
Communication Efficient Distributed Agnostic Boosting

Shang-Tse Chen
Georgia Institute of Technology

Maria-Florina Balcan
Carnegie Mellon University

Duen Horng Chau
Georgia Institute of Technology

Abstract

We consider the problem of learning from distributed data in the agnostic setting, i.e., in the presence of arbitrary forms of noise. Our main contribution is a general distributed boosting-based procedure for learning an arbitrary concept space, that is simultaneously noise tolerant, communication efficient, and computationally efficient. This improves significantly over prior works that were either communication efficient only in noise-free scenarios or computationally prohibitive. Empirical results on large synthetic and real-world datasets demonstrate the effectiveness and scalability of the proposed approach.

1 INTRODUCTION

Distributed machine learning has received increasing attention in this “big data” era (Jordan and Mitchell, 2015). The most common use case of distributed learning is when the data cannot fit into a single machine, or when one wants to speed up the training process by utilizing parallel computation of multiple machines (Li et al., 2014; Zhang et al., 2013, 2012). In these cases, one can usually freely distribute the data across entities, and an evenly distributed partition would be a natural choice.

In this paper, we consider a different setting where the data is inherently distributed across different locations or entities. Examples of this scenario include scientific data gathered by different teams, or customer information of a multinational corporation obtained in different countries. The goal is to design an efficient learning algorithm with a low generalization error over the union of the data. Note that the distribution of the data from each source may be very different. There-

fore, to deal with the worst-case situation, we assume the data can be adversarially partitioned. This scenario has been studied for different tasks, such as supervised learning (Balcan et al., 2012; Daumé et al., 2012; Bellet et al., 2015), unsupervised learning (Balcan et al., 2013, 2014), and optimization (Boyd et al., 2011; Jaggi et al., 2014).

Traditional machine learning algorithms often only care about sample complexity and computational complexity. However, since the bottleneck in the distributed setting is often the communication between machines (Balcan et al., 2012), the theoretical analysis in this paper will focus on communication complexity. A baseline approach in this setting would be to uniformly sample examples from each entity and perform centralized learning at the center. By the standard VC-theory, a sampling set of size $O(\frac{d}{\epsilon^2} \log \frac{1}{\epsilon})$ is sufficient. The communication complexity of this approach is thus $O(\frac{d}{\epsilon^2} \log \frac{1}{\epsilon})$ examples.

More advanced algorithms with better communication complexities have been proposed in recent works (Balcan et al., 2012; Daumé et al., 2012). For example, Balcan et al. (2012) propose a generic distributed boosting algorithm that achieves communication with only logarithmic dependence on $1/\epsilon$ for any concept class. Unfortunately, their method only works in the standard realizable PAC-learning setting, where the data can be perfectly classified by a function in the hypothesis set and is noiseless. This is because many boosting algorithms are vulnerable to noise (Dietterich, 2000; Long and Servedio, 2008).

The realizable case is often unrealistic in real-world problems. Therefore, we consider the more general agnostic learning setting (Kearns et al., 1994), where there is no assumption on the target function. Since it is impossible to achieve an arbitrary error rate ϵ , the goal in this setting is to find a hypothesis with error rate close to $opt(H)$, the minimum error rate achievable within the hypothesis set H . The error bound is often in the form of $O(opt(H)) + \epsilon$. Balcan et al. (2012) propose an algorithm based on the robust generalized halving algorithm with communication complexity of $\tilde{O}(k \log(|H|) \log(1/\epsilon))$ examples. However, the algo-

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

rithm works only for a finite hypothesis set H and is computationally inefficient.

We propose a new distributed boosting algorithm that works in the agnostic learning setting. While our algorithm can handle this much more difficult and more realistic scenario, it enjoys the same communication complexity as in (Balcan et al., 2012) that is logarithmic in $1/\epsilon$ and exponentially better than the natural baselines. The algorithm is computationally efficient and works for any concept class with a finite VC-dimension. The key insight, inspired by Balcan et al. (2012), is that a constant (independent of ϵ) number of examples suffice to learn a weak hypothesis, and thus if the boosting algorithm only needs $O(\log \frac{1}{\epsilon})$ iterations, we obtain the desired result.

The key challenge in this approach is that designing an agnostic boosting algorithm is highly non-trivial, even in the centralized setting. The first agnostic boosting algorithm was proposed in (Ben-David et al., 2001). Although the number of iterations is $O(\log \frac{1}{\epsilon})$ and is asymptotically optimal, their bound on the final error rate is much weaker: instead of $O(\text{opt}(H)) + \epsilon$, the bound is $O(\text{opt}(H)^{c(\beta)}) + \epsilon$, where $c(\beta) = 2(1/2 - \beta)^2 / \ln(1/\beta - 1)$. Subsequent works (Kanade and Kalai, 2009; Gavinsky, 2003; Feldman, 2010) significantly improve the bound on the error rate. However, their algorithms all require $O(1/\epsilon^2)$ iterations, which can in turn result in $O(1/\epsilon^2)$ communication in the distributed setting.

