

# Improved Quantum Boosting

Adam Izdebski<sup>1</sup> ✉

Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland

Ronald de Wolf ✉

QuSoft, CWI and University of Amsterdam, The Netherlands

---

## Abstract

Boosting is a general method to convert a weak learner (which generates hypotheses that are just slightly better than random) into a strong learner (which generates hypotheses that are much better than random). Recently, Arunachalam and Maity [5] gave the first quantum improvement for boosting, by combining Freund and Schapire’s AdaBoost algorithm with a quantum algorithm for approximate counting. Their booster is faster than classical boosting as a function of the VC-dimension of the weak learner’s hypothesis class, but worse as a function of the quality of the weak learner. In this paper we give a substantially faster and simpler quantum boosting algorithm, based on Servedio’s SmoothBoost algorithm [22].

**2012 ACM Subject Classification** Theory of computation → Quantum computation theory; Computing methodologies → Machine learning

**Keywords and phrases** Learning theory, Boosting algorithms, Quantum computing

**Digital Object Identifier** 10.4230/LIPIcs.ESA.2023.64

**Funding** *Ronald de Wolf*: Partially supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037), and through QuantERA ERA-NET Cofund project QuantAlgo (680-91-034).

**Acknowledgements** We thank Srinivasan Arunachalam for many helpful comments, Min-Hsiu Hsieh for sending us an updated version of [24] and answering some questions about this paper, Yassine Hamoudi for answering a question about [16], and the anonymous referees for some helpful pointers.

## 1 Introduction

### 1.1 Boosting

There has been tremendous growth in machine learning research and applications, both in practice (applying all sorts of methods on all sorts of data and seeing what works well) and in theory (computational learning theory). However, not very many ideas generated in theoretical machine learning have had a large impact on machine learning practice. One of the exceptions is *boosting*, which is a simple, general, and widely applicable method to improve the generalization error of a given learning method, i.e., to convert a *weak* learner into a *strong* learner.

The set-up here is binary classification: we are trying to predict binary labels  $y$  from points  $x \in \mathcal{X}$ . A typical case would be  $\mathcal{X} = \{0, 1\}^n$ . We are given  $m$  labeled examples  $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \{-1, 1\}$  where the  $x_i$ s are independent and identically distributed (i.i.d.) according to some unknown distribution  $\mathcal{D}$ , and the binary labels are determined by some unknown *target function*  $f : \mathcal{X} \rightarrow \{-1, 1\}$  that we are trying to learn, i.e.,  $y_i = f(x_i)$ . A weak learner  $\mathcal{W}$  is an algorithm that can be fed a number of examples according to a specified

---

<sup>1</sup> Work done while a student at the IILC of the University of Amsterdam.



## 64:2 Improved Quantum Boosting

distribution  $D$  (not to be confused with the unknown data-generating distribution  $\mathcal{D}$ ) over the  $m$  examples of the given sample, and that is then promised to generate a hypothesis  $h : \mathcal{X} \rightarrow \{-1, 1\}$  that is slightly better than random w.r.t. that  $D$ :

$$\Pr_{x \sim D}[h(x) \neq f(x)] \leq 1/2 - \gamma.$$

Here  $\gamma \in (0, 1/2)$  is a small but positive number that gives the quality of the weak learner. We denote the “cost” (in time complexity or whatever measure the user likes) of one run of  $\mathcal{W}$  by  $W$ , and use this number also as an upper bound on the number of examples (distributed according to  $D$ ) that the weak learner uses.

A hypothesis with a generalization error that is just slightly better than random is not very useful by itself. The goal of boosting is to convert the weak learner into a strong learner, which is one that produces hypotheses not only with small empirical error (i.e., w.r.t. the uniform distribution over the  $m$  examples), but even with small generalization error w.r.t. the unknown target function  $f : \mathcal{X} \rightarrow \{-1, 1\}$  and the unknown distribution  $\mathcal{D}$  that generated the examples:

$$\Pr_{x \sim \mathcal{D}}[h(x) \neq f(x)] \leq \varepsilon.$$

Here the desired upper bound  $\varepsilon$  on the final generalization error is a parameter of the strong learner. Unsurprisingly, achieving smaller  $\varepsilon$  requires a larger number of examples and larger runtime. For simplicity, in this introduction we focus on the case  $\varepsilon = 1/3$  (in the body of the paper we cover the general case). Similarly, the smaller the initial advantage  $\gamma$  is, the more work we will have to do find a hypothesis with small generalization error.<sup>2</sup>

The idea of boosting is to find a hypothesis with low empirical error by combining different runs of the weak learner on different distributions. Once we have a hypothesis with small empirical error on a sufficiently large set of examples, VC-theory implies that such a hypothesis will probably also have a small generalization error.

How can we find a hypothesis with small empirical error? Because empirical error is measured w.r.t. the uniform distribution over  $\{x_1, \dots, x_m\}$ , that will be our first distribution  $D^1$ . We run the weak learner on  $D^1$ , and receive a hypothesis  $h_1$  that is slightly better than random w.r.t. the uniform distribution. The next iteration then biases the distribution away from the examples that are already well-classified, by increasing the probability of misclassified examples, yielding a new distribution  $D^2$ . We then run the weak learner again, to generate a hypothesis  $h_2$  that is slightly better than random w.r.t. this new distribution, and hence hopefully better than  $h_1$  on the examples that were misclassified by  $h_1$ . Then we bias the distribution further towards the still-misclassified examples, and so on. The intuition here is that the distributions  $D^t$  “zoom in” on the hardest examples, the ones that are most difficult to classify correctly. After some  $T$  iterations, the  $T$  different weak hypotheses are combined into one hypothesis  $h$ , typically by defining the latter as the sign of a linear combination  $\sum_{t=1}^T \alpha_t h_t$  of the  $T$  weak hypotheses  $h_1, \dots, h_T$ . Surprisingly, already after a relatively small number of iterations, the resulting hypothesis will have small empirical error! Thus boosting converts the ability to generate weak hypotheses w.r.t. chosen distributions over the examples, into the ability to generate strong hypotheses, which have small error w.r.t. both the uniform distribution over the examples, and w.r.t. the unknown target function  $f$  and distribution  $\mathcal{D}$  that generated our  $m$  examples.

---

<sup>2</sup> For simplicity we will assume this  $\gamma$  is known to the strong learner we are trying to design, but this is not necessary: if it doesn't know  $\gamma$ , the strong learner can try exponentially decreasing guesses for  $\gamma$  until it finds a hypothesis with small empirical error.

A number of classical boosting algorithms exist that instantiate this meta-algorithm in different ways. The most famous of these is probably Freund and Schapire’s AdaBoost [13, 14, 19] (short for “adaptive boosting”), which biases the new distribution  $D^{t+1}$  based on the error  $\varepsilon_t$  that  $h_t$  made. It drives the empirical error all the way down to 0 (note that as soon as this error is  $< 1/m$  it must actually be 0). AdaBoost uses  $T = O(\log(m)/\gamma^2)$  iterations. Each iteration takes time  $\tilde{O}(m)$  to compute the error  $\varepsilon_t$  of  $h_t$  and to update the distribution over the  $m$  examples<sup>3</sup> and runs the weak learner  $\mathcal{W}$  once, at cost  $W$ . This gives overall complexity

$$\tilde{O}\left(\frac{W + m}{\gamma^2}\right).$$

How large should  $m$  be in order to make the inference from low empirical error to low generalization error? This depends on the hypothesis space  $\mathcal{H}_{weak}$  of the weak learner, in particular on its VC-dimension  $d$  (defined in Section 2.1). The hypothesis space  $\mathcal{H}_{strong}$  of the boosting algorithm consists of all signs of linear combinations of up to  $T$  elements of  $\mathcal{H}_{weak}$ . One can show that the VC-dimension of  $\mathcal{H}_{strong}$  is  $D = \tilde{O}(dT)$ . VC-theory implies that (for constant  $\varepsilon$ )  $m \approx D \approx dT \approx d/\gamma^2$  examples suffice to end up with generalization error  $\leq \varepsilon$  (with high probability over the choice of the sample). Accordingly, when re-expressed as a function of  $d$  rather than  $m$ , the complexity of AdaBoost is

$$\tilde{O}\left(\frac{W}{\gamma^2} + \frac{d}{\gamma^4}\right). \tag{1}$$

## 1.2 Quantum boosting

In the last few years there has been a surge in interest in possible ways in which *quantum* computers might help improve machine learning (see [9] for a survey of several algorithmic approaches and [6] for quantum learning theory).

