
Efficient Quantum Agnostic Improper Learning of Decision Trees

Sagnik Chatterjee¹

Tharrmashastha SAPV¹

Debajyoti Bera^{1,2}

¹Indraprastha Institute of Information Technology, Delhi

²Centre for Quantum Technologies, IIT-Delhi

Abstract

The agnostic setting is the hardest generalization of the PAC model since it is akin to learning with adversarial noise. In this paper, we give a poly($n, t, 1/\epsilon$) quantum algorithm for learning size t decision trees over n -bit inputs with uniform marginal over instances, in the agnostic setting, *without membership queries* (MQ). This is the first algorithm (classical or quantum) for efficiently learning decision trees without MQ. First, we construct a quantum agnostic weak learner by designing a quantum variant of the classical Goldreich-Levin algorithm that works with strongly biased function oracles. Next, we show how to quantize the agnostic boosting algorithm by Kalai and Kanade (2009) to obtain the *first* efficient quantum agnostic boosting algorithm (that has a *polynomial speedup* over existing adaptive quantum boosting algorithms). We then use the quantum agnostic boosting algorithm to boost the weak quantum agnostic learner constructed previously to obtain a quantum agnostic learner for decision trees. Using the above framework, we also give quantum decision tree learning algorithms without MQ in weaker noise models.

1 INTRODUCTION

Efficiently learning decision trees is a central problem in algorithmic learning theory since any Boolean function is learnable as a decision tree (Bshouty, 1993). There has been a large body of work (see Table 1) centered around providing theoretical guarantees for

learning decision trees under various generalizations and restrictions of the Probably Approximately Correct (PAC) model introduced by Valiant (1984). The original PAC model (Valiant, 1984) is in the noiseless setting where the learning algorithm is trained on a training set $S = \{(x_i, y_i)\}_{i \in [m]}$ consisting of m tuples of instances $x_i \in \mathbb{F}_2^n$ and their corresponding binary labels y_i . In the random classification noise (RCN) setting, the learning algorithm is trained on a set S' where each label y_i in S is flipped with a uniform probability p . In the agnostic setting (adversarial noise), each label in S is flipped with some probability which is dependent on the example. There are two types of decision tree learning algorithms: *proper learning* algorithms, where the output is a decision tree, and *improper learning* algorithms, where the output hypothesis is not necessarily required to be a decision tree. Proper learning of decision trees, *even in the noiseless setting*, is known to be computationally hard (Koch et al., 2023), and all the efficient improper learning algorithms (for different noise models) are designed to use Membership Query (MQ) oracles (see Table 1).

Downsides of MQ oracles. A MQ oracle allows a learning algorithm to fetch the label of *any desired instance* in the input space, even among the ones absent in the training set. In the famous experiment by Baum and Lang (1992), the MQ oracle was queried by the learning algorithm on instances outside the domain of the labeling function. This makes MQ oracles difficult to implement and is probably one reason that makes them unattractive to the applied machine learning community (Bshouty and Feldman, 2002; Awasthi et al., 2013), which brings us to the main question tackled in this work.

<p>Question: Does there exist a polynomial time (improper) decision tree learning algorithm without membership queries?</p>
--

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

Quantum as the silver bullet. In practice, machine learning algorithms use data in the training set to learn a hypothesis. This setup can be modeled as having query access to a random example oracle where we

sample training points according to the uniform distribution. Theoretically, it is known that the PAC+MQ model is strictly stronger than the PAC model with only random examples (Angluin, 1988; Bshouty, 1993; Feldman, 2006; Valiant, 1984). Similar to the random example oracle, access to a uniform superposition over the training set is an equivalent and a natural requirement in quantum computing. This was first demonstrated by Bshouty and Jackson (1998) where they introduced the notion of the Quantum PAC model. Many subsequent works (see Atıcı and Servedio (2007); Arunachalam and Maity (2020); Izdebski and de Wolf (2020); Chatterjee et al. (2023)) have been designed in the realizable *quantum* PAC model with access to a uniform superposition over the training examples. It is not known whether random examples are sufficient for any (quantum or classical) agnostic learning task, which was another motivation behind this work.

The query models used in our quantum algorithm for improperly learning decision trees were proposed by Bshouty and Jackson (1998) and Arunachalam and de Wolf (2017); and are generalizations of the random example oracle where the learning algorithm has query access to a superposition over all instances in the domain. A detailed description of the Quantum Example (QEX) and Quantum Agnostic Example (QAEX) oracles is given in Section 2. While the random example query model is weaker than the QEX model, the MQ model is stronger than the QEX model w.r.t. uniform marginal distribution (Bshouty and Jackson, 1998).

1.1 Our Contributions and Technical Overview

The main contribution of this work is a quantum polynomial time algorithm for improperly learning decision trees *without* MQ in the agnostic setting (and hence, in weaker noise settings). The importance is twofold.

1. To our knowledge, ours is the first *quantum algorithm* for decision tree learning (realizable or agnostic, with or without MQ).
2. Our algorithm is also the only known *efficient* agnostic PAC learning algorithm for decision trees (classical or quantum) *without* MQ¹.

We state a simplified version of our main result now.

Theorem 1. *Given m training examples, there exists a quantum algorithm for learning size- t decision trees in the agnostic setting without MQ in poly($m, t, 1/\epsilon$) time.*

¹Our result subsumes the classical realizable learning algorithm for *monotone* decision trees without MQ by O’Donnell and Servedio (2007).

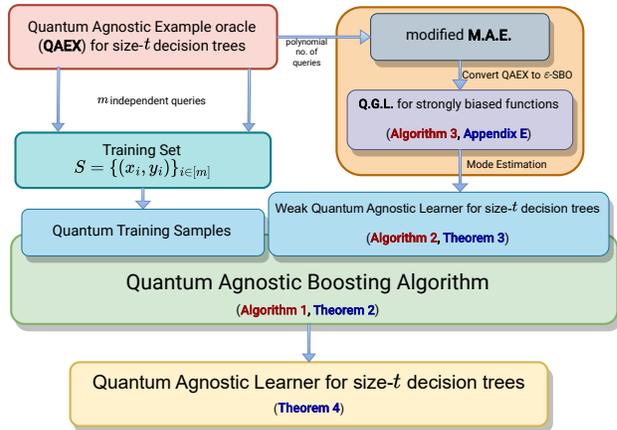


Figure 1: Agnostically learning polynomial-sized decision trees without MQ.

Here we note that the number of training samples m required for learning is polynomial w.r.t. to n where the decision trees correspond to n -bit Boolean functions. Following earlier work (see Section 1.2), we also assume a uniform marginal distribution over the instances. In Table 1, we compare our decision tree learning algorithm against existing decision tree learning algorithms. Our algorithm (see Fig. 1) follows from the existence of

- an efficient quantum agnostic boosting algorithm (see Section 3), and
- an efficient weak quantum agnostic learner for decision trees (see Section 4).

1.1.1 Quantum agnostic boosting

Quantum boosting algorithms for the realizable setting have been shown to exist (see Section 1.2), but their existence in the agnostic setting was an open question (Izdebski and de Wolf, 2020). The challenge in such algorithms is precisely estimating the margins under the presence of instance-dependent noise. Our idea was to quantize the Kalai and Kanade (2009) (KK) algorithm whose use of *relabeling* let us avoid using the amplitude amplification subroutine (a staple in the previous quantum boosting algorithms) explicitly, thereby removing a significant source of error. In Section 3, we show that given a weak quantum agnostic learner A with an associated hypothesis class \mathcal{C} and a set of m training examples, we can construct a poly($m, 1/\epsilon$) time quantum boosting algorithm to produce a hypothesis that is ϵ close to the best hypothesis in \mathcal{C} .

Table 1: Comparing different algorithms for learning size- t decision trees on n -bit Boolean functions. Note here that t and $1/\epsilon$ can be as large as $\text{poly}(n)$, which renders the running time of many of the algorithms given below as super-polynomial. Our quantum algorithms are strictly polynomial in all parameters while being the only algorithm to work in the agnostic and realizable settings and not use membership queries (denoted by MQ). Here we note that the number of training samples m required for learning is $\text{poly}(n)$. QC denotes query complexity.

Work	Setting	Type	Noise Setting	MQ	Runtime	
EH 1989	Classical	Proper	Realizable	No	$\text{poly}(n^{\log t}, 1/\epsilon)$	
KM 1991	Classical	Improper	Realizable	Yes	$\text{poly}(n, t, 1/\epsilon)$	
LMN 1993	Classical	Proper	Realizable	No	$\text{poly}(n^{\log(t/\epsilon)})$	
MR 2002	Classical	Proper	Agnostic	No	$\text{poly}(n^{\log(t/\epsilon)})$	
GKK 2008						
KK 2009	Classical	Improper	Agnostic	Yes	$\text{poly}(n, t, 1/\epsilon)$	
Feldman 2009						
BLT 2020a	Classical	Proper	Agnostic	No	$\text{poly}(n^{\log t}, 1/\epsilon)$	
This Work	Quantum	Improper	Realizable	No	$\text{poly}(n, t, 1/\epsilon)$	QC: $O(1/\epsilon^2)$
	Quantum	Improper	Agnostic	No	$\text{poly}(n, t, 1/\epsilon)$	QC: $O(n^2/\epsilon^3)$

1.1.2 Weak Quantum Agnostic Learner for Decision Trees

In Section 4, we construct a quantum weak agnostic learner for size- t decision trees using $O(n^2/\epsilon^3)$ queries to the QAEX oracle instead of the MQ oracle. The weak learner is constructed using a new quantum variant of the Goldreich-Levin algorithm (GL) (Goldreich and Levin, 1989). We use QGL to identify the monomial that best approximates the Bayes optimal predictor. This monomial serves as our weak learner. We are aware of only one prior quantum Goldreich-Levin algorithm (Adcock and Cleve, 2002); however, that algorithm involved different types of oracles and tackled a problem unrelated to ours. We now briefly touch upon the technical challenges encountered.

1. The classical GL algorithm requires obtaining $f(x)$ for specific instances x . Our QGL algorithm (see Algorithm 3), instead, was designed to work with the QAEX oracle, which generates a superposition over all $(x, f(x))$ pairs. The key step in the QGL algorithm is using a Deutsch-Jozsa-style sampler to work in tandem with the QAEX oracle. This brings us to the second technical challenge.
2. The true label $f(x)$ of any x is imperative for the classical Goldreich-Levin algorithm to work prop-

erly. However, in the agnostic scenario, both $f(x)$ (correct label) and $1 - f(x)$ (incorrect label) may be returned with non-zero probability. The probabilities could also depend on x , which makes matters worse. Thus, we designed a wrapper around QAEX denoted O_h (see Algorithm 2) employing the recent technique of multi-distribution amplitude estimation (MAE) (Bera and SAPV, 2022) to ensure a bound on the errors.

We note here that oracles in which the probability of label flips do not depend on x capture the RCN model and have been studied as *biased oracles*. To differentiate, we refer to oracles where the probability of label flips are dependent on the instance as *strongly biased oracles* — these capture the agnostic setting.

3. Our QGL algorithm is run on the wrapper oracle O_h . Unfortunately, the QGL algorithm itself uses erroneous subroutines like amplitude amplification and estimation. Such algorithms often exhibit grossly incorrect behaviors (e.g., the amplification step may amplify the amplitudes of even the undesired states due to the error arising from amplitude estimation). We meticulously ensured amplitude amplification and estimation work in tandem to keep their inherent errors in control,

particularly as the algorithm proceeds to lower levels of the prefix search tree (where the errors have a chance to accumulate).

Weaker Noise Settings. The agnostic setting generalizes the realizable and the random classification noise (RCN) settings; thus, our framework also learns decision trees in those settings, as explained in [Section 4.2](#) and [Section 4.3](#) respectively.

1.2 Related Work

Agnostic Boosting. [Kalai et al. \(2008\)](#) gave a classical agnostic boosting algorithm that achieves nearly optimal accuracy. We follow the agnostic boosting formalization of [Kalai et al. \(2008\)](#) (as opposed to earlier works like [Ben-David et al. \(2001\)](#); [Gavinsky \(2002\)](#)) in this paper. [Feldman \(2009\)](#), and [KK \(2009\)](#) came up with distribution-specific agnostic boosting algorithms to circumvent certain impossibility results on convex boosting algorithms ([Long and Servedio, 2008](#)). We give a quantum version of the KK algorithm that also achieves a quadratic speedup in the VC dimension of the weak learner.

Agnostic Learning of Decision Trees. [Ehrenfeucht and Haussler \(1989\)](#) gave the first *weakly* proper learning algorithm with quasi-polynomial running time and sample-complexity in the realizable setting using random examples. Subsequent works on properly learning decision trees ([Mehta and Raghavan, 2002](#); [Blanc et al., 2020a,b, 2022](#)) either have quasi-polynomial dependence on error parameters and intensive memory requirements or require the use of MQ (see [Table 1](#)). Recently, it was shown by [Koch et al. \(2023\)](#) that efficient proper learning of decision trees has a superpolynomial lower bound. [Bshouty \(2023\)](#) showed that the superpolynomial lower bound also holds for proper learning of monotone decision trees.

[Kushilevitz and Mansour \(1991\)](#) gave the first polynomial time improper decision tree learning algorithm (we henceforth refer to this as the KM algorithm) using MQ in the realizable setting. Their approach was later extended to the agnostic setting by [Gopalan et al. \(2008\)](#); [Kalai and Kanade \(2009\)](#); [Feldman \(2009\)](#). To our knowledge, there is no prior work on quantum agnostic learning.

Quantum Boosting. [Arunachalam and Maity \(2020\)](#) gave the first quantum adaptive boosting algorithm, which was a quantum generalization of the celebrated AdaBoost algorithm. Their approach was later extended to work on non-binary weak learners by [Chatterjee et al. \(2023\)](#). Both of the above boosting algorithms generate a quadratic speedup compared to

their classical counterparts in the VC dimension of the weak learner. This speedup is retained by our quantum agnostic boosting algorithm.