We summarize our solution to the above challenge and other major contributions as follows.

- We show an agnostic boosting algorithm with a good error bound that only requires $O(\log \frac{1}{\epsilon})$ iterations, by performing novel analysis of a boosting algorithm (Kale, 2007) previously analyzed only in the realizable case. Our approach outperforms previous results, which either use $O(1/\epsilon^2)$ iterations (Kanade and Kalai, 2009; Feldman, 2010; Gavinsky, 2003) or achieves a much weaker error bound (Ben-David et al., 2001). Note that this could be of independent interest beyond the distributed setting.
- We efficiently adapt the agnostic boosting algorithm to the distributed setting. This results in the first algorithm that is both computationally efficient *and* communication efficient to learn a general concept class in the distributed agnostic learning setting.
- Our boosting-based approach is flexible in that it can be used with various weak learners. Furthermore, the weak learner only needs to work in the traditional centralized setting rather than in the

more challenging distributed setting. This makes it much easier to design new algorithms for different concept classes in the distributed setting.

- We empirically compare our algorithm to the existing distributed boosting algorithm (Balcan et al., 2012). It does much better on the synthetic dataset and achieves promising results on real-world datasets as well.

2 PROBLEM DEFINITION

We first introduce agnostic learning as a special case of the general statistical learning problem. Then, we discuss the extension of the problem to the distributed setting, where the data is adversarially partitioned.

2.1 Statistical learning problem

In statistical learning, we have access to a sampling oracle according to some probability distribution D over $X \times \{-1, 1\}$. The goal of a learning algorithm is to output a hypothesis h with a low error rate with respect to D , defined as $\text{err}_D(h) = E_{(x,y) \sim D}(h(x) \neq y)$. Often, we compare the error rate to the minimum achievable value within a hypothesis set H , denoted by $\text{err}_D(H) = \inf_{h' \in H} \text{err}_D(h')$. More precisely, a common error bound is in the following form.

$$\text{err}_D(h) \leq c \cdot \text{err}_D(H) + \epsilon, \tag{1}$$

for some constant $c \geq 1$ and an arbitrary error parameter $\epsilon > 0$.

Many efficient learning algorithms have been proposed for the realizable case, where the target function is in H and thus $\text{err}_D(H) = 0$. In this paper, we consider the more general case where we do not have any assumption on the value of $\text{err}_D(H)$. This is often called the agnostic learning setting (Kearns et al., 1994). Ideally, we want c in the bound to be as close to one as possible. However, for some hypothesis set H , achieving such a bound with $c = 1$ is known to be NP-hard (Feldman et al., 2009).

2.2 Extension to the distributed setting

In this work, we consider the agnostic learning problem in the distributed learning framework proposed by Balcan et al. (2012). In this framework, we have k entities. Each entity $i \in [k]$ has access to a sampling oracle according to a distribution D_i over $X \times \{-1, 1\}$. There is also a *center* which can communicate with the k entities and acts as a coordinator. The goal is to learn a good hypothesis with respect to the overall distribution $D = \frac{1}{k} \sum_{i=1}^k D_i$ without too much communication

among entities. It is convenient to calculate the communication by *words*. For example, a d -dimensional vector counts as $O(d)$ words.

Main goal The problem we want to solve in this paper is to design an algorithm that achieves error bound (1) for a general concept class H . The communication complexity should depend only logarithmically on $1/\epsilon$.

3 DISTRIBUTED AGNOSTIC BOOSTING

In this work, we show a distributed boosting algorithm for any concept class with a finite VC-dimension d . In the realizable PAC setting, the boosting algorithm is assumed to have access to a γ -weak learner that, under any distribution, finds a hypothesis with error rate at most $1/2 - \gamma$. This assumption is unrealistic in the agnostic setting since even the best hypothesis in the hypothesis set can perform poorly. Instead, following the setting of (Ben-David et al., 2001), the boosting algorithm is assumed to have access to a β -weak agnostic learner defined as follows.

Definition 1. A β -weak agnostic learner, given any probability distribution D , will return a hypothesis h with error rate

$$err_D(h) \leq err_D(H) + \beta.$$

Detailed discussion of the existence of such weak learners can be found in (Ben-David et al., 2001). Since error of $1/2$ can be trivially achieved, in order for the weak learner to convey meaningful information, we assume $err_D(H) < 1/2 - \beta$. Some prior works use different definitions. For example, Kalai et al. (2008) use the definition of (α, γ) -weak learner. That definition is stronger than ours, since an (α, γ) -weak learner in that paper implies a β -weak learner in our paper with $\beta = \alpha - \gamma$. Therefore, our results still hold by using their definition. Below we show an efficient agnostic boosting algorithm in the centralized setting.

3.1 Agnostic boosting: centralized version

The main reason why many boosting algorithms (including AdaBoost (Freund and Schapire, 1997) and weight-based boosting (Mohri et al., 2012; Shalev-Shwartz and Ben-David, 2014)) fail in the agnostic setting is that they tend to update the example weights aggressively and may end up putting too much weight on noisy examples.