Recently, Arunachalam and Maity [5] gave the first speed-up for boosting on a quantum computer. Here the given sample is still the same classical sequence of  $m$  labeled examples  $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \{-1, 1\}$ , but these are now stored in a quantum-accessible classical memory, which means a quantum learner can query multiple examples in superposition.

The key insight of [5] is that the error  $\varepsilon_t$  of the base classifier in the  $t$ -th iteration of AdaBoost can be approximated faster (in time  $o(m)$ ) using a quantum counting algorithm; this approximation is subtle because it involves both multiplicative and additive error, in different regimes for  $\varepsilon_t$ . Their method works not only for boosting classical weak learners, but also for boosting *quantum* weak learners. These are fed quantum examples w.r.t. the distribution  $D$ :

$$\sum_{i=1}^m \sqrt{D(x_i)} |x_i, y_i\rangle. \tag{2}$$

If the quantum weak learner  $\mathcal{W}$  expects to receive  $W$  such examples, the quantum booster will have to prepare  $W$  copies of this state to feed into  $\mathcal{W}$ .<sup>4</sup>

<sup>3</sup> The notation  $\tilde{O}(f)$  means  $O(f \cdot \text{polylog}(f))$ .

<sup>4</sup> The classical boosting literature [19] distinguishes “boosting by resampling” and “boosting by reweighting”. Like Arunachalam and Maity [5], we follow “boosting by resampling” and explicitly prepare the  $W$  quantum or classical examples (w.r.t.  $D^t$ ) that the weak learner needs, rather than just modifying  $D^t$ .

The quantum version of AdaBoost of [5] uses the same number of iterations as classical AdaBoost, but improves the complexity of each iteration (at least as a function of  $m$  or  $d$ ). Their main complexity upper bound is:

$$\tilde{O}\left(\frac{W^{1.5}\sqrt{d}}{\gamma^{11}}\right). \quad (3)$$

Comparing with the complexity of classical AdaBoost Eq. (1), this gives a speed-up over classical boosting in terms of the dependence on the VC-dimension  $d$  of the weak learner’s hypothesis class, but at the expense of a significant deterioration in terms of the dependence on the quality of the weak learner  $\gamma$  and a milder deterioration in terms of the weak learner’s cost  $W$ .

### 1.3 Our results

In this paper we give a simpler and faster quantum boosting algorithm. Instead of AdaBoost, our starting point will be Servedio’s SmoothBoost algorithm [22], which we explain in Section 3. Servedio’s motivation for smooth boosting was to deal with malicious noise (at a rate that depends on  $\gamma$ ) in the sample better than AdaBoost. However, SmoothBoost is also very suitable for “quantization” thanks to the following advantages that it has over AdaBoost:

- SmoothBoost doesn’t need to calculate or approximate the error  $\varepsilon_t$  of  $h_t$  on the  $m$  examples, which means we don’t need to apply approximate quantum counting for this.
- The distributions  $D^t$  that it generates are “smooth” (whence its name), in the sense that no example has probability much bigger than the uniform probability  $1/m$ . Generating quantum examples as in Eq. (2) is cheaper when none of the probabilities is big.
- The weights  $\alpha_t$  in the final linear combination  $\sum_t \alpha_t h_t$  are all equal to 1 in SmoothBoost. In contrast, AdaBoost uses  $\alpha_t = \frac{1}{2} \ln((1 - \varepsilon_t)/\varepsilon_t)$ , hence the quantum algorithm’s approximation errors in  $\varepsilon_t$  lead to approximation errors in  $\alpha_t$  that need to be kept under control.

In addition to exploiting these “classical” advantages in order to obtain a simpler and faster quantum booster, we also give an improved procedure to generate quantum examples over the  $m$  examples from  $S$ . This procedure assumes access to a non-normalized version of  $D^t$  and doesn’t have to worry about the normalizing factor. As explained in Section 4.1, in a way quantum mechanics will take care of the proper normalization for us.<sup>5</sup>

We obtain the following upper bound on the complexity of our Quantum SmoothBoost:<sup>6</sup>

$$\tilde{O}\left(\frac{W}{\gamma^4} + \frac{\sqrt{d}}{\gamma^5}\right). \quad (4)$$

<sup>5</sup> We also generate the quantum examples exactly, while [5] only generate them approximately and hence has to deal with the way the errors in this process affect the other parts of their boosting procedure. Our example-generating procedure could also be used to improve the bounds of quantum AdaBoost [5], though the result won’t be as efficient as our quantum SmoothBoost. Note that if we want to use Quantum SmoothBoost with a classical weak learner  $\mathcal{W}$ , we can just measure the  $W$  quantum examples in the computational basis to obtain the  $W$  classical examples distributed according to  $D^t$  that such a  $W$  needs as input. This still gives a speed-up in terms of the desired generalization error  $\varepsilon$  compared to classically generating those  $W$  examples.

<sup>6</sup> This bound is when we aim at constant generalization error  $\varepsilon = 1/3$ . We also make explicit the complexity for much smaller  $\varepsilon$  (see Theorem 14).

This improves over the complexity of the booster of [5] (as given in Eq. (3)) in terms of the parameters  $W$  and (especially)  $\gamma$ . The  $\gamma$ -dependence is still worse than classical AdaBoost (as given in Eq. (1)), but not by large powers anymore. It is an interesting open question whether this  $\gamma$ -dependence can be improved further.

## 1.4 Related work

Our main sources of inspiration for this paper were the quantum AdaBoost of Arunachalam and Maity [5] and classical SmoothBoost of Servedio [22], and we have tried in this paper to combine the best elements of both.

Here we mention a number of related quantum papers. Wang et al. [24] (which preceded [5]) give a quantum speed-up for a specific subtask of AdaBoost, namely to compute the coefficients  $\alpha_t = \frac{1}{2} \ln((1 - \varepsilon_t)/\varepsilon_t)$  that combine given base classifiers  $h_1, \dots, h_T$  into a good hypothesis  $h = \text{sign}(\sum_t \alpha_t h_t)$ . These weights  $\alpha_t$  are approximated more efficiently than is possible classically using a version of approximate quantum counting. This, however, assumes the base classifiers have already been generated and sidesteps the most important aspect of AdaBoost, which is to generate the  $h_t$ 's adaptively by running the weak learner on a distribution  $D^t$  that depends on  $h_1, \dots, h_{t-1}$ . The even earlier paper by Schuld and Petruccione [21] considers quantum ensembles of classifiers (rather than the linear combinations used in boosting) and runs AdaBoost as a subroutine, but does not give a quantum boosting algorithm.

