has its second role: it controls the second moment of the positive part $X \triangleq \max\{Z_a, 0\}$ under the reweighted distribution P_F . Together with symmetry and $\mathbb{E}_a[Z_a^2] = 1$, it gives $\mathbb{E}_{P_F}[X] \geq 1$ and $\mathbb{E}_{P_F}[X^2] \leq 1 + \sqrt{3}/2 < 2$. The Paley-Zygmund Inequality then yields $\Pr_{P_F}[Z_a \geq c] > 1/2$ for a sufficiently small fixed $c > 0$, as desired. Thus the same moment bounds both supply useful classical candidates and turn interference into a quantum success guarantee above $1/2$.

Remark 1.7 (Limits of Parity Functions and of Pairwise Independence). Quadratic functions play an important role in the analysis of our learning algorithms. For the interested reader, in [Section C](#) we provide a concrete example to illustrate why the use of linear functions and pairwise independence is insufficient to achieve the guarantees on the success probabilities of the classical and quantum learners described above.

Amplifying Correlation and Success Probability. For [Theorem 1.2](#) Item (ii), we run the quantum learner independently on several examples and output the majority of its hypotheses. Note that a success probability above $1/2$ for the overall correlation does not by itself justify this step: hypotheses with positive correlation may still make systematic errors on the same inputs. Nevertheless, the quantum sampler provides the stronger property needed here. At *every input*, a sampled hypothesis has the same small positive bias toward the correct value, and its agreement signs at distinct inputs are independent in pairs. Repetition increases each pointwise bias, while independence across input pairs makes the overall correlation concentrate around its positive mean.

Specifically, under the sampler's output distribution $P_F(a)$, which corresponds to an output hypothesis h_a , the signs $Y_x \triangleq F(x)H_a(x)$ satisfy $\mathbb{E}_{P_F}[Y_x] = N^{-1/2}$, and Y_x, Y_y are independent whenever $x \neq y$. These properties follow from 4-wise independence of the original family $\mathcal{H}$ (see [Lemma 3.4](#)). Now take an odd number $m \leq N$ of independent runs. At a fixed input, majority turns m independent signs of mean $N^{-1/2}$ into a sign of mean $\Omega(\sqrt{m/N})$ (see [Lemma 3.5](#)). Independence across pairs of inputs can be shown to pass to their majority outputs, so the overall correlation has variance at most $1/N$. Its mean grows with m , while its variance stays bounded by the same quantity. Chebyshev's inequality consequently gives correlation $\Omega(\sqrt{m/N})$ with failure probability $O(1/m)$. Finally, choosing $m = \Theta(1/\eta)$ yields [Theorem 1.2](#) Item (ii), with time and circuit size $\eta^{-1}\text{poly}(n)$.

From Exponentially Small to Inverse Superpolynomial Advantage. We obtain [Theorem 1.5](#) by a standard padding argument and an extension of the classical lower bound. Set $r_n = 2^{k\lceil \log_2 n \rceil}$, so that $n^k \leq r_n = n^{O(1)}$. For the desired advantage $\lambda(n)$, the concept class consists of target functions $f_g(x, z) = g(x)$ that depend on only d of their n input bits, with $2^d = \Theta(r_n^3 \lambda(n)^{-2})$. The assumptions on λ ensure that $d = \omega(\log n)$ and $d < n$ eventually, while constructivity allows d to be computed efficiently.

For the quantum upper bound, it is not hard to see that quantum random examples for g can be obtained from quantum random examples for f_g . We apply the amplified quantum learner ([Theorem 1.2](#) Item (ii)) with $\eta = 1/(Lr_n^3)$, for a sufficiently large fixed constant L . Using $O(r_n^3)$ examples in polynomial time gives advantage $a_1 \sqrt{Lr_n^3/2^d} \geq \lambda(n)$, with success probability at least $1 - \eta \geq 1 - 1/n^k$.

Classically, the relevant transcript lower bound argument now concerns the 2^d independent ‘‘core’’ labels corresponding to a random function $g: \{0, 1\}^d \rightarrow \{0, 1\}$. A polynomial number of queries is negligible compared with $2^d \lambda(n)/r_n = \Theta(r_n^2/\lambda(n)) = n^{\omega(1)}$, and one can show that some target has success probability at most $1/2$ even at advantage $\lambda(n)/r_n \leq \lambda(n)/n^k$. This remains true even if the learner's hypothesis depends on all padding bits: averaging over those bits assigns each core input a weight in $[-1, 1]$, preserving symmetry of the unseen contribution.

To obtain a vanishing success probability, we also control the size of the unseen contribution. The

unknown core labels are independent and unbiased, so their contributions largely cancel, leaving fluctuations on a scale at most $2^{d/2}$. A simple calculation shows that our choice of parameters makes success at advantage $\lambda(n)/r_n$ require a favorable fluctuation larger than this scale by a factor of order $\sqrt{r_n}$. The polynomially many queried labels contribute too little to close this gap. Formally, using Chebyshev's inequality and our choice of constants it can be shown that the average success probability is at most $1/(16r_n) \leq 1/n^k$. Averaging therefore supplies a target with this bound, which also holds at the required advantage $\lambda(n)/n^k \geq \lambda(n)/r_n$.

Finally, it is possible to obtain the same quantum learning advantage for a concept class whose functions admit polynomial-size Boolean circuits, under a generic cryptographic assumption.

Transferring the Separation to Polynomial Size Concepts. It is well known that one-way functions (OWFs) yield pseudorandom function families (PRFs) with polynomial-size circuits [HILL99, GGM86]. Intuitively, a classical learner running within the security bound cannot distinguish a random function from a pseudorandom one, so its limitations on random targets should also apply to targets from the pseudorandom family. The quantum upper bounds already apply to every Boolean function, including these efficiently represented functions. The task is therefore to use the PRF family to transfer the classical lower bounds to a concept class contained in $\text{SIZE}[\text{poly}]$.

The subtlety here is *verification*. To turn a successful learner into a distinguisher, we must check whether its output has the promised positive correlation. An exponentially small correlation cannot in general be certified in polynomial time, so ordinary polynomial security does not justify this reduction. To address this issue, we use subexponential security and a change of scale dating back to the natural proofs argument [RR97], i.e., we choose the target's relevant input dimension sufficiently small relative to the cryptographic security parameter that exhaustive correlation testing fits within the security bound, while the resulting functions still have polynomial size circuits. A classical randomized learner crossing the success probability threshold would then give a distinguisher with suitable advantage. This idea and its extensions underlie Theorem 1.4 Item (ii) and the conditional part of Theorem 1.5.

1.3 Related Work

Servedio and Gortler [SG04] established polynomial relations between classical and quantum sample complexity for distribution-free PAC learning, and between their query complexities for exact learning with membership queries. For exact learning on n input bits, Q quantum queries can be simulated using $O(nQ^3)$ deterministic classical queries. Recent constructions match this bound and its sharper randomized analogue [AGM26]. For distribution-free PAC learning in the usual small error regime, Arunachalam and de Wolf [AdW18] proved that classical and quantum sample complexities coincide up to constant factors.

Under a fixed example distribution, advantages are possible. Uniform parity learning uses $O(1)$ quantum examples versus $\Omega(n)$ classical examples at fixed small error and constant confidence (see, e.g., [AdW17, Section 5.3]). With stronger oracle access, namely a quantum example preparation unitary and its inverse, a nearly quadratic improvement in the accuracy dependence is possible even for distribution-free learning [SSG24]. In contrast, our exponential sample complexity gap concerns the large error regime: under the uniform distribution, one quantum random example yields advantage $\Omega(2^{-n/2})$ with success probability above $1/2$, whereas exponentially many classical membership queries are necessary.

In the computational setting, Servedio and Gortler [SG04] proved a PAC learning separation assuming classical hardness of factoring Blum integers, i.e., they describe a concept class that is efficiently quantum learnable under arbitrary distributions but is not efficiently classically PAC learnable. Their quantum learner uses Shor's factoring algorithm and ordinary classical examples, and outputs a classical circuit. Liu,

Arunachalam, and Temme [LAT21] give a quantum kernel method that achieves 99% accuracy with success probability at least $2/3$ for a classification problem under the uniform distribution on $\mathbb{Z}_p^*$. Assuming the hardness of discrete logarithm, no classical polynomial-time learner achieves inverse polynomial advantage over random guessing with the same confidence. Jerbi et al. [JGM⁺24] also obtain a computational PAC learning advantage under the uniform distribution assuming the discrete cube root assumption, with quantum computation used during training and entirely classical evaluation of the resulting hypothesis.

These computational separations exploit arithmetic structure. Under the generic assumption that one-way functions exist, Servedio and Gortler [SG04] established a separation between quantum and classical exact learning from membership queries, where successful learning requires a hypothesis agreeing with the target on every input. This result does not establish a PAC learning separation from random examples.

Going beyond the standard PAC and exact learning frameworks, Gavinsky [Gav12] obtained an unconditional exponential separation between quantum and classical learning in a predictive model for relations, where an input may have several valid outputs. Their learner retains a quantum state to answer prediction queries, while our results concerns Boolean functions and outputs a classical hypothesis. Conditional quantum learning advantages have also been established for noisy parity and linear equations [CSS15, GKZ19], and for generative and density learning of classical distributions under cryptographic assumptions [SSHE21, PSES23]. For learning certain properties of quantum states and processes, unconditional exponential separations in sample or query complexity arise between learners with quantum memory and those retaining only classical information between samples or experiments [HKP21, CCHL22].

Our quantum sampler builds on Fourier sampling [BV97, BJ98], retaining the known branch of a labeled quantum example as a coherent reference, so that measurement probabilities contain a term linear in signed correlation. Limited independence supplies a succinct quadratic family with controlled moments, giving positive correlation with success probability strictly above $1/2$. A close precedent is van Dam’s approximate oracle interrogation [vD98], where one quantum query produces a string of N bits with expected agreement $1/2 + 1/(2\sqrt{N})$ with the oracle string. This construction already biases the output toward positive signed correlation. The main distinction is computational: van Dam’s algorithm outputs a full truth table of length $N = 2^n$, whereas our sampler produces a succinct classical hypothesis in time polynomial in n . The use of a coherent reference has conceptual precedents in phase retrieval, where a known reference makes intensity measurements sensitive to relative phase through interference terms [KH90, BSL⁺19]. A related use of interference appears in quantum classification, where measurement probabilities depend on inner products between encoded data vectors [SFP17].

The use of quadratic hypotheses connects our construction to quadratic Goldreich-Levin algorithms, initiated by Tulsiani and Wolf [TW14]. Using classical membership queries, recent algorithms find quadratic phases with absolute correlation at least $\tau - \varepsilon$ in time polynomial in n and quasi-polynomial in $1/\varepsilon$ [CSBA⁺26, Theorem 1.5], or at least τ^C in time polynomial in n and $1/\tau$, for a universal constant $C > 1$ [ADGG26, Theorem 1.3], where τ is the maximum absolute quadratic correlation. Even with classical membership queries, these bounds do not yield a polynomial-time learner at the exponentially small parameters considered here.

Acknowledgements. We thank Rahul Ilango for reminding us of the simple but insightful observation that outputting a random function gives a classical learner with a guarantee similar to the quantum learner based on Fourier sampling [IO22]. Appendix A of [AGG⁺21] and this observation served as the starting point for our investigation. We also thank Tom Gur for an email exchange on previous work related to quantum learning advantage.

Use of AI Tools. Since [AGG⁺21] and [IO22], we had believed, perhaps naively, that the uncertainty in the sign of the hypothesis produced by Fourier sampling was an inherent limitation of any efficient quantum learner, just as it is for randomized learners. Generative AI led us to recognize that this intuition was incorrect. The central idea behind the correlation-biased quantum sampler, which combines a derandomized classical learner based on anti-concentration with a variant of quantum Fourier sampling, was discovered by GPT-5.5. We also used GPT-5.5, GPT-5.6, GPT-6, and GPT-6.1 extensively to explore several related directions, prepare initial drafts of the proofs of auxiliary results, and for general assistance with writing and verification. These tools materially affected all sections of the paper. We take full responsibility for the correctness of the paper and exposition of the results.

Funding Information. This work received support from the UKRI Frontier Research Guarantee Grant EP/Y007999/1 and the Centre for Discrete Mathematics and its Applications (DIMAP) at the University of Warwick. The author would also like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme Logical Foundations of Complexity Theory (LFC), where part of the work on this paper was undertaken. The LFC programme was supported by EPSRC grant EP/Z000580/1.

2 Preliminaries

2.1 Basic Notation

For a given $n \geq 1$ we will implicitly use $N = 2^n$, unless stated otherwise.

Probabilities. We use $x \sim S$ to denote that x is sampled uniformly from the set S . We often use $\mathcal{U}_n$ to denote the uniform distribution over $\{0, 1\}^n$. We will sometimes use boldface letters to emphasise a random choice, e.g., to distinguish each function h_a in a family indexed by $a \in A$ and a random function h_a , where $a \sim A$. For an event E , we use $\mathbf{1}_E$ to denote its indicator random variable. We use E^c to denote the complement of an event E .

Function Classes. Let ALL_n denote the class of all functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. For a size bound $s: \mathbb{N} \rightarrow \mathbb{N}$, we let $\text{SIZE}[s(n)]$ denote the set of Boolean functions computed by a Boolean circuit of size at most $s(n)$.

Representation of Boolean Functions. We will often use lowercase letters for Boolean-valued functions and the corresponding uppercase letters for their sign representations. Thus a target

$$f: \{0, 1\}^n \rightarrow \{0, 1\}$$

has *sign representation*

$$F: \{0, 1\}^n \rightarrow \{-1, 1\}, \quad F(x) = (-1)^{f(x)}.$$

Similarly, a Boolean-valued hypothesis $h: \{0, 1\}^n \rightarrow \{0, 1\}$ has sign representation $H(x) = (-1)^{h(x)}$.

