

# The Cost of Parallelizing Boosting

*Xin Lyu  
Hongxun Wu  
Junzhao Yang*

Electrical Engineering and Computer Sciences  
University of California, Berkeley

Technical Report No. UCB/Eecs-2024-207

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2024/Eecs-2024-207.html>

December 4, 2024



Copyright © 2024, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

---

# The Cost of Parallelizing Boosting

by Hongxun Wu

---

## Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,  
University of California at Berkeley, in partial satisfaction of the requirements for the  
degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

### Committee:



---

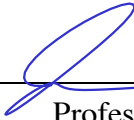
Professor Avishay Tal  
Research Advisor

12/4/2024

---

(Date)

\* \* \* \* \*



---

Professor Jelani Nelson  
Second Reader

12/2/2024

---

(Date)

# The Cost of Parallelizing Boosting

Xin Lyu \*

Hongxun Wu †

Junzhao Yang ‡

July 2023

## Abstract

We study the cost of parallelizing weak-to-strong boosting algorithms for learning, following the recent work of Karbasi and Larsen. Our main results are two-fold:

- First, we prove a tight lower bound, showing that even “slight” parallelization of boosting requires an exponential blow-up in the complexity of training.

Specifically, let  $\gamma$  be the weak learner’s advantage over random guessing. The famous AD-ABOOST algorithm produces an accurate hypothesis by interacting with the weak learner for  $\tilde{O}(1/\gamma^2)$ <sup>1</sup> rounds where each round runs in polynomial time.

Karbasi and Larsen showed that “significant” parallelization must incur exponential blow-up: Any boosting algorithm either interacts with the weak learner for  $\Omega(1/\gamma)$  rounds or incurs an  $\exp(d/\gamma)$  blow-up in the complexity of training, where  $d$  is the VC dimension of the hypothesis class. We close the gap by showing that any boosting algorithm either has  $\Omega(1/\gamma^2)$  rounds of interaction or incurs a smaller exponential blow-up of  $\exp(d)$ .

- Complementing our lower bound, we show that there exists a boosting algorithm using  $\tilde{O}(1/(t\gamma^2))$  rounds, and only suffer a blow-up of  $\exp(d \cdot t^2)$ .

Plugging in  $t = \omega(1)$ , this shows that the smaller blow-up in our lower bound is tight. More interestingly, this provides the first trade-off between the parallelism and the total work required for boosting.

Our lower bound follows from a novel interpretation of parallel boosting as a variant of “coin game”. The upper bound is inspired by the “bagging” technique in machine learning and draws a connection to differential privacy.

## 1 Introduction

Boosting is one of the most important contributions from the theory to the practice of machine learning. In the 1980s, Kearns and Valiant [14, 15] raised the fundamental question of boosting: Can every learning algorithm that, given any distribution, outputs a classifier with slightly nontrivial accuracy (i.e., weak learner) be “boosted” into an algorithm that outputs a classifier with arbitrarily high accuracy (i.e., strong learner)? This question was answered positively by Schapire [27]. Since then, boosting has found numerous applications in different fields of machine learning [1, 4, 28].

---

\*Department of EECS, UC Berkeley. Email: xinlyu@berkeley.edu. Supported by Avishay Tal’s Sloan Research Fellowship, NSF CAREER Award CCF-2145474, and Jelani Nelson’s ONR grant N00014-18-1-2562.

†Department of EECS, UC Berkeley. Email: wuhx@berkeley.edu. Supported by Avishay Tal’s Sloan Research Fellowship, NSF CAREER Award CCF-2145474, and Jelani Nelson’s ONR grant N00014-18-1-2562.

‡IIIS, Tsinghua University. Email: yang-jz20@mails.tsinghua.edu.cn

<sup>1</sup>In this paper, we use  $\tilde{O}$  to hide the terms logarithmic in  $\gamma$  and  $m$  where  $m$  is the size of the training set.

Boosting frameworks such as XGBoost [5, 6], LightGBM [13] have been popular tools in machine learning practice.

The classical boosting algorithm ADABOOST, developed by Freund and Schapire [9], follows an elegant strategy consisting of many adaptive rounds. In the initial round, it runs the weak learner on the uniform distribution over the training set to obtain the first classifier. For every following round, it adjusts the distribution by increasing the mass on the data points where the last classifier makes a mistake. Then, it obtains the next classifier by running the weak learner on this new distribution. Intuitively, this forces the weak learner to “focus” on correcting its previous mistakes. After sufficiently many rounds, ADABOOST aggregates the opinion of all classifiers by taking a majority vote. If the hypotheses in each round have an advantage  $\gamma$  over random guess,  $\tilde{O}(1/\gamma^2)$  rounds are sufficient for the aggregated classifier to achieve a 99% accuracy.

The power of the seemingly simple strategy comes from two sources: (1) running the weak learner on many different distributions and (2) the adaptivity in choosing these distributions. This is also a shared characteristic of modern gradient boosters [6, 10, 13]. However, adaptivity comes with a high cost. In reality, weak learners are usually implemented by training the model over the given distribution. One call to the weak learner could take days (especially for many modern deep neural networks). The adaptivity makes the algorithm inherently sequential: no matter how many computational resources are available, boosting will always blow up the already-long training time multiplicatively.

**Impossibility of significant parallelization.** Recently, Karbasi and Larsen [12] proved that a significant parallelization is impossible, showing that certain adaptivity is necessary for boosting.

To set up minimal notation, a  $\gamma$ -weak learner is a weak learner that always outputs classifiers with at least  $\gamma$  advantage over the random guess. We say a weak learner uses a concept class  $\mathcal{H}$  if its output is always a concept in  $\mathcal{H}$ . Then, the main theorem of Karbasi and Larsen [12] reads:

**Theorem 1** (Special Case of Theorem 1, [12], Informally Rephrased). *There is a universal constant  $\alpha > 0$  such that the following is true for any weak-to-strong learner (boosting algorithm)  $A$ . Suppose  $A$  achieves 0.99 accuracy with every valid  $\gamma$ -weak ( $0 < \gamma < \alpha$ ) learner  $\mathcal{W}$  that uses a concept set of VC dimension  $d$ . Then, either  $A$  interacts with the weak learner for at least  $p \geq 1/\gamma$  rounds, or  $A$  makes at least  $t \geq \min(\tilde{\Omega}(\exp(d/\gamma)), \exp(\exp(d)))$  oracle calls to the weak learner in total.<sup>2</sup>*

As AdaBoost takes  $\tilde{O}(1/\gamma^2)$  rounds, this shows that any algorithm that quadratically parallelizes AdaBoost must make  $\exp(d/\gamma)$  calls in each round. The authors of [12] concluded that “*The classic algorithms, such as AdaBoost, use  $O(\gamma^{-2} \ln m)$  rounds.*<sup>3</sup> Thus it is conceivable that boosting can be somewhat parallelized. We leave this as an exciting direction for future research.”, which still leaves hope for some mild parallelization of boosting.

## 1.1 Our Results

**Impossibility of slight parallelization** Answering this hope in the negative, we prove that slight parallelization is impossible without an exponential number of calls to the weak learner each round.

We note that, in the prior work of Long and Servedio [19], they studied a more restrictive class of parallel boosting algorithms. For this class of boosters, they proved a stronger lower bound

<sup>2</sup>In their paper, they proved that the loss is at least  $\ell \geq \exp(-O(p \max(\gamma, \ln(tp)\gamma^2/d)))$  when  $t \leq \exp(\exp(d))$ . Plug in  $\ell \leq 0.01$  and  $p \leq 1/\gamma$ , one gets that  $t \geq \min(\tilde{\Omega}(\exp(d/\gamma)), \exp(\exp(d)))$  here.

<sup>3</sup>Again, here  $m$  is the size of the data set.

saying that  $\Omega(1/\gamma^2)$  rounds are necessary, even with an unbounded number of calls each round. (See Section 1.2 for a more detailed discussion.)

**Theorem 2** (Special Case of Theorem 5). *There is a universal constant  $\alpha > 0$  such that the following is true for any weak-to-strong learner (boosting algorithm)  $A$ . Suppose  $A$  achieves 0.99 accuracy with every valid  $\gamma$ -weak ( $0 < \gamma < \alpha$ ) learner  $\mathcal{W}$  that uses a concept set of VC dimension  $d$ . Then, either  $A$  interacts with the weak learner for at least  $p \geq 1/\gamma^2$  rounds, or  $A$  makes at least  $t \geq \exp(\Omega(d))$  oracle calls to the weak learner in total.*

We proved this result via a connection to the coin problem, which we will discuss in detail in Section 1.3.

**Tradeoff between parallelisms and total work.** Comparing Theorem 1 and Theorem 2, one may notice that the lower bounds for  $t$  are  $\min(\exp(d/\gamma), \exp(\exp(d)))$  and  $\exp(d)$  respectively. The  $\exp(d)$  bound in Theorem 2 seems weaker. A priori, it is not clear whether this is a technicality of our proof or an inherent nature of the problem.

To answer this, we accompany our lower bound with a new boosting algorithm with fewer rounds of interaction, which shows that Theorem 2 is nearly tight regarding the dependence on  $\gamma$ .

**Theorem 3** (Informal version of Theorem 6). *There exists a boosting algorithm, that boosts any  $\gamma$ -weak learner  $\mathcal{W}$  with a hypothesis class of VC dimension  $d$  to a strong learner (with accuracy 0.99), using  $\tilde{O}(1/(\gamma^2 R))$  rounds of interaction and  $\exp(\tilde{O}(dR^2))$  calls to  $\mathcal{W}$  per round.*

*Specifically, when  $R = \omega(1)$  and  $R \ll \log(1/\gamma)$ , it uses  $o(1/\gamma^2)$  rounds and  $\exp(\tilde{O}(d))$  queries per round.*

In the work of Karbasi and Larsen [12], they gave a one-round boosting algorithm by essentially enumerating all sparsely supported distributions and training weak learners on them in parallel. In comparison, our few-round boosting algorithm is inspired by a practical technique in machine learning, bagging (also known as bootstrap aggregation). Interestingly, similar ideas of using bagging to help parallelization of boosting have been experimentally explored by machine learning researchers (e.g., Figure 4 of [31], Algorithm 2 of [26], Algorithm 2 of [21]). They gave experimental evidence that these algorithms of similar spirits work well in practice. To the best of our knowledge, our work is the first to give a theoretical analysis of such parallelization. Our analysis uses differential privacy tools, which we will discuss in more detail in Section 1.3.