2 NOTATION AND PRELIMINARIES

Fourier Analysis of Boolean Functions. Given any Boolean function $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, where $\mathbb{F}_2^n = \{0, 1\}^n$, we can uniquely express it as $f(x) = \sum_{S \in \mathbb{F}_2^n} \hat{f}(S) \chi_S(x)$. Here $\hat{f}(S) = \mathbb{E}[f(x) \chi_S(x)] = \langle f, \chi_S \rangle$ are the Fourier coefficients corresponding to every S , and $\chi_S(x) = \prod_{i \in S} (-1)^{x_i}$, where x_i are 0-1 valued. $\chi_S(x)$ is the multilinear monomial corresponding to every S (also referred to as the parity of S). For Boolean functions, the squares of the Fourier coefficients $\hat{f}^2(S)$ form a probability distribution.

In algorithmic learning, our objective is to learn an approximation of the Fourier representation² of f by finding the set of strings S that have high $\hat{f}(S)$ values. We design a quantum variant (see [Algorithm 3](#)) of the classical GL algorithm ([Goldreich and Levin, 1989](#)) to find terms with Fourier coefficients larger than a threshold τ . The QGL algorithm searches a binary tree of all possible prefixes of n -length strings; the root corresponds to the empty prefix, and the leaves correspond to complete strings, s.t. every string represents a monomial. The weight of a node a of length s is defined as $\text{PW}(a) = \sum_{b \in \{0,1\}^{n-s}} \hat{f}^2(ab)$.

Agnostic PAC Learning. Consider an n -bit function or ‘‘concept’’ $c \in \mathcal{C} : \mathbb{F}_2^n \rightarrow \{-1, 1\}$. In the agnostic setting ([Haussler, 1992](#); [Kearns et al., 1992](#)), a learning algorithm tries to learn some unknown concept w.r.t. a fixed arbitrary joint distribution \mathcal{D} over $\mathbb{F}_2^n \times \{-1, 1\}$. The agnostic setting is seen as learning with adversarial noise in the following manner: Let \mathcal{D} be a joint distribution over the examples and the labels $\mathbb{F}_2^n \times \{-1, 1\}$. We can also interpret this as a distribution \mathcal{D}' over \mathbb{F}_2^n , where the examples are labeled according to some concept c' , s.t. an adversary corrupts some η fraction of the labels given to the algorithm. In the agnostic setting, training error of a hypothesis h , $\text{err}_S(h) = \Pr_S[h(x) \neq y]$ is defined w.r.t. set S of m labeled training examples sampled from a joint distribution \mathcal{D} over $\mathbb{F}_2^n \times \{-1, 1\}$. The generalization error is defined as $\text{err}_{\mathcal{D}}(h) = \Pr_{\mathcal{D}}[h(x) \neq y]$. Correlation is defined as follows.

Definition 1 (Correlation ([Kalai and Kanade, 2009](#))). The correlation of a hypothesis $h \in \mathcal{H}$ w.r.t. \mathcal{D} over

²For a detailed survey on the connection between Fourier representation and learning theory see [Mansour \(1994\)](#).

$\mathbb{F}_2^n \times \{-1, 1\}$ is defined as $\text{cor}_{\mathcal{D}}(h) = 1 - 2\text{err}_{\mathcal{D}}(h) = \mathbb{E}_{\mathcal{D}}[h(x) \cdot y]$.

The *optimal correlation* of a class of concepts \mathcal{C} is defined as $\text{optcor}_{\mathcal{D}}(\mathcal{C}) = \text{cor}_{\mathcal{D}}(h^*) = \text{argmax}_{h \in \mathcal{C}} \text{cor}_{\mathcal{D}}(h)$.

In agnostic PAC learning, we fix some concept class \mathcal{C} (e.g., decision trees of fixed depth) and aim to learn a hypothesis h close to the best possible concept $c_{\text{opt}} \in \mathcal{C}$. Note that h may not belong to \mathcal{C} , as in improper learning. Boosting algorithms are an important class of improper learning algorithms.

Agnostic Boosting. As discussed earlier, computational hardness results for polytime proper learning led researchers to try the improper learning approach via boosting, where they would take a “weak”-agnostic learner and boost it to obtain a better (not necessarily optimal as in the realizable case) generalization performance. We make these notions precise below.

Definition 2 ((m, κ, η) -weak Agnostic Learner (Kalai and Kanade, 2009)). For some $\kappa = O(1/\text{poly}(m))$, an algorithm A learns concept class \mathcal{C} over an arbitrary distribution \mathcal{D} on $\mathcal{X} \times \{-1, 1\}$, on m examples drawn i.i.d. from \mathcal{D} , and outputs a hypothesis h s.t. $\text{cor}_{\mathcal{D}}(h) \geq \eta \cdot \text{optcor}_{\mathcal{D}}(\mathcal{C}) - \kappa$.

Definition 3 (β -optimal (ε, δ) -agnostic PAC learner (Gavinsky, 2002)). A learning algorithm A β -optimally learns a concept class \mathcal{C} if for every $\varepsilon, \delta > 0$, $0 < \beta \leq 1/2$, any arbitrary distribution \mathcal{D} over $\mathcal{X} \times \{-1, 1\}$, A takes examples drawn i.i.d. from \mathcal{D} , and outputs a hypothesis h s.t. $\text{cor}_{\mathcal{D}}(h) \geq \text{optcor}_{\mathcal{D}}(\mathcal{C}) - \beta - \varepsilon$ with probability at least $1 - \delta$.

For brevity, we shall be referring to β -optimal (ε, δ) -agnostic PAC learners as β -optimal agnostic PAC learners. The goal of Agnostic Boosting (Ben-David et al., 2001; Gavinsky, 2002) is to produce a β -optimal learner given a (m, κ, η) -weak agnostic learner.

Kalai and Kanade (2009) introduced the concept of training intermediate weak hypotheses on *randomly relabeled examples* (instead of the traditional reweighting schemes based on AdaBoost) to obtain a β -optimal agnostic learner. In the fully supervised setting (i.e., w.r.t. this paper), the semantic differences between reweighting and relabeling are negligible.

Quantum Agnostic Learning. Classical learners have access to a random example oracle $\text{EX}(f, \mathcal{D})$ for a function f w.r.t. distribution \mathcal{D} over \mathbb{F}_2^n , which samples an instance x according to \mathcal{D} , and returns a labeled example $\langle x, f(x) \rangle$. In the agnostic case, learners have access to the oracle $\text{AEX}(\mathcal{D})$ where \mathcal{D} is a joint distribution over instances and labels. An

invocation of AEX returns a labeled instance $\langle x, y \rangle$ w.r.t. \mathcal{D} . In the Quantum PAC model (Bshouty and Jackson, 1998), the quantum learners have access to a *quantum example oracle* $\text{QEX}(f, \mathcal{D})$, s.t. each invocation to $\text{QEX}(f, \mathcal{D})$ produces the quantum state $\left(\sum_{x \in \mathbb{F}_2^n} \sqrt{\mathcal{D}(x)} |x, f(x)\rangle\right)$. In the quantum agnostic setting (Arunachalam and De Wolf, 2018), quantum learners can access the oracle $\text{QAEX}(\mathcal{D})$, s.t. each invocation of the oracle produces the quantum state $\left(\sum_{(x,y) \in \mathbb{F}_2^n \times \{-1,1\}} \sqrt{\mathcal{D}(x,y)} |x,y\rangle\right)$. We now define a (m, κ, η) -weak quantum agnostic learner.

Definition 4 ((m, κ, η) -Weak Quantum Agnostic Learner). For some $\kappa = O(1/\text{poly}(m))$, a quantum algorithm A that learns a concept class \mathcal{C} over an arbitrary distribution \mathcal{D} on $\mathcal{X} \times \{-1, 1\}$, with at most m calls to a $\text{QAEX}(\mathcal{D})$ oracle, and outputs a hypothesis h s.t. $\text{cor}_{\mathcal{D}}(h) \geq \eta \cdot \text{optcor}_{\mathcal{D}}(\mathcal{C}) - \kappa$.

We can similarly define a quantum version of a β -optimal agnostic learner.

Useful Quantum Algorithms.

Lemma 2 (Amplitude Amplification (Brassard et al., 2002)). Let there be a unitary U such that $U|0\rangle = \sqrt{a}|\phi_0\rangle|0\rangle + \sqrt{1-a}|\phi_1\rangle|1\rangle$ for an unknown a such that $a \geq p > 0$ for a known p . Then there exists a quantum amplitude amplification algorithm that makes $\Theta(\sqrt{p'}/p)$ expected number of calls to U and U^{-1} and outputs the state $|\phi_0\rangle$ with a probability $p' > 0$.

Lemma 2 allows us to boost the probability of success of a marked state with a quadratic speedup compared to probabilistic amplification algorithms.

Lemma 3 (Relative Error Estimation (Brassard et al., 2002)). Given an error parameter ε , a constant $k \geq 1$, and a unitary U such that $U|0\rangle = \sqrt{a}|\phi_0\rangle|0\rangle + \sqrt{1-a}|\phi_1\rangle|1\rangle$ where either $a \geq p$ or $a = 0$. Then there exists a quantum algorithm that produces an estimate \tilde{a} of the success probability a with probability at least $1 - \frac{1}{2^k}$ such that $|a - \tilde{a}| \leq \varepsilon a$ when $a \geq p$. The expected number of calls to U and U^{-1} made by our quantum amplitude estimation algorithm is $O\left(\frac{k}{\varepsilon\sqrt{p}} \left(1 + \log \log \frac{1}{p}\right)\right)$.

We see that Lemma 3 can be used for mean estimation with a relative error by setting $p = O(1/m)$, where $|\phi_0\rangle$ is a superposition over m basis states. This lemma follows from the amplitude estimation lemma (Theorem 15 of Brassard et al. (2002)) by setting $t = \frac{1}{m}$.

Lemma 4 (Multidistribution Amplitude Estimation. Theorem 4 of Bera and SAPV (2022)). Given an oracle O that acts as $O|0\rangle = \sum_y \alpha_y |y\rangle (\eta_{0,y}|0\rangle + \eta_{1,y}|1\rangle)$, there exists an algorithm to output the quantum state

$\sum_y \alpha_y |y\rangle (\eta_{0,y} |0\rangle + \eta_{1,y} |1\rangle) |\widetilde{\eta}_{1,y}\rangle$ in $O(1/\varepsilon)$ queries with a high probability, such that $|\widetilde{\eta}_{1,y} - \eta_{1,y}| \leq \varepsilon$.

Given a joint distribution D over $\mathcal{X} \times \mathcal{Y}$, Lemma 4 allows us to estimate the conditional probability $\Pr[\mathcal{Y} = y|x]$ to within ε accuracy over all x in superposition.

3 QUANTUM AGNOSTIC BOOSTING

In this section, we describe our quantum agnostic boosting algorithm that has query access to a (m, κ, η) -weak quantum agnostic learner A , and to its corresponding QAEX (D) oracle for an unknown joint distribution D over $\mathcal{X} \times \{-1, 1\}$. As is common in quantum boosting algorithms (see Section 1.2), we also assume access to earlier hypotheses in the form of oracles $O_{H^1}, O_{H^2}, \dots, O_{H^{t-1}}$ ³. The pseudo-code for our algorithm is given in Algorithm 1 which follows the classical Kalai and Kanade (2009) algorithm⁴.

At a very high level, Algorithm 1 iteratively computes multiple hypotheses. To compute the hypothesis, say H^t in iteration t , it first randomly relabels the examples $\{(x_i, y_i)\}_i$ in a careful manner, and then obtains a hypothesis h^t from the relabeled examples. Next, it estimates the confidence margins α_t and β_t , and depending on their values, generates H^t from h^t and the hypothesis H^{t-1} from the earlier iteration. The best among $\{H^1, H^2, \dots\}$ is returned as the strong learner. We note a few key points. Firstly, the *relabeling step* in the KK algorithm can be simulated as replacing each example (x_i, y_i) with two conservatively weighted examples: (x_i, y_i) with weight $(1 + w_i^t)/2$ and (x_i, \bar{y}_i) with weight $(1 - w_i^t)/2$. We set $w_i^t = \min\left\{1, e^{\{-H^{t-1}(x_i) \cdot y_i\}}\right\}$ which can be shown to be a conservative weighting function. It is easy to show using the Chernoff-Hoeffding bounds that α_t and β_t are good estimates of $\text{cor}_{\mathcal{D}'_{w^t}}(h^t)$ and $\text{cor}_{\mathcal{D}'_{w^t}}(-H^{t-1})$, respectively, when m is large; here, \mathcal{D}'_{w^t} denotes the distribution D relabeled by the weighting function w^t .

The quantum algorithm essentially takes care of two things: creating an oracle to return a superposition of relabeled examples and estimating the confidence margins α_t and β_t . For the latter, we use quantum mean estimation with relative error (see Section 2) to obtain estimates $\tilde{\alpha}_t$ and $\tilde{\beta}_t$ respectively. The first task is accomplished by performing standard operations on

³For conciseness, we refer to $\text{sign}(H^t)$ and $-\text{sign}(H^t)$ as H^t and $-H^t$ throughout this work. This notation can be interpreted as a confidence-weighted prediction.

⁴For completeness, we give a short simplified analysis of the Kalai and Kanade (2009) algorithm (henceforth referred to as the KK algorithm) in Appendix A.

Algorithm 1: Quantum Agnostic Boosting

Input: (m, κ, η) -weak quantum agnostic learner A and its corresponding QAEX oracle.

Initialize: $H^0 = 0$, $\varepsilon > 0$, $T = O(1/\eta^2 \varepsilon^2)$. Prepare a set S of m training samples $\{(x_i, y_i)\}_{i \in [m]}$ by measuring the output of QAEX.

Output: Hypothesis $H^{\hat{t}}$ for some $\hat{t} \in \{1, 2, \dots, T\}$ such that $\text{err}_S(H^{\hat{t}}) = \min_t \text{err}_S(H^t)$.