Agreement and Correlation. We let

$$\text{corr}(H, F) \triangleq \mathbb{E}_{x \sim \mathcal{U}_n} [H(x)F(x)] \in [-1, 1]$$

denote the *correlation* between H and F . Agreement and correlation are related by

$$\Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] = \frac{1}{2} + \frac{1}{2} \text{corr}(H, F). \quad (2)$$

2.2 Classical Learning Theory

In this section, we review the standard PAC learning model, fix notation, and make our conventions explicit. We refer to a textbook such as [KV94] for additional information.

In the PAC learning framework, the learning target is an unknown Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, called a *concept*. A *concept class* $\mathcal{C}$ is a family of such functions. For each $n \in \mathbb{N}$, we write

$$\mathcal{C}_n \triangleq \mathcal{C} \cap \{f: \{0, 1\}^n \rightarrow \{0, 1\}\}$$

for the restriction of $\mathcal{C}$ to inputs of length n .

Learnability is defined relative to a class $\mathcal{D}$ of distribution ensembles. Each ensemble $D = \{D_n\}_{n \in \mathbb{N}}$ consists of probability distributions D_n supported on $\{0, 1\}^n$. Given a target concept $f \in \mathcal{C}_n$ and a distribution D_n , a *labelled example oracle* $\text{EX}(f, D_n)$ returns, on each invocation, an independent pair $(x, f(x))$, where x is sampled according to D_n .

PAC Learning Model. Let $\mathcal{C}$ be a concept class and let $\mathcal{D}$ be a class of distribution ensembles. An algorithm A *PAC learns* $\mathcal{C}$ over $\mathcal{D}$ if, for every $D \in \mathcal{D}$, every $\varepsilon, \delta \in (0, 1)$, every sufficiently large $n \in \mathbb{N}$, and every $f \in \mathcal{C}_n$,

$$\Pr_{A, \text{EX}(f, D_n)} \left[\Pr_{x \sim D_n} [h(x) = f(x)] \geq 1 - \varepsilon \text{ for } h \leftarrow A^{\text{EX}(f, D_n)}(1^n, 1^{\lceil \varepsilon^{-1} \rceil}, 1^{\lceil \delta^{-1} \rceil}) \right] \geq 1 - \delta.$$

The outer probability is taken over the internal randomness of A and the independent samples returned by $\text{EX}(f, D_n)$. In other words, except with probability at most δ , the hypothesis produced by the learner has error at most ε under D_n .

We will only consider learning algorithms over the *uniform distribution*, i.e., we assume $D_n = \mathcal{U}_n$ throughout the paper.

A PAC learner is *proper* if it always outputs a hypothesis $h \in \mathcal{C}_n$. A learner whose output need not belong to $\mathcal{C}_n$ is called *improper*. All learning algorithms considered in this paper are improper.

In some results considered in this paper, the *confidence* parameter δ and the *accuracy* parameter ε are fixed. In such cases, the parameters are not provided as inputs to the learner.

Membership Queries. In the PAC model with *membership queries*, the learner additionally receives oracle access to the target concept f . It may therefore choose any $x \in \{0, 1\}^n$ and query the oracle for the value $f(x)$. Note that a randomized learning algorithm operating under the uniform distribution can simulate each example oracle call $\text{EX}(f, \mathcal{U}_n)$ using a membership query. For this reason, when proving lower bounds on classical learning algorithms we will restrict our attention to membership queries.

2.3 Quantum Learning Algorithms

We consider quantum learning algorithms with access to *quantum random examples* under the uniform distribution. We describe the model below, and refer to [AdW17, AGS21] for more information related to quantum learning theory and quantum random examples.

Learning from Quantum Random Examples. Under the uniform distribution, a classical random example is a pair $(x, f(x))$, where x is sampled uniformly from $\{0, 1\}^n$. Its quantum analogue is the coherent state

$$|\text{QEX}_f\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |x\rangle |f(x)\rangle.$$

Measuring this state in the computational basis produces an ordinary uniform random example. A quantum learner can instead apply unitary transformations before measuring, allowing the labeled examples in the superposition to interfere. The quantum learners considered in this paper use copies of this state and make no quantum membership queries.⁷

As usual, a quantum learning algorithm is described by a uniform sequence $\{Q_n\}_{n \geq 1}$ of quantum circuits. When we say that a quantum learner runs in time $T(n)$, this means that each circuit Q_n contains at most $T(n)$ gates, and given n we can print the description of Q_n in deterministic time $T(n)$.

Quantum Notation. We write H for the single-qubit Hadamard transform, so that the letter H remains reserved for sign-valued hypotheses (Section 2.1). In our notation, for each x ,

$$H|f(x)\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle + (-1)^{f(x)} |1\rangle \right) = \frac{1}{\sqrt{2}} (|0\rangle + F(x)|1\rangle).$$

Every quantum learning algorithm described in this paper produces a classical deterministic Boolean circuit as its output hypothesis.

2.4 Probability

The following inequality will be helpful during the analysis of the success probability of the quantum learner.

The Paley-Zygmund Inequality. This is a standard second moment estimate that converts moment bounds into a lower bound on the probability that a random variable is not too small. We include here a precise formulation and its short proof for completeness.

Lemma 2.1 (Paley-Zygmund Inequality). *Let X be a random variable satisfying*

$$0 < \mathbb{E}[X] < \infty \quad \text{and} \quad 0 < \mathbb{E}[X^2] < \infty.$$

Then, for every $\lambda \in [0, 1)$,

$$\Pr[X \geq \lambda \cdot \mathbb{E}[X]] \geq (1 - \lambda)^2 \cdot \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}.$$

Proof. Consider the event

$$E_\lambda \triangleq \{X \geq \lambda \cdot \mathbb{E}[X]\}.$$

Since $X < \lambda \mathbb{E}[X]$ on the complementary event E_λ^c ,

$$\mathbb{E}[X] = \mathbb{E}[X \mathbf{1}_{E_\lambda^c}] + \mathbb{E}[X \mathbf{1}_{E_\lambda}] \leq \lambda \mathbb{E}[X] + \mathbb{E}[X \mathbf{1}_{E_\lambda}].$$

Consequently,

$$(1 - \lambda)\mathbb{E}[X] \leq \mathbb{E}[X \mathbf{1}_{E_\lambda}].$$

By the Cauchy-Schwarz inequality,

$$\mathbb{E}[X \mathbf{1}_{E_\lambda}] \leq \mathbb{E}[X^2]^{1/2} \mathbb{E}[\mathbf{1}_{E_\lambda}^2]^{1/2} = \mathbb{E}[X^2]^{1/2} \Pr[E_\lambda]^{1/2}.$$

Combining these inequalities and squaring both sides gives the desired inequality

$$\Pr[E_\lambda] \geq (1 - \lambda)^2 \cdot \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}. \quad \square$$

⁷A quantum membership oracle provides the unitary transformation $O_f|x, b\rangle = |x, b \oplus f(x)\rangle$. A quantum membership query can be used to generate a quantum random example under the uniform distribution, but a quantum random example does not in general allow the learner to implement O_f on arbitrary superpositions of its choice.

2.5 4-Wise Independence

We will need the following definition, which will be used later with $k = 4$.

k -Wise Independence. Let $\mathcal{X}$ be a (nonempty) finite domain, let $\mathcal{A}$ be a finite seed set, and let

$$\mathcal{H} = \{h_a: \mathcal{X} \rightarrow \{0, 1\}\}_{a \in \mathcal{A}}$$

be a family of Boolean functions indexed by elements in $\mathcal{A}$. For an integer $k \geq 1$, the family $\mathcal{H}$ is called *k -wise independent* if, for every $1 \leq \ell \leq k$, every set of distinct points $x^{(1)}, \dots, x^{(\ell)} \in \mathcal{X}$, and every vector $v \in \{0, 1\}^\ell$,

$$\Pr_{a \sim A} \left[(h_a(x^{(1)}), \dots, h_a(x^{(\ell)})) = v \right] = 2^{-\ell}.$$

Equivalently, the evaluations at any $\ell \leq k$ distinct inputs are independent uniform bits.

For each h_a , define its sign representation $H_a(x) = (-1)^{h_a(x)}$. Since $b \mapsto (-1)^b$ is a bijection from $\{0, 1\}$ to $\{-1, 1\}$, the family $\mathcal{H}$ is k -wise independent precisely when $(H_a(x^{(1)}), \dots, H_a(x^{(\ell)}))$, for $a \sim A$, is uniform on $\{-1, 1\}^\ell$ for every collection of $\ell \leq k$ distinct inputs.

Similarly, we will say that a collection of random variables $\{Y_x\}_{x \in \mathcal{X}}$ is k -wise independent if every tuple containing $\ell \leq k$ distinct random variables is uniformly distributed.

4-Wise Independence Bounds. We next record some consequences of 4-wise independence that will be used later. Let $N = |\mathcal{X}|$, fix an arbitrary function

$$F: \mathcal{X} \rightarrow \{-1, 1\},$$

and define, for a uniformly random seed $a \sim A$,

$$Y_x \triangleq F(x)H_a(x) \in \{-1, 1\}, \quad Z_a \triangleq \frac{1}{\sqrt{N}} \sum_{x \in \mathcal{X}} Y_x.$$

Equivalently, the random variable

$$Z_a = \sqrt{N} \cdot \mathbb{E}_{x \sim U_n} [F(x)H_a(x)]$$

is the normalized correlation between the target function F and the randomly selected hypothesis H_a .

Lemma 2.2 (Moment Bounds for Correlation with 4-Wise Independent Family). *If the family $\{H_a\}_{a \in A}$ is 4-wise independent, then for every fixed target $F: \mathcal{X} \rightarrow \{-1, 1\}$, the random variable Z_a defined above satisfies*

$$\mathbb{E}_{a \sim A}[Z_a] = 0, \quad \mathbb{E}_{a \sim A}[Z_a^2] = 1, \quad \mathbb{E}_{a \sim A}[Z_a^4] = 3 - \frac{2}{N} \leq 3.$$

Proof. Multiplication by the fixed sign $F(x)$ merely exchanges the two possible values of $H_a(x)$. Consequently, $\{Y_x\}_{x \in \mathcal{X}}$ is also 4-wise independent.

The first moment is

$$\mathbb{E}[Z_a] = \frac{1}{\sqrt{N}} \sum_x \mathbb{E}_a[Y_x] = 0.$$

For the second moment,

$$\begin{aligned}
\mathbb{E}_a[Z_a^2] &= \frac{1}{N} \sum_{x_1, x_2} \mathbb{E}_a[Y_{x_1} Y_{x_2}] \\
&= \frac{1}{N} \left(\sum_x \mathbb{E}_a[Y_x^2] + \sum_{x_1 \neq x_2} \mathbb{E}_a[Y_{x_1} Y_{x_2}] \right) \\
&= \frac{1}{N} (N + 0) = 1.
\end{aligned}$$

Here $Y_x^2 = 1$, while the mixed terms vanish by pairwise independence.

For the fourth moment,

$$\mathbb{E}_a[Z_a^4] = \frac{1}{N^2} \sum_{x_1, x_2, x_3, x_4} \mathbb{E}_a[Y_{x_1} Y_{x_2} Y_{x_3} Y_{x_4}].$$

A term of this sum has nonzero expectation only if every input occurs an even number of times. There are N tuples in which all four indices are equal. There are also⁸ $3N(N-1)$ tuples containing two distinct inputs, each appearing twice. Every such tuple contributes 1. All other tuples contribute zero by 4-wise independence. Therefore

$$\mathbb{E}_a[Z_a^4] = \frac{N + 3N(N-1)}{N^2} = 3 - \frac{2}{N} \leq 3. \quad \square$$

The Global Sign Flip Property. In our learning algorithm we will require symmetry between positive and negative correlations. We say that the family $\mathcal{H}$ admits a *global sign flip* if there is a bijection

$$\iota: \mathcal{A} \rightarrow \mathcal{A}$$

such that

$$H_{\iota(a)}(x) = -H_a(x)$$

for every seed a and every input x . Since ι is a bijection, it preserves the uniform distribution on the seed set. Moreover,

$$Z_{\iota(a)} = -Z_a,$$

so the distribution of $Z_{\mathcal{A}}$ is symmetric about zero. For a 4-wise independent family, combining this symmetry with [Lemma 2.2](#) and applying the Paley-Zygmund inequality ([Lemma 2.1](#)) with $X = Z_{\mathcal{A}}^2$ and $\lambda = 1/4$ gives the following probability bounds:

$$\Pr_{a \sim \mathcal{A}} \left[|Z_a| \geq \frac{1}{2} \right] \geq \frac{3}{16} \quad \text{and} \quad \Pr_{a \sim \mathcal{A}} \left[Z_a \geq \frac{1}{2} \right] = \frac{1}{2} \Pr_{a \sim \mathcal{A}} \left[|Z_a| \geq \frac{1}{2} \right] \geq \frac{3}{32}. \quad (3)$$

2.6 Cryptography

We will need the following standard notion from cryptography. We refer to [\[Gol01\]](#) for more background on one-way functions.