We also note that the idea from Theorem 2 can be adapted to give a smooth trade-off on the lower bound side.

**Theorem 4** (Informal version of Theorem 8). *There is a universal constant  $\alpha > 0$  such that the following is true for any weak-to-strong learner (boosting algorithm)  $A$ . Suppose  $A$  achieves 0.99 accuracy with every valid  $\gamma$ -weak ( $0 < \gamma < \alpha$ ) learner  $\mathcal{W}$  that uses a concept set of VC dimension  $d$ . For every parameter  $R \leq \frac{1}{\gamma^2}$ , either  $A$  interacts with the weak learner for at least  $p \geq \frac{1}{R} \frac{1}{\gamma^2}$  rounds, or  $A$  makes at least  $t \geq \exp(\Omega(dR))$  oracle calls to the weak learner in total.*

Note when  $p = 1/\gamma$ , Theorem 4 gives a lower bound of  $t \geq \exp(d/\gamma)$  which improves over the  $t \geq \min(\exp(d/\gamma), \exp(\exp(d)))$  lower bound in Theorem 1 by Karbasi and Larsen [12], showing that the  $\exp(\exp(d))$  term is merely an artifact of the previous approach. This is a bonus from our coin problem approach.

All together, our knowledge of the cost for parallelizing boosting can be summarized as Figure 1. Closing the gap between the lower and upper bound for the  $p \approx \frac{1}{\gamma}$  regime remains an interesting open problem.

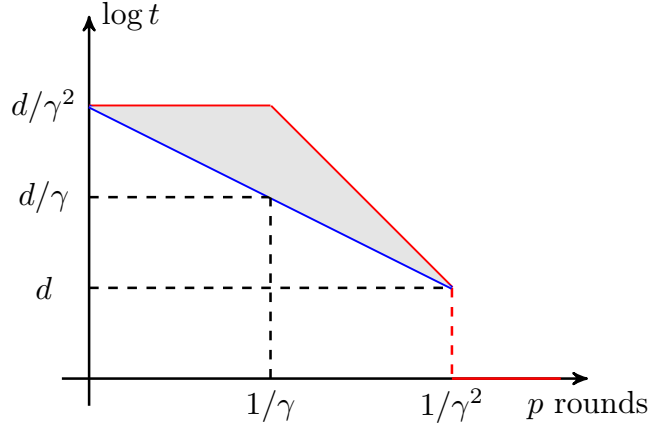


Figure 1: Tradeoff between rounds of interaction  $p$  and number of parallel queries in a single round  $t$  (from Theorem 3 and Theorem 4 (ignoring all the log factors)). The red line is the upper bound and blue line is the lower bound. There is a phase transition when  $p \approx 1/\gamma^2$ . The gray area indicates the current gap in the upper and lower bounds.

## 1.2 Related Works

**Prior works on Parallel Boosting.** Numerous works in literature have studied the amount of resources required for boosting. The earlier work of Freund [8] showed that a total number of  $\Omega(1/\gamma^2)$  calls to the weak learner is required. An important prior work of Long and Servedio [19] considered the complexity of parallel boosting and proved that  $\tilde{\Omega}(1/\gamma^2)$  rounds of interaction with the weak learner is required, *regardless* of the number of calls in each round. While this appears quantitatively stronger than our results, we note that the model considered in their work puts a significant restriction on the boosting algorithms, which makes the results weaker.

Specifically, they assume that the query distributions of the booster is always derived through “filtering”. Namely, it can only adjust the weight of any input  $x \in X$  based on the classifications of  $x$  given by the hypotheses from the weak learner in previous rounds. Under this assumption, they show that even an oblivious weak learner, whose output is independent of the query distribution, is sufficient to “fool” the booster. In contrast, our approach does not bind the booster to such constraints. Because of the assumption, their model cannot capture the boosting algorithms that query randomized distributions, such as our algorithms or the algorithms of Karbasi and Larsen [12]. Note that the ability to query randomized distributions is the key feature that allows us to achieve a tradeoff between query and round complexity and break their lower bound.

On the algorithm side, there are several boosting algorithms that make parallel calls to the weak learner, including the literature on boosting using decision tree learning [16] and branching programs [22, 11, 20, 18]. Despite making parallel queries, all these algorithms still need  $\Omega(1/\gamma^2)$  rounds of interaction with the weak learner. This is inherent because they are “filtering” algorithms, which are subject to the aforementioned lower bound by Long and Servedio [19].

**Differential Privacy and Bagging.** The idea of taking subsamples, performing computation, and finally aggregating the results is perhaps ubiquitous in computer science. For the machine learning side, the bagging technique has been investigated by numerous empirical and theoretical works and has a rich literature. For one recent example, it was shown in [17] that the bagging technique, coupled with the Empirical Risk Minimization algorithm, gives a PAC-learning algorithm

with optimal sample complexity. For the differential privacy side, the so-called “privacy amplification by sampling” [2] and the “sample-and-aggregate” [25, 29] framework have been the main workhorses behind many exciting developments.

Our work contributes a new perspective on the power of subsampling in machine learning through an information-theoretical technique first developed and popularized in the Differential Privacy community, namely the advanced composition theorem. In our work, the privacy of the data set is not a primary concern, and our result is not directly comparable to the line of DP works. We hope that our work can inspire further investigation into the power of sample-and-aggregation and its interplay with ML and DP.

### 1.3 Our Techniques

**Review of the previous lower bound.** Before introducing our construction, it is instructive to review the previous lower bound of Theorem 1, which follows from a recursive set hiding structure. The main idea of [12] is as follows: Let  $X_0$  be the set of all data points, and say we finally care about the accuracy on uniform distribution over  $X_0$ . The weak learner samples a chain of sets  $X_0 \supset X_1 \supset X_2 \supset X_3 \supset \dots \supset X_p$  with  $|X_i| = (1 - 2\gamma)|X_{i-1}|$  uniformly at random. Then, a random hypothesis  $c : X_0 \rightarrow \{\pm 1\}$  is drawn as the ground truth, and the goal of the strong learner is to approximate  $c$  as well as possible. Roughly speaking, for the  $i$ -th round of queries, the weak learner will answer randomly within  $X_i$  and truthfully outside of it (on  $X_0 \setminus X_i$ ). The idea of the construction is intuitive: before the  $i$ -th round, the boosting algorithm knows nothing about  $X_i$  and the correct labels of data points in  $X_{i-1}$ . After round  $i$ , the boosting algorithm may learn the set  $X_i$ , but it has no information about the correct labels of data points in  $X_i$ .

Since the boosting algorithm does not know what  $X_i$  is before round  $i$ , ideally, the distribution  $\mathcal{D}$  it picks in the round  $i$  should have  $1 - 2\gamma$  mass within  $X_i$  and  $2\gamma$  mass in  $X_0 \setminus X_i$ . Hence, one should expect it to have accuracy roughly  $(1 - 2\gamma) \cdot \frac{1}{2} + 2\gamma \cdot 1 = \frac{1}{2} + \gamma$ . This makes it a valid weak learner. Taking  $p = \Theta(1/\gamma)$ , one can make sure that  $|X_p| \geq 0.01|X_0|$  so that the boosting algorithm, which has no information about correct labels in  $|X_i|$ , cannot have accuracy better than 0.99.

The reason why this approach stuck at  $1/\gamma$  rounds is that, in order to ensure a  $\gamma$ -advantage over random guess, the weak learner can conceal at most  $1 - 2\gamma$  fraction of the correct labels, then naturally, after roughly  $1/\gamma$  rounds, most correct labels are released to the boosting algorithm.

**Connection to Coin Problem.** We overcome this barrier by drawing an interesting connection to the coin problem. We think of the weak learner as having  $|X_0|$  many random coins. For any  $x \in X_0$ , if the correct label is 1, then the corresponding coin is  $\gamma$ -biased towards head, and if the true label is 0, the coin is  $\gamma$ -biased towards tail. Loosely speaking, in each round  $i$ , the weak learner tosses every coin *exactly once* and makes up a hypothesis  $h^{(i)} : X_0 \rightarrow \{0, 1\}$  that records the toss results. In the ideal case, every distribution  $\mathcal{D}$  picked by the boosting algorithm is spread-out enough. Then by concentration among the coin tosses, with exponentially small probability, each query  $\mathcal{D}$  can be answered by  $h^{(i)}$  with an  $\Omega(\gamma)$  advantage.

In this way, the weak learner gives away only  $O(\gamma^2)$  bits of information on the correct label on every data point, while the prior construction reveals a  $2\gamma$  fraction of correct labels in every round. Consequently, we are able to provide a tight  $\Omega(1/\gamma^2)$  lower bound on the number of rounds.

We note that in this ideal case, this coin problem approach is similar to that of the previous lower bounds [8, 19]. The difference lies in the case that the query distribution  $\mathcal{D}$  is sparse (i.e., not spread-out enough). We construct a family of hypotheses  $\mathcal{H}^{(i)}$  consisting of  $h^{(i)}$  and a small number of random hypotheses. (This construction has the same spirit as the approach by Karbasi and Larsen [12].) In the case where  $\mathcal{D}$  is sparse, we show that the weak learner can answer the



query with  $\Omega(\gamma)$  advantage using one of the random hypotheses in  $\mathcal{H}^{(i)}$ . We also slightly simplified their approach by a more fine-grained division between the sparse and spread-out cases.