To overcome this, we consider a smoothed boosting algorithm shown in Algorithm 1, based on the a boosting

Algorithm 1 Smooth Boosting (Kale, 2007)

Initialization: Fix a $\gamma \leq \frac{1}{2}$. Let $D^{(1)}$ to be the uniform distribution over the dataset S .

for $t = 1, 2, \dots, T$ **do**

1. Call the weak learner with distribution $D^{(t)}$ and obtain a hypothesis $h^{(t)}$
2. Update the example weights

$$\hat{D}^{(t+1)}(i) = D^{(t)}(i) \cdot (1 - \gamma)^{\ell_i^{(t)}} / Z^{(t)}$$

where $\ell_i^{(t)} = \mathbf{1}[h^{(t)}(x_i) \neq y_i]$ and $Z^{(t)} = \sum_i D^{(t)}(i) \cdot (1 - \gamma)^{\ell_i^{(t)}}$ is the normalization factor.

3. Project $\hat{D}^{(t+1)}$ into the feasible set \mathcal{P} of ϵ -smooth distributions

$$D^{(t+1)} = \arg \min_{D \in \mathcal{P}} RE(D \parallel \hat{D}^{(t+1)})$$

end for

Output: The hypothesis $\text{sign} \left(\frac{1}{T} \sum_{t=1}^T h^{(t)} \right)$

technique proposed in (Kale, 2007). This algorithm uses at most $O(\log 1/\epsilon)$ iterations and enjoys a nice “smoothness” property, which is shown to be helpful in the agnostic setting (Gavinsky, 2003). While this algorithm was originally analyzed in the realizable case, in this work we prove that this algorithm actually works in the much harder agnostic setting. Below we first state their results in the realizable setting.

The boosting algorithm adjusts the example weights using the standard multiplicative weight update rule. The main difference is that it performs an additional Bregman projection step of the current example weight distribution into a convex set \mathcal{P} after each boosting iteration. The Bregman projection is a general projection technique that finds a point in the feasible set with the smallest “distance” to the original point in terms of *Bregman divergence*. Here we use a particular Bregman divergence called *relative entropy* $RE(p \parallel q) = \sum_i p_i \ln(p_i/q_i)$ for two distributions p and q . To ensure that the boosting algorithm always generates a “smooth” distribution, we set the feasible set \mathcal{P} to be the set of all ϵ -smooth distributions, which is defined as follows.

Definition 2. A distribution D on S is called ϵ -smooth if $\max_i D(i) \leq \frac{1}{\epsilon|S|}$,

It is easy to verify that \mathcal{P} is a convex set. The theoretical guarantee of Algorithm 1 is shown in Theorem 1. The proof is similar to the one in (Kale, 2007), except that they use real-valued weak learners, whereas here

we only consider binary hypotheses for simplicity. For completeness, we include the proof using our notation.

Theorem 1. *Given a sample S and access to a γ -weak learner, Algorithm 1 makes at most $T = O(\frac{\log(1/\epsilon)}{\gamma^2})$ calls to the weak learner with ϵ -smooth distributions and achieves error rate ϵ on S .*

Proof. The analysis is based on the well-studied online learning from experts problem. In each round t , the learner has to make a decision based on the advice of n experts. More precisely, the learner chooses a distribution $D^{(t)}$ from a convex feasible set \mathcal{P} and follows the advice of the i -th expert with probability $D^{(t)}(i)$. Then, the losses of each expert's suggested actions are revealed as a vector $\ell^{(t)}$. The expected loss of the learner incurred by using $D^{(t)}$ is thus $\ell^{(t)}D^{(t)}$. The goal is to achieve a total expected loss $\sum_{t=1}^T \ell^{(t)}D^{(t)}$ not too much more than $\min_{D \in \mathcal{P}} \sum_{t=1}^T \ell^{(t)}D$, the cost of always using the best fixed distribution in \mathcal{P} . Step 2 and 3 of Algorithm 1, which is also known as the multiplicative weights update algorithm, has the following regret bound (Herbster and Warmuth, 2001).

Lemma 1. *For any $0 < \gamma \leq 1/2$ and any positive integer T , the multiplicative weights update algorithm generates distributions $D^{(1)}, \dots, D^{(T)} \in \mathcal{P}$ where each $D^{(t)}$ is computed only based on $\ell^{(1)}, \dots, \ell^{(t-1)}$, such that for any $D \in \mathcal{P}$,*

$$\sum_{t=1}^T \ell^{(t)}D^{(t)} \leq (1 + \gamma) \sum_{t=1}^T \ell^{(t)}D + \frac{RE(D \parallel D^{(1)})}{\gamma}.$$

To use the above result in boosting, we can think of the n examples in sample S as the set of experts. The learner's task is thus to choose a distribution $D^{(t)} \in \mathcal{P}$ over the sample at each round. The loss $\ell_i^{(t)}$ is defined to be $\mathbf{1}[h^{(t)}(x_i) = y_i]$, where $h^{(t)}$ is the hypothesis returned by the weak learner. To ensure that the boosting algorithm always generates a "smooth" distribution, we set the feasible set \mathcal{P} to be the set of all ϵ -smooth distributions. Below we show how this can be applied in boosting, as suggested by Kale (2007).