AdaBoost may be viewed as an instance of the multiplicative weights update method, see for instance the presentation in [4, Section 3.6]. There have been several quantum speed-ups for multiplicative weights methods in other contexts, particularly the quantum SDP-solvers of Brandão et al. [11, 3, 10, 2], and the very recent quantum version of the Hedge algorithm of Reberthost et al. [18]. However, none of those speed-ups for versions of multiplicative weights seems directly applicable to our boosting setting.

## 2 Preliminaries

### 2.1 PAC learning

In this section we give a brief introduction to the PAC learning framework, which provides theoretical guarantees on learnability. The textbook by Shalev-Shwartz and Ben-David [23] provides an excellent and detailed introduction to the topic of classical PAC learning.

To formally introduce the PAC learning framework, let  $\mathcal{D}$  denote a probability distribution over the set of *points*  $\mathcal{X}$ . We want to learn an unknown *target function*  $f : \mathcal{X} \rightarrow \mathcal{Y}$ . We will assume here that the set of labels  $\mathcal{Y}$  is just  $\{-1, 1\}$ , so we are dealing with binary classification. A typical situation to keep in mind is the important special case of learning Boolean functions, where  $\mathcal{X} = \{0, 1\}^n$ , or  $\mathcal{X} = \cup_{n \geq 1} \{0, 1\}^n$ .

Learning begins by choosing a learning algorithm (a “learner”) with an associated *hypothesis class*  $\mathcal{H}$  of functions  $h : \mathcal{X} \rightarrow \{-1, 1\}$ . This hypothesis class could be any set of functions, but good examples to keep in mind are cases where  $\mathcal{X} = \{0, 1\}^n$  and  $\mathcal{H}$  consists of objects with bounded computational power, for instance all Boolean circuits of at most a certain size, all neural networks with a specific depth and number of nodes, or all decision trees of at most a certain depth. We will assume that each  $h \in \mathcal{H}$  has a succinct description and that we can efficiently evaluate a given  $h$  on a given  $x \in \mathcal{X}$ . For simplicity we assume such an evaluation has unit cost.

## 64:6 Improved Quantum Boosting

The learner is given access to a *sample*  $S = ((x_1, y_1), \dots, (x_m, y_m))$ , which is the training data. The points  $x_i$  are i.i.d. generated according to an unknown distribution  $\mathcal{D}$  on  $\mathcal{X}$ , and the labels  $y_i = f(x_i)$  are determined by the target function  $f$  that we are trying to learn. The learner's goal is to find an  $h \in \mathcal{H}$  that fits well with the given training data, in the hope that this  $h$  will generalize well to points that were not part of the data, in the sense of mostly giving the same labels as the target function. The PAC learning framework is a distribution-free setting, so we would like to design a learner that works well for every  $\mathcal{D}$ , in the sense of outputting a hypothesis with low generalization error.

► **Definition 1.** *The generalization error of  $h : \mathcal{X} \rightarrow \mathcal{Y}$  w.r.t. target function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  under distribution  $\mathcal{D}$  is*

$$\text{err}(h, f, \mathcal{D}) = \Pr_{x \sim \mathcal{D}} [h(x) \neq f(x)].$$

Generalization error is often referred to as the *true* error; it is the quantity the learner is really trying to minimize over the class  $\mathcal{H}$  of available hypotheses.

As the distribution  $\mathcal{D}$  is anyway unknown, the generalization error of a hypothesis  $h$  cannot be calculated and the learner uses the *empirical* error of a hypothesis  $h$  (the fraction of the sample that  $h$  mislabels) to measure its performance, as a proxy for the generalization error.

► **Definition 2.** *The empirical error of  $h : \mathcal{X} \rightarrow \mathcal{Y}$  w.r.t. sample  $S = ((x_1, y_1), \dots, (x_m, y_m))$  is*

$$\hat{\text{err}}(h, S) = \Pr_{i \in_R [m]} [h(x_i) \neq y_i],$$

where  $i \in_R [m]$  means that  $i$  is taken uniformly at random from  $[m] = \{1, \dots, m\}$ .

► **Definition 3.** *An  $(\epsilon, \delta)$ -PAC learner for a concept class  $\mathcal{C}$  with hypothesis class  $\mathcal{H}$  and sample complexity  $m$ , is an algorithm  $\mathcal{A}$  such that the following holds for all target functions  $f \in \mathcal{C}$  and all distributions  $\mathcal{D}$  on  $\mathcal{X}$ :*

- $\mathcal{A}$  takes as input  $m$  examples  $(x_1, f(x_1)), \dots, (x_m, f(x_m))$  where the  $x_i$  are i.i.d. according to  $\mathcal{D}$ .
- $\mathcal{A}$  outputs an  $h \in \mathcal{H}$  which is “Probably Approximately Correct” in the sense that

$$\Pr[\text{err}(h, f, \mathcal{D}) \leq \epsilon] \geq 1 - \delta,$$

where the probability is taken over the sample and over the learner's internal randomness.

The end goal is to find a learner with small sample complexity  $m$ , small error probability  $\delta$ , and (most important of all) small generalization error  $\epsilon$ . Often we will start, however, with a “weak” learner, one whose generalization error is only slightly better than random. Since we restricted to binary labels ( $\mathcal{Y} = \{-1, 1\}$ ), generalization error  $\epsilon = 1/2$  is no better than random guessing. A weak learner is a learner that does slightly better than that:

► **Definition 4 (Weak learning).** *A  $\gamma$ -weak learner  $\mathcal{W}$  for concept class  $\mathcal{C}$  with hypothesis class  $\mathcal{H}_{\text{weak}}$  is a  $(1/2 - \gamma, 0)$ -PAC learner. Hypotheses returned by a weak learner are called base classifiers.*

Following Servedio [22], we assume the weak learner  $\mathcal{W}$  has error probability  $\delta = 0$ , so it always outputs a hypothesis with generalization error  $\leq 1/2 - \gamma$ . If instead we start with a  $\mathcal{W}$  that has non-zero error probability, say  $1/3$ , then we can reduce this error probability

to small  $\delta > 0$  as follows. Run  $\mathcal{W}$  a total number of  $r = \lceil \log_3(1/\delta) \rceil$  times, each time with fresh independent examples. Then, except with probability  $\leq (1/3)^r \leq \delta$ , at least one of the returned hypotheses  $h_1, \dots, h_r \in \mathcal{H}_{weak}$  will have generalization error  $\leq 1/2 - \gamma$ . However, finding (with success probability  $\geq 1 - \delta$ ) among these  $r$  hypotheses one with such low error has a cost. Deciding (with success probability  $\geq 2/3$ ) for a given hypothesis  $h$  whether it has error  $\leq 1/2 - \gamma$  under a given distribution can be done by sampling  $O(1/\gamma^2)$  examples according to that distribution and estimating the fraction of examples where  $h$  predicts the label correctly. Searching over the  $r$  hypotheses to find a good one adds a factor of  $O(r)$  to the classical cost, and reducing the overall error probability from  $1/3$  to  $\delta$  adds another factor of  $O(\log(1/\delta))$ .<sup>7</sup>

Suppose we ideally want to run an errorless weak learner  $T$  times, namely once in each of  $T$  iterations. But instead we start with a weak learner with error probability  $1/3$ . Reducing  $1/3$  to  $\delta \ll 1/T$  allows us to take a union bound over all  $T$  iterations, and conclude that with high probability each of the  $T$  iterations produces a base classifier with generalization error  $\leq 1/2 - \gamma$  (w.r.t. the distribution  $D^t$  of that iteration). Because here we assumed our weak learner has no error probability from the start, we do not have to factor in the additional cost for this error reduction, but it anyway doesn't significantly affect the complexities of classical or quantum boosting (Eqs. (1) and (4) respectively).