**Connection to Bagging and Differential Privacy** Bagging is the technique of creating many data sets by independently sampling the training set. Our algorithm itself is very simple (See Algorithm 1). The key idea is to simulate a weak learner with a small set of hypotheses  $\mathcal{H}_k$  obtained by bagging.

---

**Algorithm 1:** Sketch of the Parallel Boosting Algorithm

---

```

1 Initialize the distribution  $D$  to be uniform over the training set.
2 for  $k \leftarrow 0$  to  $\lceil K/R \rceil - 1$  do
    // Bagging Step.
3   for  $q \leftarrow 1$  to  $Q$  do
4     Subsample  $n$  elements from the training set according to  $D$ .
5      $D_{k,q} \leftarrow$  the uniform distribution over these  $n$  elements.
6    $\mathcal{H}_k \leftarrow$  hypotheses obtained by running weak learner on  $D_{k,q}$  for each  $q \in [Q]$  separately.
    // Boosting Step.
7   for  $r \leftarrow kR$  to  $(k+1)R - 1$  do
8      $h_r \leftarrow \arg \min_{h \in \mathcal{H}_k} \mathcal{L}_D(h)$ . // Here  $h_r$  is easy to find because  $|\mathcal{H}_k| \leq Q$ .
9     Update  $D$  as if the weak learner has outputted  $h_r$  on  $D$ .
10 return Aggregation of all  $K$  different  $h_r$ 's.
```

---

For its analysis, we want to argue that bagging helps the algorithm be more robust against a few distribution updates. This is when *advanced composition*, a tool from differential privacy (see, e.g. [7]), comes in handy. Recall that bagging generates many new datasets by subsampling. Loosely speaking, we will use it to show that the mechanism for generating one such new dataset satisfies approximate DP. (Looking ahead, the input to the mechanism will be the accumulated loss on every data point.) Hence intuitively, for just a few updates, they will not affect the result of bagging much. With some effort, we can show that at least one of the hypotheses in  $\mathcal{H}_k$  will have  $\gamma$  advantage over random guess. This guarantees that at Line 9, we successfully simulated a valid weak learner by our parallel calls to  $\mathcal{W}$  at Line 6.

## 2 Preliminaries

**Notation.**  $\mathbb{R}_+^n$  denotes the set of all non-negative real numbers.  $\text{Ber}(p)$  denotes the Bernoulli distribution with parameter  $p$ . Namely,  $X \in \text{Ber}(p)$  takes value 1 with probability  $p$  and value 0 otherwise. For a finite set  $A$ , we use  $x \sim A$  to denote a random variable  $x$  distributed uniformly at random from  $A$ . We usually use  $\mathcal{X}$  to denote the universe. Then, for every subset  $S \subseteq \mathcal{X}$ , we denote by  $\bar{S} := \mathcal{X} \setminus S$  the complementary set of  $S$ .

### 2.1 PAC Learning and Boosting

This paper studies *boosting algorithms*. Namely, algorithms that make oracle queries to a *weak learner* and convert it into a *strong learner*.

**PAC Learning.** We consider the task of *binary classification* and employ the PAC learning framework, which we briefly review now. We use  $\mathcal{X}$  to denote a finite domain of inputs. The binary labels are denoted by  $\{\pm 1\}$ . Let  $\mathcal{H} \subseteq \{h : \mathcal{X} \rightarrow \{\pm 1\}\}$  be a class of hypotheses. The VC dimension of  $\mathcal{H}$ , denoted by  $\text{VCDim}(\mathcal{H})$ , is the maximum  $d \in \mathbb{N}$  such that there exists a subset  $S \subseteq \mathcal{X}$  of size  $d$  that is *shattered*<sup>4</sup> by  $\mathcal{H}$ . A learning task is usually described by a distribution  $\mathcal{D}$  over  $\mathcal{X} \times \{\pm 1\}$ , and the goal of learning is to find a hypothesis  $\hat{h}$  (not necessarily from  $\mathcal{H}$ ) minimizing the loss function on  $\mathcal{D}$ , which is defined as  $\mathcal{L}_{\mathcal{D}}(\hat{h}) := \mathbf{E}_{(x,y) \sim \mathcal{D}}[h(x) \neq y]$ . The distribution  $\mathcal{D}$  is called *realizable*, if there is a hypothesis  $h \in \mathcal{H}$  such that  $\mathcal{L}_{\mathcal{D}}(h) = 0$ .

**Weak learner.** Let  $\gamma > 0$  be a parameter. For our purpose, a  $\gamma$ -*weak learner* with hypothesis set  $\mathcal{H}$  is an oracle that can be queried with distributions supported on  $\mathcal{X} \times \{\pm 1\}$ . Given a query distribution  $\mathcal{D}$ , the weak learner returns a hypothesis  $h \in \mathcal{H}$  such that

$$\mathcal{L}_{\mathcal{D}}(h) := \mathbf{E}_{(x,y) \sim \mathcal{D}}[h(x) \neq y] \leq \frac{1}{2} - \gamma.$$

If there are multiple  $h \in \mathcal{H}$  with loss bounded by  $\frac{1}{2} - \gamma$ , the weak learner can return an arbitrary one. The weak learner may declare “failure” if there is no valid hypothesis. We call the oracle a “weak learner” because the learner produces hypotheses “slightly” better than random guess (note that we think of  $\gamma$  as small).

**Strong learner and Boosting.** A strong learner is a learning algorithm that, given  $n$  i.i.d. samples from a *realizable* distribution  $\mathcal{D}$ , with probability  $1 - \delta$ , produces a hypothesis  $\hat{h}$  with  $\mathcal{L}_{\mathcal{D}}(\hat{h}) \leq \varepsilon$ . Usually, in a strong learning algorithm, the accuracy  $\varepsilon$  and confidence  $\delta$  can be arbitrarily small with only a  $\text{poly}(1/\varepsilon, \log(1/\delta))$  overhead in the sample and computational complexity.

In the boosting framework, the goal is to design a general weak-to-strong learner (a.k.a. a boosting algorithm) *without* the knowledge of the hypothesis class  $\mathcal{H}$ . Instead, the boosting algorithm is given oracle access to a weak learner for  $\mathcal{H}$ , and the boosting procedure should work well for every valid weak learner oracle.

More precisely, let  $\mathcal{A}$  be a boosting algorithm. There is a hypothesis class  $\mathcal{H}$  and a realizable distribution  $\mathcal{D}$ , both unknown to the boosting algorithm  $\mathcal{A}$ . First,  $\mathcal{A}$  receives  $m$  i.i.d. samples from  $\mathcal{D}$ , denoted by  $\{(x_i, y_i)\}_{i \in [m]}$ . Then,  $\mathcal{A}$  interacts with a weak learner  $\mathcal{W}$  of  $\mathcal{H}$  for  $p \geq 1$  rounds. In each round,  $\mathcal{A}$  sends a set of at most  $t$  queries to  $\mathcal{W}$ , where each query is a distribution  $\mathcal{D}'$  over  $S$ . The weak learner then answers all the queries in parallel. Namely, for each  $\mathcal{D}'$ , the weak learner reports a hypothesis  $h \in \mathcal{H}$  such that  $\mathcal{L}_h(\mathcal{D}') \leq \frac{1}{2} - \gamma$ . Finally, after the  $p$  rounds of interaction,  $\mathcal{A}$  returns a hypothesis  $h^*$  (not necessarily from  $\mathcal{H}$ ) that tries to minimize the loss  $\mathcal{L}_{\mathcal{D}}(h^*)$ . We use  $\mathcal{A}^{\mathcal{W}}(S)$  to denote the output of  $\mathcal{A}$  with oracle  $\mathcal{W}$  and input data set  $S$ .

We can see that the complexity of the boosting algorithm is parameterized by three parameters  $m, p, t$ . We call  $m$  the *sample complexity* of  $\mathcal{A}$ , and say the algorithm has *parallel complexity*  $(p, t)$ .

### 3 Lower Bound Against Slight Parallelization

Our negative result shows that, given oracle access to a  $\gamma$ -weak learner, any boosting algorithm with  $o(\frac{1}{\gamma^2})$  rounds of interaction cannot learn the target hypothesis with  $1 - o(1)$  accuracy unless it makes exponential (in VC dimension) many calls to the weak learner in each round. We formalize our claim into the following main theorem.

<sup>4</sup>This means the projection of  $\mathcal{H}$  onto  $S$  consists of all  $2^{|S|}$  possible hypotheses.

**Theorem 5.** *There is a universal constant  $c_{\text{lb}} > 0$  for which the following is true. For every  $0 < \gamma < 1/2$ , every  $d, m \geq 1$ , let  $\mathcal{A}$  be a boosting algorithm that uses  $m$  samples and has parallel complexity  $(p, t)$  where  $p \leq \min(\frac{c_{\text{lb}}d}{\gamma^2}, \exp(c_{\text{lb}}d))$  and  $t \leq \exp(c_{\text{lb}}d)$ . Then, there exists a domain  $\mathcal{X}$  of size  $2m$ , a hypothesis class  $\mathcal{H} \subseteq \{\pm 1\}^{|\mathcal{X}|}$  of VC dimension  $\frac{d}{c_{\text{lb}}}$ , a realizable distribution over  $\mathcal{X} \times \{\pm 1\}$ , and a  $\gamma$ -weak learner  $\mathcal{W}$  for  $\mathcal{H}$  such that*

$$\mathbf{E}_{S \sim \mathcal{D}^m, \mathcal{A}} [\mathcal{L}_{\mathcal{D}}(\mathcal{A}^{\mathcal{W}}(S))] \geq \exp(-O(p\gamma^2 + 1)). \quad (1)$$

The rest of the section is devoted to the proof of Theorem 5.