⁸There are $\binom{N}{2}$ choices of the unordered pair $\{x, y\}$. Once this pair is fixed, an ordered tuple (x_1, x_2, x_3, x_4) is determined by choosing the two positions occupied by x ; the remaining two positions are occupied by y . Thus the number of tuples is precisely

$$\binom{N}{2} \binom{4}{2} = 6 \binom{N}{2} = 3N(N-1).$$

Subexponentially Secure One-Way Functions. A uniform family of functions $f_n: \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ computable in deterministic polynomial time is a *subexponentially secure one-way function* if there is a constant $\varepsilon > 0$ such that, for every uniform randomized algorithm $\mathcal{I}$ running in time at most 2^{n^ε} , every positive polynomial p , and every sufficiently large n ,

$$\Pr_{x \sim \mathcal{U}_n, y = f_n(x), \mathcal{I}}[\mathcal{I}(1^n, y) \in f_n^{-1}(y)] < \frac{1}{p(n)}.$$

3 Quantum Learning Algorithms: Proof of Theorem 1.2

The proof follows the outline presented in Section 1.2. We first describe a classical learner with a weaker guarantee on its success probability. We then explain how to improve it beyond $1/2$ using a quantum algorithm. Both constructions rely on a suitable 4-wise independent function family, which we describe next.

3.1 Technical Background

3.1.1 Quadratic Functions: A Suitable 4-Wise Independent Family

As in Section 2.5, let $\mathcal{X}$ be a finite domain with $N = 2^n = |\mathcal{X}|$, let $\mathcal{A}$ be a finite seed set, and let

$$\mathcal{H} \triangleq \{h_a: \mathcal{X} \rightarrow \{0, 1\}\}_{a \in \mathcal{A}}$$

be a family of Boolean functions indexed by elements in $\mathcal{A}$. Recall that for each h_a , we define its *sign representation* $H_a(x) = (-1)^{h_a(x)}$.

The Quadratic Boolean Family. We now construct an explicit family having all required properties. Let

$$r \triangleq 1 + n + \binom{n}{2}.$$

A seed $a \in \{0, 1\}^r$ is interpreted as

$$a = (a_\emptyset, \{a_i\}_{i=1}^n, \{a_{ij}\}_{1 \leq i < j \leq n}).$$

Associate with this seed the multilinear quadratic polynomial

$$q_a(x) = a_\emptyset \oplus \bigoplus_{i=1}^n a_i x_i \oplus \bigoplus_{1 \leq i < j \leq n} a_{ij} x_i x_j,$$

where all arithmetic is over $\mathbb{F}_2$. Finally, define

$$h_a(x) = q_a(x), \quad H_a(x) = (-1)^{q_a(x)}.$$

We will need the following result about the 4-wise independence of quadratic Boolean functions.⁹

Lemma 3.1 (Properties of the Quadratic Boolean Family). *For every $n \geq 1$, the family*

$$\mathcal{H}_n = \{H_a: \{0, 1\}^n \rightarrow \{-1, 1\}\}_{a \in \{0, 1\}^r}$$

has the following properties.

⁹For a more general version of this construction, we refer to [CGL⁺05].

1. The family is 4-wise independent under a uniform choice of the seed a .
2. The family admits a global sign flip $\iota : \mathcal{A} \rightarrow \mathcal{A}$.
3. There is an injective map

$$J : \{\perp\} \cup \{0, 1\}^n \rightarrow \{0, 1\}^r$$

such that $J(\perp) = 0^r$ and

$$H_a(x) = (-1)^{a \cdot J(x)}.$$

Both J and its inverse on its image are efficiently computable.

4. For every seed a , the Boolean function h_a is computable by a deterministic classical Boolean circuit of size $O(n^2)$.

Proof. We provide a detailed proof for completeness. Define the feature map

$$J(\perp) = 0^r, \quad J(x) = (1, \{x_i\}_{i=1}^n, \{x_i x_j\}_{1 \leq i < j \leq n}).$$

The ordinary binary inner product satisfies $a \cdot J(x) = q_a(x)$ over $\mathbb{F}_2$, and hence $H_a(x) = (-1)^{a \cdot J(x)}$.

To prove 4-wise independence, fix distinct points

$$x^{(1)}, \dots, x^{(k)} \in \{0, 1\}^n, \quad 1 \leq k \leq 4.$$

Their output bits are obtained from the seed through the linear map

$$L : \mathbb{F}_2^r \longrightarrow \mathbb{F}_2^k, \\ a \longmapsto \left(a \cdot J(x^{(1)}), \dots, a \cdot J(x^{(k)}) \right).$$

We argue next that the vectors

$$J(x^{(1)}), \dots, J(x^{(k)})$$

are linearly independent over $\mathbb{F}_2$. Suppose, toward a contradiction, that these feature vectors are linearly dependent. Since the coefficients in a binary linear dependence belong to $\{0, 1\}$, there is a nonempty set

$$T \subseteq \{x^{(1)}, \dots, x^{(k)}\}$$

such that

$$\sum_{x \in T} J(x) = 0.$$

The constant coordinate of every $J(x)$ is 1, so the constant coordinate of this sum is $|T|$ modulo 2. Thus $|T|$ must be even. Since $1 \leq |T| \leq 4$, it remains to consider the cases $|T| = 2$ and $|T| = 4$.

If $T = \{x, y\}$, the linear coordinates of $J(x) + J(y) = 0$ imply $x + y = 0$, and hence $x = y$. This contradicts the assumption that the selected points are distinct.

Suppose now that

$$T = \{x^{(1)}, x^{(2)}, x^{(3)}, x^{(4)}\}.$$

The linear coordinates of the assumed relation give

$$x^{(1)} + x^{(2)} + x^{(3)} + x^{(4)} = 0.$$

Set

$$u = x^{(1)}, \quad v = x^{(2)} + x^{(1)}, \quad w = x^{(3)} + x^{(1)}.$$

Then the four points are

$$u, \quad u + v, \quad u + w, \quad u + v + w.$$

Because the four points are distinct, v and w are nonzero and distinct. Two nonzero distinct vectors over $\mathbb{F}_2$ are linearly independent. The matrix whose rows are v and w therefore has rank 2, so it has a nonzero 2 by 2 minor. Consequently, there are coordinates $i < j$ such that

$$v_i w_j + v_j w_i = 1.$$

For these coordinates, the sum of the quadratic feature $x_i x_j$ over the four points is

$$\begin{aligned} & u_i u_j + (u_i + v_i)(u_j + v_j) + (u_i + w_i)(u_j + w_j) \\ & \quad + (u_i + v_i + w_i)(u_j + v_j + w_j) \\ & = v_i w_j + v_j w_i \\ & = 1. \end{aligned}$$

This contradicts the assumption that every coordinate of $\sum_{x \in T} J(x)$ is zero. Hence every collection of at most four distinct feature vectors is linearly independent.

From linear independence, it follows that L has rank k . The rank nullity theorem gives $\dim(\ker L) = r - k$, and a standard linear algebra argument shows that every element of $\mathbb{F}_2^k$ has exactly 2^{r-k} preimages under L . A uniform seed $a \in \mathbb{F}_2^r$ consequently produces a uniform evaluation vector

$$(h_a(x^{(1)}), \dots, h_a(x^{(k)})) \in \mathbb{F}_2^k.$$

This proves 4-wise independence.

For the global sign flip, define ι by replacing the constant coefficient $a_\emptyset$ with $a_\emptyset \oplus 1$ and leaving every other coefficient unchanged. Then $q_{\iota(a)}(x) = q_a(x) \oplus 1$ for every x , and therefore $H_{\iota(a)}(x) = -H_a(x)$.

Note that J is injective on $\{\perp\} \cup \{0, 1\}^n$. Its constant coordinate separates $J(\perp) = 0^r$ from every $J(x)$, while its linear coordinates recover x and therefore distinguish any two Boolean inputs. Its inverse on its image returns $\perp$ on 0^r and otherwise reads x from the linear coordinates.

Finally, computing $J(x)$, evaluating $q_a(x)$, and implementing the corresponding hypothesis circuit require $O(n^2)$ elementary Boolean operations. $\square$

3.1.2 Derandomized Classical Learner with $p^* \leq 1/2$

In this section, we describe a zero-query polynomial-time classical learner in the large error regime that succeeds with probability $\Omega(1)$. The quantum learner presented in [Section 3.2](#) builds on the construction described here.

Proposition 3.2 (Classical Zero-Query Learner). *There is a randomized zero-query algorithm $\mathcal{R}$ that runs in $\text{poly}(n)$ time and outputs a deterministic classical Boolean circuit h such that, for every target $f: \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\Pr_{h \leftarrow \mathcal{R}} \left[\Pr_x [h(x) = f(x)] \geq \frac{1}{2} + \frac{1}{4} 2^{-n/2} \right] \geq \frac{3}{32}.$$

Proof. We use the efficiently computable 4-wise independent family $\mathcal{H}$ from [Section 3.1.1](#). The learner $\mathcal{R}$ samples a uniform seed a and outputs the corresponding deterministic circuit h_a , whose sign representation is H_a . By [Lemma 3.1](#), $\mathcal{R}$ runs in polynomial time and its output hypothesis is a deterministic Boolean circuit.

For a fixed target $f: \{0, 1\}^n \rightarrow \{0, 1\}$ with sign representation $F(x) = (-1)^{f(x)}$, set

$$Z_a \triangleq \frac{1}{\sqrt{N}} \sum_x F(x) H_a(x).$$

All probabilities and expectations in the remainder of this proof are over the uniformly sampled seed a . The target F remains fixed. By [Lemma 2.2](#),

$$\mathbb{E}[Z_a] = 0, \quad \mathbb{E}[Z_a^2] = 1, \quad \mathbb{E}[Z_a^4] \leq 3.$$

Recall that the family $\mathcal{H}$ admits a global sign flip ([Lemma 3.1](#)). Consequently, applying the Paley-Zygmund Inequality ([Lemma 2.1](#)) as in the derivation of [Equation \(3\)](#) from [Section 2.5](#), we get that

$$\Pr\left[Z_a \geq \frac{1}{2}\right] \geq \frac{3}{32}.$$

On this event, $\text{corr}(H_a, F) = Z_a/\sqrt{N} \geq (1/2) \cdot (1/\sqrt{N})$, which provides the guarantee in the statement of [Proposition 3.2](#) via [Equation \(2\)](#). $\square$

3.2 The Correlation-Biased Quantum Sampler: Proof of Item 1

In this section, we describe a quantum learning algorithm that uses a single quantum random example and achieves success probability strictly larger than $1/2$. Note that the learning procedure is quantum, but the hypothesis it returns is classical.

3.2.1 Hypothesis Family $\mathcal{H}$ and Isometry W

We will need the following notions to describe the quantum algorithm $\mathcal{Q}$ from [Theorem 1.2](#).

The Hypothesis Family $\mathcal{H}$. Let $\mathcal{H} = \{H_a: \{0, 1\}^n \rightarrow \{-1, 1\}\}_{a \in \mathcal{A}}$ be an efficiently computable 4-wise independent family of sign-valued functions. As in [Lemma 3.1](#), we assume that the seed set $\mathcal{A}$ is identified with $\{0, 1\}^r$. In particular,

$$M \triangleq |\mathcal{A}| = 2^r,$$

and expressions such as $a \cdot J(x)$ below are ordinary binary inner products in this chosen encoding. We also assume a *global sign flip*: there is a bijection $\iota: \mathcal{A} \rightarrow \mathcal{A}$ preserving the uniform seed distribution such that

$$H_{\iota(a)}(x) = -H_a(x) \quad \text{for every } a \in \mathcal{A} \text{ and every } x.$$

Finally, assume that there is an efficiently computable injective feature map

$$J: \{\perp\} \cup \{0, 1\}^n \rightarrow \{0, 1\}^r,$$

whose inverse on its image is also efficiently computable. We require

$$J(\perp) = 0^r \quad \text{and} \quad H_a(x) = (-1)^{a \cdot J(x)}.$$

The Isometry W . We now introduce a map W needed by the quantum algorithm. Let $\perp \notin \{0, 1\}^n$ be an auxiliary reference index. We use the following $N + 1$ basis states of an $n + 1$ qubit index register:

$$|\perp\rangle \triangleq |0\rangle|0^n\rangle, \quad |x\rangle_{\text{idx}} \triangleq |1\rangle|x\rangle \quad (x \in \{0, 1\}^n).$$

Writing $|\ell\rangle$ for the encoded state of an index $\ell \in \{\perp\} \cup \{0, 1\}^n$, we define

$$W|\ell\rangle \triangleq H^{\otimes r}|J(\ell)\rangle.$$

The standard Hadamard identity

$$H^{\otimes r}|y\rangle = \frac{1}{\sqrt{M}} \sum_{a \in \{0,1\}^r} (-1)^{a \cdot y} |a\rangle$$

and the defining properties of J give

$$W|\perp\rangle = \frac{1}{\sqrt{M}} \sum_{a \in \mathcal{A}} |a\rangle, \quad W|x\rangle_{\text{idx}} = \frac{1}{\sqrt{M}} \sum_{a \in \mathcal{A}} H_a(x) |a\rangle. \quad (4)$$

Since J is injective and the Hadamard transform is unitary, it is not hard to check that W is an isometry.¹⁰

To use W in our polynomial-time quantum learner, we must also show that it can be implemented using polynomially many quantum gates. For completeness, we describe this implementation next. Using a fresh r qubit register, reversibly compute

$$|\ell\rangle|0^r\rangle \mapsto |\ell\rangle|J(\ell)\rangle.$$

Next, evaluate J^{-1} on the second register and XOR the recovered binary encoding of ℓ into the first register, obtaining

$$|\ell\rangle|J(\ell)\rangle \mapsto |0^{n+1}\rangle|J(\ell)\rangle.$$

This resets the original index register while retaining its information in the second register. Temporary work qubits are returned to zero after each evaluation. Finally, apply $H^{\otimes r}$ to the second register:

$$|0^{n+1}\rangle|J(\ell)\rangle \mapsto |0^{n+1}\rangle W|\ell\rangle.$$

The first register is now in a fixed zero state and is unentangled with the second, so we can ignore it. By linearity, the second register contains $W|\psi\rangle$ for every input state $|\psi\rangle$ in the domain of W . (This is the seed register measured by the learner.) The assumed efficient evaluation and inversion of J admit reversible implementations with polynomial overhead. Together with $r = \text{poly}(n)$ for the concrete family, this gives an implementation using $\text{poly}(n)$ gates.