## 2.2 How many examples suffice to ensure small generalization error?

The number of examples that are necessary and sufficient for learning is governed by the VC-dimension of the relevant hypothesis class and by the desired generalization error, as follows. A set  $S \subseteq \mathcal{X}$  of  $d$  points is said to be *shattered* by  $\mathcal{H}$  if for each of the  $2^d$  labelings  $\ell : S \rightarrow \{0, 1\}$ , there exists an  $h \in \mathcal{H}$  that agrees with  $\ell$  on the points in  $S$ . The *VC-dimension* of a hypothesis class  $\mathcal{H}$  is the size of a largest  $S$  that is shattered by  $\mathcal{H}$ . Intuitively, if the VC-dimension of  $\mathcal{H}$  is small, then it should be relatively simple to find a good hypothesis in it, i.e., one that minimizes empirical error.

The following theorem implies that for sufficiently large  $m$ , every  $h \in \mathcal{H}$  has a generalization error that is only slightly worse than its empirical error. Such a result means it suffices to look for a hypothesis with small empirical error.

► **Theorem 5** (Theorem 2.5 in [19]). *Let  $\mathcal{H}$  be a hypothesis class of finite VC-dimension  $d$ . Assume that a sample  $S$  of size  $m$  is chosen for some target function  $f$ , i.i.d. according to some distribution  $\mathcal{D}$ . Then for every  $\eta > 0$  it holds that*

$$\Pr[\exists h \in \mathcal{H} : \text{err}(h, f, \mathcal{D}) > \hat{\text{err}}(h, S) + \eta] \leq 8 \left( \frac{em}{d} \right)^d \exp \left\{ \frac{-m\eta^2}{32} \right\}.$$

If we set  $\eta = \varepsilon/2$  and

$$m = O \left( \frac{d \log(d/(\delta\varepsilon)) + \log(1/\delta)}{\varepsilon^2} \right),$$

with a sufficiently large constant in the  $O(\cdot)$ , then (except with probability  $\delta$ ), each  $h \in \mathcal{H}$  has a generalization error that is at most  $\varepsilon/2$  bigger than its empirical error. Accordingly, if a learner now outputs any hypothesis  $h \in \mathcal{H}$  whose empirical error is  $\leq \varepsilon/2$ , then its generalization error will be  $\leq \varepsilon$ , as desired.

<sup>7</sup> In the quantum case the  $O(1/\gamma^2)$  can be replaced by  $O(1/\gamma)$  using quantum approximate counting (Theorem 7 below), and the  $O(r)$  can be replaced by  $O(\sqrt{r})$  using Grover's algorithm [15].

### 2.3 Quantum PAC learning and helpful quantum subroutines

In order to introduce the quantum boosting algorithm in Section 4, we explain the query model that the quantum algorithm works with. We say that an algorithm has *query access* to a string  $z \in Z_a^N$  over alphabet  $Z_a = \{0, \dots, a-1\}$  if it can apply a unitary  $O_z$  such that

$$O_z : |i, b\rangle \mapsto |i, b \oplus z_i\rangle,$$

where  $i \in \{0, 1\}^{\lceil \log N \rceil}$ ,  $b \in Z_a$ , and  $\oplus$  denotes addition modulo  $a$ . Naturally, a quantum algorithm can apply  $O_z$  on a superposition of distinct inputs  $i$ .

In the classical setting we assumed a learner is given a sample  $S = ((x_1, y_1), \dots, (x_m, y_m))$  of  $m$  labeled examples. Here points  $x_i \in \mathcal{X}$  are independently drawn from an unknown distribution  $\mathcal{D}$ , and labeled  $y_i = f(x_i)$  according to an unknown target function  $f$ . Such a classical sample will still be the starting point of our quantum boosting algorithm; we assume the learner has query access to the sample (viewed as a string  $z \in (\mathcal{X} \times \{-1, 1\})^m$ ). One may think of the sample as being stored in a quantum-accessible classical memory, sometimes called QRAM, which may be queried throughout the algorithm (incl. by the weak learner). However, our setting also encompassed the case of *synthetic* data, where we would have an efficient procedure which, on input  $i$ , computes the example  $(x_i, y_i)$ .

Even though the initially given sample is classical, like Arunachalam and Maity [5] we will set up our quantum booster so that it can work to improve a classical weak learner but also to improve a *quantum* weak learner. The latter is given *quantum examples* w.r.t. distribution  $D$ :

$$\sum_{x \in \mathcal{X}} \sqrt{D(x)} |x, f(x)\rangle.$$

One can think of a quantum example as the coherent version of a random example  $(x, f(x))$  where  $x \sim D$ . A quantum learner is given access to several copies of the quantum example and performs a POVM measurement, where each outcome is associated with a hypothesis  $h$  in its hypothesis class. It won't matter for the purposes of this paper, but [7] proved that in the general PAC and agnostic learning settings, the required number of classical and quantum examples are the same up to constant factor.

In the case of boosting, the weak learner will be fed quantum examples w.r.t. a distribution  $D$  that only has support on the  $m$  given examples. Since our initially given sample is classical, our boosting algorithm will itself have to prepare the quantum examples that it wants to feed into the weak learner in each iteration, and we have to (and will) account for the cost of this.

We also assume that we can evaluate a given  $h$  (in the weak learner's hypothesis class  $\mathcal{H}_{weak}$ ) in superposition, meaning we can apply a unitary that maps  $|h\rangle|x\rangle|b\rangle \mapsto |h\rangle|x\rangle|h(x) \cdot b\rangle$ ; here the basis states of the first space are the names of the  $h \in \mathcal{H}_{weak}$ , the basis states of the second space are the elements of  $\mathcal{X}$ , and the basis states of the third space are the labels in  $\mathcal{Y} = \{-1, 1\}$ .

The definitions of PAC learning and weak learning straightforwardly generalize to the quantum setting. We refer to the survey [6] for more on this model. The following basic quantum subroutines can be derived from Brassard et al. [12] (or from [1] if one wants to avoid use of the quantum Fourier transform):



► **Theorem 6** (Amplitude amplification). *Suppose we have an  $m$ -qubit unitary  $U$  such that*

$$U|0^m\rangle = \sqrt{a}|\phi_0\rangle|0\rangle + \sqrt{1-a}|\phi_1\rangle|1\rangle,$$

*and we know a lower bound  $a'$  on  $a$ . Then there exists a quantum algorithm  $V$  using  $O(1/\sqrt{a'})$  applications of  $U$  and  $U^\dagger$ , and  $\tilde{O}(1/\sqrt{a'})$  other gates, such that*

$$V|0^m\rangle = \sqrt{b}|\phi_0\rangle|0\rangle + \sqrt{1-b}|\phi_1\rangle|1\rangle,$$

*where  $b \in [1/2, 1]$ .*

► **Theorem 7** (Approximate counting). *Suppose we have query access to a string  $z \in [0, 1]^N$ , with sum  $s = \sum_{i=1}^N z_i \geq 1$ . There exists a quantum algorithm that uses  $O(\frac{1}{\varepsilon}\sqrt{N} \log(1/\delta))$  queries and  $\tilde{O}(\frac{1}{\varepsilon}\sqrt{N} \log(1/\delta))$  other operations, and that outputs (except with probability  $\leq \delta$ ) an  $\tilde{s}$  such that  $(1 - \varepsilon)s \leq \tilde{s} \leq (1 + \varepsilon)s$ .*

### 3 SmoothBoost

We first consider the classical SmoothBoost algorithm of Servedio [22]. It generates only smooth distributions, in the sense that none of the examples get too much weight. In the next section we will introduce a quantum version of SmoothBoost.