### 3.1 Construction of Hard Instances

In this subsection, we fix parameters  $d, m$  in Theorem 5, and describe our construction of the hard hypotheses class  $\mathcal{H}$  and weak learner  $\mathcal{W}$ . We will give a randomized construction that is *independent* of the boosting algorithm. Then, for any fixed boosting algorithm  $\mathcal{A}$ , we will show that our construction incurs a noticeable loss in expectation, which implies that there exists an instantiation of the construction on which  $\mathcal{A}$  fails to boost.

**Defining the learning task.** Let  $m > 0$  be the number of samples requested by the strong learner, which can be arbitrarily large (e.g., it can be as large as  $d^{d^d}$ ). We aim to construct a learning task with a hypothesis class of VC dimension  $d$ , on which the strong learner fails even with  $m$  samples.

Now, given  $m$ , we define the input domain to be  $\mathcal{X} = \{x_1, x_2, \dots, x_{2m}\}$ . We sample a concept  $c$  over  $\mathcal{X}$  by setting each  $c(x_i)$  to  $\pm 1$  randomly and independently. In the lower bound proof, we will think of  $c$  as the “ground truth”, which the strong learner attempts to learn. Hence, we define the target distribution  $\mathcal{D}$  as the uniform distribution over  $\{(x_i, c(x_i))\}_{1 \leq i \leq 2m}$ .

**Weak learner and hypothesis class.** We now construct the hypothesis class  $\mathcal{H}$ . We conduct the construction in  $p$  stages. In each stage  $i \in [p]$ , we first create a hypothesis class  $\mathcal{H}^{(i)}$  of size  $2^{O(d)}$  where big- $O$  hides an absolute constant. The construction of  $\mathcal{H}^{(i)}$  is as follows.

- First, we sample a hypothesis  $a^{(i)} \in \{1, -1\}^{2m}$  by drawing each  $a^{(i)}(x_j) \sim (-1)^{\text{Ber}(1/2 - C_{\text{bias}} \cdot \gamma \cdot c(x_j))}$  independently. Considering  $c$  as the “ground truth”, the hypothesis  $a^{(i)}$  gives an expected advantage of  $C_{\text{bias}}\gamma$  over random guess.
- For a parameter  $\hat{d} = \Theta(d)$  (Here,  $\Theta$  hides an absolute constant that will be specified later), draw  $r_1^{(i)}, \dots, r_{\hat{d}}^{(i)}$  where each  $r_j^{(i)} \sim \{\pm 1\}^{2m}$ . Namely, we draw  $2^{\hat{d}}$  hypotheses uniformly at random.
- Finally, we define  $\mathcal{H}^{(i)} = \{a^{(i)}\} \cup \{r_j^{(i)} : 1 \leq j \leq 2^{\hat{d}}\}$ .

We set our final hypothesis class as  $\mathcal{H} = \{c\} \cup \bigcup_{i=1}^p \mathcal{H}^{(i)}$ . It is easy to see that  $|\mathcal{H}| \leq 2^{O(d)}$  and consequently  $\text{VCDim}(\mathcal{H}) \leq O(d)$ .

Next, we define the weak learner  $\mathcal{W}$ . We sort all hypotheses in the order of  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}, \{c\}$  (the order inside a sub-class can be arbitrary). Then, on a given query  $\mathcal{D}$ ,  $\mathcal{W}$  outputs the *first* hypothesis  $h$  from the list such that  $\mathcal{L}_{\mathcal{D}}(h) \leq \frac{1}{2} - \gamma$ . This defines a valid weak learner for  $\mathcal{H}$ .

**Proof Outline.** We present the overall proof structure and state two key claims. Assuming them, we quickly finish the proof of Theorem 5. We prove the claims in subsequent subsections.

Our first claim is that, with high probability, the queries made by  $\mathcal{A}$  in the first  $i$  rounds can be answered by hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(i)}$  only.

**Claim 1.** *Let  $\mathcal{A}$  be a boosting algorithm. For every  $i \leq p$  and every fixed realization of  $c \sim \{\pm 1\}^{2m}$  and  $S \sim \mathcal{D}^m$ , with probability  $1 - i \cdot \exp(-\Omega(d))$  (over the construction of  $\mathcal{H}$  and the interaction between  $\mathcal{A}$  and  $\mathcal{W}$ ), the queries made by  $\mathcal{A}$  in the first  $i$  rounds can be answered by hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(i)}$ .*

Claim 1 is proved in Section 3.2. Assuming it for now, the following corollary is immediate.

**Corollary 1.** *With probability  $1 - \exp(-\Omega(d))$ , the interaction between  $\mathcal{A}$  and  $\mathcal{W}$  can be simulated given  $S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}$ . In particular, the simulation does not require the knowledge of  $c$ .*

*Proof.* By Claim 1 and the condition that  $p \leq \exp(c_{\text{lb}} \cdot d)$  where  $c_{\text{lb}} > 0$  is sufficiently small, we have with probability  $1 - \exp(-\Omega(d))$  that the queries of  $\mathcal{A}$  can be answered with hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}$  only.

We now describe the simulation. We can start by simulating  $\mathcal{A}$  given the data  $S$ . In each of the  $p$  rounds,  $\mathcal{A}$  issues a set of queries. Under the aforementioned event, each query can be answered by some hypothesis from  $\mathcal{H} \setminus \{c\}$ , which means we can compute the responses of the weak learner given  $\mathcal{H} \setminus \{c\}$ . Finally, after the  $p$  rounds of interaction, we finish by simulating the remaining pieces of  $\mathcal{A}$  and reporting its output.  $\square$

Next, we would like to show that, given only  $S$  and  $\mathcal{H} \setminus \{c\}$ , the learner cannot produce a hypothesis with low generalization error.

**Claim 2.** *Let  $\mathcal{B} = \mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})$  be an arbitrary aggregation procedure that outputs a hypothesis  $\hat{h}$ . We have*

$$\mathbf{E}_{c, S, \mathcal{H}} [\mathcal{L}_{\mathcal{D}}(\hat{h} = \mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}))] \geq \exp(-O(p\gamma^2 + 1)).$$

Claim 2 is proved in Section 3.3. Assuming it, we prove our main theorem.

**Proof of Theorem 5.** Let  $\mathcal{E}$  be the event defined in Corollary 1. It follows that  $\Pr[\bar{\mathcal{E}}] \leq \exp(-\Omega(d))$ . Under the event  $\mathcal{E}$ , the behavior of the boosting algorithm can be written as a procedure  $\mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})$  that depends only on  $S$  and  $\mathcal{H} \setminus \{c\}$ . By Claim 2, we obtain

$$\begin{aligned} & \mathbf{E}_{c \sim \{\pm 1\}^{2m}, S \sim \mathcal{D}^m, \mathcal{H}, \mathcal{A}} [\mathcal{L}_{\mathcal{D}}(\mathcal{A}^{\mathcal{W}}(S))] \\ & \geq \Pr[\mathcal{E}] \cdot \mathbf{E}[\mathcal{L}_{\mathcal{D}}(\mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})) \mid \mathcal{E}] \\ & \geq \mathbf{E}[\mathbb{1}_{\mathcal{E} \text{ happens}} \cdot \mathcal{L}_{\mathcal{D}}(\mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}))] \\ & \geq \mathbf{E}[\mathcal{L}_{\mathcal{D}}(\mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)}))] - \Pr[\bar{\mathcal{E}}] \quad (\text{loss is always bounded by 1}) \\ & \geq \exp(-O(p\gamma^2 + 1)). \end{aligned}$$

This shows that a random construction of  $\mathcal{H}$  incurs a loss of  $\exp(-O(p\gamma^2 + 1))$  in expectation. By averaging, there exists an instantiation of the construction satisfying the proposition of Theorem 5, which completes the proof.

### 3.2 Proof of Claim 1

We prove Claim 1 by inducting on  $i \leq p$ . Assume the claim is true for the first  $i - 1$  rounds. Fix the data set  $S$ . Under the event that queries of the first  $i - 1$  rounds were answered using hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(i-1)}$ , the current set of at most  $t$  queries, denoted by  $\mathcal{D}_1^{(i)}, \dots, \mathcal{D}_{t'}^{(i)}$ , are independent of  $\mathcal{H}^{(i)}$ : they only depend on the internal randomness of  $\mathcal{A}$  after conditioning on  $S$  and  $\mathcal{D}^{(j)}$  for  $j < i$ .

Let  $\mathcal{D}'$  be one query from  $\{\mathcal{D}_1^{(i)}, \dots, \mathcal{D}_{t'}^{(i)}\}$ . We will show that the probability  $\mathcal{D}'$  fails to be answered by  $\mathcal{H}^{(j)}, j \leq i$ , is bounded by  $\exp(-\Omega(d))$ . Having established the claim, we may union bound over all  $t' \leq \exp(c_{\text{lb}} \cdot d)$  queries and finish the proof for the  $i$ -th round.

Now, assume  $\mathcal{D}'$  is a query distribution that is independent of  $\mathcal{H}^{(i)}$ . We will further divide the query into two cases. Namely the *spread* case and the *concentrated* case. Let  $\alpha_{\text{thr}} > 0$  be a (large) constant to be specified later. We begin by formalizing the definition of spread distribution.

**Definition 1** (Spread distribution). Let  $\mathcal{D}'$  be a distribution over  $\mathcal{X}$ . Sort elements in  $\mathcal{X}$  by their probability mass under  $\mathcal{D}'$  and let  $w_{(i)}$  denote the  $i$ -th largest probability mass. We define the spreadness of  $\mathcal{D}'$  as

$$F(\mathcal{D}') = \sum_{i=1}^d w_{(i)} + \sqrt{d} \left( \sum_{i=d+1}^{|\mathcal{X}|} w_{(i)}^2 \right)^{1/2}.$$

We say that  $\mathcal{D}'$  is *spread* if  $F(\mathcal{D}')$  is smaller than  $\alpha_{\text{thr}}\gamma$ . Otherwise, we say  $\mathcal{D}'$  is *concentrated*.