By the assumption of the γ -weak learner, we have

$$\ell^{(t)}D^{(t)} = \sum_i D^{(t)}(i) \mathbf{1}[h^{(t)}(x_i) = y_i] \geq 1/2 + \gamma.$$

After $T = \lceil \frac{2 \ln(1/\epsilon)}{\gamma^2} \rceil + 1$ rounds, we set the final hypothesis $f = \text{sign}(\frac{1}{T} \sum_{t=1}^T h^{(t)})$. Let $E \subseteq S$ be the set of examples where f predicts incorrectly. Suppose $|E| \geq \epsilon n$. Let $D = u_E$, the uniform distribution on E and 0 elsewhere. It is easy to see that $u_E \in \mathcal{P}$, since $|E| \geq \epsilon n$. For each example $(x_i, y_i) \in E$, we have

$$\sum_{t=1}^T \ell_i^{(t)} = \sum_{t=1}^T \mathbf{1}[h^{(t)}(x_i) = y_i] \leq T/2,$$

since f misclassifies (x_i, y_i) . Therefore, $\sum_{t=1}^T \ell^{(t)}u_E \leq T/2$. Furthermore, since $|E| \geq \epsilon n$, we have

$$RE(D \parallel D^{(1)}) = RE(u_E \parallel D^{(1)}) \leq \ln(1/\epsilon).$$

Plugging these facts into in Lemma 1, we get

$$(1/2 + \gamma)T \leq (1 + \gamma)T/2 + \frac{\ln(1/\epsilon)}{\gamma},$$

which implies $T \leq \frac{2 \ln(1/\epsilon)}{\gamma^2}$, a contradiction. \square

Note that in Theorem 1, it is not explicitly assumed to be in the realizable case. In other words, If we have a γ -weak learner in the agnostic setting, we can achieve the same guarantee. However, in the agnostic setting, we only have the much weaker and more realistic assumption of β -weak agnostic learner. The next theorem shows the error bound we get under this usual assumption in the agnostic setting.

Theorem 2. *Given a sample S and access to a β -weak agnostic learner, Algorithm 1 uses at most $O(\frac{\log(1/\epsilon)}{(1/2-\beta)^2})$ iterations and achieves an error rate $\frac{2err_S(H)}{1/2-\beta} + \epsilon$ on S , where $err_S(H)$ is the optimal error rate on S achievable using the hypothesis class H .*

Proof. The idea is to show that as long as the boosting algorithm always generates some ϵ' -smooth distributions, the β -weak agnostic learner is actually a γ -weak learner for some $\gamma > 0$, i.e., it achieves error rate $1/2 - \gamma$ for any ϵ' -smooth distributions. In each iteration t , the β -weak agnostic learner, given S with distribution $D^{(t)}$, returns a hypothesis $h^{(t)}$ such that

$$err_{D^{(t)}}(h^{(t)}) \leq err_{D^{(t)}}(H) + \beta \leq \frac{1}{\epsilon'} err_S(H) + \beta.$$

The second inequality utilizes the ϵ' -smoothness property of $D^{(t)}$. The reason is that if h is the optimal hypothesis on S , we have

$$\begin{aligned} err_{D^{(t)}}(H) &\leq err_{D^{(t)}}(h) \leq \frac{\#\text{mistakes on } S}{\epsilon'|S|} \\ &= \frac{1}{\epsilon'} err_S(h) = \frac{1}{\epsilon'} err_S(H). \end{aligned}$$

Let $\frac{1}{\epsilon'} err_S(H) + \beta = \frac{1}{2} - \gamma$, or equivalently $\gamma = (\frac{1}{2} - \beta) - \frac{1}{\epsilon'} err_S(H)$. Then, if $\epsilon' \geq \frac{2err_S(H)}{1/2-\beta}$, we have $\gamma \geq \frac{1}{2}(1/2 - \beta) > 0$. Therefore, we can use Theorem 1, and achieves error rate ϵ' on S by using $O(\frac{\log(1/\epsilon')}{(1/2-\beta)^2})$ iterations. Alternatively, it achieves error rate $\frac{2err_S(H)}{1/2-\beta} + \epsilon$ by using $O(\frac{\log(1/\epsilon)}{(1/2-\beta)^2})$ iterations. \square

Next, we show how to adapt this algorithm to the distributed setting.

3.2 Agnostic boosting: distributed version

The technique of adapting a boosting algorithm to the distributed setting is inspired by Balcan et al. (2012). They claim that any weight-based boosting algorithm can be turned into a distributed boosting algorithm with communication complexity that depends linearly on the number of iterations in the original boosting algorithm. However, their result is not directly applicable to our boosting algorithm due to the additional projection step. We will describe our distributed boosting algorithm by showing how to simulate the three steps in each iteration of Algorithm 1 in the distributed setting with $O(d)$ words of communication. Then, since there are at most $O(\log(1/\epsilon))$ iterations, the desired result follows.