We give the pseudocode of SmoothBoost in Algorithm 1. There are a few cosmetic changes compared to the pseudocode of [22] that will make it easier for us to quantize it later. The algorithm takes four inputs. First, a  $\gamma$ -weak learner  $\mathcal{W}$  with associated hypothesis class  $\mathcal{H}_{weak}$  and cost and sample complexity  $W$ . Second, a sample  $S = ((x_1, y_1), (x_2, y_2), \dots, (x_m, y_m)) \in (\mathcal{X} \times \{-1, 1\})^m$ , for some sample size  $m$  that we will choose later. Lastly, a parameter  $\kappa \in (0, 1)$  which controls the empirical error of SmoothBoost and a parameter  $\theta \in [0, \frac{1}{2}]$  which controls the desired margin of the output hypothesis  $h$ . The goal of SmoothBoost is to output a hypothesis  $h : \mathcal{X} \rightarrow \{-1, 1\}$  with small empirical error (and as we shall see later, for sufficiently large  $m$ , this  $h$  will also have large generalization error). The final  $h$  is going to be the sign of a sum of elements of  $\mathcal{H}_{weak}$ , so the strong learner's hypothesis class is larger than that of the weak learner.

The central objects in this algorithm are the vectors  $M^1, M^2, \dots, M^T \in [0, 1]^m$ , which are unnormalized distributions over the  $m$  examples. The distribution  $D^t$  is the normalized version of  $M^t$ . SmoothBoost starts by initializing weights to  $N_i^0 = 0$  and  $M_i^1 = 1$ , for all  $i \in [m]$ , so  $D^1$  is uniform. In each iteration, Step 6 checks whether the sum of  $M_i^t$  is below  $\kappa m$ , and if so it terminates. Otherwise it runs the weak learner on  $W$  i.i.d. examples sampled (from  $S$ ) according to  $D^t$ , producing a base classifier  $h_t : \mathcal{X} \rightarrow \{-1, 1\}$ . Step 9 updates  $N^{t-1}$  to  $N^t$  and  $M^t$  to  $M^{t+1}$ . For each  $i \in [m]$ ,  $N_i^t$  is the cumulative amount by which hypotheses  $h_1, \dots, h_t$  beat the desired margin  $\theta$ . If  $x_i$  got correctly classified by  $h_t$ , then it will get higher weight  $N_i^t$ , which results in smaller weight  $M_i^{t+1}$  and smaller probability  $D_i^{t+1}$  in the next round of boosting. This mechanism forces the next run of the weak learner  $\mathcal{W}$  to “zoom in” (i.e., assign higher probabilities) to systematically misclassified instances. The procedure terminates if the sum of all weights  $\sum_{i \in [m]} M_i^t$  gets sufficiently small, as controlled by the parameter  $\kappa$ .

We now state three results from Servedio [22] that show, respectively, that the intermediate distributions are smooth, that SmoothBoost terminates after a small number of iterations, and that it returns a hypothesis with low empirical error.

► **Claim 8** (Lemma 1 of [22]). For each  $1 \leq t \leq T$ , it holds that  $\max_{i \in [m]} |D_i^t| \leq \frac{1}{\kappa m}$ .

Proof. This follows from the condition of Step 6: before termination we have  $\sum_{i=1}^m M_i^t \geq \kappa m$  and  $M_i^t \in [0, 1]$ , hence  $D_i^t = M_i^t / \sum_{i \in [m]} M_i^t \leq 1/\kappa m$  for all  $i \in [m]$ . ◁

## 64:10 Improved Quantum Boosting

### ■ Algorithm 1 SmoothBoost.

**Input:** A  $\gamma$ -weak learner  $\mathcal{W}$  with complexity  $W$ .

A sample  $S = ((x_1, y_1), \dots, (x_m, y_m)) \in (\mathcal{X} \times \{-1, 1\})^m$ .

Parameters  $\kappa \in (0, 1), \theta \in [0, \frac{1}{2}]$ .

**Output:** Hypothesis  $h : \mathcal{X} \rightarrow \{-1, 1\}$ .

- 1: **function** SMOOTHBOOST( $\mathcal{W}, S, \kappa, \theta$ )
- 2: For all  $i \in [m]$  initialize:  $N_i^0 \leftarrow 0, M_i^1 \leftarrow 1$ .
- 3:  $t \leftarrow 1$
- 4: **while** true **do**
- 5:     Compute  $s = \sum_{i=1}^m M_i^t$ .
- 6:     If  $s < \kappa m$  then  $T \leftarrow t - 1$ , **return**  $h \leftarrow \text{sign}(\sum_{t=1}^T h_t)$ , and **terminate**.
- 7:     Prepare  $W$  i.i.d. examples w.r.t. distribution  $D^t = M^t / \sum_i M_i^t$  (see Footnote 8).
- 8:     Feed those  $W$  examples into the weak learner  $\mathcal{W}$  to obtain base classifier  $h_t$ .
- 9:     For all  $i \in [m]$  set

$$N_i^t \leftarrow N_i^{t-1} + h_t(x_i)y_i - \theta$$

and

$$M_i^{t+1} \leftarrow \begin{cases} 1 & \text{for } N_i^t < 0 \\ (1 - \gamma)^{\frac{N_i^t}{2}} & \text{for } N_i^t \geq 0 \end{cases}$$

- 10:  $t \leftarrow t + 1$ .

▷ **Claim 9** (Theorem 3 of [22]). If  $\theta = \frac{\gamma}{2+\gamma}$  and for all  $t$  it holds that  $\Pr_{i \sim D^t}[h_t(x_i) \neq y_i] \leq \frac{1}{2} - \gamma$ , then SmoothBoost terminates with  $T < \frac{2}{\kappa\gamma^2\sqrt{1-\gamma}}$  iterations.

Proof. [22, Lemmas 4 and 5] imply  $\frac{2m}{\gamma\sqrt{1-\gamma}} > \gamma \sum_{t=1}^T \sum_{i=1}^m M_i^t$  (this is the hard part of Servedio's correctness proof of SmoothBoost). We have  $\sum_{i=1}^m M_i^t \geq \kappa m$  for all  $t$  until termination. Hence  $\frac{2m}{\gamma\sqrt{1-\gamma}} > \gamma T \kappa m$ , which implies the claim. ◁

▷ **Claim 10** (Theorem 2 of [22]). After  $t$  iterations of SmoothBoost, the hypothesis  $h = \text{sign}(\sum_{t=1}^t h_t)$  has empirical error  $e\hat{r}(h) \leq \sum_{i=1}^m M_i^{t+1}/m$ .

Proof. Note that  $N_i^T = \sum_{t=1}^T (h_t(x_i)y_i - \theta)$ . Hence if  $i$  is such that  $\sum_{t=1}^T h_t(x_i)y_i < \theta T$ , then  $N_i^T < 0$  and  $M_i^{T+1} = 1$ . The final hypothesis  $h$  errs on the  $i$ th example iff  $\sum_{t=1}^T h_t(x_i)y_i < 0$ . We upper bound the number of  $i \in [m]$  for which this happens:

$$\begin{aligned} \left| \left\{ i \mid \sum_{t=1}^T h_t(x_i)y_i < 0 \right\} \right| &\leq \left| \left\{ i \mid \sum_{t=1}^T h_t(x_i)y_i < \theta T \right\} \right| = \sum_{i: \sum_{t=1}^T h_t(x_i)y_i < \theta T} M_i^{T+1} \\ &\leq \sum_{i=1}^m M_i^{T+1}. \end{aligned} \quad \triangleleft$$

Since SmoothBoost terminates if  $\sum_{i=1}^m M_i^{T+1} < \kappa m$ , Claim 10 implies that the empirical error of the final hypothesis is  $< \kappa$ .