**Remark 1.** *The idea of considering spread and concentrated queries separately is directly inspired by [12]. However, we use a different definition of “spread” than the prior work, which turns out to be much simpler and facilitate our proof significantly.*

**Concentrated Queries.** Suppose  $\mathcal{D}'$  is a concentrated distribution. In this case, we first observe that the expected loss of a truly random hypothesis  $r \sim \{1, -1\}^{2m}$  is  $\frac{1}{2}$ . For concentrated query distribution  $\mathcal{D}'$ , we will use an anti-concentration inequality to lower bound the probability that  $\mathcal{L}_{\mathcal{D}'}(r) < 1/2 - \gamma$ . To start, define a function  $F' : \mathbb{R}_+^n \times \mathbb{R}_+ \rightarrow \mathbb{R}$  for every non-negative vector  $w \in \mathbb{R}_+^n$  and real  $t > 0$ :

$$F'(w, t) := \sum_{i=1}^{\lfloor t^2 \rfloor} |w_{(i)}| + t \left( \sum_{j=\lfloor t^2 \rfloor + 1}^n w_{(j)}^2 \right)^{1/2}$$

where  $w_{(i)}$  denotes the  $i$ -th largest entry of  $w$ .

**Lemma 1** ([24]). *There are universal constants  $\alpha_1, \alpha_2 > 0$  such that the following holds: For any vector  $w \in \mathbb{R}_+^n$  and  $t > 0$ , it holds that  $\Pr_{x \sim \{\pm 1\}^n}[\langle w, x \rangle > \alpha_1 F'(w, t)] \geq \alpha_2^{-1} \exp(-\alpha_2 t^2)$ .*

By our definition, note that  $F(\mathcal{D}') = F'(\mathcal{D}', \sqrt{d})$ . Setting  $t = \sqrt{d}$  and applying Lemma 1, we obtain

$$\Pr[1/2 - \mathcal{L}_{\mathcal{D}'}(r) > \gamma] \geq \Pr[\langle \mathcal{D}', r \cdot c \rangle > \alpha_1 F(\mathcal{D}')] \geq \alpha_2^{-1} \exp(-\alpha_2 d),$$

provided that  $\alpha_{\text{thr}} > \frac{2}{\alpha_1}$ . Here,  $r \cdot c$  denotes the pointwise product of  $r$  and  $c$ . Setting  $\hat{d} = \log_2(\alpha_2 \exp(\alpha_2 d)) + 2d = \Theta(d)$ , the probability that exists one hypothesis  $r \in \mathcal{H}^{(i)}$  such that  $r$  has  $\gamma$  advantage over  $\mathcal{D}'$  is at least  $1 - (1 - \alpha_2^{-1} \exp(-\alpha_2 d))^{2^{\hat{d}}} \geq 1 - \exp(-\exp(d))$ .

**Spread Queries.** For a fixed round  $i$ , let  $\mathcal{D}'$  be a spread distribution queried in the  $i$ -th round. Recall we chose  $a^{(i)}$  as a random hypothesis with an expected advantage of  $4C_{\text{bias}} \cdot \gamma$ .

Assume the claim is true for the first  $i - 1$  rounds (that is, all the queries in the first  $i - 1$  round were answered using hypotheses from  $\mathcal{H}^{(1)} \cup \dots \cup \mathcal{H}^{(i-1)}$ ). Then, the query  $\mathcal{D}'$  is independent of  $a^{(i)}$ . We now argue that, for a spread query  $\mathcal{D}'$ , the probability that  $a^{(i)}$  fails to achieve  $\gamma$  advantage on  $\mathcal{D}'$  is exponentially small.

For each  $j \in [2m]$ , let  $y_j := \mathcal{D}'(x_j) \cdot \mathbb{1}[a^{(i)}(x_j) \neq c(x_j)]$  be the random variable denoting the contribution to the loss from the input  $x_j$ . Let  $B \subseteq [2m]$  be the set of indices of the  $d$  largest probability mass. We observe that

$$\Pr \left[ \mathcal{L}_{\mathcal{D}}(a^{(i)}) > \frac{1}{2} - \gamma \right] = \Pr \left[ \sum_{j \in B} y_j + \sum_{j \in \bar{B}} y_j > \frac{1}{2} - \gamma \right].$$

By definition, with probability one, we have

$$\sum_{j \in B} y_j \leq \sum_{j \in B} \mathcal{D}'(x_j) \leq F(\mathcal{D}') \leq \alpha_{\text{thr}} \gamma.$$

Therefore, it suffices to prove that

$$\Pr \left[ \sum_{j \in \bar{B}} y_j > \frac{1}{2} - (\alpha_{\text{thr}} + 1)\gamma \right] \leq \exp(-d).$$

We wish to apply Hoeffding's inequality. To begin with, we calculate:

$$\left( \sum_{j \in \bar{B}} \mathcal{D}'(x_j)^2 \right)^{1/2} \leq \frac{F(\mathcal{D}')}{\sqrt{d}} \leq \frac{\alpha_{\text{thr}} \cdot \gamma}{\sqrt{d}}.$$

Denote  $S := \sum_{j \in \bar{B}} y_j$ . We have

$$\mathbf{E}[S] \leq \frac{1}{2} - C_{\text{bias}} \gamma.$$

Hence, choosing  $C_{\text{bias}}$  to be larger than  $3\alpha_{\text{thr}} + 1$ , we obtain

$$\Pr[S > \frac{1}{2} - (\alpha_{\text{thr}} + 1)\gamma] \leq \Pr[S - \mathbf{E}[S] > 2\alpha_{\text{thr}} \cdot \gamma] \leq \exp \left( \frac{(2\gamma)^2}{\sum_{j \in \bar{B}_j} (2\mathcal{D}'(x_j))^2} \right) \leq \exp(-d).$$

To wrap up the two cases, given a query  $\mathcal{D}'$ , being spread or not, with probability  $1 - \exp(-\Omega(d))$  over the sampling of  $\mathcal{H}^{(i)}$ , there is an  $h \in \mathcal{H}^{(i)}$  such that  $\mathcal{L}_{\mathcal{D}'}(h) < 1/2 - \gamma$ , which means that  $\mathcal{D}'$  can be answered by a hypothesis from  $\bigcup_{j \leq i} \mathcal{H}^{(j)}$ . We conclude the proof by union-bounding over all queries in the  $i$ -th round.

### 3.3 Proof of Claim 2

In the coin problem, there is a hidden coin obeying either  $\text{Ber}(\frac{1+\varepsilon}{2})$  or  $\text{Ber}(\frac{1-\varepsilon}{2})$ . The algorithm is given as input  $n$  tosses  $x_1, x_2, \dots, x_n$  of the coin, and is asked to distinguish these two cases. In our proof, we use the following information-theoretic lower bound for the coin problem.

**Lemma 2** (Coin problem. See e.g. Exercise 3.19, [23]). *Let  $\varepsilon \in (0, 1/2)$ . For every (possibly randomized) algorithm  $f : \{0, 1\}^n \rightarrow \{\pm 1\}$ , we have*

$$\mathbf{E}_{b \sim \{\pm 1\}} \left[ \mathbf{E}_{x \sim \text{Ber}(\frac{1+b \cdot \varepsilon}{2})^n} [f(x) \neq b] \right] \geq \exp(-O(\varepsilon^2 n + 1)).$$

*Proof of Claim 2.* Let  $\mathcal{B} = \mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})$  be an arbitrary procedure that outputs a hypothesis. Let  $S_0 \subseteq \mathcal{X}$  be the subset of inputs that appeared in  $S$  (namely,  $S_0$  is the projection of  $S$  onto  $\mathcal{X}$ ). Conditioning on  $S_0$ , we observe that  $\mathcal{B}$  is essentially playing  $2m - |S_0|$  independent instances of coin games in parallel.

In more detail, for each  $\mathcal{H}^{(i)}$ , all the random hypotheses from  $\mathcal{H}^{(i)}$  are independent of the ground truth  $c$  (which means  $\mathcal{B}$  could have generated them by itself). Next, each  $a^{(i)}$  is obtained by, for each input  $x \in \mathcal{X}$ , flipping a biased coin from the distribution  $\text{Ber}(\frac{1}{2} + C_{\text{bias}} \cdot \gamma \cdot c(x))$ . Fix an input  $x \in \mathcal{X} \setminus S_0$ . We further observe that the value of  $a^{(i)}(x')$  for  $x' \neq x$  is independent of  $c(x)$ . Thus, by Lemma 2, we obtain

$$\mathbf{E}_{c \sim \{\pm 1\}^{2m}} \left[ \mathbf{Pr}_{\substack{\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)} \\ h = \mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})}} [h(x) \neq c(x)] \right] \geq \exp(-O(\gamma^2 p + 1)).$$

By linearity of expectation, we thus obtain

$$\mathbf{E}_{\substack{c, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)} \\ h = \mathcal{B}(S, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p)})}} [\mathcal{L}_{\mathcal{D}}(h)] \geq \frac{2m - |S_0|}{2m} \cdot \exp(-O(\gamma^2 p)) \geq \exp(-O(\gamma^2 p + 1)).$$

Finally, averaging over  $S_0$  finishes the proof. □

## 4 Trade-off between Parallelism and Total Work

In this section, we study the round-query trade-off of parallel boosting. In Section 4.1, we will first present a boosting algorithm based on a variant ADABOOST and inspired by bagging. The key lemma in its analysis will be proved via differential privacy in Section 4.2. Finally, we generalize our lower bound to give the round-query trade-off in Section 4.3.

### 4.1 The Few Rounds Boosting Algorithm

This subsection is devoted to the following theorem.