3.2.2 Quantum Algorithm $\mathcal{Q}$

Using the concepts and notation introduced above, the quantum learner can be described in four steps. As usual, let $N = 2^n$, and let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be the target function. Moreover, let $\mathcal{H}$, W , $\mathcal{A}$, and $M = |\mathcal{A}| = 2^r$ be as in [Section 3.2.1](#).

Quantum Procedure $\mathcal{Q}$:

¹⁰Recall that an isometry is a linear map W that preserves inner products: $\langle W\psi | W\phi \rangle = \langle \psi | \phi \rangle$ for all vectors in its domain. In particular, it maps an orthonormal basis to an orthonormal set in the output space, which may have larger dimension.

1. Receive one uniform quantum random example

$$|\text{QEX}_f\rangle = \frac{1}{\sqrt{N}} \sum_x |x\rangle |f(x)\rangle.$$

Apply a Hadamard gate to the label qubit. Using $F(x) = (-1)^{f(x)}$ and linearity,

$$\begin{aligned} |\text{QEX}_f\rangle &\mapsto \frac{1}{\sqrt{2N}} \sum_x |x\rangle (|0\rangle + F(x)|1\rangle) \\ &= \frac{1}{\sqrt{2N}} \sum_x |x\rangle |0\rangle + \frac{1}{\sqrt{2N}} \sum_x F(x) |x\rangle |1\rangle. \end{aligned}$$

Regard the rightmost qubit as a flag and place it first in the notation. The same state can be rewritten as

$$\frac{1}{\sqrt{2}} |0\rangle \left(\frac{1}{\sqrt{N}} \sum_x |x\rangle \right) + \frac{1}{\sqrt{2}} |1\rangle \left(\frac{1}{\sqrt{N}} \sum_x F(x) |x\rangle \right).$$

(The flag 0 branch is a uniform state, while the flag 1 branch contains the target values as phases. The flag is left unmeasured, allowing the “reference” branch and the “target” branch to interfere in Step 3.)

2. Apply $H^{\otimes n}$ to the input register only on the flag 0 branch. Because

$$H^{\otimes n} \left(\frac{1}{\sqrt{N}} \sum_x |x\rangle \right) = |0^n\rangle,$$

the resulting state can be written as

$$\frac{1}{\sqrt{2}} |0\rangle |0^n\rangle + \frac{1}{\sqrt{2N}} \sum_x F(x) |1\rangle |x\rangle.$$

As in [Section 3.2.1](#), encode the reference label by $|\perp\rangle = |0\rangle |0^n\rangle$ and encode an input label x by $|x\rangle_{\text{lab}} = |1\rangle |x\rangle$. In this label space the state is

$$|\psi_F\rangle \triangleq \frac{1}{\sqrt{2}} |\perp\rangle + \frac{1}{\sqrt{2N}} \sum_x F(x) |x\rangle_{\text{lab}}.$$

(Note that this preparation is exact. It uses the given quantum example, one Hadamard gate, and n Hadamard gates controlled on the flag being 0. It makes no membership query and uses no measurement or postselection.)

3. Apply the isometry W . By [Equation \(4\)](#) and linearity,

$$\begin{aligned} W|\psi_F\rangle &= \frac{1}{\sqrt{2}} W|\perp\rangle + \frac{1}{\sqrt{2N}} \sum_x F(x) W|x\rangle_{\text{lab}} \\ &= \frac{1}{\sqrt{2M}} \sum_a |a\rangle + \frac{1}{\sqrt{2NM}} \sum_x \sum_a F(x) H_a(x) |a\rangle \\ &= \frac{1}{\sqrt{2M}} \sum_{a \in \mathcal{A}} (1 + Z_a) |a\rangle. \end{aligned}$$

where

$$Z_a = \frac{1}{\sqrt{N}} \sum_x F(x) H_a(x) = \sqrt{N} \text{corr}(H_a, F).$$

(The amplitude of the basis state $|a\rangle$ is therefore proportional to $1 + Z_a$. In this sense, the normalized correlation between H_a and F has been encoded into a quantum amplitude.)

4. Measure the seed register. The probability of observing a basis state is the squared absolute value of its amplitude, so the resulting seed is distributed according to

$$P_F(a) \triangleq \frac{(1 + Z_a)^2}{2M}. \quad (5)$$

Once a has been measured, output the corresponding deterministic classical circuit h_a .

Intuition. It is helpful to isolate the role of the reference branch. Without $|\perp\rangle$, the amplitude of $|a\rangle$ would be proportional to Z_a , so measurement would sample a with probability proportional to Z_a^2 . Squaring would erase the sign of the correlation, just as in ordinary Fourier sampling (Section A). In contrast, the reference branch adds the constant amplitude 1, producing

$$(1 + Z_a)^2 = 1 + 2Z_a + Z_a^2.$$

The linear term $2Z_a$ favors positive correlations over negative ones. More concretely, the global sign flip pairs a seed with normalized correlation z with a seed having normalized correlation $-z$. For $z > 0$, their measurement weights differ by

$$(1 + z)^2 - (1 - z)^2 = 4z > 0.$$

Interference changes the sampling distribution before measurement so that the positively correlated orientation is intrinsically more likely.

3.2.3 Correctness

We now prove that the quantum procedure $\mathcal{Q}$ from Section 3.2.2 outputs a positively correlated hypothesis with probability strictly greater than $1/2$. This completes the proof of Theorem 1.2 Item (i).

Theorem 3.3 (Correctness of the Quantum Learner $\mathcal{Q}$). *For every $n \geq 1$ and every target function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, $\mathcal{Q}$ uses one uniform quantum random example and runs in time $\text{poly}(n)$. With probability $p^* \geq 0.53$, it outputs a deterministic classical Boolean circuit h such that*

$$\Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] \geq \frac{1}{2} + a_0 \cdot 2^{-n/2},$$

where $a_0 = 0.001$.

Proof. Fix the target function $F(x) = (-1)^{f(x)}$. Let U be the uniform distribution on the seed set $\mathcal{A}$, and let P_F be the measurement distribution of $\mathcal{Q}$ (Section 3.2.2). Recall that $Z_a \triangleq \sqrt{N} \cdot \text{corr}(H_a, F)$. We write Z for this function of the seed and indicate the seed distribution (U or P_F) by the subscript on each probability or expectation.

Set $c \triangleq 2a_0 = 0.002$. By Equation (2) and $N = 2^n$, a measured seed a yields the desired agreement precisely when $Z_a \geq c$. Our goal is therefore to show that

$$\Pr_{P_F}[Z \geq c] \geq 0.53.$$

We first recall the properties of Z under the uniform seed distribution. By [Lemma 3.1](#) and [Lemma 2.2](#),

$$\mathbb{E}_U[Z^2] = 1, \quad \mathbb{E}_U[Z^4] \leq 3.$$

The global sign flip also makes Z symmetric about zero under U , as explained in [Section 2.5](#). In particular, its odd moments vanish.

To use these facts under the quantum measurement distribution, recall from Step 4 of [Section 3.2.2](#) that $P_F(a) = (1 + Z_a)^2/(2M)$. Since $U(a) = 1/M$, for every function $g: \mathbb{R} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}_{P_F}[g(Z)] = \frac{1}{2} \mathbb{E}_U[(1 + Z)^2 g(Z)].$$

This identity lets us express moments under P_F in terms of moments under U .

The classical learner in [Proposition 3.2](#) applies Paley-Zygmund ([Lemma 2.1](#)) under the uniform seed distribution. Here we apply the same inequality under the quantum measurement distribution, to the non-negative random variable

$$X \triangleq \max\{Z, 0\}.$$

Since $c > 0$, the events $X \geq c$ and $Z \geq c$ coincide. We will therefore bound the first two moments of X under P_F .

For the first moment, the conversion identity and the discussion above on the (uniform) moments of Z give

$$\begin{aligned} \mathbb{E}_{P_F}[Z] &= \frac{1}{2} \mathbb{E}_U[Z(1 + Z)^2] \\ &= \frac{1}{2} (\mathbb{E}_U[Z] + 2\mathbb{E}_U[Z^2] + \mathbb{E}_U[Z^3]) = 1, \end{aligned}$$

where the last equality uses symmetry and $\mathbb{E}_U[Z^2] = 1$. Since $X \geq Z$, it follows that

$$\mu \triangleq \mathbb{E}_{P_F}[X] \geq 1.$$

For the second moment, observe that $X^2 = Z^2 \mathbf{1}_{\{Z > 0\}}$. The function $z^2(1 + |z|)^2$ is even and vanishes at zero, so symmetry under U assigns half of its expectation to positive values of Z . Applying the conversion identity and using that Z is symmetric about zero under U , we obtain

$$\begin{aligned} \nu \triangleq \mathbb{E}_{P_F}[X^2] &= \frac{1}{2} \mathbb{E}_U[Z^2(1 + Z)^2 \mathbf{1}_{\{Z > 0\}}] \\ &= \frac{1}{4} \mathbb{E}_U[Z^2(1 + |Z|)^2] \\ &= \frac{1}{4} (1 + 2\mathbb{E}_U[|Z|^3] + \mathbb{E}_U[Z^4]). \end{aligned}$$

By the Cauchy-Schwarz inequality and the uniform moment bounds,

$$\mathbb{E}_U[|Z|^3] = \mathbb{E}_U[|Z| Z^2] \leq \sqrt{\mathbb{E}_U[Z^2] \mathbb{E}_U[Z^4]} \leq \sqrt{3}.$$

Consequently,

$$\nu \leq \frac{1 + 2\sqrt{3} + 3}{4} = 1 + \frac{\sqrt{3}}{2}.$$

We can now apply Paley-Zygmund. Since $\mu \geq 1$ and $c = 0.002$, the parameter $\lambda \triangleq c/\mu$ belongs to $(0, 1)$. Applying [Lemma 2.1](#) to X under P_F with this parameter gives

$$\begin{aligned} \Pr_{P_F}[Z \geq c] &= \Pr_{P_F}[X \geq c] = \Pr_{P_F}[X \geq \lambda \cdot \mu] \\ &\geq \left(1 - \frac{c}{\mu}\right)^2 \frac{\mu^2}{\nu} = \frac{(\mu - c)^2}{\nu} \\ &\geq \frac{(1 - c)^2}{1 + \sqrt{3}/2} = \frac{(0.998)^2}{1 + \sqrt{3}/2} \\ &> 0.53. \end{aligned}$$

As noted at the beginning of the proof, this establishes the required agreement guarantee.

Finally, $\mathcal{Q}$ uses one uniform quantum random example. The state preparation in Steps 1 and 2 and the implementation of W in [Section 3.2.1](#) give the polynomial running time. By [Lemma 3.1](#), the measured seed specifies a deterministic classical Boolean circuit of size $O(n^2)$. This proves the theorem. $\square$

3.3 Boosting Correlation and Success Probability: Proof of Item 2

We now prove Item (ii) of [Theorem 1.2](#). The quantum algorithm $\mathcal{Q}^+$ runs $\mathcal{Q}$ independently on several quantum random examples and takes the majority of the resulting hypotheses. The key property is that, at each input, a sampled hypothesis has a small positive bias toward the target. Moreover, these agreement indicators are independent across distinct input pairs. Majority increases the pointwise bias, and independence among pairs then gives concentration of the overall correlation. We provide the details next.

Fix a target $f: \{0, 1\}^n \rightarrow \{0, 1\}$, write $F = (-1)^f$, and let $N = 2^n$. We use the quadratic Boolean family from [Lemma 3.1](#). Let U be the uniform distribution on its seed set $\mathcal{A}$, let $M = |\mathcal{A}|$, and recall that the measurement distribution of $\mathcal{Q}$ is

$$P_F(a) = \frac{(1 + Z_a)^2}{2M}, \text{ where } Z_a = \frac{1}{\sqrt{N}} \sum_x F(x) H_a(x).$$

We will need the following lemmas to prove the correctness of $\mathcal{Q}^+$.

Lemma 3.4 (Pointwise Bias and Independence). *For $a \sim P_F$, define $Y_x \triangleq F(x) H_a(x)$. The following statements hold:*

- $\mathbb{E}_{P_F}[Y_x] = \frac{1}{\sqrt{N}}$ for every x .
- $\mathbb{E}_{P_F}[Y_x Y_y] = \frac{1}{N}$ for every $x \neq y$.
- For $x \neq y$, the random variables Y_x and Y_y are independent under P_F .

Proof. Regard Y_x and $Z = Z_a$ as functions of the seed. As above, under U , the signs Y_x are unbiased and 4-wise independent. Similarly to the proof of [Theorem 3.3](#), for any function G of the seed, the change of distribution gives

$$\mathbb{E}_{P_F}[G] = \frac{1}{2} \mathbb{E}_U[(1 + 2Z + Z^2)G]. \quad (6)$$

For a fixed input x , the uniform moments needed here are

$$\mathbb{E}_U[Y_x] = 0, \quad \mathbb{E}_U[Z Y_x] = \frac{1}{\sqrt{N}}, \quad \mathbb{E}_U[Z^2 Y_x] = 0.$$

Indeed, in $\mathbb{E}_{\mathcal{U}}[ZY_x] = N^{-1/2} \sum_u \mathbb{E}_{\mathcal{U}}[Y_u Y_x]$, only $u = x$ contributes. In the expansion of $Z^2 Y_x$, cancellation of repeated factors leaves either one unbiased sign or three distinct unbiased signs, so every term has expectation zero due to the 4-wise independence of $\mathcal{H}$. Consequently, Equation (6) with $G = Y_x$ yields

$$\mathbb{E}_{P_F}[Y_x] = \frac{1}{2} \left(0 + \frac{2}{\sqrt{N}} + 0 \right) = \frac{1}{\sqrt{N}}.$$

Now fix distinct inputs x, y . We will rely on Equation (6) again. The same cancellation argument gives

$$\mathbb{E}_{\mathcal{U}}[Y_x Y_y] = 0, \quad \mathbb{E}_{\mathcal{U}}[ZY_x Y_y] = 0.$$

For the quadratic term,

$$\mathbb{E}_{\mathcal{U}}[Z^2 Y_x Y_y] = \frac{1}{N} \sum_{u,v} \mathbb{E}_{\mathcal{U}}[Y_u Y_v Y_x Y_y] = \frac{2}{N}.$$

Only $(u, v) = (x, y)$ and $(u, v) = (y, x)$ contribute, each with value 1. Every other term reduces to a product of two or four distinct unbiased signs and vanishes by 4-wise independence. Thus $\mathbb{E}_{P_F}[Y_x Y_y] = 1/N$.