- 1 **for** $t = 1$ **to** T **do**
 - 2 Prepare $2 + m$ copies of
 $|\psi_0\rangle = 1/\sqrt{m} \sum_{i \in [m]} |x_i, y_i\rangle$.
 - 3 Query the oracle $O_{H^{t-1}}$ to obtain $2 + m$ copies of the state $1/\sqrt{m} \sum_{i \in [m]} |x_i, y_i\rangle |w_i^t\rangle$.
 - 4 On the last m copies, perform arithmetic operations to obtain $1/\sqrt{m} \sum_{i \in [m]} |x_i, y_i\rangle |z_i\rangle$.
 \triangleright Let $z_i = (1 + w_i^t)/2$, $z'_i = (1 - w_i^t)/2$.
 - 5 Obtain $|\phi_3\rangle$ by a conditional rotation on $|z_i\rangle$.
 $|\phi_3\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle (\sqrt{z_i} |0\rangle + \sqrt{z'_i} |1\rangle)$.
 - 6 Perform a CNOT operation on $|y_i\rangle$ with the last register as control to obtain m copies of
 $|\phi_4\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i\rangle (\sqrt{z_i} |y_i, 0\rangle + \sqrt{z'_i} |\bar{y}_i, 1\rangle)$.
 \triangleright Denote the unitary for obtaining $|\phi_4\rangle$ by QAEX $_t$.
 - 7 Obtain oracle O_{h^t} corresponding to hypothesis h^t produced by weak learner A using QAEX $_t$ as the quantum example oracle instead of QAEX.
 - 8 Invoke O_{h^t} on the 1st copy of $|\phi_0\rangle$ to obtain
 $\frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i, w_i^t, h^t(x_i)\rangle$.
 \triangleright Let $\alpha_t = \frac{1}{m} \sum_{i \in [m]} (w_i^t y_i h^t(x_i))$.
 - 9 Prepare the state

$$\sqrt{1 - \alpha_t} |\psi_0, 0\rangle + \sqrt{\alpha_t} |\psi_1, 1\rangle$$
.
 Estimate α_t as $\tilde{\alpha}_t$.
 Invoke O_{h^t} on the 2nd copy of $|\phi_0\rangle$ to obtain
 $\frac{1}{\sqrt{m}} \sum_i |x_i, y_i, w_i^t\rangle |-H^{t-1}(x_i)\rangle$.
 \triangleright Let $\beta_t = 1/m \sum_{i \in [m]} (w_i^t \cdot y_i \cdot -H^{t-1}(x_i))$.
 - 10 Prepare the state

$$\sqrt{1 - \beta_t} |\psi_0, 0\rangle + \sqrt{\beta_t} |\psi_1, 1\rangle$$
.
 Estimate β_t as $\tilde{\beta}_t$.
 If $\tilde{\alpha}_t > \tilde{\beta}_t$, $H^t = H^{t-1} + \tilde{\alpha}_t \cdot h^t$. Otherwise,
 $H^t = (1 - \tilde{\beta}_t) H^{t-1}$. Construct the oracle O_{H^t} .
 - 11 Prepare the state

$$\sqrt{1 - \beta_t} |\psi_0, 0\rangle + \sqrt{\beta_t} |\psi_1, 1\rangle$$
.
 Estimate β_t as $\tilde{\beta}_t$.
 If $\tilde{\alpha}_t > \tilde{\beta}_t$, $H^t = H^{t-1} + \tilde{\alpha}_t \cdot h^t$. Otherwise,
 $H^t = (1 - \tilde{\beta}_t) H^{t-1}$. Construct the oracle O_{H^t} .
 - 12 Return the $H^{\hat{t}}$ with the least training error on S for $\hat{t} \in [T]$.
-

a superposition state. We now state the main theorem w.r.t. the complexity and correctness of Algorithm 1, and provide further exposition and detailed proofs of Algorithm 1 in Appendix B and Appendix C.

Theorem 5 (Quantum Agnostic Boosting). *Given a (m, κ, η) -weak quantum agnostic learner A with a VC dimension of d , Algorithm 1 makes at most $\tilde{O}\left(\frac{1}{\eta^4 \varepsilon^3} \sqrt{d} \log \frac{1}{\delta}\right)$ queries to A and runs for an ad-*

ditional $\tilde{O}\left(\frac{n^2 T}{\eta^2 \varepsilon} \sqrt{d} \log \frac{1}{\delta}\right)$ time, to obtain a quantum (κ/η) optimal agnostic learner with a probability of failure of at most $5\delta T$ for any $\varepsilon > 0$ and $T = O\left(\frac{1}{\eta^2 \varepsilon^2}\right)$.

Proof Sketch. Almost all steps follow the KK algorithm. The first major source of error arises from the estimation of the margins. So, let's focus on [Line 12](#) where H^t is generated. At this point, [Algorithm 1](#) needs to determine the combined classifier for the next step. Accordingly, we pick a classifier among h^t and $-H^{t-1}$ that best correlates to the optimal classifier in the relabeled distribution and add a weighted version of it to the earlier hypothesis H^{t-1} . We denote this classifier as g^t , and observe that it's corresponding estimated confidence margin is $\gamma_t = \max(\alpha_t, \beta_t)$. Of course, the algorithm has computed only $\tilde{\alpha}_t$ and $\tilde{\beta}_t$, and we denote $\max(\tilde{\alpha}_t, \tilde{\beta}_t)$ as $\tilde{\gamma}_t$.

Observe that

$$\begin{aligned} \left| \tilde{\gamma}_t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| &\leq \left| \tilde{\gamma}_t - \gamma^t \right| + \left| \gamma^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| \\ &\leq \varepsilon \gamma^t + \eta \varepsilon / 20 \\ &\leq \eta \varepsilon / 10. \end{aligned}$$

The first inequality follows from the triangle inequality. The second inequality follows from relative estimation ([Lemma 3](#) in [Section 2](#)) and Chernoff-Hoeffding bounds. The final inequality stems from observing that $\gamma^t \leq 1$ and setting $\varepsilon = \eta \varepsilon / 20$. This shows that $\tilde{\gamma}_t$ is a good estimate of the correlation of g^t on to the relabeled distribution. Therefore, [Algorithm 1](#) chooses the right hypothesis with high probability in every iteration. Each iteration of [Algorithm 1](#) makes $\tilde{O}\left(\frac{1}{\eta \varepsilon} \sqrt{m} \log \frac{1}{\delta}\right)$ queries for estimating various quantities using [Lemma 3](#). This gives us the required query complexity.

Finally, we note that there are three points of failure in every iteration of [Algorithm 1](#): (a) Estimation of $\tilde{\gamma}_t$ fails w.p. $\leq 3\delta$, (b) weak learner A fails to produce a hypothesis w.p. $\leq \delta$, and (c) estimating the correlation of g^t fails w.p. $\leq \delta$. Therefore the entire algorithm fails w.p. at most $5\delta T$. ■

4 QUANTUM DECISION TREE LEARNING WITHOUT MEMBERSHIP QUERIES

This section shows how to obtain efficient decision tree learning algorithms in the agnostic setting without membership queries, in particular, using only states that are superpositions of pairs of random examples and labels provided by the QAEX oracle. We use

an improper learning approach with two main steps: Obtain a weak learner and then use an appropriate boosting algorithm to obtain a strong learner (see [Fig. 1](#)). Since the agnostic/adversarial noise setting is the hardest generalization of PAC learning, it follows that the above blueprint would also work for designing efficient learning algorithms for decision trees without MQ for more restricted noise models such as the random classification noise model, and the realizable/noiseless model. In fact, there exist simpler algorithms for both of these restricted settings as discussed in [Section 4.2](#) and [Section 4.3](#).

4.1 The Adversarial Noise (Agnostic) setting

Here the task is to learn an unknown concept $f(x)$ (represented by a decision tree) given a QAEX oracle. There were several difficulties in using the existing techniques to construct a weak learner for the agnostic settings. For example, in the technique proposed by [Gopalan et al. \(2008\)](#), a function is implicitly constructed from the AEX oracle whose samples are used as an approximation of the true labeling function; this is, however, not possible due to the inherent differences between AEX and QAEX oracles.

The approach taken in [Section 4.2](#) would also not work since it is not entirely clear how to obtain a Fourier sampling state⁵ directly from the QAEX oracle without an explicit oracle O_f where $f(x)$ is the unknown concept we are trying to learn. Finally, the techniques of [Iwama et al. \(2005\)](#) that we used for the random classification noise setting also do not apply here since it acts only with oracles O_f^γ where the bias γ is the same for all x . However, in the agnostic setting, the bias is dependent on x .

[Algorithm 2](#) constructs a weak quantum agnostic learner for decision trees from the QAEX oracle⁶. We can then use [Algorithm 1](#) to boost this weak learner into a quantum agnostic learner for decision trees. [Algorithm 2](#) first constructs an operator O_h using the QAEX oracle that can act as a biased oracle for some predictor f such that f is an approximation of the Bayes optimal predictor f_B .

Internally, the algorithm checks for each x , in superposition, which amongst $\alpha_{0|x}$ and $\alpha_{1|x}$ is the largest and sets $h(x)$ accordingly. Further, it performs multiple such checks over multiple independent copies to reduce any error arising from the amplitude estimation of the $\alpha_{y|x}$ states. Next, [Algorithm 2](#) offloads the bulk of its work to a quantum rendering of the GL

⁵This is the state $\left(\frac{1}{\sqrt{2^{n/2}}}\right) \sum_x \hat{f}(x) |x\rangle |\dots\rangle$.

⁶[Algorithm 2](#) is detailed in [Appendix E](#).

Algorithm 2: Weak Quantum Agnostic Learner

Input: The QAEX oracle, t . **Initialize** $\delta = 1/10$, $\gamma \leq \min\{\delta/4nt^2, \varepsilon^2/8\}$.

Output: A $(m, 1/t, \varepsilon)$ WL for size- t decision trees.

- 1 Query the QAEX oracle to obtain $|\psi_1\rangle = \sum_{x,y} \sqrt{D_{x,y}} |x\rangle |y\rangle = \frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\sum_y \alpha_{y|x} |y\rangle \right)$.
- 2 Perform ℓ independent estimations of $\text{MAE}(\varepsilon, 1 - 8/\pi^2)$ conditioned on the second register to obtain $\frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\sum_y \alpha_{y|x} |y\rangle \right) (\beta_{gx} |\tilde{\alpha}_{1|x}\rangle + \beta_{bx} |\text{Err}\rangle)^{\otimes \ell}$. \triangleright Let $\ell = O(\log 1/\gamma)$.
- 3 On each of the ℓ registers, perform thresholding on 3rd register to obtain $\frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\hat{\beta}_{gx} |h(x)\rangle |\phi'(x)\rangle + \hat{\beta}_{bx} |\overline{h(x)}\rangle |\phi''(x)\rangle \right)^{\otimes \ell}$. \triangleright Let $h(x) = \mathbb{I} \left[\tilde{\alpha}_{1|x} > \frac{1}{\sqrt{2}} \right]$.
- 4 Perform majority on $|h(x)\rangle$ registers over all ℓ copies to create $\frac{1}{\sqrt{2^n}} \sum_x |x\rangle |\xi(x)\rangle |h^*(x)\rangle$.
- 5 Let O_h be the combined unitary from steps 1 to 4.
- 6 Perform a binary search over the intervals $(\tau', \tau]$ of size $\varepsilon/16$ on $(0, 1]$, to find the largest τ such that $\text{QGL}(O_h, n, \tau, \varepsilon, \delta/4 \log(\varepsilon))$ outputs a tuple (l, \tilde{S}) with $l = 1$. The search terminates if $\tau \leq 1/t$.
- 7 Return the parity monomial $\chi_{\tilde{S}}$ as our weak learner.

algorithm denoted QGL^7 .

QGL tries to approximate a decision tree with a monomial, and it's operations are motivated by the classical GL algorithm (see Section 2). The technical difficulty was to generalize it to take as input a strongly biased oracle instead of an (error-free) oracle for a Boolean function and further enhance it to contain three kinds of errors: (a) errors from the biased oracle, (b) errors arising from amplitude estimation, and (c) errors from amplitude amplification (the state that we will amplify may contain false positives arising due to the first two errors, and those will now be incorrectly amplified).

We now state the main theorem for obtaining Quantum Weak Learners, with a detailed proof in Appendix E.1.

Theorem 6 (Weak Agnostic Learner for size- t Decision Trees). *Let $\eta = 1/t$, and let $\kappa \in [0, 1/2)$. Given access to a QAEX oracle, Algorithm 2 makes $m = \tilde{O}\left(\frac{n}{\eta\kappa^3} \cdot \log \frac{1}{\kappa}\right)$ calls to the QAEX oracle and runs for an additional $\tilde{O}\left(\frac{n}{\eta\kappa^3} \cdot \log \frac{1}{\kappa}\right)$ time to obtain a (m, κ, η) -weak quantum agnostic learner for size- t decision trees w.h.p.*

Proof Sketch. Let \mathcal{C} be a family of size- t decision trees with $c \in \mathcal{C}$ as the optimal classifier. Using the Fourier expansion of c and applying Definition 1 we have $\text{cor}_D(c(x)) = \sum_{S \subseteq [n]} \hat{c}(S) \text{cor}_D(\chi_S(x))$. Kushilevitz and Mansour showed that $\sum_{S \subseteq [n]} |\hat{c}(S)| \leq t$. Using an averaging argument, we have $\max_S |\text{cor}_D(\chi_S(x))| \geq \frac{1}{t} \text{cor}_D(c(x))$.

We now claim that Algorithm 2 produces \tilde{S} s.t. $|\max_S \text{cor}_D(\chi_S(x)) - \text{cor}_D(\chi_{\tilde{S}}(x))| \leq \varepsilon$. This claim follows from the detailed analysis of the QGL algorithm (Algorithm 3; see Appendix D.1 for details).

⁷For details refer to Algorithm 3 detailed in Appendix D.

Given \tilde{S} , we have $\text{cor}_D(\chi_{\tilde{S}}(x)) \geq \frac{1}{t} \text{cor}_D(c(x)) - \varepsilon$. This is an $(m, \varepsilon, \frac{1}{t})$ -weak quantum agnostic learner w.r.t c (from Definition 4). \blacksquare

We state the main result of this work now.