**Theorem 6.** *For any data distribution  $\mathcal{D}$  over  $\mathcal{X}$ , any unknown concept  $c : \mathcal{X} \rightarrow \{1, -1\}$ , any  $\gamma$ -weak learner  $\mathcal{W}$  that produces hypothesis class of VC dimension  $d$ , for any  $1 \leq R \leq 1/2\gamma$ , let training set  $S \sim \mathcal{D}^m$  and setting  $m = \tilde{O}(d\gamma^{-2})$ , there exists a (randomized) boosting algorithm  $\mathcal{A}$  satisfying  $\mathbf{Pr}_{h \leftarrow \mathcal{A}}[\mathcal{L}_{\mathcal{D}}(h) < 0.1] \geq 0.9$ , such that  $\mathcal{A}$  runs in  $O(\gamma^{-2} \ln m/R)$  rounds and makes  $\exp(O(dR^2)) \ln(1/\gamma)$  queries each round.*

**Algorithm Description.** Our algorithm contains  $K/R$  rounds. In the  $k$ -th round ( $k \in [K/R]$ ), it contains a bagging step followed by a boosting step.



- Bagging step: It samples  $Q$  different training sets of size  $n$  from the current distribution  $D_{kR}$ . For each of these training sets, it makes a parallel call to the weak learner  $\mathcal{W}$ . Let  $\mathcal{H}_k$  be the set of output hypotheses of the weak learner calls.
- Boosting step: It constructs a simulated weak learner  $\mathcal{W}'$  which, for any distribution, simply outputs the best-performing hypotheses among these  $Q$  hypotheses in  $\mathcal{H}_k$ . It then performs  $R$  steps of ADABOOST by interacting  $R$  rounds with simulated weak learner  $\mathcal{W}'$ . ADABOOST updates the current distribution  $D_{kR}$  times, generating  $D_{kR+1}, D_{kR+2}, \dots, D_{(k+1)R}$ .

---

**Algorithm 2:** Round-query trade-off boosting

---

**Input** : Labelled Training Set  $(S, c(S)) = (x_1, c(x_1)), (x_2, c(x_2)), \dots, (x_m, c(x_m))$ ,  
number of calls to weak learner per round  $Q$

**Output:** Aggregated voting classifier  $\hat{h}$

```

1  $w \leftarrow \frac{1}{2} \ln\left(\frac{1/2+\gamma/4}{1/2-\gamma/4}\right)$  // Learning Rate
2  $n \leftarrow c'd\gamma^{-2}$ 
3  $\mathcal{D}_0 \leftarrow (\frac{1}{m}, \frac{1}{m}, \dots, \frac{1}{m})$ 
4 for  $k \leftarrow 0$  to  $\lceil K/R \rceil - 1$  do
    // Bagging Step
5    $\mathcal{Q}_k \leftarrow \emptyset$ 
6   for  $q \leftarrow 1$  to  $Q$  do
7     Sample a multiset  $T_{k,q} \sim (\mathcal{D}_{kR})^n$ 
8      $D_{k,q} \leftarrow$  uniform distribution over  $T_{k,q}$ 
9      $\mathcal{Q}_k \leftarrow \mathcal{Q}_k \cup \{D_{k,q}\}$ 
10   $\mathcal{H}_k \leftarrow$  hypotheses from  $\mathcal{W}$  after querying  $\mathcal{Q}_k$  in parallel
    // Boosting Step
11  for  $r \leftarrow kR$  to  $\min((k+1)R, K) - 1$  do
12    if exists  $h \in \mathcal{H}_k$  s.t.  $\mathcal{L}_{\mathcal{D}_r}(h) \leq 1/2 - \gamma/4$  then
13       $h_r \leftarrow$  such  $h$  that  $\mathcal{L}_{\mathcal{D}_r}(h) \leq 1/2 - \gamma/4$ 
14      for  $i \leftarrow 1$  to  $m$  do
15         $\mathcal{D}_{r+1}(i) \leftarrow \mathcal{D}_r(i) \exp(-c(x_i)h_r(x_i) \cdot w)$ 
16      Normalize  $\mathcal{D}_{r+1}$ 
17    else
18      return failed
19 for  $x \in \mathcal{X}$  do
20    $g(x) \leftarrow \frac{1}{K} \sum_{k=0}^{K-1} h_k(x)$ 
21    $\hat{h}(x) \leftarrow \text{sign}(g(x))$ 
22 return  $\hat{h}$ 

```

---

Specifically, we will set  $K = 16\gamma^{-2} \ln m$  and  $Q = \exp(16c'dR^2) \ln(1/\gamma)$  where  $m$  is the size of the training set. See Algorithm 2 for the detailed pseudocode of this algorithm.

**Analysis.** By the standard analysis of ADABOOST, as long as simulated  $\mathcal{W}'$  always succeeds in outputting a hypotheses with  $\gamma$  advantage (i.e. Algorithm 2 never fails and reaches Line Line 18), the aggregated hypotheses  $\hat{h}$  will satisfy  $\mathcal{L}_D(h) < 0.1$  with good probability. Hence the key of the analysis is to prove the following lemma:



**Lemma 3** (Key lemma). *There exists a universal constant  $c' > 1$ , such that for any  $r \in [K]$ , if we set  $R \leq 1/2\gamma$ ,  $Q \geq \exp(16c'dR^2) \ln(1/\gamma)$ ,  $K = 16\gamma^{-2} \ln m$  and  $m = \ln(1/\gamma) \ln(d/\gamma) c'd/\gamma^2$  in Algorithm 2, the probability that  $h$  does not exist for that  $\mathcal{D}_r$  at Line 12 is less than  $0.01/K$ .*

We defer the proof of Lemma 3 to the next subsection. Once we have Lemma 3, a union bound over all  $K$  many  $\mathcal{D}_r$ 's shows that the algorithm fails with probability less than 0.01. When it does not fail, the following theorem guarantees generalization.

**Theorem 7.** *Let  $c : \mathcal{X} \rightarrow \{1, -1\}$  be an unknown concept,  $\mathcal{H}$  be a hypothesis class of VC dimension  $d$ , and  $\mathcal{D}$  be an arbitrary distribution over  $\mathcal{X}$ . Let  $S$  be a training set of size  $m$  randomly sampled from  $\mathcal{D}^m$ . Let  $\hat{h}$  be the voting classifier generated by Algorithm 2 ( $g(x) = \frac{1}{k} \sum_{i=1}^k h_i(x)$  and  $\hat{h}(x) = \text{sign}(g(x))$ ). It must have large margins  $c(x)g(x) \geq \gamma$  for all  $x \in S$ . It then follows that, with probability  $1 - \delta$ ,*

$$\mathcal{L}_{\mathcal{D}}(\hat{h}) \leq \alpha_{\text{gen}} \cdot \frac{d \ln(m) \ln(m/d) + \ln(1/\delta)}{\gamma^2 m}$$

where  $\alpha_{\text{gen}}$  is a universal constant.

*Proof Sketch.* We will prove in Appendix A that by standard analysis of ADABOOST,  $\hat{h}$  will satisfy the large margin property  $c(x)g(x) \geq \gamma/16$  for all  $x \in S$ . The rest follows from Breiman's min-margin bound [3].  $\square$

*Proof of Theorem 6.* Conditioning on the algorithm does not fail, it follows from Theorem 7 that  $\mathcal{L}_{\mathcal{D}}(\hat{h}) \leq 0.01$  with probability 0.99 for  $m = \ln(1/\gamma) \ln(d/\gamma) c'd/\gamma^2$  and a large constant  $c'$ . Since our algorithm only fails with probability 0.01 from Lemma 3, this proves Theorem 6.  $\square$

## 4.2 Proof of Lemma 3

To prove that the algorithm fails with small probability, we have to show that, at Line 12, such good  $h \in \mathcal{H}_k$  always exists. Note  $\mathcal{H}_k$  is generated by running weak learners on the multiset  $T_{k,q}$ 's. We will prove that at least one of the  $T_{k,q}$ 's is an " $\varepsilon$ -approximation" of  $\mathcal{D}_r$ , so that the hypothesis generated from running weak learner on  $T_{k,q}$  will also have  $O(\gamma)$  for  $\mathcal{D}_r$ .

**Definition 2** ( $\varepsilon$ -approximation). A multiset  $T$  is an  $\varepsilon$ -approximation for a hypothesis class  $\mathcal{H}$  if for any  $h \in \mathcal{H}$ , we have  $|\mathcal{L}_{\mathcal{D}}(h) - \mathcal{L}_{\mathcal{D}_T}(h)| \leq \varepsilon$  where  $\mathcal{D}_T$  is the uniform distribution over  $T$ .

Note that  $T_{k,q}$  is drawn from  $(\mathcal{D}_{kR})^n$  not from  $(\mathcal{D}_r)^n$ . For it to be a good approximation for  $\mathcal{D}_r$ , we have to show that  $\mathcal{D}_{kR}$  and  $\mathcal{D}_r$  are close. Namely, the distribution does not change much after a few (less than  $R$ ) updates of ADABOOST. Specifically, the only property we need is  $(\varepsilon, \delta)$ -indistinguishability:

**Definition 3** ( $(\varepsilon, \delta)$ -indistinguishable distributions). For any two distributions  $\mathcal{D}$  and  $\mathcal{D}'$  over  $S$ , they are  $(\varepsilon, \delta)$ -indistinguishable if for any event  $E \subseteq S$ , it holds that  $\Pr_{\mathcal{D}}[E] \leq \Pr_{\mathcal{D}'}[E] \cdot e^\varepsilon + \delta$  and  $\Pr_{\mathcal{D}'}[E] \leq \Pr_{\mathcal{D}}[E] \cdot e^\varepsilon + \delta$  where  $\Pr_{\mathcal{D}}[E]$  denotes  $\Pr_{x \in \mathcal{D}}[x \in E]$ .

To establish that  $(\mathcal{D}_{kR})^n$  not  $(\mathcal{D}_r)^n$  are  $(\varepsilon, \delta)$ -indistinguishable, we start by upper bounding the max-divergence between  $\mathcal{D}_{kR}$  and  $\mathcal{D}_r$ .