In step 1, in order to obtain a 2β -weak hypothesis (we use 2β instead of β for convenience, which only affects the constant terms), the center calls the β -weak agnostic learner on a dataset sampled from $D^{(t)} = \frac{1}{k} \sum_{i=1}^k D_i^{(t)}$. The sampling procedure is as follows. Each entity first sends its sum of weights to the center. Then, the center samples $O(\frac{d}{\beta^2} \log \frac{1}{\beta})$ examples in total across the k entities proportional to their sum of weights. By the standard VC-theory, the error rate of any hypothesis on the the sample is within β to the true error rate with respect to the underlying distribution, with high probability. It is thus sufficient to find a hypothesis with error within β to the best hypothesis, which can be done thanks to the assumed β -weak learner.

Step 2 is relatively straightforward. The center broadcasts $h^{(t)}$ and each entity updates its own internal weights independently. Each entity then sends the summation of internal weights to the center for the calculation of the normalization factor. The communication in this step is $O(kd)$ for sending $h^{(t)}$ and some numbers. What is left is to show that the projection in step 3 can be done in a communication efficient way. As shown in (Herbster and Warmuth, 2001), the projection using relative entropy as the distance into \mathcal{P} , the set of all ϵ -smooth distributions, can be done by the following simple algorithm.

For a fixed index m , we first clip the largest m coordinates of p to $\frac{1}{\epsilon n}$, and then rescale the rest of the coordinates to sum up to $1 - \frac{m}{\epsilon n}$. We find the least index m such that the resulting distribution is in \mathcal{P} , i.e. all the coordinates are at most $\frac{1}{\epsilon n}$. A naive algorithm by first sorting the coordinates takes $O(n \log n)$ time, but it is communicationally inefficient.

Fortunately, Herbster and Warmuth (2001) also propose a more advanced algorithm by recursively finding the median. The idea is to use the median as the

Algorithm 2 Distributed Bregman projection

Input:

Each entity i :

a disjoint subset \mathcal{W}_i of $\mathcal{W} = \{w_1, \dots, w_n\}$

Center:

$n_0 = n; \quad C = 0; \quad C^w = 0$

while $n_0 \neq 0$ do

distributedly find the median θ of $(\mathcal{W}_1, \dots, \mathcal{W}_k)$

Each entity i :

$\mathcal{L}_i = \{w : w < \theta, w \in \mathcal{W}_i\}; \quad L_i^w = \sum_{w \in \mathcal{L}_i} w$

$\mathcal{M}_i = \{w : w = \theta, w \in \mathcal{W}_i\}; \quad M_i^w = \sum_{w \in \mathcal{M}_i} w$

$\mathcal{H}_i = \{w : w > \theta, w \in \mathcal{W}_i\}; \quad H_i^w = \sum_{w \in \mathcal{H}_i} w$

Center:

$L = \sum_i |\mathcal{L}_i|; \quad L^w = \sum_i L_i^w$

$M = \sum_i |\mathcal{M}_i|; \quad M^w = \sum_i M_i^w$

$H = \sum_i |\mathcal{H}_i|; \quad H^w = \sum_i H_i^w$

$m_0 = \frac{1 - (C+H)\frac{1}{\epsilon n}}{1 - (C^w + H^w)}$ and broadcasts it

if $\theta m_0 > \frac{1}{\epsilon n}$ then

$C = C + H + M; \quad C^w = C^w + H^w + M^w$

if $L = 0$ then $\theta = \max(w : w < \theta, w \in \mathcal{W})$

set $n_0 = L$

notify each entity i to set $\mathcal{W}_i = \mathcal{L}_i$

else

set $n_0 = H$

notify each entity i to set $\mathcal{W}_i = \mathcal{H}_i$

end while

Center: $m_0 = \frac{1 - C\frac{1}{\epsilon n}}{1 - C^w}$ and broadcasts it

Each entity i : set each coordinate as

$$w'_i = \begin{cases} \frac{1}{\epsilon n} & \text{if } w_i > \theta \\ w_i m_0 & \text{if } w_i \leq \theta \end{cases}$$

threshold, which corresponds to a potential index m , i.e., m is the number of coordinates larger than the median. We then use a binary search to find the least index m . The distributed version of the algorithm, shown in Algorithm 2, has the following communication complexity bound.