Theorem 7 (Restating Theorem 1). *For any $\delta > 0$, $\varepsilon \in (0, 1/2)$, there exists a quantum learning algorithm with VC dimension d that makes $\tilde{O}\left(\frac{nt^5\sqrt{d}}{\varepsilon^6} \log(1/\delta)\right)$ queries to the QAEX oracle and takes an additional $\tilde{O}\left(\frac{n^3t^5\sqrt{d}}{\varepsilon^6} \log(1/\delta)\right)$ time for $(t\varepsilon)$ -optimal agnostic PAC learning size- t decision trees on n -bits.*

Proof Sketch. We use the weak quantum agnostic learner for size- t decision trees constructed in Algorithm 2 (set $\kappa = \varepsilon$ and $\eta = \frac{1}{t}$ in Theorem 6) as a weak learner for the quantum agnostic boosting algorithm as described in Algorithm 1. By Theorem 5, the output of Algorithm 1 is a $(t\varepsilon)$ -optimal agnostic learner for size- t decision trees. \blacksquare

4.2 The Noiseless (Realizable) Setting

Many quantum algorithms use the Fourier sampling oracle to obtain speedups over their classical counterparts. A Fourier sampling oracle (Bernstein and Vazirani, 1993a) yields the state $\sum_{S \subseteq [n]} \hat{f}(S) |S\rangle$, given access to an oracle for the function f , and upon measurement, returns S such that $\hat{f}^2(S)$ is the largest with high probability. It is, therefore, natural to use as f the labeling function in a realizable setting. Further, we can use the majority of several Fourier samples, from multiple copies of the above state, as a realizable weak learner for size- t decision trees from O_f without using membership queries (see Appendix E.2 for details). This weak learner can be fed into quantum realizable boosting algorithms Arunachalam and Maity (2020); Izdebski and de Wolf (2020) to obtain a strong PAC learner for size- t decision trees.

Algorithm 3: QGL algorithm

Input: Oracle O_h , n , threshold τ , accuracy $\epsilon \in (0, \tau)$, error $\delta' \in (0, 1/2)$

Output: A tuple (l, \tilde{S}) such that if $l = 1$ then $\hat{h}(\tilde{S}) \geq \tau - \epsilon$ and if $l = 0$ then $\hat{h}(S) < \tau$, $\forall S$

Initialize: Set L_F to be the first level that has the number of nodes r to be at least $\frac{1}{\tau^2}$ nodes. Set $L_{i,g} = L_F$.

Initialize: Set $q = \lceil \log(1/\epsilon) \rceil + 5$ and $\delta' = \delta\tau^2/8n$.

- 1 Prepare the state $|\psi_1\rangle = \frac{1}{\sqrt{r}} \sum_{p' \in L_{i,g}} |p'\rangle$
 - 2 **do**
 - 3 Append $|+\rangle |0\rangle |0^n\rangle |0^n\rangle |0^q, 0\rangle^{\otimes O(\log(1/\delta'))} |0\rangle$ to the state.
 - 4 Prepare the state $|\nu_1\rangle$ in the fourth register and the state $|\nu_2^p\rangle$ in the fifth register.
 - ▷ $|\nu_1\rangle$ and $|\nu_2^p\rangle$ are defined in [Appendix D.1](#).
 - 5 Perform the swap test with 3^{rd} register as the control state and 4^{th} and 5^{th} registers as the target state to obtain

$$\frac{1}{\sqrt{2^{|L_{i,g}|}}} \sum_{p \in L_{i,g} \times \{0,1\}} |P\rangle (|0^q\rangle |0\rangle)^{\otimes \ell} |0\rangle.$$
 - ▷ $L_{i,g}$ is the set of “good” prefixes of level $i - 1$.
 - ▷ $|P\rangle = |p\rangle (\sigma_{0,p} |0\rangle |\phi_{0,p}\rangle + \sigma_{1,p} |1\rangle |\phi_{1,p}\rangle)$.
 - ▷ ▷ Let $\ell = O(\log 1/\delta')$.
 - 6
 - 7 Perform M.A.E. $(\epsilon/2, 1 - \frac{8}{\pi^2})$ to estimate $\sigma_{0,p}$ and obtain a state of the form

$$\frac{1}{\sqrt{2^{|L_{i,g}|}}} \sum_{p \in L_{i,g} \times \{0,1\}} |P\rangle |W\rangle^{\otimes \ell} |0\rangle.$$
 - ▷ $|W\rangle = (v_{p,g} |\widetilde{\sigma_{0,p}}\rangle + v_{p,b} |E_p\rangle) |0\rangle$.
 - 8 In each of the $O(\log(1/\delta'))$ estimate registers, mark all the estimates that are at least $\frac{1}{2} + \frac{1}{2}(\tau - \epsilon)$.
 - 9 Perform a majority over all the $O(\log(1/\delta'))$ indicator registers and store the majority in the last register.
 - 10 Amplify the probability of measuring the last register as $|1\rangle$ to obtain the following state of the form with error at most $\delta'/2n$

$$\alpha_{i,g} \sum_{p \in L_{i+1,g}} |p, \xi_p, 1\rangle + \alpha_{i,b} \sum_{p \in L_{i+1,b}} |p, \xi_p, 0\rangle.$$

Measure the last qubit as m . If $m = 0$ and $i \neq n$, return to step 1.
 - 11 **while** $i \neq n + 1$
 - 12 Measure the first register as \tilde{S} and return (m, \tilde{S})
-

4.3 The Random Classification Noise setting

In this model, the labels associated with instances suffer from an independent random noise, and we can

model it as a biased oracle O_f^ϵ for the true labeling function $f(x)$ s.t. $O_f^\epsilon |x\rangle |0^{m-1}\rangle |0\rangle$ gives us the state $|x\rangle \left(\alpha |u_x\rangle |f(x)\rangle + \beta |w_x\rangle |\overline{f(x)}\rangle \right)$ with $|\alpha|^2 \geq \frac{1}{2} + \epsilon$. [Iwama et al. \(2005\)](#) showed that for any $O(T)$ query quantum algorithm that solves a problem with high probability using access to a perfect oracle, there exists an $O(T/\epsilon)$ query quantum algorithm that solves the same problem with high probability but using access to an ϵ -biased oracle⁸. Thus, to obtain a weak learner in the RCN setting, we only need to design a QGL variant using an unbiased oracle, and then use the result by [Iwama et al. \(2005\)](#) to adapt it for a biased oracle. It suffices to state that the QGL algorithm in [Algorithm 3](#) also works for unbiased oracles.

5 DISCUSSION

[Rudin et al. \(2022\)](#) lists decision tree learning as one of ten grand challenges in interpretable machine learning. Current state-of-the-art decision tree learning algorithms ([Table 1](#)) make use of membership queries that detract from human explainability. Therefore, there is a well-motivated need to move away from MQ and towards weaker query models. We give such an algorithm using QAEX queries in this work. We also remark here that the agnostic setting is particularly suitable for NISQ devices. However, since the boosting algorithms proposed in this work appear too complex to be implemented on NISQ hardware, simpler alternatives may be appealing, particularly to the practitioners of quantum ML. The ultimate goal is to obtain efficient learning algorithms for decision trees in the agnostic setting by only using random examples (from the training set). Another immediate follow-up would be obtaining lower bounds for improper learning of decision trees without MQ in the agnostic setting.

6 ACKNOWLEDGEMENTS

The authors would like to thank Marcel Hinsche for pointing out an error in an earlier version of this work.

References

N.H. Bshouty. Exact learning via the monotone theory. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 302–311, 1993. doi: 10.1109/SFCS.1993.366857.

Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

Caleb Koch, Carmen Strassle, and Li-Yang Tan. Superpolynomial lower bounds for decision tree learn-

⁸We provide a small discussion on [Iwama et al. \(2005\)](#) in [Appendix F](#) for completeness.

- ing and testing. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1962–1994. SIAM, 2023.
- Eric B Baum and Kenneth Lang. Query learning can work poorly when a human oracle is used. In *International joint conference on neural networks*, volume 8, page 8. Beijing China, 1992.
- Nader H Bshouty and Vitaly Feldman. On using extended statistical queries to avoid membership queries. *Journal of Machine Learning Research*, 2 (Feb):359–395, 2002.
- Pranjal Awasthi, Vitaly Feldman, and Varun Kanade. Learning using local membership queries. In *Conference on Learning Theory*, pages 398–431. PMLR, 2013.
- Dana Angluin. Queries and concept learning. *Machine learning*, 2(4):319–342, 1988.
- Vitaly Feldman. Optimal hardness results for maximizing agreements with monomials. In *21st Annual IEEE Conference on Computational Complexity (CCC’06)*, pages 9–pp. IEEE, 2006.
- 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.
- Alp Atıcı and Rocco A Servedio. Quantum algorithms for learning and testing juntas. *Quantum Information Processing*, 6(5):323–348, 2007.
- Srinivasan Arunachalam and Reevu Maity. Quantum boosting. In *37th Int. Conf. Mach. Learn. ICML 2020*, volume PartF16814, pages 354–364. PMLR, nov 2020. ISBN 9781713821120.
- Adam Izdebski and Ronald de Wolf. Improved Quantum Boosting, sep 2020.
- Sagnik Chatterjee, Rohan Bhatia, Parmeet Singh Chani, and Debajyoti Bera. Quantum boosting using domain-partitioning hypotheses. *Quantum Machine Intelligence*, 5(2):33, Jul 2023. ISSN 2524-4914. doi: 10.1007/s42484-023-00122-3. URL <https://doi.org/10.1007/s42484-023-00122-3>.
- Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. *ACM SIGACT News*, 48(2):41–67, 2017.
- Andrzej Ehrenfeucht and David Haussler. Learning decision trees from random examples. *Information and Computation*, 82(3):231–246, 1989.
- Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the fourier spectrum. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 455–464, 1991.
- Nathan Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, fourier transform, and learnability. *J. ACM*, 40(3):607–620, jul 1993. ISSN 0004-5411. doi: 10.1145/174130.174138. URL <https://doi.org/10.1145/174130.174138>.
- Dinesh Mehta and Vijay Raghavan. Decision tree approximations of boolean functions. *Theoretical Computer Science*, 270(1):609–623, 2002. ISSN 0304-3975. doi: [https://doi.org/10.1016/S0304-3975\(01\)00011-1](https://doi.org/10.1016/S0304-3975(01)00011-1).
- Parikshit Gopalan, Adam Tauman Kalai, and Adam R Klivans. Agnostically learning decision trees. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 527–536, 2008.
- Adam Tauman Kalai and Varun Kanade. Potential-Based Agnostic Boosting. *Proceedings of the 22nd International Conference on Neural Information Processing Systems*, pages 880–888, December 2009.
- Vitaly Feldman. Distribution-Specific Agnostic Boosting, September 2009.
- Guy Blanc, Jane Lange, and Li-Yang Tan. Provable guarantees for decision tree induction: the agnostic setting. In *International Conference on Machine Learning*, pages 941–949. PMLR, 2020a.
- Ryan O’Donnell and Rocco A Servedio. Learning monotone decision trees in polynomial time. *SIAM Journal on Computing*, 37(3):827–844, 2007.
- Oded Goldreich and Leonid A Levin. A hard-core predicate for all one-way functions. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 25–32, 1989.
- Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In *STACS 2002: 19th Annual Symposium on Theoretical Aspects of Computer Science Antibes-Juan les Pins, France, March 14–16, 2002 Proceedings*, pages 323–334. Springer, 2002.
- Debajyoti Bera and Tharmashastha SAPV. Few quantum algorithms on amplitude distribution. *arXiv preprint arXiv:2208.00162*, 2022.
- Adam Tauman Kalai, Yishay Mansour, and Elad Verbin. On agnostic boosting and parity learning. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, pages 629–638, Victoria British Columbia Canada, May 2008. ACM. ISBN 978-1-60558-047-0. doi: 10.1145/1374376.1374466.
- Shai Ben-David, Philip M. Long, and Yishay Mansour. Agnostic Boosting. *Proceedings of the 14th Annual Conference on Computational Learning Theory and and 5th European Conference on Computational Learning Theory*, pages 507–516, July 2001.

- Dmitry Gavinsky. Optimally-Smooth Adaptive Boosting and Application to Agnostic Learning. In G. Goos, J. Hartmanis, J. van Leeuwen, Nicolò Cesa-Bianchi, Masayuki Numao, and Rüdiger Reichuk, editors, *Algorithmic Learning Theory*, volume 2533, pages 98–112. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002. ISBN 978-3-540-00170-6 978-3-540-36169-5.
- Philip M. Long and Rocco A. Servedio. Random classification noise defeats all convex potential boosters. In *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, page 608–615, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605582054. doi: 10.1145/1390156.1390233. URL <https://doi.org/10.1145/1390156.1390233>.
- Guy Blanc, Neha Gupta, Jane Lange, and Li-Yang Tan. Universal guarantees for decision tree induction via a higher-order splitting criterion. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9475–9484. Curran Associates, Inc., 2020b.
- Guy Blanc, Jane Lange, Mingda Qiao, and Li-Yang Tan. Properly learning decision trees in almost polynomial time. *J. ACM*, 69(6), nov 2022. ISSN 0004-5411. doi: 10.1145/3561047. URL <https://doi.org/10.1145/3561047>.
- Nader H. Bshouty. Superpolynomial lower bounds for learning monotone classes, 2023.
- Yishay Mansour. Learning boolean functions via the fourier transform. In *Theoretical advances in neural computation and learning*, pages 391–424. Springer, 1994.
- David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and computation*, 100(1):78–150, 1992.
- Michael J Kearns, Robert E Schapire, and Linda M Sellie. Toward efficient agnostic learning. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 341–352, 1992.
- Srinivasan Arunachalam and Ronald De Wolf. Optimal quantum sample complexity of learning algorithms. *The Journal of Machine Learning Research*, 19(1):2879–2878, 2018.
- Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum Comput. Inf. ({W}ashington, {DC}, 2000)*, volume 305 of *Contemp. Math.*, pages 53–74. Amer. Math. Soc., Providence, RI, may 2002. doi: 10.1090/conm/305/05215.
- Kazuo Iwama, Rudy Raymond, and Shigeru Yamashita. General bounds for quantum biased oracles. *IPSJ Digital Courier*, 1:415–425, 2005.
- Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '93, page 11–20, New York, NY, USA, 1993a. Association for Computing Machinery. ISBN 0897915917. doi: 10.1145/167088.167097. URL <https://doi.org/10.1145/167088.167097>.
- Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistic Surveys*, 16:1–85, 2022.
- Debajyoti Bera and Sapv Tharrmashastha. Quantum and randomised algorithms for non-linearity estimation. *ACM Transactions on Quantum Computing*, 2(2):1–27, 2021.
- Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 11–20, 1993b.