Combining the previous two claims, we see that the empirical error decreases like  $O(1/(T\gamma^2))$ . This contrasts with AdaBoost, where the empirical error goes down exponentially fast in  $T$  and hence can be driven down to  $< 1/m$  (and hence to 0) quite cheaply.

SmoothBoost does not drive the empirical error down all the way to 0, because that would require setting  $\kappa < 1/m$  which implies a very large number of iterations,  $T = O(m/\gamma^2)$ . However, small but non-zero empirical error is good enough for our purposes, because (with sufficiently large sample size  $m$ ) that already implies small generalization error.

We will choose  $\kappa$  to be  $\varepsilon/2$ , which sets the above upper bound on the empirical error of the final hypothesis  $h$  to half of the allowed generalization error. By the discussion following Theorem 5, if the sample size  $m$  is large enough, the final hypothesis will have generalization error  $\leq \varepsilon$ , as desired. The required  $m$  depends on the VC-dimension of the hypothesis class  $\mathcal{H}_{strong}$  of SmoothBoost, which consists of signs of sums of  $T$  elements of the hypothesis class  $\mathcal{H}_{weak}$  of the weak learner. The VC-dimensions of these two classes are related as follows:

▷ **Claim 11** (Shalev-Shwartz & Ben-David, p. 109 [23]). Let  $\mathcal{H}_{weak}$  be a hypothesis class of VC-dimension  $d$  and  $\mathcal{H}_{strong} = \{\text{sign}(\sum_{i=1}^T h_i) \mid h_1, \dots, h_T \in \mathcal{H}_{weak}\}$ . Then the VC-dimension of  $\mathcal{H}_{strong}$  is  $D = O(Td \log(Td))$ .

By Theorem 5 and the fact that  $T = O(\frac{1}{\varepsilon\gamma^2})$  it thus suffices to take

$$m = O\left(\frac{D \log(D/(\delta\varepsilon)) + \log(1/\delta)}{\varepsilon^2}\right) = O\left(\frac{d \log(d/(\delta\varepsilon\gamma))^2}{\varepsilon^3\gamma^2} + \frac{\log(1/\delta)}{\varepsilon^2}\right) \quad (5)$$

examples in order to be able to infer (with success probability  $\geq 1 - \delta$ ) generalization error  $\leq \varepsilon$  from empirical error  $\leq \varepsilon/2$ .

Finally, let us determine the complexity of SmoothBoost, in terms of the overall number of elementary operations and queries to the sample and to the  $h_t$ . There are  $T = O(\frac{1}{\varepsilon\gamma^2})$  iterations. Each iteration involves one application of the weak learner  $\mathcal{W}$ , and  $\tilde{O}(m)$  other operations. The weak learner needs to be fed  $W$  examples sampled according to distribution  $D^t$ . Using rejection sampling, we can generate  $W$  such examples at cost  $O(W/\kappa) = O(W/\varepsilon)$ .<sup>8</sup>

▶ **Theorem 12.** *Let  $\mathcal{W}$  be a  $\gamma$ -weak learner of complexity  $W$  for concept class  $\mathcal{C}$ , with hypothesis class  $\mathcal{H}_{weak}$  of VC-dimension  $d$ . Then given  $m$  examples according to Eq. (5), SmoothBoost is an  $(\varepsilon, \delta)$ -PAC learner for  $\mathcal{C}$ , with hypothesis class  $\mathcal{H}_{strong}$ . It runs the weak learner  $O(\frac{1}{\varepsilon\gamma^2})$  times and uses*

$$\tilde{O}(T(W/\varepsilon + m)) = \tilde{O}\left(\frac{W}{\varepsilon^2\gamma^2} + \frac{m}{\varepsilon\gamma^2}\right) = \tilde{O}\left(\frac{W}{\varepsilon^2\gamma^2} + \frac{d}{\varepsilon^4\gamma^4}\right)$$

*other operations (elementary computational steps, queries to the sample, and evaluations of base classifiers).*

## 4 Quantum Smooth Boosting

In this section we introduce our quantum version of SmoothBoost. The algorithm is given query access to a quantum (or classical) weak learner  $\mathcal{W}$  with sample complexity  $W$ , and to a sample  $S$  of size  $m$ . The quantum weak learner needs to be fed *quantum* examples according to the distribution  $D^t$  obtained by normalizing the weight-vector  $M^t$ . We will start with that.

<sup>8</sup> Specifically, Step 7 of SmoothBoost can be implemented as follows. Sample  $i \in [m]$  uniformly. With probability  $M_i^t$  output  $(x_i, y_i)$ , and otherwise repeat. Since the probability to output  $(x_i, y_i)$  is proportional to  $M_i^t$ , the example (if we indeed output an example) is sampled according to the desired probability distribution  $D^t$ . Note that the probability that we output an example in one try is  $\frac{1}{m} \sum_i M_i^t \geq \kappa$ , because of the condition of Step 6. Hence the expected number of repetitions before we output an example is  $\leq 1/\kappa$ .

### 4.1 Preparing quantum examples

Here we show how we can efficiently prepare quantum examples w.r.t. the distribution  $D^t$  induced by the non-normalized  $M^t$ , thanks to its smoothness. This may be viewed as a quantum analogue of the classical rejection sampling sketched in Footnote 8.

► **Theorem 13.** *Suppose we have query access to the  $m$  numbers  $M_1, \dots, M_m \in [0, 1]$ . Let  $s = \sum_{i=1}^m M_i$  be their (unknown) sum, which has a known lower bound of  $\kappa m$ . Define a probability distribution  $D$  on  $[m]$  by  $D_i = M_i/s$ . Then using an expected number of  $O(1/\sqrt{\kappa})$  queries and  $\tilde{O}(1/\sqrt{\kappa})$  other gates, we can prepare the state*

$$\sum_{i=1}^m \sqrt{D_i} |i\rangle.$$

**Proof.** Start by preparing the uniform state

$$\frac{1}{\sqrt{m}} \sum_{i=1}^m |i\rangle |0\rangle.$$

Using two queries (the second to uncompute the value  $M_i$ ), and a few other gates to implement a conditional rotation by angle  $\arcsin(\sqrt{M_i})$ , prepare

$$\frac{1}{\sqrt{m}} \sum_{i=1}^m |i\rangle (\sqrt{M_i} |0\rangle + \sqrt{1 - M_i} |1\rangle).$$

The squared norm of the part of the state ending in  $|0\rangle$  is  $s/m \geq \kappa$ . Now use  $O(1/\sqrt{\kappa})$  rounds of amplitude amplification (Theorem 6) to increase that squared norm to  $\geq 1/2$ . This costs  $O(1/\sqrt{\kappa})$  queries and  $\tilde{O}(1/\sqrt{\kappa})$  other gates.