Theorem 3. *Algorithm 2 projects an n -dimensional distribution into the set of all ϵ -smooth distributions \mathcal{P} with $O(k \log^2(n))$ words of communication complexity.*

Proof. Since Algorithm 2 is a direct adaptation of the centralized projection algorithm (Herbster and War-

Algorithm 3 Distributedly finding the median

Input:

Each entity i :

a disjoint subset \mathcal{W}_i of $\mathcal{W} = \{w_1, \dots, w_n\}$

Each entity i :

Send the median m_i of \mathcal{W}_i to the center

while $|\mathcal{W}_i| > 1$ for some $i \in [k]$ **do**

Center:

Send $\bar{m} = \max_i m_i$ to entity A

Send $\underline{m} = \min_i m_i$ to entity B

Entity A :

Send $\bar{n} = |i : w_i \in \mathcal{W}_A, w_i \geq \bar{m}|$ to the center

Entity B :

Send $\underline{n} = |i : w_i \in \mathcal{W}_B, w_i < \bar{m}|$ to the center

Center: Send $r = \min\{\bar{n}, \underline{n}\}$ to entity A and B

Entity A: Remove the largest r elements in \mathcal{W}_A

Entity B: Remove the smallest r elements in \mathcal{W}_B

Entity A and B:

Send the new median to the center

end while

math, 2001), we omit the proof of its correctness. Because we use a binary search over possible thresholds, the algorithm runs at most $O(\log(n))$ iterations. Therefore, it suffices to show that the communication complexity of finding the median is at most $O(k \log n)$, which can be achieved by Algorithm 3.

Each entity first sends its own median to the center. The center identifies the maximum and minimum local medians, denoted as \bar{m} and \underline{m} , respectively. The global median must be between \bar{m} and \underline{m} , and removing the same number of elements larger than or equal to \bar{m} and less than \underline{m} will not change the median. Therefore, the center can notify the two corresponding entities and let them remove the same number of elements. At least one entity will reduce its size by half, so the algorithm stops after $O(k \log n)$ iterations. Note that except for the first round, we only need to communicate the updated medians of two entities, so the overall communication complexity is $O(k \log n)$ words.

In practice, it is often easier and more efficient to use a *quickselect*-based distributed algorithm to find the median. The idea is to randomly select and broadcast a weight at each iteration. This, in expectation, can remove half of the possible median candidates. This approach achieves the same communication complexity in expectation. \square

The complete distributed agnostic boosting algorithm is shown in Algorithm 4. We summarize our theoretical results in the next Theorem.

Theorem 4. *Given access to a β -weak agnostic learner, Algorithm 4 achieves error rate $\frac{2err_D(H)}{1/2-\beta} + \epsilon$ by using at most $O(\frac{\log(1/\epsilon)}{(1/2-\beta)^2})$ rounds, each involving $O((d/\beta^2)\log(1/\beta))$ examples and an additional $O(kd \log^2(\frac{d \log(1/\epsilon)}{(1/2-\beta)\epsilon}))$ words of communication per round.*

Proof. The boosting algorithm starts by drawing from D a sample S of size $n = \tilde{O}(\frac{\log(1/\epsilon)d}{(1/2-\beta)^2\epsilon^2})$ across the k entities without communicating them. If S is a centralized dataset, then by Theorem 2 we know that Algorithm 1 achieves error rate $\frac{2err_S(H)}{1/2-\beta} + \frac{\epsilon}{2}$ on S using $O(\frac{\log(1/\epsilon)}{(1/2-\beta)^2})$ iterations. We have shown that Algorithm 4 is a correct simulation of Algorithm 1 in the distributed setting, and thus we achieve the same error bound on S . The number of communication rounds is the same as the number of iterations of the boosting algorithm. And in each round, the communication includes $O(d/\beta^2 \log(1/\beta))$ examples for finding the β -weak hypothesis, $O(kd)$ words for broadcasting the hypothesis and some numbers, and $O(k \log^2(n))$ words for the distributed Bregman projection.

So far we only have the error bound of $\frac{2err_S(H)}{1/2-\beta} + \frac{\epsilon}{2}$ on S . To obtain the generalization error bound, note that with $n = \tilde{O}(\frac{\log(1/\epsilon)d}{(1/2-\beta)^2\epsilon^2})$ and by the standard VC-dimension argument, we have that with high probability $err_S(H) \leq err_D(H) + \frac{(1/2-\beta)\epsilon}{8}$, and the generalization error of our final hypothesis deviates from the empirical error by at most $\epsilon/4$, which completes the proof with the desired generalization error bound. \square

4 EXPERIMENTS

We compare the empirical performance of the proposed distributed boosting algorithm with two other algorithms on synthetic and real-world datasets. The first one is distributed AdaBoost (Balcan et al., 2012), which is similar to our algorithm but without the projection step. The second one is the distributed logistic regression algorithm available in the MPI implementation of the Liblinear package (Zhuang et al., 2015). We choose it as a comparison to a non-boosting approach. Note that Liblinear is a highly-optimized package while our implementation is not, so the run time comparison is not absolutely fair. However, we show that our approach, grounded in a rigorous framework, is comparable to this leading method in practice.