CHECKLIST

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
Detailed Exposition for [Algorithm 1](#), [Algorithm 2](#), and [Algorithm 3](#) are given in [Appendix A.1](#), [Appendix E.1](#), and [Appendix D.1](#) respectively.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
Detailed analysis and proofs for [Algorithm 1](#), [Algorithm 2](#), and [Algorithm 3](#) are given in [Appendix A.1](#), [Appendix E.1](#), and [Appendix D.1](#) respectively.
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
See [Theorem 1](#), [Theorem 5](#), and [Theorem 6](#)
 - (b) Complete proofs of all theoretical results. [Yes]
See [Appendix A.1](#), [Appendix E.1](#), and [Appendix D.1](#).
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix

A The Kalai-Kanade Algorithm

Algorithm 4: The Kalai-Kanade algorithm

Input: (m, κ, η) -weak agnostic learner A with complexity R , and m labeled training samples $S = \{(x_i, y_i)\}_{i \in [m]}$.

Output: (κ/η) -optimal hypothesis $H^{\hat{t}}$ for $1 \leq \hat{t} \leq T$ such that $\text{err}_S(H^{\hat{t}}) = \mathbf{argmin}_t \text{err}_S(H^t)$.

Data: Initialize $H^0 = 0$, and a worst-case guess for T .

1 **for** $t = 1$ **to** T **do**

2 Define $w_i^t = -\phi'(z_i) = \min \left\{ 1, e^{\{-H^{t-1}(x_i) \cdot y_i\}} \right\}$.

3 **Relabeling Step:** Set $\tilde{y}_i = y_i$ w.p. $(1 + w_i^t)/2$, and w.p. $(1 - w_i^t)/2$ set $\tilde{y}_i = \bar{y}_i$.

4 Pass the set of relabeled samples $\tilde{S} = \{(x_i, \tilde{y}_i)\}_{i \in [m]}$ to A to obtain intermediate hypothesis h^t .

5 Let $\alpha_t = \frac{1}{m} \sum_{i \in [m]} (w_i^t \cdot y_i \cdot h^t(x_i))$, and $\beta_t = \frac{1}{m} \sum_{i \in [m]} (w_i^t \cdot y_i \cdot -H^{t-1}(x_i))$.

6 If $\alpha_t > \beta_t$, set $H^t = H^{t-1} + \alpha_t \cdot h^t$. Otherwise, set $H^t = (1 - \beta_t) H^{t-1}$.

We first define the *conservative* weighting function used to relabel training samples.

Definition 5 (Conservative weighting function). A function $w : \mathcal{X} \times \{-1, 1\} \rightarrow [0, 1]$ is conservative for any function $h : \mathcal{X} \rightarrow \{-1, 1\}$ if $w(x, -h(x)) = 1$ for all $x \in \mathcal{X}$.

Consider the potential function

$$\phi(z) = \begin{cases} 1 - z & \text{if } z \leq 0 \\ e^{-z} & \text{if } z > 0 \end{cases}.$$

Observe that the weights in the Kalai-Kanade algorithm are set to the negative gradients of ϕ whose argument contains a combined hypothesis from the previous iterations; therefore, we try to use the weak learner to form a combined hypothesis that lowers the potential function in gradient descent like fashion. We note here that $\phi(z)$ is differentiable everywhere and $-\phi'(z) \in \{1/e, 1\}$. We state the following lemma using this fact and Taylor's expansion.

Claim 1 (Lemma 2 of [Kalai and Kanade \(2009\)](#)). $\phi(z) - \phi(z + \varepsilon) \geq -\phi'(z) \cdot \varepsilon - \frac{\varepsilon^2}{2}$.

The Kalai-Kanade algorithm produces a combined classifier H^t on round t , which has a lower potential than H^{t-1} until the potential eventually drops from 1 in iteration $t = 1$ to (or gets arbitrarily close to) 0 for some iteration \hat{t} . Since there is a lower bound on how much the potential can drop every round, this gives us an upper bound on the number of iterations until the Kalai-Kanade algorithm converges. Finally, we see that when the potential drops to its lowest value, the combined classifier $H^{\hat{t}}$ qualifies as an agnostic learner. Let \mathcal{D} be any arbitrary joint distribution over $\mathcal{X} \times \{-1, 1\}$. We denote the resulting relabeled distribution⁹ (relabelled using any weighting function $w : \mathcal{X} \times \{-1, 1\} \rightarrow [0, 1]$) by \mathcal{D}'_w .

Claim 2 (Lemma 1 of [Kalai and Kanade \(2009\)](#)). *Given any arbitrary distribution \mathcal{D} over $\mathcal{X} \times \{-1, 1\}$, an optimal classifier c and a classifier h s.t. $c, h : \mathcal{X} \rightarrow [-1, 1]$, and a weighting function $w : \mathcal{X} \times \{-1, 1\} \rightarrow [0, 1]$ which is conservative for h , we can show that $\text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h) \geq \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h)$.*

Proof. From [Algorithm 4](#), we can see that $\mathbb{E}_{\{x, y\} \in \mathcal{D}'_w} [h(x) \cdot y] = \mathbb{E}_{\{x, y\} \in \mathcal{D}} [h(x) \cdot y \cdot w(x, y)]$. We now evaluate the quantity $\text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h)$ using [Definition 1](#).

$$\begin{aligned} \text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h) &= \text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h) + \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h) - \text{cor}_{\mathcal{D}}(c) + \text{cor}_{\mathcal{D}}(h) \\ &= \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h) + \text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}'_w}(h) + \text{cor}_{\mathcal{D}}(h) \\ &= \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h) + \mathbb{E}_{\{x, y\} \in \mathcal{D}} [c(x) \cdot y \cdot (1 - w(x, y))] - \mathbb{E}_{\{x, y\} \in \mathcal{D}} [h(x) \cdot y \cdot (1 - w(x, y))] \\ &= \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h) - \mathbb{E}_{\{x, y\} \in \mathcal{D}} [(c(x) - h(x)) \cdot y \cdot (1 - w(x, y))]. \end{aligned}$$

⁹Technically, this is $\mathcal{D}'_{w,t}$, but the usage should be apparent from the context.

The proof follows from [Definition 5](#), and the fact that $w(x, y) = -\phi'(h(x)y)$. When $h(x) = y$, we have $w(x, y) = 1/e \implies 1 - w(x, y) > 0$, and $c(x) \cdot y \leq 1$ (true for any classifier). Therefore $\text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h) \geq \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h)$. Alternatively, when $h(x) = -y$, we have $w(x, y) = 1 \implies 1 - w(x, y) = 0$ which implies $\text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(h) = \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(h)$. \blacksquare

Consider the case when $\text{cor}_{\mathcal{D}'_w}(\mathcal{C}) = 0$. In this case, the optimal classifier behaves like a random guesser under the relabeled distribution. Therefore, either the combined classifier H^{t-1} is worse than random guessing (since it was used to set the weights for relabeling), and we should use its negation as a weak agnostic learner, or the hypothesis returned by the weak learner trained on the relabeled distribution is close to optimal. Therefore, we need to pick either of these to add to the combined classifier for the next iteration. The selected hypothesis is denoted by g^t . The Kalai-Kanade algorithm combines the existing combined classifier and g^t (weighted by its correlation γ^t) to form the combined classifier for the t^{th} iteration. Next we state a result that lower bounds the drop in potential in every iteration.

Claim 3 (Lemma 3 of [Kalai and Kanade \(2009\)](#)). *Given any function $H : \mathcal{X} \rightarrow \mathbb{R}$, hypothesis $h : \mathcal{X} \rightarrow [-1, 1]$, a weight $\gamma \in \mathbb{R}$, an arbitrary joint distribution $\mathcal{D} \sim \mathcal{X} \times \{-1, 1\}$, a weighting function $w(x, y) = -\phi'(y \cdot H(x))$, and a relabeled distribution \mathcal{D}'_w , we have $\mathbb{E}_{\{x, y\} \sim \mathcal{D}}[\phi(y \cdot H(x))] - \mathbb{E}_{\{x, y\} \sim \mathcal{D}}[\phi(y \cdot (H + \gamma h)(x))] \geq \text{cor}_{\mathcal{D}'_w}(h) - \frac{\gamma^2}{2}$.*

The proof follows directly by plugging in appropriate values for z and ε in [Claim 1](#) and taking an expectation over both sides. The main result of [Kalai and Kanade \(2009\)](#), which shows that the combined classifier output by the Kalai-Kanade algorithm is an agnostic learner, is as follows.

Lemma 8 (Theorem 1 of [Kalai and Kanade \(2009\)](#)). *Let A be an (m, κ, η) -weak agnostic learner w.r.t. some concept class \mathcal{C} s.t. $\text{VCdim}(\mathcal{C}) = d$. Then, for any $\varepsilon, \delta > 0$, there exists an agnostic boosting algorithm that uses $m = O\left(\frac{1}{\eta^2 \varepsilon^2} \log \frac{1}{\delta}\right)$ examples and $T = O\left(\frac{1}{\eta^2 \varepsilon^2}\right)$ iterations, makes $\tilde{O}\left(\frac{T \cdot d}{\eta^2 \varepsilon} \log \frac{1}{\delta}\right)$ queries to A and runs for an additional $\tilde{O}\left(\frac{n^2 \cdot T \cdot d}{\eta^2 \varepsilon} \log \frac{1}{\delta}\right)$ to output a hypothesis h with probability at least $1 - O(\delta T)$, such that $\text{cor}_{\mathcal{D}}(h) \geq \text{optcor}_{\mathcal{D}}(\mathcal{C}) - \frac{\kappa}{\eta} - \varepsilon$.*

For a large enough training set size, we can give a tight enough estimate for the correlation of the new classifier g^t , which is an (m, κ, η) -weak agnostic learner. We also see from [Claim 3](#) that a confidence-based weighted combination drops the potential, and we can lower bound this drop in potential. Therefore, we can obtain an upper bound on the number of iterations of [Algorithm 4](#), such that the potential function eventually reaches the minimum possible value. The proof follows from the fact that when the potential function reaches the minimum possible value, the corresponding combined classifier is a κ/η -optimal agnostic learner.

A.1 Proof of [Lemma 8](#)

Claim 4. *Either the weak hypothesis produced by [Algorithm 4](#) on the t^{th} iteration, or the negation of the combined hypotheses up to the $t - 1^{\text{th}}$ step has a correlation greater than $\frac{\eta \varepsilon}{3}$.*

Proof. Consider the optimal hypothesis $c \in \mathcal{C}$, and the combined hypothesis produced by [Algorithm 4](#) at iteration $t - 1$ to be H^{t-1} . If H^{t-1} is not a β -optimal agnostic learner, then we have $\text{cor}_{\mathcal{D}}(c) > \text{cor}_{\mathcal{D}}(H^{t-1}) + \beta + \varepsilon$. Plugging in [Claim 2](#), we have $\text{cor}_{\mathcal{D}'_w}(c) > \text{cor}_{\mathcal{D}'_w}(H^{t-1}) + \beta + \varepsilon$.

First consider the case $\text{cor}_{\mathcal{D}'_w}(c) > \beta + \frac{\varepsilon}{2}$, where $\beta = \frac{\kappa}{\eta}$. Consider the hypothesis h^t produced by the weak learner at the t^{th} iteration in [Algorithm 4](#). By the weak learning assumption, we have $\text{cor}_{\mathcal{D}'_w}(h^t) \geq \eta \cdot \text{cor}_{\mathcal{D}'_w}(c) - \kappa \implies \text{cor}_{\mathcal{D}'_w}(h^t) \geq \frac{\eta \varepsilon}{2}$.

Now consider the other case $\beta + \frac{\varepsilon}{2} > \text{cor}_{\mathcal{D}'_w}(c)$. This implies that $\text{cor}_{\mathcal{D}'_w}(-H^{t-1}) > \frac{\varepsilon}{2}$. \blacksquare

Lemma 8 (Theorem 1 of [Kalai and Kanade \(2009\)](#)). *Let A be an (m, κ, η) -weak agnostic learner w.r.t. some concept class \mathcal{C} s.t. $\text{VCdim}(\mathcal{C}) = d$. Then, for any $\varepsilon, \delta > 0$, there exists an agnostic boosting algorithm that uses $m = O\left(\frac{1}{\eta^2 \varepsilon^2} \log \frac{1}{\delta}\right)$ examples and $T = O\left(\frac{1}{\eta^2 \varepsilon^2}\right)$ iterations, makes $\tilde{O}\left(\frac{T \cdot d}{\eta^2 \varepsilon} \log \frac{1}{\delta}\right)$ queries to A and runs for an additional $\tilde{O}\left(\frac{n^2 \cdot T \cdot d}{\eta^2 \varepsilon} \log \frac{1}{\delta}\right)$ to output a hypothesis h with probability at least $1 - O(\delta T)$, such that $\text{cor}_{\mathcal{D}}(h) \geq \text{optcor}_{\mathcal{D}}(\mathcal{C}) - \frac{\kappa}{\eta} - \varepsilon$.*

Proof. Since $\eta \in [0, \frac{1}{2})$, we have from [Claim 4](#) that $\text{cor}_{\mathcal{D}'_w}(g^t) \geq \frac{\eta\varepsilon}{3}$, where g^t is the better of the two candidate hypotheses at iteration t . Now, consider the margin γ^t of the best classifier g^t at iteration t obtained using m training samples.

$$\gamma^t = \frac{1}{m} \sum_{i \in [m]} g^t(x_i) \cdot y_i \cdot w^t(x_i, y_i).$$

This margin is simply the estimated correlation of g^t . Using Chernoff-Hoeffding bounds and setting $m = O\left(\frac{1}{\eta^2\varepsilon^2} \log \frac{1}{\delta}\right)$, we have $|\gamma^t - \text{cor}_{\mathcal{D}'_w}(g^t)| \leq O(\eta\varepsilon)$ with high probability. Setting the appropriate values for $\text{cor}_{\mathcal{D}'_w}(g^t)$ allows us to lower bound the potential drop to at least $O(\eta^2\varepsilon^2)$ in iteration $t > 0$ using [Claim 3](#).