If we measure the last qubit of the resulting state, then we obtain outcome 0 with probability  $\geq 1/2$  and the state collapses to the state that we want to prepare (with an extra  $|0\rangle$ -qubit that we can remove). Note that we know when we succeed to produce the desired state. Since the probability of success is  $\geq 1/2$ , the expected number of repetitions before success is  $\leq 2$ . ◀

Once we have produced a copy of the state

$$\sum_{i=1}^m \sqrt{D_i} |i\rangle,$$

we can easily convert this into a quantum example

$$\sum_{i=1}^m \sqrt{D_i} |x_i, y_i\rangle$$

by querying the sample  $S$ .

### 4.2 Quantizing SmoothBoost

The pseudocode of Quantum Smooth Boosting is given in Algorithm 2. The algorithm receives as input a weak quantum learner  $\mathcal{W}$  with sample complexity  $W$ , and query access to a sample  $S$  of  $m$  examples. Additionally, it receives two parameters  $\kappa, \theta$ .

The algorithm looks a bit different from classical SmoothBoost because it doesn't update the  $m$ -dimensional vectors  $N_i^t$  and  $M_i^t$  explicitly anymore; the  $O(m)$  that this costs is more than we are willing to spend in the quantum case. Instead, we will store the earlier base

classifiers  $h_1, \dots, h_t$ . Queries to these classifiers together with queries to the sample  $S$  allow us to calculate each entry  $N_i^t$  and  $M_i^t$  on demand in time  $\tilde{O}(t)$ , via the formulas of Step 9 of SmoothBoost.

The algorithm begins by initializing  $N^0, M^1$  and by setting  $t = 1$ . Like in the classical case, we iterate until the sum of the weights  $\sum_{i \in [m]} M_i^t$  becomes small enough. In contrast to the classical case, we do not have the time to sum these  $m$  numbers exactly, so we will instead estimate the sum with small approximation error using quantum counting (Theorem 7).

■ **Algorithm 2** Quantum SmoothBoost.

**Input:** A  $\gamma$ -weak quantum learner  $\mathcal{W}$  with complexity  $W$ .

A sample  $S = ((x_1, y_1), \dots, (x_m, y_m)) \in (\mathcal{X} \times \{-1, 1\})^m$ .

Parameters  $\kappa \in (0, 1), \theta \in [0, \frac{1}{2})$ .

**Output:** Hypothesis  $h : \mathcal{X} \rightarrow \{-1, 1\}$ .

---

```

1: function QUANTUMSMOOTHBOOST( $\mathcal{W}, S, \kappa, \theta$ )
2:    $t \leftarrow 1$ 
3:   while true do
4:     Compute an estimate  $\tilde{s}$  of  $s = \sum_{i=1}^m M_i^t$  with multiplicative error 1.1 (using
       Theorem 7), where the  $M_i^t$  are as defined in Step 9 of SmoothBoost (and only
       computed on demand).
5:     If  $\tilde{s} < \kappa m$  then  $T \leftarrow t - 1$ , return  $h \leftarrow \text{sign}(\sum_{t=1}^T h_t)$ , and terminate.
6:     Prepare  $W$  copies of example  $|D^t\rangle$  w.r.t. distribution  $D_i^t = \frac{M_i^t}{\sum_{i \in [m]} M_i^t}$  (using
       Theorem 13).
7:     Feed those  $W$  examples into the weak learner  $\mathcal{W}$  to obtain base classifier  $h_t$ .
8:      $t \leftarrow t + 1$ .

```

---

Quantum Smoothboost runs  $O(TW)$  quantum subroutines that each have some error probability. By setting this error probability to be  $\ll 1/TW$ , the union bound implies that the probability that at least one of them will fail, is very small. The extra cost-factor  $\log(TW)$  that this error-reduction incurs will be absorbed by our  $\tilde{O}(\cdot)$  notation.

If we condition on the very-high-probability event that the various quantum subroutines involved all succeed, then the weights  $N^t$  and  $M^t$  are just equal to the weights as they would be in classical SmoothBoost with the same number of iterations. Because our approximation  $\tilde{s}$  of  $s$  for the stopping criterion has small multiplicative error, the smoothness of the intermediate distributions  $D^t$  before termination can be marginally worse than in classical SmoothBoost (Claim 8): For each  $1 \leq t \leq T$ , it holds that  $\max_{i \in [m]} |D_i^t| \leq \frac{1.1}{\kappa m}$ .

Quantum SmoothBoost terminates if  $\tilde{s} < \kappa m$ . Because  $\tilde{s}$  might underestimate the true  $s$  by at most a factor 1.1, upon termination we have  $s < 1.1\kappa m$  and hence Claim 10 implies empirical error  $e\hat{r}(h) < 1.1\kappa$ . We set  $\kappa = \varepsilon/2.2$  in order to ensure  $e\hat{r}(h) \leq \varepsilon/2$ . Like before, we choose the sample size  $m$  given by Eq. (5) to ensure generalization error  $\leq \varepsilon$ .

The total number of iterations is still  $O(\frac{1}{\varepsilon\gamma^2})$ .<sup>9</sup> It remains to determine the complexity of one iteration. The most costly steps in one iteration are Steps 4 and 6. As a subroutine these will use the fact that we can compute  $M_i^t$  and  $N_i^t$  using  $O(t) = O(T)$  calls to the earlier base classifiers and the sample  $S$ .

---

<sup>9</sup> There is one small change in the proof of Claim 9: since we condition on all runs of quantum counting in Step 4 giving an estimate of  $\sum_i M_i^t$  up to multiplicative error 1.1, we now have  $\sum_{i=1}^m M_i^t \geq \kappa m/1.1$  for all  $t$  until termination.

## 64:14 Improved Quantum Boosting

- Step 4. The approximation of  $s = \sum_{i=1}^m M_i^t$  up to multiplicative error 1.1 using Theorem 7 costs  $\tilde{O}(T\sqrt{m})$  (for simplicity assume  $\varepsilon \gg 1/m$  to ensure the condition  $s \geq 1$  in Theorem 7 holds.)
- Step 6. Preparing one copy of  $|D^t\rangle$  costs  $\tilde{O}(T/\sqrt{\varepsilon})$  by Section 4.1 (using our setting of  $\kappa = \varepsilon/2.2$ ), so overall this step costs  $\tilde{O}(WT/\sqrt{\varepsilon})$ .

Adding these costs shows that one iteration costs  $\tilde{O}(T(W/\sqrt{\varepsilon} + \sqrt{m}))$ . Plugging in  $T = O(\frac{1}{\varepsilon\gamma^2})$ , and the same sample size  $m = \tilde{O}(d/\varepsilon^3\gamma^2)$  as for classical SmoothBoost (from Eq. (5)), gives our main result:

► **Theorem 14.** *Let  $\mathcal{W}$  be a  $\gamma$ -weak quantum learner of complexity  $W$  for concept class  $\mathcal{C}$ , with hypothesis class  $\mathcal{H}_{\text{weak}}$  of VC-dimension  $d$ . Then given  $m$  examples according to Eq. (5), QuantumSmoothBoost is an  $(\varepsilon, \delta)$ -PAC learner for  $\mathcal{C}$ , with hypothesis class  $\mathcal{H}_{\text{strong}}$ . It runs the weak learner  $O(\frac{1}{\varepsilon\gamma^2})$  times and uses*

$$\tilde{O}(T^2(W/\sqrt{\varepsilon} + \sqrt{m})) = \tilde{O}\left(\frac{W}{\varepsilon^{2.5}\gamma^4} + \frac{\sqrt{m}}{\varepsilon^{2}\gamma^4}\right) = \tilde{O}\left(\frac{W}{\varepsilon^{2.5}\gamma^4} + \frac{\sqrt{d}}{\varepsilon^{3.5}\gamma^5}\right)$$

other operations (elementary computational steps, queries to the sample, and evaluations of base classifiers).