Algorithm 4 Distributed agnostic boosting algorithm**Initialization:****Center:** Access to a β -agnostic weak learner. Set $\gamma = \frac{1}{2}(\frac{1}{2} - \beta)$ **Each entity i :**Sample S_i drawn from D_i such that $S = \cup_i S_i$ with size $n = \tilde{O}(\frac{\log(1/\epsilon)d}{(1/2-\beta)^2\epsilon^2})$ Set weights $v_{i,x}^{(1)} = 1/|S_i|$ for each $(x, y) \in S_i$ **for $t = 1, 2, \dots, T$ do****Each entity i :** Send $w_i^{(t)} = \sum_{x \in S_i} v_{i,x}^{(t)}$ to the center**Center:** Let $W^{(t)} = \sum_i w_i^{(t)}$. Determine the number of examples $n_i^{(t)}$ to request from each entity i by sampling $O(\frac{d}{\beta^2} \log \frac{1}{\beta})$ times from the multinomial distribution $w_i^{(t)}/W^{(t)}$, and then send each number $n_i^{(t)}$ to entity i .**Each entity i :** sample $n_i^{(t)}$ times from S_i proportional to $v_{i,x}^{(t)}$ and send them to the center**Center:** run the β -agnostic weak learner on the union of the received $O(\frac{d}{\beta^2} \log \frac{1}{\beta})$ examples, and then broadcast the returned hypothesis $h^{(t)}$ **Each entity i :** update the weight of each example (x, y)

$$\hat{v}_{i,x}^{(t+1)} = \begin{cases} v_{i,x}^{(t)}(1 - \gamma) & \text{if } h^{(t)}(x) = y \\ v_{i,x}^{(t)} & \text{otherwise} \end{cases}$$

Distributedly normalize and then project the weights by Algorithm 2

end for**Output:** The hypothesis $\text{sign}\left(\frac{1}{T} \sum_{t=1}^T h^{(t)}\right)$

Table 1: Average (over 10 trials) error rate (%) and standard deviation on the synthetic dataset

Noise	Dist. AdaBoost	Dist. SmoothBoost	Liblinear -LR
0.1%	11.64 \pm 3.82	4.28 \pm 0.66	0.00
1%	25.97 \pm 1.56	13.38 \pm 4.66	0.00
10%	28.04 \pm 0.94	27.07 \pm 1.60	37.67

4.1 Experiment setup

All three algorithms are implemented in C using MPI, and all the experiments are run on Amazon EC2 with 16 m3.large machines. The data is uniformly partitioned across 16 machines. All the results are averaged over 10 independent trials. Logistic regression is a deterministic algorithm, so we do not show the standard deviation of the error rate. We however still run it for 10 times to get the average running time. Since each algorithm has different number of parameters, for fairness, we do not tune the parameters. For the two boosting algorithms, we use $T = 100$ decision stumps as our weak learners and set $\beta = 0.2$ and $\epsilon = 0.1$ in all experiments. For logistic regression, we use the default parameter $C = 1$.

4.2 Synthetic dataset

We use the synthetic dataset from (Long and Servedio, 2008). This dataset has an interesting theoretical property that although it is linearly separable, by randomly flipping a tiny fraction of labels, all convex potential boosting algorithms, including AdaBoost, fail to learn well. A random example is generated as follows. The label y is randomly chosen from $\{-1, +1\}$ with equal odds. The feature $x = \langle x_1, \dots, x_{21} \rangle$, where $x_i \in \{-1, +1\}$, is sampled from a mixture distribution: **1)** With probability 1/4, set all x_i to be equal to y . **2)** With probability 1/4, set $x_1 = x_2 = \dots = x_{11} = y$ and $x_{12} = x_{13} = \dots = x_{21} = -y$. **3)** With probability 1/2, randomly set 5 coordinates from the first 11 and 6 coordinates from the last 10 to be equal to y . Set the remaining coordinates to $-y$.

We generate 1,600,000 examples in total for training on 16 machines and test on a separate set of size 100,000. The results are shown in Table 1. One can see that our approach (Dist.SmoothBoost), is more resistant to noise than Dist.AdaBoost and significantly outperforms it for having upto 1% noise. In high noise setting (10%), Liblinear performs poorly, while our approach achieves the best error rate.

Table 2: Average (over 10 trials) error rate (%) and standard deviation on real-world datasets

Dataset	#examples	# features	Dist.AdaBoost	Dist.SmoothBoost	Liblinear-LR
Adult	48,842	123	15.71 \pm 0.16	15.07 \pm 2.32	15.36
Ijcnn1	141,691	22	5.90 \pm 0.10	4.33 \pm 0.18	7.57
Cod-RNA	488,565	8	6.12 \pm 0.09	6.51 \pm 0.11	11.79
Covtype	581,012	54	24.98 \pm 0.22	24.68 \pm 0.30	24.52
Yahoo	3,251,378	10	37.08 \pm 0.15	36.86 \pm 0.27	39.15

Table 3: Average run time (s) on real-world datasets

Dataset	Dist. AdaBoost	Dist. SmoothBoost	Liblinear-LR
Adult	5.02	15.54	0.06
Ijcnn1	0.76	9.19	0.10
Cod-RNA	1.08	10.11	0.12
Covtype	3.71	6.48	0.31
Yahoo	3.37	3.79	1.37

4.3 Real-world datasets

We run the experiments on 5 real-world datasets with sizes ranging from 50 thousands to over 3 millions: ADULT, IJCNN1, COD-RNA, and COVTYPE from the LibSVM data repository ¹; YAHOO from the Yahoo! WebScope dataset (Chu et al., 2009). The Yahoo dataset contains user click logs and is extremely imbalanced. We trim down this dataset so that the number of positive and negative examples are the same. The detailed information of the datasets are summarized in Table 2. Each dataset is randomly split into 4/5 for the training set and 1/5 for the testing set.