For the last item, let $x \neq y$. For any sign $Y \in \{-1, 1\}$ and any $s \in \{-1, 1\}$, we have the pointwise identity

$$\mathbf{1}_{\{Y=s\}} = \frac{1 + sY}{2}.$$

Indeed, the right side equals 1 when $Y = s$ and 0 otherwise. Write $\mu = 1/\sqrt{N}$, and take all probabilities and expectations below under P_F . By the first two items,

$$\mathbb{E}[Y_x] = \mathbb{E}[Y_y] = \mu, \quad \mathbb{E}[Y_x Y_y] = \mu^2.$$

Taking expectations in the indicator identity gives

$$\Pr[Y_x = s] = \frac{1 + s\mu}{2}, \quad \Pr[Y_y = t] = \frac{1 + t\mu}{2}.$$

Now let $s, t \in \{-1, 1\}$ be arbitrary. Applying the same identity to both signs,

$$\begin{aligned} \Pr[Y_x = s, Y_y = t] &= \mathbb{E}[\mathbf{1}_{\{Y_x=s\}} \mathbf{1}_{\{Y_y=t\}}] \\ &= \frac{1}{4} \mathbb{E}[(1 + sY_x)(1 + tY_y)] \\ &= \frac{1}{4} (1 + s\mathbb{E}[Y_x] + t\mathbb{E}[Y_y] + st\mathbb{E}[Y_x Y_y]) \\ &= \frac{1}{4} (1 + s\mu + t\mu + st\mu^2) \\ &= \frac{1 + s\mu}{2} \frac{1 + t\mu}{2} \\ &= \Pr[Y_x = s] \Pr[Y_y = t]. \end{aligned}$$

This establishes the required factorization for every possible pair of values, so Y_x and Y_y are independent. $\square$

Lemma 3.5 (Majority Bias; see [Wat20, Appendix A]). *There is a universal constant $0 < b \leq 1$ such that, for every odd integer m with $1 \leq m \leq N$, independent signs $B_1, \dots, B_m \in \{-1, +1\}$ of mean $1/\sqrt{N}$ satisfy*

$$\mathbb{E}[\text{Sign}(B_1, \dots, B_m)] \geq b \sqrt{\frac{m}{N}},$$

where $\text{Sign}(B_1, \dots, B_m) \in \{-1, +1\}$ denotes the sign of their sum.

Next, we provide a precise description of $\mathcal{Q}^+$.

The Amplified Learner $\mathcal{Q}^+(n, m)$. For an odd integer $m \leq N$, run $\mathcal{Q}$ independently m times, using one fresh quantum random example each time. Let $a_1, \dots, a_m$ be the measured seeds. Output the Boolean circuit

$$h(x) = \text{Maj}(h_{a_1}(x), \dots, h_{a_m}(x)),$$

where $\text{Maj}(y_1, \dots, y_m)$ is the ordinary bit majority, and all seeds are hardwired into h . Its sign representation is

$$H(x) = \text{Sign}(H_{a_1}(x), \dots, H_{a_m}(x)).$$

Theorem 3.6 (Guarantee for the Amplified Learner). *Let b be the constant in [Lemma 3.5](#), and set $C_0 = 4/b^2$. For every target $f \in \text{ALL}_n$ and every odd integer $1 \leq m \leq N$, $\mathcal{Q}^+(n, m)$ uses exactly m independent quantum random examples, runs in time $m \cdot \text{poly}(n)$, and outputs a deterministic classical Boolean circuit of size $m \cdot \text{poly}(n)$ such that*

$$\Pr_{\mathcal{Q}^+} \left[\Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] \geq \frac{1}{2} + \frac{b}{4} \sqrt{\frac{m}{N}} \right] \geq 1 - \frac{C_0}{m}.$$

Proof. For each input x , set $X_x = F(x)H(x)$. Multiplication of all arguments of an odd majority by the same sign multiplies its output by that sign, so

$$X_x = \text{Sign}(F(x)H_{a_1}(x), \dots, F(x)H_{a_m}(x)).$$

The arguments are independent signs of mean $1/\sqrt{N}$ by [Lemma 3.4](#). Hence by [Lemma 3.5](#) their signed majority has a common mean

$$\beta_m \triangleq \mathbb{E}[X_x] \geq b \sqrt{\frac{m}{N}}.$$

For distinct inputs x, y , [Lemma 3.4](#) gives independence of $F(x)H_{a_i}(x)$ and $F(y)H_{a_i}(y)$ for each i . Independence of the runs then implies independence of the two full sequences, and therefore of X_x and X_y .

It follows by a simple calculation that the output correlation satisfies

$$\mathbb{E}[\text{corr}(H, F)] = \beta_m, \quad \text{Var}(\text{corr}(H, F)) = \frac{1 - \beta_m^2}{N} \leq \frac{1}{N}.$$

Chebyshev's inequality gives

$$\Pr \left[\text{corr}(H, F) < \frac{\beta_m}{2} \right] \leq \frac{4}{N\beta_m^2} \leq \frac{C_0}{m}.$$

On the complementary event, [Equation \(2\)](#) yields

$$\Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] \geq \frac{1}{2} + \frac{\beta_m}{4} \geq \frac{1}{2} + \frac{b}{4} \sqrt{\frac{m}{N}}.$$

Each run of $\mathcal{Q}$ takes $\text{poly}(n)$ time and produces a circuit of size $O(n^2)$. A majority circuit can sum the m output bits and compare the sum with $(m+1)/2$, using $O(m \log(m+1))$ gates. Since $m \leq 2^n$, this is $O(mn)$. Thus both the complete computation and the construction of its output circuit take $m \cdot \text{poly}(n)$ time, and the circuit has the stated size. In particular, it is a deterministic classical circuit built from the quadratic Boolean hypotheses in [Lemma 3.1](#). $\square$

Choosing the Failure Probability. Set $C = C_0 + 1$, and suppose $\eta \in (0, 1/2)$ satisfies $\eta \geq C/N$. Choose m to be the smallest odd integer at least C_0/η . Then

$$\frac{C_0}{\eta} \leq m < \frac{C_0}{\eta} + 2 \leq \frac{C_0 + 1}{\eta} \leq N.$$

Thus $m = \Theta(1/\eta)$ and $C_0/m \leq \eta$. Moreover,

$$\frac{b}{4} \sqrt{\frac{m}{N}} \geq \frac{b\sqrt{C_0}}{4\sqrt{\eta}} 2^{-n/2}.$$

Taking $a_1 = b\sqrt{C_0}/4$ in [Theorem 3.6](#) proves Item (ii) of [Theorem 1.2](#), including the claimed time bound $\text{poly}(n)/\eta$, circuit size, and number of quantum random examples.

4 Limitations of Classical Learners: Proof of [Theorem 1.4](#)

As alluded to in [Section 1.2](#), different techniques are needed to establish each item of [Theorem 1.4](#). We first discuss the unconditional query complexity lower bound for classical learners for the concept class $\mathcal{C} = \text{ALL}$.

4.1 Unconditional Information Theoretic Lower Bound for $\mathcal{C} = \text{ALL}$: Proof of Item 1

The polynomial-time classical learner in [Proposition 3.2](#) achieves correlation $\Omega(2^{-n/2})$ with constant success probability (without making any queries). We now show that guaranteeing success probability strictly greater than $1/2$ at this correlation requires $\Omega(2^{n/2})$ classical membership queries. We note that the argument applies to adaptive queries and places no restriction on the learner's running time or the size of its output circuit.

Theorem 4.1 (Classical Query Lower Bound). *Let $\mathcal{A}$ be a randomized classical learner that makes at most q membership queries to a target $f \in \text{ALL}_n$ and outputs a deterministic Boolean circuit h . For every $\gamma > q/(2N)$,*

$$\mathbb{E}_{f \sim \text{ALL}_n} \left[\Pr_{\mathcal{A}} \left[\Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] \geq \frac{1}{2} + \gamma \right] \right] \leq \frac{1}{2}.$$

Consequently, $\mathcal{A}$ has success probability at most $1/2$ at this correlation regime for some target $f \in \text{ALL}_n$.

Proof. Choose f uniformly from ALL_n , so the signs $F(x) = (-1)^{f(x)}$ are independent and uniform. Fix all random coins of $\mathcal{A}$ and condition on a feasible complete query transcript. Let Q be the set of distinct queried points, let $k = |Q| \leq q$, and let $H = (-1)^h$ be the resulting fixed hypothesis.

Every target agreeing with the recorded answers on Q produces the same transcript. Thus, even for adaptive queries, the unqueried signs remain independent and uniform under this conditioning. Consequently,

$$S \triangleq \sum_{x \notin Q} H(x)F(x)$$

is a sum of $N - k$ independent uniform signs, and

$$N \cdot \text{corr}(H, F) = \sum_{x \in Q} H(x)F(x) + S \leq k + S.$$

By Equation (2), the success event

$$\mathcal{E} \triangleq \left\{ \Pr_{x \sim \mathcal{U}_n} [h(x) = f(x)] \geq \frac{1}{2} + \gamma \right\}$$

therefore implies

$$S \geq 2\gamma N - k \geq 2\gamma N - q > 0,$$

where the last inequality relied on the assumption that $\gamma > q/(2N)$. Since the conditional distribution of S is symmetric about zero, its probability of being strictly positive is at most $1/2$. Hence the conditional probability of $\mathcal{E}$ is at most $1/2$.

Averaging over transcripts and random coins gives

$$\mathbb{E}_{f \sim \text{ALL}_n} \left[\Pr_{\mathcal{A}} [\mathcal{E}] \right] \leq \frac{1}{2}.$$

In particular, at least one target function has success probability no greater than the expected value. $\square$

This proves Item (i) of Theorem 1.4. In particular, for $\gamma = c \cdot 2^{-n/2}$ with fixed $c > 0$, the condition is $q < 2c2^{n/2}$, which rules out classical learners making $q = o(2^{n/2})$ queries.

4.2 Conditional Computational Lower Bound for $\mathcal{C} = \text{SIZE}[\text{poly}]$: Proof of Item 2

4.2.1 Strong Pseudorandom Functions

We first state a standard pseudorandomness consequence of the assumption that subexponentially secure one-way functions exist (Section 2.6). All adversaries in this subsection are uniform randomized classical algorithms. Their oracle queries are classical and may be adaptive.

Definition (Strong PRFs). Let $\ell: \mathbb{N} \rightarrow \mathbb{N}$ be polynomially bounded and computable in polynomial time. Consider a family

$$\mathcal{G}_\lambda = \{g_{\lambda,s}: \{0,1\}^\lambda \rightarrow \{0,1\}\}_{s \in \{0,1\}^{\ell(\lambda)}}$$

with a deterministic evaluation algorithm running in $\text{poly}(\lambda)$ time on input $(1^\lambda, s, x)$. For a randomized oracle algorithm D , define

$$\text{Adv}_{\mathcal{G}}(D, \lambda) \triangleq \left| \Pr_{s \sim \mathcal{U}_{\ell(\lambda)}, D} [D^{g_{\lambda,s}}(1^\lambda) = 1] - \Pr_{R \sim \text{ALL}_{\lambda}, D} [D^R(1^\lambda) = 1] \right|,$$

where R is a uniformly random Boolean function on λ bits. We call $\mathcal{G}$ a *strong pseudorandom function family* (strong PRF) if there is a rational constant $\varepsilon \in (0, 1)$ such that every D running in time at most 2^{λ^ε} satisfies

$$\text{Adv}_{\mathcal{G}}(D, \lambda) \leq 2^{-\lambda^\varepsilon} \tag{7}$$

for every sufficiently large λ .

We will need the following well-known cryptographic result.

Lemma 4.2 (Strong PRFs from Subexponentially Secure OWFs). *If subexponentially secure one-way functions exist, then a strong pseudorandom function family exists. In particular, there are a rational constant $\varepsilon \in (0, 1)$ and a fixed integer $d \geq 1$ such that the family satisfies Equation (7) and every member of $\mathcal{G}_\lambda$ has a Boolean circuit of size at most λ^d for all sufficiently large λ .*

For completeness, we provide a proof sketch and verify the parameters in the statement of [Lemma 4.2](#) needed in our application.

Proof Sketch. Let f_λ be the assumed OWF on λ input bits, with security exponent $\varepsilon_0 > 0$. We retain λ as the security parameter throughout.