For direct comparison with the quantum boosting result of Arunachalam and Maity [5], we instantiate this by setting  $\varepsilon = \delta = 1/3$ , in which case the complexity of Quantum SmoothBoost is

$$\tilde{O}\left(\frac{W}{\gamma^4} + \frac{\sqrt{d}}{\gamma^5}\right).$$

This polynomially improves over the time complexity  $\tilde{O}\left(\frac{W^{1.5}\sqrt{d}}{\gamma^{11}}\right)$  of the quantum version of AdaBoost of [5], in the  $W$ -dependence but especially in the  $\gamma$ -dependence.

### 5 Future work

This work leaves open many questions for future work:

- The  $\gamma$ -dependence of Quantum Smoothboost is still slightly worse than in classical boosting ( $1/\gamma^5$  vs  $1/\gamma^4$ ). Is there a way to improve this further, or can we prove a lower bound on the  $\gamma$ -dependence for every quantum boosting algorithm that has  $\sqrt{d}$ -dependence on the VC-dimension of the weak learner's hypothesis class?
- Our quantum version of SmoothBoost improves the cost per iteration but not the number of iterations, which remains  $T = O(1/\varepsilon\gamma^2)$ . Can we reduce the number of iterations by quantizing SmoothBoost differently, or by quantizing some other boosting approach? There are classical boosting-type algorithms with fewer iterations than SmoothBoost, for instance [8], but it's not clear there how to significantly reduce the cost per iteration on a quantum computer.
- Boosting has many applications in theory and practice. Can we find applications where quantum Smoothboost is particularly suitable – some problem where the weak learner has relatively large advantage  $\gamma$  and large VC-dimension  $d$ , so that the square-root improvement in  $d$  dominates the worse dependence on  $1/\gamma$ ?
- Can we do boosting for *agnostic* learning, where the label  $y$  of an example  $(x, y)$  is not determined by  $x$  but  $(x, y)$  is jointly generated by some distribution  $\mathcal{D}$  on  $\mathcal{X} \times \mathcal{Y}$ ?

- What about learning with various kinds of noise in the sample: random classification noise, or Massart noise, or Tsybakov noise, or malicious noise? Servedio [22] designed SmoothBoost motivated by its ability to deal with malicious noise on the labels of the sample: if  $1/100$  of the  $m$  given examples have their label flipped, then a distribution that puts probability  $\leq c/m$  on each  $i$  only puts total probability  $\leq c/100$  on the erroneous examples. Servedio used this to give a learning algorithm for linear threshold functions that is robust against small,  $\gamma$ -dependent amounts of malicious noise (see [17] and references therein for follow-up work). This result straightforwardly carries over to the quantum case using our Quantum SmoothBoost (we omit the details), but we don't know much more about quantum learning with malicious noise.
- What about learning functions that have a larger range than just  $\{-1, 1\}$ ?
- Our booster relies on a VC-dimension-based analysis of boosting. However, there are non-VC-based analyses of boosting [20], which to some extent can explain why boosting avoids overfitting even when the VC-dimension of the booster's hypothesis class gets large (i.e., when  $T$  gets large). We might use this to give an alternative analysis of quantum boosters.

---

## References

- 1 Scott Aaronson and Patrick Rall. Quantum approximate counting, simplified. In *Symposium on Simplicity in Algorithms*, pages 24–32, 2020. [arXiv:1908.10846](#).
- 2 Joran van Apeldoorn and András Gilyén. Improvements in quantum SDP-solving with applications. In *Proceedings of 46th ICALP*, pages 99:1–99:15, 2019. [arXiv:1804.05058](#). [doi:10.4230/LIPIcs.ICALP.2019.99](#).
- 3 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-solvers: Better upper and lower bounds. *Quantum*, 4(230), 2020. Earlier version in FOCS'17. [arXiv:1705.01843](#).
- 4 Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012. [doi:10.4086/toc.2012.v008a006](#).
- 5 Srinivasan Arunachalam and Reevu Maity. Quantum boosting. In *Proceedings of 37th International Conference on Machine Learning (ICML'20)*, pages 377–387, 2020. [arXiv:2002.05056](#).
- 6 Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. *SIGACT News*, 48(2):41–67, 2017. [arXiv:1701.06806](#).
- 7 Srinivasan Arunachalam and Ronald de Wolf. Optimal quantum sample complexity of learning algorithms. *Journal of Machine Learning Research*, 19, 2018. Earlier version in CCC'17. [arXiv:1607.00932](#).
- 8 Boaz Barak, Moritz Hardt, and Satyen Kale. The uniform hardcore lemma via approximate Bregman projections. In *Proceedings of 20th ACM-SIAM SODA*, pages 1193–1200, 2009.
- 9 Jacob Biamonte, Peter Wittek, Nicola Pancotti, P. Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671), 2017. [arXiv:1611.09347](#).
- 10 Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. In *Proceedings of 46th ICALP*, pages 27:1–27:14, 2019. [arXiv:1710.02581](#). [doi:10.4230/LIPIcs.ICALP.2019.27](#).
- 11 Fernando G. S. L. Brandão and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *Proceedings of 58th IEEE FOCS*, pages 415–426, 2017. [arXiv:1609.05537](#). [doi:10.1109/FOCS.2017.45](#).
- 12 Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74. AMS, 2002. [arXiv:quant-ph/0005055](#).

## 64:16 Improved Quantum Boosting

- 13 Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 5:119–139, 1997.
- 14 Yoav Freund, Robert E. Schapire, and Naoki Abe. A short introduction to boosting. *Journal of the Japanese Society For Artificial Intelligence*, 14(771-780):1612, 1999.
- 15 Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th ACM STOC*, pages 212–219, 1996. [quant-ph/9605043](#).
- 16 Yassine Hamoudi, Maharshi Ray, Patrick Reberntrost, Miklos Santha, Xin Wang, and Siyi Yang. Quantum algorithms for hedging and the Sparsitron. [arXiv:2002.06003](#), 2020.
- 17 Phil M. Long and Rocco A. Servedio. Learning large-margin halfspaces with more malicious noise. In *Proceedings of NIPS*, pages 91–99, 2011.
- 18 Patrick Reberntrost, Yassine Hamoudi, Maharshi Ray, Xin Wang, Siyi Yang, and Miklos Santha. Quantum algorithms for hedging and the learning of ising models. *Phys. Rev. A*, 103:012418, January 2021. [doi:10.1103/PhysRevA.103.012418](#).
- 19 Robert E. Schapire and Yoav Freund. Boosting: Foundations and algorithms. *Kybernetes*, 2013.
- 20 Robert E. Schapire, Yoav Freund, Peter Bartlett, and Wee Sun Lee. Boosting the margin: a new explanation for the effectiveness of voting methods. In *Proceedings of 14th International Conference on Machine Learning (ICML'97)*, pages 322–330, 1997.
- 21 Maria Schuld and Francesco Petruccione. Quantum ensembles of quantum classifiers. *Scientific reports*, 8(1):1–12, 2018. [arXiv:1704.02146](#).
- 22 Rocco A. Servedio. Smooth boosting and learning with malicious noise. *Journal of Machine Learning Research*, 4:633–648, 2003. Earlier version in COLT/EuroCOLT'01.
- 23 Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014.
- 24 Ximing Wang, Yuechi Ma, Min-Hsiu Hsieh, and Manhong Yung. Quantum speedup in adaptive boosting of binary classification. *Science China Physics, Mechanics & Astronomy*, 64(2):220311, 2021. [arXiv:1902.00869](#).