Since the potential function is bounded in the range $[0, 1]$, and the potential drops by at least $O(\eta^2\varepsilon^2)$, in $O\left(\frac{1}{\eta^2\varepsilon^2}\right)$ iterations, [Algorithm 4](#) must produce a hypothesis such that the potential function drops to its lowest value. Consider the iteration τ in which potential drops to its lowest. From [Claim 2](#) we have

$$\begin{aligned} \text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(g^\tau) &\geq \text{cor}_{\mathcal{D}}(c) - \text{cor}_{\mathcal{D}}(g^\tau) \\ \implies \text{cor}_{\mathcal{D}}(g^\tau) &\geq \text{cor}_{\mathcal{D}}(c) - [\text{cor}_{\mathcal{D}'_w}(c) - \text{cor}_{\mathcal{D}'_w}(g^\tau)]. \end{aligned}$$

Substituting $\text{cor}_{\mathcal{D}'_w}(c) > \frac{\kappa}{\eta} + \frac{\varepsilon}{2}$ (since the potential is lowest at this iteration) and $\text{cor}_{\mathcal{D}'_w}(g^\tau) \geq \frac{\eta\varepsilon}{3}$, we have $\text{cor}_{\mathcal{D}}(H^\tau) \geq \text{cor}_{\mathcal{D}}(c) - \frac{\kappa}{\eta} - \varepsilon$. Therefore, we have that in $O\left(\frac{1}{\eta^2\varepsilon^2}\right)$ iterations, [Algorithm 4](#) produces a $\left(\frac{\kappa}{\eta}\right)$ -optimal agnostic learner. \blacksquare

B Details of Quantum Agnostic Boosting Algorithm ([Algorithm 1](#))

Prepare a set S of m training samples $\{(x_i, y_i)\}_{i \in [m]}$ by measuring the output of QAEX. At the start of every iteration, we prepare $2 + m$ copies of the uniform state

$$|\psi_0\rangle = |\phi_0\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle.$$

Then, we query the $t - 1^{\text{th}}$ oracle $O_{H^{t-1}}$.

$$\begin{aligned} \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle |0\rangle |0\rangle &\xrightarrow{O_{H^{t-1}}} \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle \left| \underbrace{-H^{t-1}(x_i) \cdot y_i}_{z_i} \right\rangle |0\rangle \\ &\rightarrow \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle |z_i\rangle |w_i^t\rangle. \end{aligned}$$

The second step uses arithmetic operations to compute $w_i^t = \min\{1, e^{\{-H^{t-1}(x_i) \cdot y_i\}}\}$. We uncompute the $|z_i\rangle$ register using one query to the $O_{H^{t-1}}$ oracle to obtain $2 + m$ copies of the state

$$|\phi_2\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle |w_i^t\rangle.$$

Take the first m copies of $|\phi_2\rangle$, and perform arithmetic operations to obtain m copies of the state

$$\frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle \left| \frac{1 + w_i^t}{2} \right\rangle.$$

Perform a conditional rotation on the third register to obtain the state $|\phi_3\rangle$ as shown in [Line 5](#).

$$|\phi_3\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle \left(\sqrt{\frac{1 + w_i^t}{2}} |0\rangle + \sqrt{\frac{1 - w_i^t}{2}} |1\rangle \right).$$

After we perform the C-NOT, we get Q copies of a state $|\phi_4\rangle$ with *conservatively* relabeled samples, as shown in Line 6.

$$|\phi_4\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i\rangle \left(\sqrt{\frac{1+w_i^t}{2}} |y_i, 0\rangle + \sqrt{\frac{1-w_i^t}{2}} |\bar{y}_i, 1\rangle \right).$$

We denote the unitary for obtaining $|\phi_4\rangle$ as QAEX_t . Now, we pass QAEX_t to the (m, κ, η) -weak quantum agnostic learner A , to obtain query access to the t^{th} intermediate hypothesis h^t . Note that the weak learner A obtains the intermediate hypothesis using QAEX_t as the quantum example oracle instead of QAEX .

At this point, we have two copies of $|\phi_2\rangle$ left over. On the first copy, use the O_{h^t} oracle to obtain

$$|\psi_3^1\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle |w_i^t\rangle |w_i^t \cdot y_i \cdot h^t(x_i)\rangle.$$

Perform a conditional rotation on the last register to obtain

$$\begin{aligned} |\psi_4^1\rangle &= \frac{1}{\sqrt{m}} \sum_{i \in [m]} \sqrt{\kappa_i} |x_i, y_i\rangle |w_i^t\rangle |\kappa_i\rangle |1\rangle \\ &\quad + \frac{1}{\sqrt{m}} \sum_{i \in [m]} \sqrt{1-\kappa_i} |x_i, y_i\rangle |w_i^t\rangle |\kappa_i\rangle |0\rangle \end{aligned}$$

where $\kappa_i = w_i^t \cdot y_i \cdot h^t(x_i)$. We can rewrite the first part as

$$\sqrt{\alpha_t} \sum_{i \in [m]} \sqrt{\frac{\kappa_i}{\sum_{i \in [m]} \kappa_i}} |x_i, y_i\rangle |w_i^t, \kappa_i, 1\rangle.$$

We perform quantum amplitude estimation with relative error ε , conditioned on the $|1\rangle$ register, to obtain an estimate $\tilde{\alpha}_t$. On the second copy, use the $O_{H^{t-1}}$ oracle to obtain the state

$$|\psi_3^2\rangle = \frac{1}{\sqrt{m}} \sum_{i \in [m]} |x_i, y_i\rangle |w_i^t\rangle |w_i^t \cdot y_i \cdot -H^{t-1}(x_i)\rangle.$$

Let $\kappa_i = w_i^t \cdot y_i \cdot -H^{t-1}(x_i)$. Perform a conditional rotation on the last register to obtain the state

$$\begin{aligned} |\psi_4^2\rangle &= \frac{1}{\sqrt{m}} \sum_{i \in [m]} \sqrt{\kappa_i} |x_i, y_i\rangle |w_i^t\rangle |\kappa_i\rangle |1\rangle \\ &\quad + \frac{1}{\sqrt{m}} \sum_{i \in [m]} \sqrt{1-\kappa_i} |x_i, y_i\rangle |w_i^t\rangle |\kappa_i\rangle |0\rangle. \end{aligned}$$

We can rewrite the first part as

$$\sqrt{\beta_t} \sum_{i \in [m]} \sqrt{\frac{\kappa_i}{\sum_{i \in [m]} \kappa_i}} |x_i, y_i\rangle |w_i^t, \kappa_i, 1\rangle.$$

Again, we perform quantum amplitude estimation with relative error ε to obtain an estimate for $\tilde{\beta}_t$. We now state the following claims.

Claim 5. *Algorithm 1* computes estimates of margins $\tilde{\alpha}_t$ and $\tilde{\beta}_t$ s.t. $\tilde{\gamma}_t = \max(\tilde{\alpha}_t, \tilde{\beta}_t)$ using $\tilde{O}\left(\frac{1}{\eta\varepsilon}\sqrt{m}\log\frac{1}{\delta}\right)$ queries. $\left|\tilde{\gamma}^t - \text{cor}_{\mathcal{D}'_w}(g^t)\right| \leq \eta\varepsilon/10$ with probability $\geq 1 - 3\delta T$.

Claim 5 shows that we can estimate the correlation of the best classifier g^t at every step $t > 0$ with a high probability.

Claim 6. *Algorithm 1* takes as input an (m, κ, η) -weak quantum agnostic learner and outputs a (κ/η) -quantum agnostic learner with a probability of failure of at most $5\delta T$.

Claim 6 shows that our algorithm succeeds with high probability.

Claim 7. *Given a weak (m, κ, η) -weak quantum agnostic learner A with a VC dimension of d , Algorithm 1 makes at most $\tilde{O}\left(\frac{1}{\eta^4 \varepsilon^3} \sqrt{d} \log \frac{1}{\delta}\right)$ queries to A .*

Claim 7 gives an upper bound on the query complexity of our boosting algorithm.

Combining the three claims, we get Theorem 5, which states that the hypothesis h produced by our agnostic boosting algorithm is very close to the accuracy of the best hypothesis in the concept class \mathcal{C} with high probability, essentially guaranteeing that our boosting algorithm agnostically learns \mathcal{C} . All the proofs are given in Appendix C.

C Analysis of Algorithm 1

The analysis of Algorithm 1 relies heavily on the analysis of the classical Kalai-Kanade algorithm as presented in Appendix A and Appendix A.1.

C.1 Proof of Correctness

The following claim shows us that the estimated quantity $\tilde{\gamma}_t$ in every iteration of Algorithm 1 is good.

Claim 5. *Algorithm 1 computes estimates of margins $\tilde{\alpha}_t$ and $\tilde{\beta}_t$ s.t. $\tilde{\gamma}_t = \max(\tilde{\alpha}_t, \tilde{\beta}_t)$ using $\tilde{O}\left(\frac{1}{\eta \varepsilon} \sqrt{m} \log \frac{1}{\delta}\right)$ queries. $|\tilde{\gamma}^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t)| \leq \eta \varepsilon / 10$ with probability $\geq 1 - 3\delta T$.*

Proof. Let g^t be the classifier chosen by Algorithm 1 at the t^{th} iteration. We denote the correlation of g^t w.r.t. the relabeled distribution as $\text{cor}_{\mathcal{D}'_w}(g^t)$. Using Definition 1, we can restate this as

$$\text{cor}_{\mathcal{D}'_{w^t}}(g^t) = \mathbb{E}_{x_i, y_i \sim \mathcal{D}} [w_i^t \cdot y_i \cdot g^t(x_i)]. \quad (1)$$

Let $X_i = w_i^t \cdot y_i \cdot g^t(x_i)$ be a random variable. Applying Definition 5, we get that $X_i \in [-\frac{1}{e}, 1]$. Let $\gamma^t = \frac{1}{m} \sum_{i \in [m]} X_i$. Then by applying Chernoff-Hoeffding bounds, we have

$$\begin{aligned} \Pr \left[\left| \gamma^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| \geq \frac{\eta \varepsilon}{20} \right] \\ \leq 2 \cdot \exp \left(\frac{-2 \eta^2 \varepsilon^2}{\sum_{i=1}^m \left(1 + \frac{1}{e}\right)^2} \right) \\ \leq 2\delta. \end{aligned}$$

Therefore by setting $m = \frac{200}{\eta^2 \varepsilon^2} \log \frac{1}{\delta}$, we can obtain with probability at least $1 - 2\delta$,

$$\left| \gamma^t - \text{cor}_{\mathcal{D}'_w}(g^t) \right| \leq \frac{\eta \varepsilon}{20}. \quad (2)$$

We can obtain an estimate $\tilde{\gamma}^t$ of γ^t using Lemma 3 with probability at least $1 - \delta$, such that

$$\left| \tilde{\gamma}^t - \gamma^t \right| \leq \epsilon \cdot \gamma^t. \quad (3)$$

We note here that Eq. (3) and Claim 4 together make it impossible for the estimate $\tilde{\gamma}^t$ to be so far from the actual margin γ^t , that we end up choosing the classifier with the worse correlation.

Use triangle inequality on Eq. (2) and Eq. (3), we obtain with probability at least $1 - 3\delta$,

$$\begin{aligned} \left| \tilde{\gamma}^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| &\leq \left| \tilde{\gamma}^t - \gamma^t + \gamma^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| \\ &\leq \left| \tilde{\gamma}^t - \gamma^t \right| + \left| \gamma^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t) \right| \\ &\leq \epsilon \cdot \gamma^t + \frac{\eta \varepsilon}{20}. \end{aligned} \quad (4)$$

In the last step we observe that γ^t can be at most 1. Setting $\epsilon = \frac{\eta\epsilon}{20}$ gives us the required upper-bound on $|\tilde{\gamma}^t - \text{cor}_{\mathcal{D}'_{w^t}}(g^t)|$. As a final point, we get the required query complexity by plugging the terms of Eq. (3) into Lemma 3. ■

We now show that our boosting algorithm actually boosts the given weak learner to produce an agnostic learner.

Claim 6. *Algorithm 1 takes as input an (m, κ, η) -weak quantum agnostic learner and outputs a (κ/η) -quantum agnostic learner with a probability of failure of at most $5\delta T$.*

Proof. Using Claim 5 and Claim 3, we obtain that the drop in potential for Algorithm 1 at every iteration is bounded by at most $O(\eta^2\epsilon^2)$. We now follow the proof for Lemma 8 given in Appendix A.1 to show that Algorithm 1 produces a (κ/η) -agnostic learner in at most $O\left(\frac{1}{\eta^2\epsilon^2}\right)$ iterations.

We allow the algorithm to fail with probability 3δ during estimation of $\tilde{\gamma}^t$ (see Claim 5). We allow the algorithm to fail with another δ probability while invoking the weak learner to produce a hypothesis h^t at the t^{th} iteration. Finally, estimating the correlation of the constructed hypothesis g^t can fail with an additional probability of δ at every iteration. ■

C.2 Complexity Analysis

Claim 7. *Given a weak (m, κ, η) -weak quantum agnostic learner A with a VC dimension of d , Algorithm 1 makes at most $\tilde{O}\left(\frac{1}{\eta^4\epsilon^3}\sqrt{d}\log\frac{1}{\delta}\right)$ queries to A .*

Proof. The quantum algorithm runs for $O\left(\frac{1}{\eta^2\epsilon^2}\right)$ iterations (see Claim 6). From Claim 5, we see that each iteration makes $\tilde{O}\left(\frac{\sqrt{m}}{\eta\epsilon}\log\frac{1}{\delta}\right)$ queries. Plugging in sample complexity upper bounds from Arunachalam and De Wolf (2018), we have $m = \tilde{\Theta}\left(\frac{d}{\eta^2}\right)$ for both the classical and quantum case¹⁰, where d is the VC-dimension of the $\left(\frac{\kappa}{\eta}\right)$ -optimal agnostic learner. This gives us a total of $\tilde{O}\left(\frac{\sqrt{d}}{\eta^4\epsilon^3}\log\frac{1}{\delta}\right)$ queries made by Algorithm 1. ■

We note here that the classical algorithm has a query complexity of $\tilde{O}\left(\frac{d}{\epsilon^2}\log\frac{1}{\delta}\right)$ (Arunachalam and De Wolf, 2018). Therefore, we have a polynomial blowup in the given parameters, while we have a quadratic speedup in the VC dimension of the agnostic learner. We restate the main theorem here for completeness.