First, we amplify the security of the assumed one-way function. Fix a rational $0 < \gamma < \varepsilon_0$ and take the direct product of λ independent copies of f_λ , with total input length λ^2 . The classical amplification argument [[Gol01](#), Theorem 2.3.2] embeds a challenge in each coordinate among independently generated instances and repeats the inversion attempt. Tracking that proof, an inverter for the direct product function of time T and success $\eta \geq 2^{-\lambda^\gamma}$ yields a uniform inverter for f_λ with constant success probability in time $\text{poly}(T, \lambda, 1/\eta)$. For $T \leq 2^{\lambda^\gamma}$, the resulting inverter's running time is $2^{O(\lambda^\gamma)} < 2^{\lambda^{\varepsilon_0}}$, contradicting the assumed negligible success probability condition of f_λ . Thus every inverter for the direct product function running in time at most 2^{λ^γ} has success probability at most $2^{-\lambda^\gamma}$, for all sufficiently large λ .

Next, we construct a subexponentially secure pseudorandom generator (strong PRG) from the amplified one-way function. In more detail, the HILL construction [[HILL99](#)] gives a polynomial-time pseudorandom generator

$$P_\lambda: \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{2\ell(\lambda)}, \quad \ell(\lambda) = \text{poly}(\lambda),$$

with advantage at most $2^{-\lambda^\beta}$ against time 2^{λ^β} , for some rational $0 < \beta < \gamma$. The polynomial security loss is absorbed by choosing $\beta > 0$ sufficiently small.

Finally, we obtain a strong PRF from this pseudorandom generator. For this, we apply the GGM binary tree construction [[GGM86](#)] with depth λ and a uniform root seed in $\{0, 1\}^{\ell(\lambda)}$. The input bits select a path, and the function outputs one bit of the resulting leaf seed. The standard hybrid argument turns a PRF distinguisher of time T and advantage δ into a PRG distinguisher of time $\text{poly}(T, \lambda)$ and advantage at least $\text{poly}(\delta, 1/\lambda, 1/T)$. Fix a rational $0 < \varepsilon < \beta$. As a consequence of the preceding discussion, a violation of [Equation \(7\)](#) would give a PRG distinguisher with time $2^{O(\lambda^\varepsilon)}$ and advantage at least $2^{-O(\lambda^\varepsilon)}$, contradicting the security of the pseudorandom generator for all sufficiently large λ . Finally, each function evaluation uses λ generator evaluations, so fixing the key gives a circuit of size at most λ^d for a fixed d . $\square$

4.2.2 Lower Bound

The proof of [Theorem 4.1](#) uses a uniformly random target, which typically has no polynomial size circuit. To obtain a lower bound for small circuits, we replace that target by a member of the family from [Lemma 4.2](#), restricted to a smaller subcube. The subcube is small enough that a distinguisher can evaluate the learner's correlation exactly within the PRF security bound. Our use of a smaller input dimension for this purpose is inspired by a similar scaling argument in the natural proofs barrier construction [[RR97](#)].

Theorem 4.3 (Conditional Classical Hardness for Polynomial Size Concepts). *Assume the existence of subexponentially secure one-way functions. There is a fixed integer $K \geq 1$ such that, for every $\rho > 0$ and $p > 1/2$, no uniform randomized classical algorithm $\mathcal{A}$ has the following property: for every sufficiently large input length m and every target $f \in \text{SIZE}[m^K]$ on m bits, it runs in $o(2^{m/2})$ time with membership queries and outputs a deterministic Boolean circuit h satisfying*

$$\Pr_{\mathcal{A}}[\text{corr}(H, F) \geq \rho 2^{-m/2}] \geq p,$$

where $F(x) = (-1)^{f(x)}$ and $H(x) = (-1)^{h(x)}$.

Proof. Take $\mathcal{G}, d, \varepsilon$ from [Lemma 4.2](#), and set $K = \lceil d/\varepsilon \rceil + 1$. Suppose a learner $\mathcal{A}$ with the stated guarantee exists. Truncate it after its promised running time, and output a fixed constant circuit if it has not produced a valid circuit by then. This preserves its guarantee on the promised targets and bounds its running time, query count, and output size on every oracle.

The Distinguisher. Let $\rho_* \triangleq \rho/2$. In addition, at security parameter λ , let

$$m \triangleq \left\lfloor \frac{\lambda^\varepsilon}{10} \right\rfloor, \quad M \triangleq 2^m.$$

Given an oracle $\mathcal{O}: \{0, 1\}^\lambda \rightarrow \{0, 1\}$, restrict it to

$$f'(x) \triangleq \mathcal{O}(x0^{\lambda-m}), \quad F'(x) \triangleq (-1)^{f'(x)}, \quad x \in \{0, 1\}^m.$$

Run $\mathcal{A}$ with membership access to f' , obtaining a circuit h' with sign representation H' . Next enumerate the entire subcube and accept precisely when

$$\text{corr}(H', F') \geq \rho_* \cdot 2^{-m/2}. \quad (8)$$

Acceptance on a Pseudorandom Target. If $\mathcal{O} = g_{\lambda,s}$, then its restriction f' has a circuit of size $O(\lambda^d) = O(m^{d/\varepsilon})$. Since $K > d/\varepsilon$, this is at most m^K for all sufficiently large λ . The learner's guarantee at threshold $\rho \cdot 2^{-m/2}$ implies acceptance at the smaller threshold in Equation (8) with probability at least $p > 1/2$, for every key s and hence for a uniform key.

Acceptance on a Random Target. If $\mathcal{O}$ is uniformly random, f' is a uniformly random Boolean function on m bits. Let $q_m = o(2^{m/2})$ bound the learner's query complexity. For all sufficiently large m we have $\rho_* \sqrt{M} > q_m$. Apply Theorem 4.1 with agreement advantage $\gamma = \rho_*/(2\sqrt{M}) > q_m/(2M)$. It follows that

$$\Pr[\text{accept} \mid \mathcal{O} \text{ is uniformly random}] \leq \frac{1}{2}.$$

As a consequence of the last two estimates, the distinguishing advantage is at least the fixed positive constant $p - 1/2$.

Running Time. The distinguisher runs $\mathcal{A}$ as a subroutine and makes M oracle queries. Each query takes time $\text{poly}(\lambda)$, since we need to pad the input. The correlation can be computed in time $2^{2m} \cdot \text{poly}(\lambda)$, since the output hypothesis h' has size at most $2^{m/2}$. Thus, by computing all queries in advance, running the learner, then computing the correlation, the total time is at most

$$2^m \cdot \text{poly}(\lambda) + o(2^{m/2}) + 2^{2m} \cdot \text{poly}(\lambda) \leq 2^{\lambda^\varepsilon/2} \cdot \text{poly}(\lambda) \leq 2^{\lambda^\varepsilon},$$

for all sufficiently large λ . The distinguisher is a uniform randomized oracle algorithm within the allowed time bound, and its constant advantage contradicts Equation (7). $\square$

By Equation (2), taking $\rho = 2a_0$ and $p = p^*$ proves Item (ii) of Theorem 1.4. The exponent K depends only on the PRF family, so the same polynomial size concept class works for every fixed $a_0 > 0$ and $p^* > 1/2$.

5 Quantum PAC Learning Advantage: Proof of Theorem 1.5

Throughout this section, fix $k \geq 1$. Assume without loss of generality that k is an integer. Write $\lambda_n = \lambda(n)$, and set

$$r_n \triangleq 2^{k \lceil \log_2 n \rceil},$$

so that $n^k \leq r_n = n^{O(1)}$. We use a standard padding argument to adjust the effective input dimension to the desired advantage. We choose targets that depend on only m of their n input bits, where

$$M = 2^m = \Theta(r_n^3 \lambda_n^{-2}).$$

With this choice of parameters, $O(r_n^3)$ quantum random examples suffice for the amplified learner from [Theorem 1.2](#) Item (ii) to achieve advantage λ_n and success probability at least $1 - 1/n^k$. Thus, for sufficiently large n , it meets the required guarantee in polynomial time.

For the classical lower bound, we adapt the earlier transcript argument ([Theorem 4.1](#)) to padded random functions, allowing the learner's hypothesis to depend arbitrarily on all n input bits. Bounding its success probability, averaged over the random target, by $1/n^k$ at advantage λ_n/n^k suffices for the unconditional assertion of [Theorem 1.5](#). In contrast, the conditional assertion concerning polynomial size concepts requires a stronger quantitative bound. We establish that the average classical success probability is bounded above by $1/(16r_n)$, leaving a noticeable gap below $1/n^k$. This gap is needed when replacing random functions by pseudorandom functions under the parameters of [Theorem 1.5](#).

We first specify how “constructivity” lets us choose m , then establish the necessary query lower bound and prove the assertions of the corollary.

Constructive Advantage. We say that a positive sequence $(\lambda_n)_{n \geq 1}$ is *constructive* if there is a deterministic algorithm that, on input 1^n , runs in time $\text{poly}(n)$ and outputs a nonnegative integer $s(n)$ such that, for every sufficiently large n ,

$$2^{-s(n)} \leq \lambda_n \leq 4 \cdot 2^{-s(n)}. \quad (9)$$

Thus only an approximation to the magnitude of λ_n is required. Note that, under [Equation \(9\)](#), the assumptions $\lambda_n 2^{n/2} = n^{\omega(1)}$ and $\lambda_n = n^{-\omega(1)}$ yield

$$s(n) = \omega(\log n), \quad \frac{n}{2} - s(n) = \omega(\log n). \quad (10)$$

5.1 Refined Query Lower Bound for Padded Functions

We will need the following second moment version of the transcript argument ([Theorem 4.1](#)).

Lemma 5.1 (Padded Query Lower Bound). *Let $1 \leq m < n$, write $M = 2^m$, and consider the target class*

$$\mathcal{D}_{n,m} = \{f_g(x, z) = g(x) : g \in \text{ALL}_m\}, \text{ where } (x, z) \in \{0, 1\}^m \times \{0, 1\}^{n-m}.$$

Suppose a randomized classical learner makes at most q membership queries and outputs a deterministic Boolean circuit h on n bits. Then, for every $\gamma > q/(2M)$,

$$\mathbb{E}_{g \sim \text{ALL}_m} \left[\Pr_{\mathcal{A}} \left[\Pr_{x,z} [h(x, z) = g(x)] \geq \frac{1}{2} + \gamma \right] \right] \leq \frac{M}{(2M\gamma - q)^2}.$$

Proof. Choose g uniformly from ALL_m , fix the learner's random coins, and condition on a feasible complete transcript. The *core input* of a query $y \in \{0, 1\}^n$ is its m -bit prefix. Let Q be the distinct queried core inputs, so $|Q| \leq q$. As in the proof of [Theorem 4.1](#), the signs $G(x) = (-1)^{g(x)}$ at unqueried inputs remain independent and uniform, even for adaptive queries. The output circuit h is fixed under this conditioning.

Define

$$w_x \triangleq \mathbb{E}_z [(-1)^{h(x,z)}] \in [-1, 1] \quad \text{and} \quad S \triangleq \sum_{x \notin Q} w_x \cdot G(x).$$

These weights are used only in the analysis, allowing h to depend arbitrarily on the padding bits. Since

$$\Pr_{x,z}[h(x,z) = g(x)] = \frac{1}{2} + \frac{1}{2M} \sum_x w_x \cdot G(x)$$

success implies

$$\sum_{x \in Q} w_x \cdot G(x) + S \geq 2M\gamma.$$

Each queried input $y = xz$ satisfies $w_x \cdot G(x) \leq |w_x| \leq 1$, so

$$\sum_{x \in Q} w_x \cdot G(x) \leq \sum_{x \in Q} |w_x| \leq |Q| \leq q.$$

Consequently, on the success event,

$$S \geq 2M\gamma - \sum_{x \in Q} w_x \cdot G(x) \geq 2M\gamma - q > 0,$$

where the strict positivity follows from $\gamma > q/(2M)$.

Recall that we choose g uniformly from ALL_m , fix the learner's random coins, and condition on a feasible complete transcript. For the following calculation, all expectations, variances, and probabilities are conditional on the transcript and coins. Under this conditioning, the weights w_x are fixed, while the unqueried signs $G(x)$ are independent and uniform in $\{-1, 1\}$. Hence $\mathbb{E}[G(x)] = 0$ and $\mathbb{E}[G(x)^2] = 1$, giving

$$\mathbb{E}[S] = \sum_{x \notin Q} w_x \cdot \mathbb{E}[G(x)] = 0.$$

For distinct unqueried inputs x_1, x_2 , independence gives $\mathbb{E}[G(x_1)G(x_2)] = \mathbb{E}[G(x_1)]\mathbb{E}[G(x_2)] = 0$. Thus all cross terms vanish when we expand S^2 , and

$$\text{Var}(S) = \mathbb{E}[S^2] = \sum_{x \notin Q} w_x^2 \leq M - |Q| \leq M.$$

The first equality uses $\mathbb{E}[S] = 0$, and the inequality uses $|w_x| \leq 1$ and the fact that there are $M - |Q|$ unqueried core inputs.

Set $a \triangleq 2M\gamma - q > 0$. Since success implies $S \geq a$ and $\mathbb{E}[S] = 0$, the success event is contained in $\{|S - \mathbb{E}[S]| \geq a\}$. Chebyshev's inequality therefore gives

$$\begin{aligned} \Pr[\text{success}] &\leq \Pr[S \geq a] \\ &\leq \Pr[|S - \mathbb{E}[S]| \geq a] \\ &\leq \frac{\text{Var}(S)}{a^2} \\ &\leq \frac{M}{(2M\gamma - q)^2}. \end{aligned}$$

Averaging over transcripts and coins proves the lemma. □

5.2 Unconditional Separation

In this section, we prove the unconditional part of [Theorem 1.5](#).