**Lemma 4.** *The max-divergence between  $\mathcal{D}_{kR}$  and  $\mathcal{D}_r$ , defined as  $D_\infty(\mathcal{D}, \mathcal{D}') := \ln(\sup_{x \in S} \mathcal{D}(x)/\mathcal{D}'(x))$ , satisfies  $D_\infty(\mathcal{D}_{kR}, \mathcal{D}_r) \leq 2\gamma R$ . When two distributions are switched, we also have  $D_\infty(\mathcal{D}_r, \mathcal{D}_{kR}) \leq 2\gamma R$ .*

*Specifically, this implies that  $\mathcal{D}_{kR}$  and  $\mathcal{D}_r$  are  $(2\gamma R, 0)$ -indistinguishable.*

*Proof.* First of all, since  $|kR - r| \leq R$  and max-divergence satisfies triangle inequality, it suffices to prove  $D_\infty(\mathcal{D}_i, \mathcal{D}_{i+1}) \leq 2\gamma$  and  $D_\infty(\mathcal{D}_{i+1}, \mathcal{D}_i) \leq 2\gamma$  for any  $kR \leq i \leq r - 1$ .

Define exponential accumulated loss

$$Z_{i,j} := \exp\left(-\sum_{r=0}^{i-1} c(x_j)h_r(x_j) \cdot w\right) \text{ for all } j \in [m],$$

$$Z_i := \sum_{j=1}^m Z_{i,j}.$$

We can then express these two distributions as

$$\mathcal{D}_i(x_j) = Z_{i,j}/Z_i,$$

$$\mathcal{D}_{i+1}(x_j) = Z_{i+1,j}/Z_{i+1}.$$

Note that  $Z_{i+1,j}$  will either be  $Z_{i,j}e^w$  or  $Z_{i,j}e^{-w}$ , thus  $\left|\ln \frac{Z_{i+1,j}}{Z_{i,j}}\right| \leq w$ . By definition

$$\frac{Z_{i+1}}{Z_i} = \frac{\sum_j Z_{i+1,j}}{\sum_j Z_{i,j}},$$

hence we also have  $\left|\ln \frac{Z_{i+1}}{Z_i}\right| \leq w$ . Note  $\gamma \in [0, 0.5]$  and

$$w = \ln((1/2 + \gamma/4)/(1/2 - \gamma/4))/2 \leq \ln(1 + 2\gamma)/2 \leq \gamma, \quad (2)$$

thus  $\left|\ln \frac{\mathcal{D}_i(x_j)}{\mathcal{D}_{i+1}(x_j)}\right| \leq \left|\ln \frac{Z_{i+1,j}}{Z_{i,j}}\right| + \left|\ln \frac{Z_{i+1}}{Z_i}\right| \leq 2\gamma$  for any  $j \in [m]$ . Therefore we have  $D_\infty(\mathcal{D}_i, \mathcal{D}_{i+1}) \leq 2\gamma$  and  $D_\infty(\mathcal{D}_{i+1}, \mathcal{D}_i) \leq 2\gamma$ .  $\square$

Then, loosely speaking, we will view  $(\mathcal{D}_{kR})^n$  and  $(\mathcal{D}_r)^n$  as  $\mathcal{D}_{kR}$  and  $\mathcal{D}_r$  ‘‘composed’’  $n$  times. This allows us to apply the advanced composition theorem from the differentiating privacy literature.

**Lemma 5** (Advanced Composition [7]). *For any integer  $n > 0$ , any two  $(\varepsilon, \delta)$ -indistinguishable distributions  $\mathcal{D}$  and  $\mathcal{D}'$ , and any  $\delta' > 0$ , the product distributions  $\mathcal{D}^n$  and  $(\mathcal{D}')^n$  are  $(\widehat{\varepsilon}, \widehat{\delta})$ -indistinguishable for*

$$\widehat{\varepsilon} = n\varepsilon(e^\varepsilon - 1) + \varepsilon\sqrt{2n \ln(1/\delta')},$$

$$\widehat{\delta} = n\delta + \delta'.$$

**Corollary 2.**  $D_r$  and  $D_{kR}$  are  $(12c'dR^{-2}, 1/4)$ -indistinguishable.

*Proof.* Using advanced composition (Lemma 5) and  $n = c'd\gamma^{-2}$ , setting  $\delta' = 1/4$ , we have

$$\widehat{\varepsilon} \leq 2n\varepsilon^2 + \varepsilon\sqrt{2n \ln(4)} \leq 8\gamma^2 R^2 n + 4\gamma R\sqrt{n} = 8c'dR^2 + 4R\sqrt{c'd} \leq 12c'dR^2$$

for large enough  $c' > 1$ . Here we used  $2\gamma R \leq 1$  and  $e^x - 1 \leq 2x$  for  $x \in [0, 1]$ .  $\square$

The rest of the proof proceeds as follows. First, By standard VC theory (Lemma 6), for large enough  $n$ , a random subset  $T$  drawn from  $(\mathcal{D}_r)^n$  is  $\varepsilon$ -approximation of  $\mathcal{D}_r$  with good probability. Together with the fact that  $(\mathcal{D}_{kR})^n$  and  $(\mathcal{D}_r)^n$  are  $(\varepsilon, \delta)$  indistinguishable, we can argue that at least of the  $T_{k,q}$ 's will  $\varepsilon$ -approximate  $(\mathcal{D}_r)^n$ . This finishes the proof.

**Lemma 6** ([30]). *There exists a universal constant  $c' > 0$  for which the following is true. For any  $0 < \varepsilon, \delta < 1$ , any distribution  $\mathcal{D}$  over  $\mathcal{X}$  and any hypothesis class  $\mathcal{H}$  of VC dimension  $d$ , a random subset  $T \sim \mathcal{D}^n$  is an  $\varepsilon$ -approximation for  $\mathcal{H}$  with probability  $1 - \delta$ , provided that  $n \geq c'(d + \ln(1/\delta))\varepsilon^{-2}$ .*

*Proof of Lemma 3.* Let  $c$  be a large enough constant. We denote by  $G \subseteq \mathcal{X}^n$  the set of multisets that  $\gamma/2$ -approximates  $\mathcal{D}_r$ . This is the set of “good” multisets. For a randomly sampled  $T' \sim \mathcal{D}_r^n$ , it follows from Lemma 6 that  $\Pr[T' \in G] \geq 3/4$ .

For a randomly sampled  $T \sim \mathcal{D}_r^n$ , it follows from Corollary 2 that  $\Pr[T' \in G] \leq \Pr[T \in G] \cdot \exp(8c'dR^2) + 1/4$ . Together with  $\Pr[T' \in G] \geq 3/4$ , it implies that  $\Pr[T \in G] \geq \exp(-8c'dR^2)/2$ .

Moreover,  $T \in G$  means  $\mathcal{L}_{\mathcal{D}_r}(\mathcal{W}(\mathcal{D}_T)) \leq \mathcal{L}_{\mathcal{D}_r}(\mathcal{W}(\mathcal{D}_T)) + \gamma/2 \leq \frac{1}{2} - \gamma/2$  where  $\mathcal{W}(\mathcal{D}_T)$  is the hypothesis returned by the weak learner  $\mathcal{W}$  when it is called with distribution  $\mathcal{D}_T$ . Hence, as long as one of  $T_{k,q}$ 's is in  $G$ , the algorithm will not fail for  $\mathcal{D}_r$  at Line 7. When  $Q \geq \exp(16c'dR^2) \ln(1/\gamma)$ , the failure probability for  $\mathcal{D}_r$  is

$$\begin{aligned} \Pr[\text{fail for } \mathcal{D}_r] &\leq \Pr[\forall j \in [Q], T_{k,j} \notin G] \\ &\leq (1 - \exp(-12c'dR^2)/2)^Q \\ &\leq \exp(-\exp(-12c'dR^2)Q/2) \\ &\leq \exp(-\exp(4c'dR^2) \ln(1/\gamma)/2) \\ &\leq 0.01/(16\gamma^{-2} \ln m) = 0.01/K, \end{aligned}$$

where we used the condition that  $m = \ln(1/\gamma) \ln(d/\gamma) c'd/\gamma^2$  in the last inequality.  $\square$

### 4.3 Generalizing Lower Bound to Smooth Trade-Off

In this subsection, we will generalize our lower bound to rounds fewer than  $O(1/\gamma^2)$  by a slight modification.

**Theorem 8.** *There is a universal constant  $c_{\text{lb}} > 0$  for which the following is true. For every  $R \geq 1, \gamma \in (0, 1/2)$  and every  $d, m \geq 1$ , let  $\mathcal{A}$  be a boosting algorithm that uses  $m$  samples and has parallel complexity  $(p, t)$  where  $p \leq \min(\frac{c_{\text{lb}}}{\gamma^2}, \exp(c_{\text{lb}} \cdot d))/R$  and  $t \leq \exp(c_{\text{lb}} R \cdot d)$ . Then, there exists a domain  $\mathcal{X}$  of size  $2m$ , a hypothesis class  $\mathcal{H} \subseteq \{\pm 1\}^{|\mathcal{X}|}$  of VC dimension  $\frac{d}{c_{\text{lb}}}$ , a realizable distribution over  $\mathcal{X} \times \{\pm 1\}$ , and a  $\gamma$ -weak learner  $\mathcal{W}$  for  $\mathcal{H}$  such that*

$$\mathbf{E}_{S \sim \mathcal{D}^m, \mathcal{A}} [\mathcal{L}_{\mathcal{D}}(\mathcal{A}^{\mathcal{W}}(S))] \geq \exp(-O(p\gamma^2 R)). \quad (3)$$

To prove Theorem 8, we use the construction in Section 3.1 but with  $p \cdot R$  stages (namely, the hypothesis class will be  $\mathcal{H} = \bigcup_{j=1}^{pR} \mathcal{H}^{(j)} \cup \{c\}$  where each  $\mathcal{H}^{(j)}$  contains a special hypothesis  $a^{(j)}$  and many purely random hypotheses).