Theorem 5 (Quantum Agnostic Boosting). *Given a (m, κ, η) -weak quantum agnostic learner A with a VC dimension of d , Algorithm 1 makes at most $\tilde{O}\left(\frac{1}{\eta^4\epsilon^3}\sqrt{d}\log\frac{1}{\delta}\right)$ queries to A and runs for an additional $\tilde{O}\left(\frac{n^2 \cdot T}{\eta^2\epsilon}\sqrt{d}\log\frac{1}{\delta}\right)$ time, to obtain a quantum (κ/η) optimal agnostic learner with a probability of failure of at most $5\delta T$ for any $\epsilon > 0$ and $T = O\left(\frac{1}{\eta^2\epsilon^2}\right)$.*

D Quantum Goldreich-Levin Algorithm

Claim 8. *Given an oracle O_h , threshold τ , accuracy ϵ and error parameter δ , Algorithm 3 performs $O\left(\frac{n}{\epsilon^2\tau}\log\left(\frac{\delta\tau^2}{n}\right)\right)$ queries to O_h and outputs a pair (l, \tilde{S}) such that if $l = 1$, then $\hat{h}(\tilde{S}) \geq \tau - \epsilon$, else if $l = 0$, then $\nexists S$ such that $\hat{h}(S) \geq \tau$, both w.p. $\geq 1 - \delta$.*

D.1 Proof of correctness of Algorithm 3:

We first present how the state evolves at each level of the quantum Goldreich-Levin algorithm. Consider the i^{th} level. Let L_i denote the i^{th} level of the Goldreich-Levin tree. Also, let $L_{i,g}$ be the set of "good" prefixes of level

¹⁰Refer Theorem 14 of Arunachalam and De Wolf (2018) for the optimal quantum agnostic sample complexity.

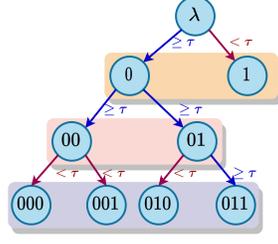


Figure 2: A partial QGL tree (up to 3 levels) indicating the level ordered traversal of “good” prefixes, that have PW above threshold τ . Bad prefixes are indicated using red arrows. The sub-trees of bad prefixes are not explored further. The set of good prefixes for level $i + 1$ is decided in superposition over the set of good prefixes of level i . Shaded boxes indicate nodes evaluated in superposition.

Algorithm 5: IGL: Interval-Search Goldreich Levin (Bera and Tharmashastha, 2021)

Input: Oracle O_h , n , accuracy ϵ and probability of error δ .

- 1 Set $k = \lceil \log_2 \frac{1}{\epsilon} \rceil + 1$ ▷ k is the smallest integer s.t. $\frac{1}{2^k} \leq \frac{\epsilon}{2}$; thus, $\frac{\epsilon}{4} < \frac{1}{2^k} \leq \frac{\epsilon}{2}$
 - 2 Set gap $g = \frac{1}{8} (\epsilon - \frac{1}{2^k})$ and threshold $\tau = \frac{1}{2}$. ▷ $8g + \frac{1}{2^k} = \epsilon \implies \frac{3}{2} \frac{\epsilon}{16} \geq g \geq \frac{\epsilon}{16}$
 - 3 **for** $i = 1 \dots k$ **do**
 - 4 Invoke $\text{QGL}(O_h, n, \tau, g, \frac{\delta}{k}) \rightarrow (l, \tilde{S})$.
 - 5 **If** $l = 0$, set $\tau = \tau + \frac{1}{2^{i+1}}$. **Else** set $\tau = \tau - \frac{1}{2^{i+1}}$.
 - 6 **If** $\frac{1}{2^{i+1}} < g/2$ or $\tau \leq \frac{1}{i}$ return \tilde{S} .
-

$i - 1$. (By “good,” we mean the prefixes p such that $\text{PW}(p)$ is greater than the threshold.) The state obtained at the end of the $i - 1^{\text{th}}$ level will be of the form

$$|\psi_i\rangle = \begin{cases} \frac{1}{|L_{i,g}|} \sum_{p' \in L_{i,g}} |p'\rangle |\phi_{p'}\rangle, & i \neq F \\ \frac{1}{\sqrt{r}} \sum_{p' \in L_F} |p'\rangle, & i = F. \end{cases}$$

Let $k = O(\log(1/\delta'))$. We append the state $|+\rangle |0\rangle^{1+2n} (|0^q\rangle |0\rangle)^{\otimes k} |0\rangle$ to $|\psi_i\rangle$ to get

$$\begin{aligned} |\psi_i^{(1)}\rangle &= \frac{1}{\sqrt{|L_{i,g}|}} \sum_{p' \in L_{i,g}} |p'\rangle |+\rangle |\phi_{p'}\rangle |0\rangle |0^n\rangle |0^n\rangle (|0^q\rangle |0\rangle)^{\otimes k} |0\rangle \\ &= \frac{1}{\sqrt{2|L_{i,g}|}} \sum_{p' \in L_{i,g}} |p'\hat{\ }0\rangle |\phi_{p'}\rangle |0\rangle |\rho\rangle |0^n\rangle |0^n\rangle (|0^q\rangle |0\rangle)^{\otimes k} |0\rangle \\ &\quad + |p'\hat{\ }1\rangle |\phi_{p'}\rangle |0\rangle |\rho\rangle |0^n\rangle |0^n\rangle (|0^q\rangle |0\rangle)^{\otimes k} |0\rangle \\ &= \frac{1}{\sqrt{2|L_{i,g}|}} \sum_{p \in L_{i,g} \times \{0,1\}} |p\rangle |\phi_p\rangle |0\rangle |0^n\rangle |0^n\rangle (|0^q\rangle |0\rangle)^{\otimes k} |0\rangle \\ &= R_1 R_2 R_3 R_4 R_5 R_6 R_7 R_8 \quad (\text{say}) \end{aligned}$$

where $p'\hat{\ }0$ and $p'\hat{\ }1$ are p' concatenated with 0 and 1 respectively, $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, $R_6 R_7 = R_{6,1} R_{7,1} \dots R_{6,k} R_{7,k}$.

Notice that the first register contains an equal superposition of all the immediate children of the “good” prefixes of the previous level. In the next step, we prepare the state $|\nu_1\rangle$ in R_4 where

$$|\nu_1\rangle = \frac{1}{2^{n/2}} \sum_x |x\rangle \left[\eta_{g,x} (-1)^{h(x)} |h(x)\rangle |\psi_{x,g}\rangle + \eta_{x,b} (-1)^{\overline{h(x)}} |\overline{h(x)}\rangle |\psi_{x,b}\rangle \right].$$

We also prepare the state $|\nu_2^p\rangle$ in R_5 where

$$|\nu_2^p\rangle = \frac{1}{2^{n/2}} \sum_x |x\rangle (-1)^{x_1 \cdot p} \left[\eta_{g,x} |h(x)\rangle |\psi_{x,g}\rangle + \eta_{x,b} |\overline{h(x)}\rangle |\psi_{x,b}\rangle \right].$$

Then, we perform the swap test with R_3 as the control qubit and R_4 and R_5 as the target qubits. This gives us,

$$\sigma_{0,p} |0\rangle |\phi_{0,p}\rangle + \sigma_{1,p} |1\rangle |\phi_{1,p}\rangle$$

as the state of the registers R_3, R_4 and R_5 for each p where $|\sigma_{0,p}|^2 = \frac{1}{2} + \frac{1}{2} |\langle \nu_1 | \nu_2^p \rangle|^2$.

Next, for each $j = 1, \dots, O(\log(1/\delta'))$, we use M.A.E to $\epsilon/2$ -estimate $|\sigma_{0,p}|^2$ in $R_{6,i}$ with error at most $1 - \frac{8}{\pi^2}$ and flip the state in $R_{7,i}$ to $|1\rangle$ if the estimate is at least $\frac{1}{2} - \frac{1}{2}(\tau - \epsilon)$. Notice that this essentially marks all the "good" states but with an error $1 - \frac{8}{\pi^2}$, i.e., the algorithm acts as a biased oracle to mark the "good" states.

As the next step, we perform a majority over $O(\log(1/\delta'))$ $R_{7,i}$ copies and store the result in R_8 . This is followed by an amplitude amplification to obtain the "good" states with high probability. For the correctness of majority followed by amplitude amplification, we direct the reader to Appendix H of [Bera and SAPV \(2022\)](#). As the last step for this level, we measure R_8 . If the measurement outcome is 1, then the post-measurement state would contain an equal superposition of all the "good" prefixes of that level.

Now, we analyze the quality of the estimate returned by the algorithm. Recall that the sum of the squares of the Fourier coefficients of a function h at all points with prefix p can be given as $\text{PW}_h(p)$

$$\begin{aligned} &= \sum_{s \in \{0,1\}^{n-|p|}} \hat{h}^2(p \frown s) \\ &= \mathbb{E}_{X_1, X_2, Z_1, Z_2} [(-1)^{f(X_1 \frown X_2) \oplus f(Z_1 \frown Z_2) \oplus p \cdot X_1 \oplus p \cdot Z_1}] \\ &= \frac{1}{2^{2n}} \sum_{x_1, x_2, z_1, z_2} (-1)^{f(x_1 \frown x_2) \oplus f(z_1 \frown z_2) \oplus p \cdot x_1 \oplus p \cdot z_1} \end{aligned}$$

where the random variables X_1 and Z_1 are samples uniformly from $\{0, 1\}^{|p|}$ and X_2 and Z_2 are samples uniformly from $\{0, 1\}^{n-|p|}$.

Now, consider the following states

$$|\nu_1\rangle = \frac{1}{2^{n/2}} \sum_x |x\rangle \left(\eta_{g,x} (-1)^{h(x)} |h(x)\rangle |\psi_{x,g}\rangle + \eta_{x,b} (-1)^{\overline{h(x)}} |\overline{h(x)}\rangle |\psi_{x,b}\rangle \right)$$

and

$$|\nu_2^p\rangle = \frac{1}{2^{n/2}} \sum_x |x\rangle (-1)^{x_1 \cdot p} \left(\eta_{g,x} |h(x)\rangle |\psi_{x,g}\rangle + \eta_{x,b} |\overline{h(x)}\rangle |\psi_{x,b}\rangle \right)$$

where x_1 is the first $|p|$ bits of x . Let $\mathcal{W} = |\langle \nu_1 | \nu_2^p \rangle|^2$. Naturally, if $\eta_{x,b} = 0$, \mathcal{W} directly yields us $\text{PW}_h(p)$. i.e,

$$\begin{aligned} \mathcal{W} &= \left(\frac{1}{2^n} \sum_x (-1)^{h(x) \oplus x_1 \cdot p} \right)^2 \\ &= \frac{1}{2^{2n}} \sum_{x, z \in \{0,1\}^n} (-1)^{h(x) \oplus h(z) \oplus p \cdot x_1 \oplus p \cdot z_1} = \text{PW}_h(p). \end{aligned}$$

However, if $\eta_{x,b} \neq 0$ for some x , then the cross terms would push the inner product away from $\text{PW}_h(p)$. Here, we show that if one is interested only in an ϵ -estimate of $\text{PW}_h(p)$, then under certain conditions on $\eta_{x,b}$, an ϵ -estimate of the inner product is not too far away from $\text{PW}_h(p)$. More concretely, we show that for an $\epsilon/2$ -estimate $\widehat{\mathcal{W}}$ of \mathcal{W} ,

$$\left| \widehat{\mathcal{W}} - \text{PW}_h(p) \right| \leq \epsilon$$

with probability at least $1 - \delta$ if $\gamma = \max_x \{\eta_{x,b}\} \leq \epsilon/8$.

Let $\eta_{x,b} \neq 0$ for some x 's. Then, we have

$$\begin{aligned} \mathcal{W} &= \frac{1}{2^{2n}} \sum_{x, z \in \{0,1\}^n} (-1)^{p \cdot x_1 \oplus p \cdot z_1} \left[\eta_{x,g}^2 \eta_{z,g}^2 (-1)^{h(x) \oplus h(z)} + \eta_{x,b}^2 \eta_{z,g}^2 (-1)^{\overline{h(x)} \oplus h(z)} + \right. \\ &\quad \left. \eta_{x,g}^2 \eta_{z,b}^2 (-1)^{h(x) \oplus \overline{h(z)}} + \eta_{x,b}^2 \eta_{z,b}^2 (-1)^{\overline{h(x)} \oplus \overline{h(z)}} \right] \end{aligned}$$

This implies

$$\begin{aligned}
 & \mathcal{W} - \text{PW}_h^2(p) \\
 &= \frac{1}{2^{2n}} \sum_{x,z \in \{0,1\}^n} (-1)^{p \cdot x_1 \oplus p \cdot z_1} \left[\eta_{x,g}^2 \eta_{z,g}^2 (-1)^{h(x) \oplus h(z)} + \eta_{x,b}^2 \eta_{z,g}^2 (-1)^{\overline{h(x)} \oplus h(z)} + \eta_{x,g}^2 \eta_{z,b}^2 (-1)^{h(x) \oplus \overline{h(z)}} \right. \\
 &\quad \left. + \eta_{x,b}^2 \eta_{z,b}^2 (-1)^{\overline{h(x)} \oplus \overline{h(z)}} \right] - \frac{1}{2^{2n}} \sum_{x,z \in \{0,1\}^n} (-1)^{h(x) \oplus h(z) \oplus p \cdot x_1 \oplus p \cdot z_1} \\
 &= \frac{1}{2^{2n}} \sum_{x,z \in \{0,1\}^n} (-1)^{p \cdot x_1 \oplus p \cdot z_1} \left[\eta_{x,g}^2 \eta_{z,g}^2 (-1)^{h(x) \oplus h(z)} + \eta_{x,b}^2 \eta_{z,g}^2 (-1)^{\overline{h(x)} \oplus h(z)} + \eta_{x,g}^2 \eta_{z,b}^2 (-1)^{h(x) \oplus \overline{h(z)}} \right. \\
 &\quad \left. + \eta_{x,b}^2 \eta_{z,b}^2 (-1)^{\overline{h(x)} \oplus \overline{h(z)}} - (-1)^{h(x) \oplus h(z)} \right]
 \end{aligned}$$

For any fixed $x, z \in \{0, 1\}^n$, let

$$\Delta_{x,z} = \eta_{x,g}^2 \eta_{z,g}^2 (-1)^{h(x) \oplus h(z)} + \eta_{x,b}^2 \eta_{z,g}^2 (-1)^{\overline{h(x)} \oplus h(z)} + \eta_{x,g}^2 \eta_{z,b}^2 (-1)^{h(x) \oplus \overline{h(z)}} + \eta_{x,b}^2 \eta_{z,b}^2 (-1)^{\overline{h(x)} \oplus \overline{h(z)}} - (-1)^{h(x) \oplus h(z)}.$$