Proof. First, we introduce some parameters and establish a few useful estimates. Let $s(n)$ be given by constructivity ([Equation \(9\)](#)), and set

$$\theta_n \triangleq 2^{-s(n)}, \quad m_n \triangleq 2s(n) + 3k \lceil \log_2 n \rceil + 4, \quad M_n \triangleq 2^{m_n} = \frac{16r_n^3}{\theta_n^2}.$$

The two asymptotic conditions in [Equation \(10\)](#) give $m_n = \omega(\log n)$ and $n - m_n = \omega(\log n)$. Thus $1 \leq m_n < n$ eventually. The parameter choice has three useful consequences:

$$M_n \theta_n^2 = 16r_n^3, \quad M_n \lambda_n^2 \leq 256r_n^3, \quad \frac{M_n \theta_n}{r_n} = \frac{16r_n^2}{\theta_n} = n^{\omega(1)}. \quad (11)$$

Next, define $\mathcal{C}_n \triangleq \mathcal{D}_{n, m_n}$ (as in [Lemma 5.1](#)) for all sufficiently large n . At the finitely many remaining lengths, take the singleton containing the constant zero function and let the quantum learner output that function.

Quantum Learner. First, we observe that quantum random examples for f_g supply quantum random examples for g . In more detail, compute m_n using the algorithm for $s(n)$. After permuting registers, a quantum random example for $f_g(x, z) = g(x)$ factors as

$$|\text{QEX}_{f_g}\rangle = \left(2^{-m_n/2} \sum_x |x\rangle |g(x)\rangle \right) \otimes |+\rangle^{\otimes (n-m_n)},$$

where we have used that, for $b = n - m_n$,

$$|+\rangle^{\otimes b} = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right)^{\otimes b} = 2^{-b/2} \sum_{z \in \{0,1\}^b} |z\rangle.$$

Discard the padding registers and apply [Theorem 1.2](#) Item (ii) to the resulting independent examples for g , with $\eta_n \triangleq 1/(Lr_n^3)$, where $L \geq \max\{1, 256/a_1^2\}$ is a fixed integer. Its condition $C2^{-m_n} \leq \eta_n < 1/2$ holds eventually because $\eta_n M_n = 16/(L\theta_n^2) \rightarrow \infty$. With probability at least $1 - \eta_n \geq 1 - 1/n^k$, the output h_0 satisfies

$$\Pr_x[h_0(x) = g(x)] \geq \frac{1}{2} + a_1 \sqrt{\frac{Lr_n^3}{M_n}} \geq \frac{1}{2} + \lambda_n,$$

where the last inequality follows from $\lambda_n \sqrt{M_n} \leq 16r_n^{3/2}$ and $a_1 \sqrt{L} \geq 16$. Output $h(x, z) = h_0(x)$, preserving agreement. Computing $s(n)$ takes polynomial time, and processing the $O(r_n^3)$ examples and constructing the output circuit costs $\text{poly}(n)$ time in total due to our choice of $\eta_n = 1/(Lr_n^3)$.

Classical Lower Bound. Fix a classical polynomial time learner $\mathcal{B}$, with polynomial query bound $q(n)$. By [Equation \(11\)](#), $q(n) \leq M_n \theta_n / r_n$ eventually. Apply [Lemma 5.1](#) at advantage θ_n / r_n to obtain

$$\mathbb{E}_{g \sim \text{ALL}_{m_n}} \left[\Pr_{\mathcal{B}} \left[\Pr_{x,z} [h(x, z) = g(x)] \geq \frac{1}{2} + \frac{\theta_n}{r_n} \right] \right] \leq \frac{r_n^2}{M_n \theta_n^2} = \frac{1}{16r_n}. \quad (12)$$

Since $\theta_n / r_n \leq \lambda_n / n^k$, the same bound applies at advantage λ_n / n^k . Averaging provides a target function with success probability at most $1/(16r_n)$, and hence at most $1/n^k$, at every sufficiently large length. $\square$

5.3 Conditional Separation for $\mathfrak{C} \subseteq \text{SIZE}[\text{poly}]$

In this section, we prove the conditional part of [Theorem 1.5](#). For this, we replace the random (core) functions by restrictions of strong PRFs. We enlarge the security parameter so that exhaustive agreement testing on all 2^n inputs remains within the PRF security bound, while the resulting concepts still have polynomial size.

Proof. Define $s(n)$, θ_n , r_n , and m_n as in the unconditional construction ([Section 5.2](#)), and let $\mathcal{G}, \varepsilon, d$ be as in [Lemma 4.2](#). Choose a fixed integer A with $A > 1/\varepsilon$ and use security parameter $t_n = n^A$. Replace the concept class above, at sufficiently large input lengths, by

$$\mathcal{C}_n^{\text{PRF}} = \left\{ (x, z) \mapsto g_{t_n, \sigma}(0^{t_n - m_n} \circ x) : \sigma \in \{0, 1\}^{\ell(t_n)} \right\},$$

where $\ell(t_n)$ is the PRF key length, $|x| = m_n$, and $|x| + |z| = n$. Hardwiring the key and zero prefix gives circuits of size at most $t_n^d = n^{Ad} = \text{poly}(n)$.

The quantum upper bound argument applies unchanged. This follows from our choice of parameters (as in the proof given in [Section 5.2](#)) and the fact that the quantum learner from [Theorem 1.2](#) Item (ii) applies to every target Boolean function.

Next, we obtain a conditional lower bound against classical learners. This part depends on the security of the strong PRF, which in turn relies on the assumption that subexponentially secure OWFs exist.

Distinguisher Construction and Advantage. Fix a uniform randomized polynomial time learner $\mathcal{B}$, and suppose that at infinitely many lengths n it succeeds with probability greater than $1/n^k$ on every target in $\mathcal{C}_n^{\text{PRF}}$ at advantage λ_n/n^k . Let $\mathcal{I}$ denote the set of such input lengths.

The distinguisher computes as follows. For an oracle $\mathcal{O}: \{0, 1\}^{t_n} \rightarrow \{0, 1\}$, simulate $\mathcal{B}$ on

$$f^{\mathcal{O}}(x, z) = \mathcal{O}(0^{t_n - m_n} \circ x).$$

After the learner outputs h , accept exactly when

$$\Pr_{x, z}[h(x, z) = f^{\mathcal{O}}(x, z)] \geq \frac{1}{2} + \frac{\theta_n}{r_n}.$$

For a random oracle, the selected restriction is uniform in ALL_{m_n} , so [Equation \(12\)](#) bounds the acceptance probability by $1/(16r_n)$. On the other hand, for an input length $n \in \mathcal{I}$, every PRF key gives acceptance probability greater than $1/n^k \geq 1/r_n$, because success at advantage λ_n/n^k implies acceptance at the smaller threshold θ_n/r_n . Thus this test distinguishes the two oracle distributions with advantage at least $15/(16r_n)$.

Distinguisher Uniformity and Running Time. To implement the test by a single uniform distinguisher D , on input 1^t reject unless $t = n^A$ for an integer n above a sufficiently large fixed threshold, and then compute $s(n)$, r_n , and m_n . Clock $\mathcal{B}$ at a sufficiently large polynomial bound on every oracle, with a fixed circuit hypothesis as fallback if the clock expires. This does not change its behavior on the PRF targets, and ensures that the random oracle simulation also has polynomial running time.

Evaluate h and $f^{\mathcal{O}}$ on all 2^n inputs to test their agreement exactly. Including the simulation, circuit evaluations, and oracle queries of length $t_n = \text{poly}(n)$, the running time is

$$2^n \cdot \text{poly}(n, t_n) \leq 2^{n^{A\varepsilon}} = 2^{t_n^\varepsilon}$$

for all sufficiently large n , by the choice of $A > 1/\varepsilon$ and $t_n = n^A$.

Since $r_n = n^{O(1)}$, the distinguishing advantage $15/(16r_n)$ exceeds $2^{-t_n^\varepsilon}$ for every sufficiently large n . Together with the running time bound above, this violates Equation (7) at infinitely many security lengths t_n . This contradicts the security of the PRF family. Hence every sufficiently large n has a target in $\mathcal{C}_n^{\text{PRF}}$ on which $\mathcal{B}$ succeeds at advantage λ_n/n^k with probability at most $1/n^k$. $\square$

References

- [AC17] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Computational Complexity Conference (CCC)*, pages 22:1–22:67, 2017.
- [ACL⁺19] Srinivasan Arunachalam, Sourav Chakraborty, Troy Lee, Manaswi Paraashar, and Ronald de Wolf. Two new results about quantum exact learning. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 16:1–16:15, 2019.
- [ADGG26] Srinivasan Arunachalam, Arkopal Dutt, Sabee Grewal, and Aparna Gupte. Marton’s conjecture in polynomial time. arXiv:2609.20771, 2026.
- [AdW17] Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. *ACM SIGACT News*, 48(2):41–67, 2017.
- [AdW18] Srinivasan Arunachalam and Ronald de Wolf. Optimal quantum sample complexity of learning algorithms. *Journal of Machine Learning Research*, 19(71):1–36, 2018.
- [AGG⁺21] Srinivasan Arunachalam, Alex Bredariol Grilo, Tom Gur, Igor C. Oliveira, and Aarthi Sundaram. Quantum learning algorithms imply circuit lower bounds. In *Symposium on Foundations of Computer Science (FOCS)*, pages 562–573, 2021.
- [AGM26] Srinivasan Arunachalam, Amin Shiraz Gilani, and Nikhil S. Mande. Optimal quantum-classical separations for exact learning. arXiv:2609.38073, 2026.
- [AGS21] Srinivasan Arunachalam, Alex Bredariol Grilo, and Aarthi Sundaram. Quantum hardness of learning shallow classical circuits. *SIAM Journal on Computing*, 50(3):972–1013, 2021.
- [AS07] Alp Atıcı and Rocco A. Servedio. Quantum algorithms for learning and testing juntas. *Quantum Information Processing*, 6(5):323–348, 2007.
- [BJ98] Nader H. Bshouty and Jeffrey C. Jackson. Learning DNF over the uniform distribution using a quantum example oracle. *SIAM Journal on Computing*, 28(3):1136–1153, 1998.
- [BSL⁺19] David A. Barmherzig, Ju Sun, Po-Nan Li, T. J. Lane, and Emmanuel J. Candès. Holographic phase retrieval and reference design. *Inverse Problems*, 35(9):094001, 2019.
- [BV97] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.
- [CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. Exponential separations between learning with and without quantum memory. In *Symposium on Foundations of Computer Science (FOCS)*, pages 574–585, 2022.
- [CGL⁺05] A. Robert Calderbank, Anna C. Gilbert, Kirill Levchenko, S. Muthukrishnan, and Martin Strauss. Improved range-summable random variable construction algorithms. In *Symposium on Discrete Algorithms (SODA)*, pages 840–849, 2005.
- [CSBA⁺26] Davi Castro-Silva, Jop Briët, Srinivasan Arunachalam, Arkopal Dutt, and Tom Gur. An algorithmic polynomial Freiman–Ruzsa theorem. arXiv:2604.04547, 2026.

- [CSS15] Andrew W. Cross, Graeme Smith, and John A. Smolin. Quantum learning robust against noise. *Physical Review A*, 92(1):012327, 2015.
- [Gav12] Dmitry Gavinsky. Quantum predictive learning and communication complexity with single input. *Quantum Information and Computation*, 12(7&8):575–588, 2012.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.
- [GKLO22] Halley Goldberg, Valentine Kabanets, Zhenjian Lu, and Igor C. Oliveira. Probabilistic Kolmogorov complexity with applications to average-case complexity. In *Computational Complexity Conference CCC*, pages 16:1–16:60, 2022.
- [GKZ19] Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Physical Review A*, 99(3):032314, 2019.
- [Gol01] Oded Goldreich. *The Foundations of Cryptography - Volume 1: Basic Tools*. Cambridge University Press, 2001.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HKP21] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Information-theoretic bounds on quantum advantage in machine learning. *Physical Review Letters*, 126(19):190505, 2021.
- [IO22] Rahul Ilango and Igor C. Oliveira. Email thread “Trivial Quantum Learner”. Personal communication, February 20, 2022.
- [JGM⁺24] Sofiene Jerbi, Casper Gyurik, Simon C. Marshall, Riccardo Molteni, and Vedran Dunjko. Shadows of quantum machine learning. *Nature Communications*, 15:5676, 2024.
- [KH90] Wooshik Kim and Monson H. Hayes. Phase retrieval using two Fourier-transform intensities. *Journal of the Optical Society of America A*, 7(3):441–449, 1990.
- [KRS19] Varun Kanade, Andrea Rocchetto, and Simone Severini. Learning DNFs under product distributions via μ -biased quantum Fourier sampling. *Quantum Information and Computation*, 19(15&16):1261–1278, 2019.
- [KV94] M. Kearns and U. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, 1994.
- [LAT21] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 17(9):1013–1017, 2021.
- [NW99] Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Symposium on Theory of Computing (STOC)*, pages 384–393, 1999.
- [O’D14] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.
- [OS17] Igor C. Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds, and pseudorandomness. In *Computational Complexity Conference (CCC)*, pages 18:1–18:49, 2017.
- [PSES23] Niklas Pirnay, Ryan Sweke, Jens Eisert, and Jean-Pierre Seifert. Superpolynomial quantum-classical separation for density modeling. *Physical Review A*, 107(4):042416, 2023.
- [RR97] Alexander A. Razborov and Steven Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24–35, 1997.
- [SFP17] Maria Schuld, Mark Fingerhuth, and Francesco Petruccione. Implementing a distance-based classifier with a quantum interference circuit. *EPL*, 119(6):60002, 2017.