Then, we have the following observation.

**Claim 3.** *Consider the setup as in Theorem 8. With probability  $1 - \exp(-\Omega(d))$ , all the queries from  $\mathcal{A}$  can be answered by hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(p \cdot R)}$ .*

*Proof.* Fix a realization of  $c \sim \{\pm 1\}^{2m}$  and  $S \sim \mathcal{D}^m$ . We prove by induction on  $i \in [p]$  that, the queries from the first  $i$  rounds can be answered by hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(iR)}$  with probability  $1 - i \cdot \exp(-\Omega(d))$ .

Assume the claim has been established for  $i - 1$ . Under the event that queries of the first  $(i - 1)$  rounds were answered by hypotheses from  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(iR)}$ , we get that each query  $\mathcal{D}'$  in

the  $i$ -th round is independent of  $\mathcal{H}^{(iR+1)}, \dots, \mathcal{H}^{((i+1)R)}$ , and that  $\mathcal{H}^{(iR+1)}, \dots, \mathcal{H}^{((i+1)R)}$  are all mutually independent. By the argument in Section 3.2, the probability that  $\mathcal{D}'$  can be answered by a hypothesis from  $\mathcal{H}^{(iR+j)}$  is  $1 - \exp(-\Omega(d))$  for every fixed  $j \in [R]$ . It follows that the probability that  $\mathcal{D}'$  cannot be answered by any hypothesis from  $\mathcal{H}^{(iR+1)}, \dots, \mathcal{H}^{((i+1)R)}$  is at most  $\exp(-R \cdot \Omega(d))$ . We can then union-bound over at most  $\exp(cRd)$  queries and finish the proof for the  $i$ -th round.  $\square$

*Proof of Theorem 8.* Given Claim 3, with probability  $1 - \exp(-\Omega(d))$ , the interaction between  $\mathcal{A}$  and the weak learner can be simulated given only  $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(pR)}$ . Then, the same ‘‘coin game’’ argument gives the desired lower bound for the loss function.  $\square$

## Acknowledgements

We thank the SODA reviewers for pointing out the important references and various helpful comments. We are grateful to Yuzhou Gu for helpful discussion.

## References

- [1] Jacob Abernethy, Pranjal Awasthi, and Satyen Kale. A multiclass boosting framework for achieving fast and provable adversarial robustness. *arXiv preprint arXiv:2103.01276*, 2021.
- [2] Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *NeurIPS*, pages 6280–6290, 2018.
- [3] Leo Breiman. Prediction games and arcing algorithms. *Neural computation*, 11(7):1493–1517, 1999.
- [4] Nataly Brukhim, Elad Hazan, and Karan Singh. A boosting approach to reinforcement learning. *Advances in Neural Information Processing Systems*, 35:33806–33817, 2022.
- [5] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [6] Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, Rory Mitchell, Ignacio Cano, Tianyi Zhou, et al. Xgboost: extreme gradient boosting. *R package version 0.4-2*, 1(4):1–4, 2015.
- [7] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60, 2010.
- [8] Yoav Freund. Boosting a weak learning algorithm by majority. *Inf. Comput.*, 121(2):256–285, 1995.
- [9] 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*, 55(1):119–139, 1997.
- [10] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [11] Adam Kalai and Rocco A Servedio. Boosting in the presence of noise. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 195–205, 2003.

- [12] Amin Karbasi and Kasper Green Larsen. The impossibility of parallelizing boosting. *arXiv preprint arXiv:2301.09627*, 2023.
- [13] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.
- [14] M Kearns. Thoughts on hypothesis boosting, ml class project. 1988.
- [15] M Kearns and LG Valiant. Cryptographic limitations on learning boolean formulae and finite automata. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 433–444, 1989.
- [16] Michael Kearns and Yishay Mansour. On the boosting ability of top-down decision tree learning algorithms. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 459–468, 1996.
- [17] Kasper Green Larsen. Bagging is an optimal PAC learner. In *COLT*, volume 195 of *Proceedings of Machine Learning Research*, pages 450–468. PMLR, 2023.
- [18] Phil Long and Rocco Servedio. Adaptive martingale boosting. *Advances in Neural Information Processing Systems*, 21, 2008.
- [19] Phil Long and Rocco Servedio. Algorithms and hardness results for parallel large margin learning. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.
- [20] Philip M Long and Rocco A Servedio. Martingale boosting. In *Learning Theory: 18th Annual Conference on Learning Theory, COLT 2005, Bertinoro, Italy, June 27-30, 2005. Proceedings 18*, pages 79–94. Springer, 2005.
- [21] Fernando Lozano and Pedro Rangel. Algorithms for parallel boosting. In *Fourth International Conference on Machine Learning and Applications (ICMLA'05)*, pages 6–pp. IEEE, 2005.
- [22] Yishay Mansour and David McAllester. Boosting using branching programs. *Journal of Computer and System Sciences*, 64(1):103–112, 2002.
- [23] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [24] Stephen J Montgomery-Smith. The distribution of rademacher sums. *Proceedings of the American Mathematical Society*, 109(2):517–522, 1990.
- [25] Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In *STOC*, pages 75–84. ACM, 2007.
- [26] Indranil Palit and Chandan K Reddy. Scalable and parallel boosting with mapreduce. *IEEE Transactions on Knowledge and Data Engineering*, 24(10):1904–1916, 2011.
- [27] Robert E Schapire. The strength of weak learnability. *Machine learning*, 5:197–227, 1990.
- [28] Zebang Shen, Hamed Hassani, Satyen Kale, and Amin Karbasi. Federated functional gradient boosting. In *International Conference on Artificial Intelligence and Statistics*, pages 7814–7840. PMLR, 2022.

- [29] Abhradeep Thakurta and Adam D. Smith. Differentially private feature selection via stability arguments, and the robustness of the lasso. In *COLT*, volume 30 of *JMLR Workshop and Conference Proceedings*, pages 819–850. JMLR.org, 2013.
- [30] Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity: festschrift for alexey chervonenkis*, pages 11–30. Springer, 2015.
- [31] C Yu and DB Skillicorn. Parallelizing boosting and bagging. *Queen’s University, Kingston, Canada, Tech. Rep*, 2001.

## A Generalization of ADABOOST

In this appendix, we provide a proof for Theorem 7 following the standard analysis of ADABOOST. First we prove that  $\hat{h}$  satisfies the large margin property  $c(x)g(x) \geq \gamma/16$  for all training samples  $x \in S$ . Then the generalization guarantee follows from Breiman's min-margin bound.

First, let exponential loss for round  $k$  (which is also defined in the proof of Lemma 4) be

$$Z_{i,j} := \exp\left(-\sum_{r=0}^{i-1} c(x_j)h_r(x_j) \cdot w\right) \text{ for all } j \in [m],$$

$$Z_i := \sum_{j=1}^m Z_{i,j}.$$

The main observation is that the exponential loss decreases exponentially.

**Claim 4.**  $Z_{k+1} \leq Z_k \cdot \sqrt{1 - \gamma^2/4}$ .

*Proof.* First observe that if we only normalize  $\mathcal{D}_k$  at last, we have  $\mathcal{D}_k(i) = Z_{k,i}/Z_k$ . To update  $Z_{k+1}$  from  $Z_k$ , we have  $Z_{k+1,i} = Z_{k,i} \cdot \exp(-c(x_i)h_k(x_i)w)$ , and thus

$$\begin{aligned} Z_{k+1} &= \sum_{c(x_i)=h_k(x_i)} Z_{k+1,i} + \sum_{c(x_i) \neq h_k(x_i)} Z_{k+1,i} \\ &= Z_k \left( \sum_{i:c(x_i)=h_k(x_i)} \mathcal{D}_k(i)e^{-w} + \sum_{i:c(x_i) \neq h_k(x_i)} \mathcal{D}_k(i)e^w \right) \\ &= Z_k((1 - \mathcal{L}_{\mathcal{D}_k}(h_k))e^{-w} + \mathcal{L}_{\mathcal{D}_k}(h_k)e^w) \\ &\leq Z_k((1/2 + \gamma/4)e^{-w} + (1/2 - \gamma/4)e^w) \\ &= Z_k \cdot 2\sqrt{(1/2 + \gamma/4)(1/2 - \gamma/4)} = Z_k \sqrt{1 - \gamma^2/4} \end{aligned}$$

where  $\mathcal{L}_{\mathcal{D}_k}(h_k) \leq 1/2 - \gamma/4$  and setting it to  $1/2 - \gamma/4$  maximizes the term.  $\square$

Now, since the exponential loss is exponentially small, the voting classifier will have large margins on every sample.

**Claim 5.** Let  $g$  be produced by Line 20 of Algorithm 2. Setting  $K = 16 \ln m \cdot \gamma^{-2}$ , for every training sample  $(x_i, c(x_i)) \in (S, c(S))$ , it holds that  $c(x_i)g(x_i) \geq \gamma/16$ .

*Proof.* We will prove by contradiction. Suppose exists some  $x_{i^*}$  that  $c(x_{i^*})g(x_{i^*}) < \gamma/16$ , we have

$$Z_K \geq Z_{k,i^*} > \exp(-wK\gamma/16) \geq \exp(-K\gamma^2/16) = 1/m$$

where we used the fact  $\gamma \in [0, 0.5]$  and

$$w = \ln((1/2 + \gamma/4)/(1/2 - \gamma/4))/2 \leq \ln(1 + 2\gamma)/2 \leq \gamma. \quad (4)$$

On the other hand, from Claim 4, we have

$$Z_K \leq Z_0 \cdot (1 - \gamma^2/4)^{K/2} \leq m \cdot \exp(-K\gamma^2/8) = 1/m,$$

which leads to a contradiction.  $\square$

Assuming the large margin property, the rest of Theorem 7 follows from Breiman's min-margin bound [3].