Using the equality $1 - \eta_{x,g}^2 \eta_{z,g}^2 = \eta_{x,g}^2 \eta_{z,b}^2 + \eta_{x,b}^2 \eta_{z,g}^2 + \eta_{x,b}^2 \eta_{z,b}^2$ in the above equation, we get

$$\begin{aligned}
 \Delta_{x,z} &= \eta_{x,b}^2 \eta_{z,g}^2 [(-1)^{\overline{h(x)} \oplus h(z)} - (-1)^{h(x) \oplus h(z)}] \\
 &\quad + \eta_{x,g}^2 \eta_{z,b}^2 [(-1)^{h(x) \oplus \overline{h(z)}} - (-1)^{h(x) \oplus h(z)}] \\
 &\quad + \eta_{x,b}^2 \eta_{z,b}^2 [(-1)^{\overline{h(x)} \oplus \overline{h(z)}} - (-1)^{h(x) \oplus h(z)}]
 \end{aligned}$$

giving the equation

$$\mathcal{W} - \text{PW}_h^2(p) = \frac{1}{2^{2n}} \sum_{x,z \in \{0,1\}^n} \Delta_{x,z}.$$

Notice that for any $a, b, c, d \in \{0, 1\}$, $-2 \leq [(-1)^{a \oplus b} - (-1)^{c \oplus d}] \leq 2$. From this observation, we get that

$$\Delta_{x,z} \geq -2 \left(\eta_{x,b}^2 \eta_{z,g}^2 + \eta_{x,g}^2 \eta_{z,b}^2 + \eta_{x,b}^2 \eta_{z,b}^2 \right)$$

and

$$\Delta_{x,z} \leq 2 \left(\eta_{x,b}^2 \eta_{z,g}^2 + \eta_{x,g}^2 \eta_{z,b}^2 + \eta_{x,b}^2 \eta_{z,b}^2 \right).$$

Now,

$$\begin{aligned}
 & \eta_{x,b}^2 \eta_{z,g}^2 + \eta_{x,g}^2 \eta_{z,b}^2 + \eta_{x,b}^2 \eta_{z,b}^2 \\
 &= \eta_{x,b}^2 \eta_{z,g}^2 + \eta_{z,b}^2 (\eta_{x,g}^2 + \eta_{x,b}^2) \\
 &= \eta_{x,b}^2 \eta_{z,g}^2 + \eta_{z,b}^2 \\
 &\leq \eta_{x,b}^2 + \eta_{z,b}^2 \\
 &\leq 2\gamma.
 \end{aligned}$$

The second-last inequality follows since $\eta_{z,g}^2 \leq 1$ and the last equality follows because $\eta_{x,b}^2 \leq \gamma$ and $\eta_{z,b}^2 \leq \gamma$. This gives us that $-4\gamma \leq \Delta_{x,z} \leq 4\gamma$ implying

$$-4\gamma \leq \frac{1}{2^{2n}} \sum_{x,z \in \{0,1\}^n} \Delta_{x,z} \leq 4\gamma.$$

Or,

$$|\mathcal{W} - \text{PW}_h^2(p)| \leq 4\gamma.$$

Now, if $\gamma \leq \epsilon/8$, then $4\gamma \leq \epsilon/2$. Then, for any $\epsilon/2$ -estimate of \mathcal{W} , we have,

$$|\widehat{\mathcal{W}} - \text{PW}_h^2(p)| \leq |\widehat{\mathcal{W}} - \mathcal{W}| + |\mathcal{W} - \text{PW}_h^2(p)| \leq \epsilon.$$

Now, we show that if $\delta' < \delta\tau^2/4n$, then the probability that this algorithm fails is at most δ . The error induced due to estimation is at most δ' . The number of candidate prefixes at any level for which estimates are obtained is at most $2/\tau^2$. Using union bound on errors, the error at any level is at most the sum of errors due to the estimation and the amplification routines. This gives us $\delta_{level} \leq \frac{2\delta'}{\tau^2} + \frac{\delta}{2n}$. Hence, the total error of the algorithm is at most $n \cdot \left(\frac{2\delta'}{\tau^2} + \frac{\delta}{2n}\right) = \frac{2n\delta'}{\tau^2} + \frac{\delta}{2}$. Setting $\delta' \leq \frac{\delta\tau^2}{4n}$, the upper bound on the total error is δ .

E Quantum Decision Tree Learning: Agnostic Setting

We detail the steps of [Algorithm 2](#) as follows:

1. We start with the state $\sum_{x,y} \sqrt{\mathcal{D}_{x,y}} |x\rangle |y\rangle$. Assuming a uniform marginal distribution over \mathcal{X} , this can be written as $\frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\sum_y \alpha_{y|x} |y\rangle \right)$.
2. We make $k = O\left(\log \frac{1}{\gamma}\right)$ independent estimations using [Lemma 4](#) (M.A.E.) with parameters $(\epsilon, 1 - 8/\pi^2)$ to obtain the state

$$\frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\sum_y \alpha_{y|x} |y\rangle \right) (\beta_{gx} |\tilde{\alpha}_{1|x}\rangle + \beta_{bx} |\text{Err}\rangle)^{\otimes k}.$$

We note here that we want to set the value of $h(x)$ as the label in the third register with the larger conditional probability.

3. On each of the $k = O\left(\log \frac{1}{\gamma}\right)$ registers, perform thresholding to obtain

$$\frac{1}{\sqrt{2^n}} \sum_x |x\rangle \left(\hat{\beta}_{gx} |h(x)\rangle |\psi'(x)\rangle + \hat{\beta}_{bx} |\overline{h(x)}\rangle |\psi''(x)\rangle \right)^{\otimes k}.$$

4. Perform majority on $O\left(\log \frac{1}{\gamma}\right)$ copies of $|h(x)\rangle$.
5. Let the product of unitaries from steps 1 to 5 be denoted as O_h . Run [Algorithm 5](#) with the oracle O_h and accuracy and error parameters as ϵ and δ to obtain a string \tilde{S} .
6. Return $\chi_{\tilde{S}(x)}$ as our desired weak learner.

E.1 Proofs of Agnostic Setting

We now state the following claims, which prove the correctness and give us the query and time complexity of [Algorithm 2](#). First, we restate [Claim 8](#), which is proven in [Appendix D](#).

Claim 8. *Given an oracle O_h , threshold τ , accuracy ϵ and error parameter δ , [Algorithm 3](#) performs $O\left(\frac{n}{\epsilon^2\tau} \log\left(\frac{\delta\tau^2}{n}\right)\right)$ queries to O_h and outputs a pair (l, \tilde{S}) such that if $l = 1$, then $\hat{h}(\tilde{S}) \geq \tau - \epsilon$, else if $l = 0$, then $\nexists S$ such that $\hat{h}(S) \geq \tau$, both w.p. $\geq 1 - \delta$.*

Claim 9. *[Algorithm 2](#) performs $\tilde{O}\left(\frac{n}{\eta\kappa^3} \cdot \log \frac{1}{\kappa}\right)$ queries to QAEEX using the QGL algorithm ([Algorithm 3](#)) where κ is the accuracy parameter. The time complexity for [Algorithm 2](#) is the same as its query complexity with a logarithmic overhead.*

Claim 10. *IGL ([Algorithm 5](#)) produces \tilde{S} such that $|\text{cor}_D(\chi_{\tilde{S}(x)}) - \max_S \text{cor}_D(\chi_S(x))| \leq \kappa$.*

The proofs for [Claim 9](#) and [Claim 10](#) follows directly from [Claim 8](#) and [Algorithm 5](#).

Lemma 9 (Kushilevitz and Mansour (1991)). *Given a size- t decision tree f , the L_1 norm of its support is upper-bounded by t , i.e., $\sum_S |\hat{f}(S)| \leq t$. Such a function f is said to be t -sparse.*

Claim 11. *The parity monomial $\chi_{\tilde{S}}$ produced by Algorithm 2 is a weak agnostic learner.*

Proof. Let \mathcal{C} be a family of size- t decision trees, and let $c \in \mathcal{C}$ be the optimal classifier. Using the Fourier expansion of c and applying Definition 1 we have

$$\text{cor}_D(c(x)) = \sum_{S \subseteq [n]} \hat{c}(S) \text{cor}_D(\chi_S(x)).$$

From Lemma 9 we have $\sum_{S \subseteq [n]} |\hat{c}(S)| \leq t$. Using an averaging argument, we have

$$\max_S |\text{cor}_D(\chi_S(x))| \geq \frac{1}{t} \text{cor}_D(c(x)). \quad (5)$$

Given any estimated mode \tilde{S} such that

$$|\text{cor}_D(\chi_{\tilde{S}}(x)) - \max_S \text{cor}_D(\chi_S(x))| \leq \kappa$$

using Eq. (5), we have

$$\text{cor}_D(\chi_{\tilde{S}}(x)) \geq \frac{1}{t} \text{cor}_D(c(x)) - \kappa.$$

From Definition 4, we see that this is indeed an $(m, \kappa, \frac{1}{t})$ -weak quantum agnostic learner w.r.t. c . \blacksquare

Claim 9 gives us the final query complexity and runtime for Algorithm 2 as stated in Theorem 6. Claim 10, and Claim 11 guarantee that Algorithm 2 produces a weak learner for size- t decision trees in polynomial running time. We restate Theorem 6 below for completeness.

Theorem 6 (Weak Agnostic Learner for size- t Decision Trees). *Let $\eta = 1/t$, and let $\kappa \in [0, 1/2)$. Given access to a QAEX oracle, Algorithm 2 makes $m = \tilde{O}\left(\frac{n}{\eta\kappa^3} \cdot \log \frac{1}{\kappa}\right)$ calls to the QAEX oracle and runs for an additional $\tilde{O}\left(\frac{n}{\eta\kappa^3} \cdot \log \frac{1}{\kappa}\right)$ time to obtain a (m, κ, η) -weak quantum agnostic learner for size- t decision trees w.h.p.*

E.2 Proofs of Realizable Setting

It is well known that the output state of the Fourier Sampling algorithm can be given as $|\psi\rangle = \sum_S \hat{f}(S) |S\rangle$. Measuring the state $|\psi\rangle$ yields subset S with probability $\hat{f}(S)^2$. We use $1/\varepsilon^2$ queries to the Fourier sampling oracle O_f to estimate the mode \tilde{S} of the output distribution with ε error. This yields the $\chi_{\tilde{S}}$ term of Claim 11.

Claim 12. *Any weak agnostic learner w.r.t. h obtained by Algorithm 2 is also a weak agnostic learner w.r.t. to the Bayes optimal predictor $f_{\mathcal{B}}$.*

Proof. Using Lemma 4, we have that $|\alpha_{1|x} - \tilde{\alpha}_{1|x}| \leq \varepsilon$, for some $\varepsilon > 0$. In Algorithm 2, we set $h(x) = \mathbb{I}[\tilde{\alpha}_{1|x} > 1/\sqrt{2}]$. Therefore, we have $|\text{err}_D(h) - \text{err}_D(f_{\mathcal{B}})| \leq 2\varepsilon$. This implies that $|\text{cor}_D(f) - \text{cor}_D(f_{\mathcal{B}})| \leq 4\varepsilon$ or $\text{cor}_D(f) \in [\text{cor}_D(f_{\mathcal{B}}) - 4\varepsilon, 1]$. The upper bound is 1 since the Bayes predictor is the optimal predictor. Therefore given h s.t., $\text{cor}_D(h) \geq \eta \cdot \text{cor}_D(f) - \kappa'$, we have $\text{cor}_D(h) \geq \eta \cdot \text{cor}_D(f_{\mathcal{B}}) - \kappa$ for appropriate $\kappa', \kappa > 0$. \blacksquare

The Bayes predictor $f_{\mathcal{B}}$ is the optimal predictor on a joint distribution \mathcal{D} over $\mathcal{X} \times \{0, 1\}$, and defined as $f_{\mathcal{B}}(x) = \text{argmax}_{y \in \{0, 1\}} \Pr_{\mathcal{D}}[y|x], \forall x \in \mathcal{X}$.

Claim 13. $\chi_{\tilde{S}}$ is a weak realizable learner for size- t decision trees.

Proof. From Claim 12, we know that $\text{cor}_D(\chi_{\tilde{S}}) \geq \frac{1}{t} \text{cor}_D(f) - \kappa$. For the realizable setting, $\text{cor}_D(f) = 1$. Therefore by setting $\varepsilon = \frac{1}{2} - O\left(\frac{1}{n}\right)$ we prove that $\chi_{\tilde{S}}$ is a weak realizable learner for size- t decision trees. \blacksquare

F Discussion on Iwama et al. (2005)

Iwama et al. (2005) showed that for any T query quantum algorithm A that solves a problem with error at most δ using a *perfect oracle*, there exists an $O(T/\varepsilon)$ query algorithm A' that solves the same problem with error at most $\delta/6$ using an ε -biased oracle. Note that here we are not referring to strongly-biased oracles.

Let us assume that the oracle invoked by A is perfect. Then if a T -query algorithm A solves a problem with error at most $\delta < 1/2$, then it is possible to construct an algorithm to solve the same problem with error at most δ' by taking the majority of $O\left(\frac{8(1-\delta)}{(1-2\delta)^2} \log(1/\delta')\right)$ invocations of A .

In the case of an ε -biased oracle, the oracle outputs the correct value with probability $1/2 + \varepsilon$. If one tries to directly use A , since errors add up linearly in quantum (Bernstein and Vazirani, 1993b), the errors at each step of A will add up to $O(T\varepsilon)$.

Alternatively, one can perform some k many invocations of the biased oracle, obtain the majority, and use the value of the majority as the oracle output. This will serve as an “almost” perfect oracle. If the error at each step is bounded to at most δ/T , then we obtain an algorithm that solves the problem with error at most δ . If we were to bound the error due to the oracle at each step to at most δ , then we need to find the right value of k . Since the oracle outputs the correct value with probability $1/2 + \varepsilon$, using Hoeffding’s inequality, we can obtain the right value of k as $k = \Omega(\log(1/\delta)/\varepsilon^2)$. This would increase the query complexity of the algorithm to $\tilde{O}(T/\varepsilon^2)$. On the other hand, Iwama et al. (2005) showed that the same problem can be solved using just $\tilde{O}(T/\varepsilon)$ queries.