- [SG04] Rocco A. Servedio and Steven J. Gortler. Equivalences and separations between quantum and classical learnability. *SIAM Journal on Computing*, 33(5):1067–1092, 2004.
- [SSG24] Wilfred Salmon, Sergii Strelchuk, and Tom Gur. Provable advantage in quantum PAC learning. In *Conference on Learning Theory (COLT)*, pages 4487–4510, 2024.
- [SSHE21] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert. On the quantum versus classical learnability of discrete distributions. *Quantum*, 5:417, 2021.
- [TW14] Madhur Tulsiani and Julia Wolf. Quadratic Goldreich–Levin theorems. *SIAM Journal on Computing*, 43(2):730–766, 2014.
- [Vad12] Salil P. Vadhan. Pseudorandomness. *Found. Trends Theor. Comput. Sci.*, 7(1-3):1–336, 2012.
- [vD98] Wim van Dam. Quantum oracle interrogation: Getting all information for almost half the price. In *Symposium on Foundations of Computer Science (FOCS)*, pages 362–367, 1998.
- [Wat20] Thomas Watson. Communication complexity with small advantage. *Comput. Complex.*, 29(1):2, 2020.
- [Wil14] Ryan Williams. Nonuniform ACC circuit lower bounds. *Journal of the ACM*, 61(1):2:1–2:32, 2014.

A Fourier-Based Quantum Learner for $\mathcal{C} = \text{ALL}$ with $p^* \leq 1/2$

In this section, we review the Fourier-based quantum learner for $\mathcal{C} = \text{ALL}$ described in [AGG⁺21, Appendix A]. This learner achieves correlation $\geq \gamma \cdot 2^{-n/2}$ and its success probability $p^* \rightarrow 1/2$ from below as $\gamma \rightarrow 0$.

The Key Idea. Recall that, for a sign-valued function $F: \{0, 1\}^n \rightarrow \{-1, 1\}$, its Fourier expansion is¹¹

$$F(x) = \sum_{S \subseteq [n]} \widehat{F}(S) \chi_S(x), \quad \chi_S(x) = (-1)^{S \cdot x}, \quad \widehat{F}(S) = \mathbb{E}_{x \sim \mathcal{U}_n} [\chi_S(x) \cdot F(x)],$$

where $S \cdot x$ denotes $\sum_{i \in S} x_i \pmod{2}$. Crucially for us, $\widehat{F}(S)$ is the *correlation* between F and the parity function χ_S .

Parseval’s identity gives

$$\sum_S \widehat{F}(S)^2 = 1,$$

so $\{\widehat{F}(S)^2\}_S$ is a probability distribution. We use $S \sim \widehat{F}^2$ to denote a set S sampled from this distribution. The main observation is that with constant probability we obtain a nontrivial correlation.

Claim A.1. *For every $0 \leq \gamma \leq 1$ and every $F: \{0, 1\}^n \rightarrow \{-1, 1\}$,*

$$\Pr_{S \sim \widehat{F}^2} \left[|\widehat{F}(S)| \geq \gamma \cdot 2^{-n/2} \right] \geq 1 - \gamma^2.$$

Proof. Let $N = 2^n$, and consider the set

$$B = \left\{ S : |\widehat{F}(S)| < \frac{\gamma}{\sqrt{N}} \right\}.$$

¹¹The necessary background in Fourier analysis of Boolean functions can be found in a standard textbook such as [O’D14].

Since Fourier sampling assigns probability $\hat{F}(S)^2$ to S ,

$$\Pr_{S \sim \hat{F}^2} [S \in B] = \sum_{S \in B} \hat{F}(S)^2 \leq |B| \frac{\gamma^2}{N} \leq \gamma^2.$$

Taking the complementary event proves the claim. $\square$

Consequently, a Fourier-sampled parity χ_S or its negation has correlation $\Omega(2^{-n/2})$ with F with constant probability. The obstruction is the *sign*: Fourier sampling reveals S , but not whether χ_S or $-\chi_S$ is positively correlated with F . If the learner simply guesses the sign, its orientation succeeds with probability $1/2$ conditional on having sampled a sufficiently large coefficient. For this reason, the learner achieves constant success probability, but this probability is not larger than $1/2$.

The Standard Fourier Quantum Sampler. It remains to explain how to sample from the distribution $\hat{F}^2$. We use the standard Fourier sampling procedure, implemented using access to uniform quantum examples. Recall that

$$|\text{QEX}_f\rangle = \frac{1}{\sqrt{N}} \sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle, \quad F(x) = (-1)^{f(x)}.$$

The procedure consists of three steps. First, we convert the quantum example into a phase state. Applying a Hadamard gate to the label qubit gives

$$\begin{aligned} (I^{\otimes n} \otimes H)|\text{QEX}_f\rangle &= \frac{1}{\sqrt{N}} \sum_x |x\rangle H|f(x)\rangle \\ &= \frac{1}{\sqrt{2N}} \sum_x |x\rangle (|0\rangle + (-1)^{f(x)}|1\rangle) \\ &= \frac{1}{\sqrt{2}} \left(\frac{1}{\sqrt{N}} \sum_x |x\rangle \right) |0\rangle + \frac{1}{\sqrt{2}} \left(\frac{1}{\sqrt{N}} \sum_x F(x)|x\rangle \right) |1\rangle. \end{aligned}$$

Define the phase state

$$|F\rangle \triangleq \frac{1}{\sqrt{N}} \sum_x F(x)|x\rangle.$$

The state above can then be written as

$$\frac{1}{\sqrt{2}} \left(\frac{1}{\sqrt{N}} \sum_x |x\rangle \right) |0\rangle + \frac{1}{\sqrt{2}} |F\rangle |1\rangle.$$

Second, we measure the label qubit in the computational basis. Since the two displayed branches each have squared norm $1/2$, the outcome 1 occurs with probability exactly $1/2$. Conditional on this outcome, the input register is left in the state $|F\rangle$. If the outcome is 0, we discard the input register and repeat the procedure using a fresh quantum example.

Third, after obtaining $|F\rangle$, we apply $H^{\otimes n}$ to the input register. Since

$$H^{\otimes n}|x\rangle = \frac{1}{\sqrt{N}} \sum_{S \in \{0,1\}^n} (-1)^{S \cdot x} |S\rangle,$$

we obtain

$$\begin{aligned} H^{\otimes n}|F\rangle &= \frac{1}{N} \sum_{S \in \{0,1\}^n} \sum_{x \in \{0,1\}^n} F(x) (-1)^{S \cdot x} |S\rangle \\ &= \sum_{S \in \{0,1\}^n} \hat{F}(S) |S\rangle, \end{aligned}$$

where

$$\widehat{F}(S) = \frac{1}{N} \sum_{x \in \{0,1\}^n} F(x)(-1)^{S \cdot x}.$$

A measurement of the input register in the computational basis therefore returns S with probability

$$|\widehat{F}(S)|^2 = \widehat{F}(S)^2.$$

Hence the measurement outcome after this step is distributed exactly according to $\widehat{F}^2$, as desired.

B Orienting a Weak Hypothesis is Hard

Let $N = 2^n$ and let $H: \{0, 1\}^n \rightarrow \{-1, 1\}$ be a known hypothesis for an unknown target $F: \{0, 1\}^n \rightarrow \{-1, 1\}$. Under the promise

$$|\mathbb{E}_x[F(x)H(x)]| \geq c/\sqrt{N}$$

for fixed $c > 0$ and uniform x , *orienting* H means choosing H or $-H$ to obtain *positive correlation*.

Consider a string $z \in \{0, 1\}^N$ indexed by $x \in \{0, 1\}^n$, where

$$z(x) \triangleq (1 - F(x)H(x))/2.$$

Since H is known, query access to F and to z are equivalent. Since $\mathbb{E}_x[F(x)H(x)] = 1 - 2|z|_1/N$, where $|z|_1 \triangleq \sum_x z(x)$, orientation is the “gap majority problem”, i.e., we need to decide whether

$$|z|_1 \leq N/2 - c\sqrt{N}/2 \quad \text{or} \quad |z|_1 \geq N/2 + c\sqrt{N}/2.$$

Nayak and Wu [NW99, Corollary 1.2] show that distinguishing the two boundary weights requires

$$\Omega(\sqrt{N}) = \Omega(2^{n/2})$$

*quantum membership queries*¹² for any fixed success probability greater than $1/2$ on every promised input.

Consequently, any quantum algorithm that orients an arbitrary known hypothesis H with fixed success probability greater than $1/2$ for every promised target F requires $\Omega(2^{n/2})$ membership queries to the unknown target F in the worst case. The quantum learner in Section 3.2.2 avoids this obstacle. Its measurement distribution favors positively correlated hypotheses, without first fixing a candidate and then determining its sign.

C Limits of Parity Functions and of Pairwise Independence

Both the classical learner in Proposition 3.2 and the quantum learner in Theorem 3.3 use the quadratic Boolean family from Lemma 3.1. We examine what happens when this family is replaced by affine parities. The resulting family is efficiently computable, pairwise independent, and admits a global sign flip. Nevertheless, as explained below, this replacement fails to preserve the success guarantee of either learner.

¹²A quantum query is defined by a unitary map O_z acting on $n+1$ qubits whose action on basis states is defined as $O_z: |x, b\rangle \rightarrow |x, b \oplus z(x)\rangle$, for every $x \in \{0, 1\}^n$ and $b \in \{0, 1\}$. A q -query quantum algorithm makes at most q such calls. Computation between calls is unrestricted.

The Classical Learner: Pairwise Independence Does Not Guarantee Constant Success Probability.

The classical learner in [Proposition 3.2](#) samples a uniform seed and outputs the corresponding hypothesis. With the quadratic family, it achieves agreement at least $1/2 + 1/(4\sqrt{N})$ with probability at least $3/32$ for every target, where $N = 2^n$.

Consider instead the family of affine parities

$$h_{u,b}(x) \triangleq (u \cdot x) \oplus b, \quad \chi_u(x) \triangleq (-1)^{u \cdot x}, \quad H_{u,b}(x) \triangleq (-1)^b \chi_u(x) = (-1)^{h_{u,b}(x)},$$

where $u \in \{0, 1\}^n$ and $b \in \{0, 1\}$ are independent and uniform. Write U for this seed distribution. The signs $H_{u,b}(x)$ are unbiased and pairwise independent. Changing b to $b \oplus 1$ also gives a global sign flip.

Fix a parity target $f(x) = s \cdot x$, with sign representation $F = \chi_s$. Orthogonality of the Fourier characters gives

$$Z_{u,b} \triangleq \sqrt{N} \text{corr}(H_{u,b}, F) = \begin{cases} \sqrt{N}, & (u, b) = (s, 0), \\ -\sqrt{N}, & (u, b) = (s, 1), \\ 0, & u \neq s. \end{cases}$$

Consequently,

$$\mathbb{E}_U[Z] = 0, \quad \mathbb{E}_U[Z^2] = 1, \quad \mathbb{E}_U[Z^4] = N,$$

but

$$\Pr_U[|Z| \geq 1/2] = \frac{1}{N}.$$

In particular, the success probability of the classical learner using this family is exactly

$$\begin{aligned} & \Pr_{(u,b) \sim U} \left[\Pr_{x \sim \mathcal{U}_n} [h_{u,b}(x) = f(x)] \geq \frac{1}{2} + \frac{1}{4\sqrt{N}} \right] \\ &= \Pr_U[Z \geq 1/2] = \frac{1}{2N}. \end{aligned}$$

This example identifies the role of the fourth moment bound in the analysis of [Proposition 3.2](#). Applying the Paley-Zygmund inequality ([Lemma 2.1](#)) to $X = Z^2$, using $\mathbb{E}_U[Z^2] = 1$ and $\lambda = 1/4$, gives

$$\Pr_U[|Z| \geq 1/2] \geq \frac{9}{16\mathbb{E}_U[Z^4]}.$$

The global sign flip makes the distribution of Z symmetric, and hence

$$\Pr_U[Z \geq 1/2] = \frac{1}{2} \Pr_U[|Z| \geq 1/2] \geq \frac{9}{32\mathbb{E}_U[Z^4]}.$$

For the quadratic family, [Lemma 2.2](#) gives $\mathbb{E}_U[Z^4] \leq 3$, yielding the bound $3/32$ used in [Equation \(3\)](#) and in the classical learner's analysis. For affine parities, the same argument gives only $9/(32N)$, consistent with the exact success probability $1/(2N)$. Thus pairwise independence does not suffice to guarantee $\Omega(1)$ success probability for this classical sampling procedure.

The Quantum Learner: Affine Parities Do Not Preserve the Success Guarantee Above $1/2$. The state preparation and family transform in [Section 3.2.2](#) remain well defined for affine parities. With seeds encoded as (u, b) , the feature map

$$J(\perp) = 0^{n+1}, \quad J(x) = (x, 1)$$

is injective, has an efficiently computable inverse on its image, and satisfies $H_{u,b}(x) = (-1)^{(u,b) \cdot J(x)}$. Since there are $2N$ seeds, the quantum measurement distribution for an arbitrary target F is

$$P_F(u, b) = \frac{(1 + Z_{u,b})^2}{4N}, \quad Z_{u,b} = (-1)^b \sqrt{N} \hat{F}(u),$$

where $\hat{F}(u) = \mathbb{E}_{x \sim \mathcal{U}_n}[F(x)\chi_u(x)]$.

For the parity target $F = \chi_s$ considered above, the only output with positive correlation is $H_{s,0} = F$. Substituting the three possible values of Z gives

$$\begin{aligned} P_F(s, 0) &= \frac{1}{4} + \frac{1}{2\sqrt{N}} + \frac{1}{4N}, \\ P_F(s, 1) &= \frac{1}{4} - \frac{1}{2\sqrt{N}} + \frac{1}{4N}, \\ \sum_{u \neq s, b} P_F(u, b) &= \frac{1}{2} - \frac{1}{2N}. \end{aligned}$$

Thus the probability of meeting the agreement threshold in [Theorem 3.3](#) is exactly

$$\frac{1}{4} + \frac{1}{2\sqrt{N}} + \frac{1}{4N},$$

which tends to $1/4$. In other words, while the quantum procedure admits $\Omega(1)$ success probability on this target, it fails to achieve success probability $> 1/2$.