The average error rates and the total running times are summarized in Table 2 and Table 3, respectively. The bold entries indicates the best error rate. Our approach outperforms the other two on 3 datasets and performs competitively on the other 2 datasets.

In terms of running time, Liblinear is the fastest on all datasets. However, the communication of our algorithm only depends on the dimension d , so even for the largest dataset (YAHOO), it can still finish within 4 seconds. Therefore, our algorithm is suitable for many real-world situations where the number of examples is much larger than the dimension of the data.

Furthermore, it is possible to replace the centralized weak learner (decision stump here) with distributed weak learners, such as distributed logistic regression, to further reduce the running time.

¹<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets>.

5 CONCLUSION

Our main contribution has two parts. The first part is a novel analysis of a centralized boosting algorithm in the agnostic setting, which gives the first agnostic boosting algorithm to achieve an error bound of $O(\text{opt}(H)) + \epsilon$ by using only $O(\log 1/\epsilon)$ iterations. The second part is the adaptation of the algorithm to the distributed setting. The resulting distributed boosting algorithm enjoys strong performance guarantees, being simultaneously noise tolerant, communication efficient, and computationally efficient; furthermore, it is quite flexible in that it can be used with a variety of weak learners. This improves significantly over the prior work of (Balcan et al., 2012; Daumé et al., 2012) that were either communication efficient only in noise-free scenarios or computationally prohibitive. While enjoying nice theoretical guarantees, our algorithm also shows promising empirical results on large synthetic and real-world datasets.

Finally, we raise some related open questions. In this work we assumed a star topology, i.e., the center can communicate with all players directly. An interesting open question is to extend our results to general communication topologies. Another concrete open question is reducing the constant in our error bound while maintaining good communication complexity. Finally, our approach uses centralized weak learners for learning general concept classes, so the computation is mostly done in the center. Are there efficient distributed weak learners for some specific concept classes? That could provide a more computation balanced distributed learning procedure that enjoys strong communication complexity as well.

Acknowledgments

This work was supported in part by NSF grants CCF-1101283, CCF-1451177, CCF 1422910, TWC-1526254, IIS-1217559, IIS-1563816, ONR grant N00014-09-1-0751, and AFOSR grant FA9550-09-1-0538. We thank the Amazon’s AWS Education grant program for providing the Amazon Web Services.

References

- Maria-Florina Balcan, Avrim Blum, Shai Fine, and Yishay Mansour. Distributed learning, communication complexity and privacy. In *Proceedings of COLT*, 2012.
- Maria-Florina Balcan, Steven Ehrlich, and Yingyu Liang. Distributed k-means and k-median clustering on general communication topologies. In *Proceedings of NIPS*, 2013.
- Maria-Florina Balcan, Yingyu Liang, Vandana Kanchanapally, and David Woodruff. Improved distributed principal component analysis. In *Proceedings of NIPS*, pages 3113–3121, 2014.
- Aurlien Bellet, Yingyu Liang, Alireza Bagheri Garakani, Maria-Florina Balcan, and Fei Sha. Distributed frank-wolfe algorithm: A unified framework for communication-efficient sparse learning. In *Proceedings of SDM*, 2015.
- Shai Ben-David, Philip M Long, and Yishay Mansour. Agnostic boosting. In *Computational Learning Theory*, pages 507–516. Springer, 2001.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011.
- Wei Chu, Seung-Taek Park, Todd Beaupre, Nitin Motgi, Amit Phadke, Seinjuti Chakraborty, and Joe Zachariah. A case study of behavior-driven conjoint analysis on yahoo!: Front page today module. In *Proceedings of KDD*, pages 1097–1104. ACM, 2009.
- Hal Daumé, Jeff M. Phillips, Avishek Saha, and Suresh Venkatasubramanian. Efficient protocols for distributed classification and optimization. In *Proceedings of the 23rd International Conference on Algorithmic Learning Theory, ALT’12*, pages 154–168, 2012.
- Thomas G. Dietterich. An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Mach. Learn.*, 40(2):139–157, 2000.
- Vitaly Feldman. Distribution-specific agnostic boosting. In *ICS*, pages 241–250, 2010.
- Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS ’09*, pages 385–394, 2009.
- Yoav Freund and Rob 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.
- Dmitry Gavinsky. Optimally-smooth adaptive boosting and application to agnostic learning. *J. Mach. Learn. Res.*, 4:101–117, 2003.
- Mark Herbster and Manfred K. Warmuth. Tracking the best linear predictor. *J. Mach. Learn. Res.*, 1: 281–309, 2001. ISSN 1532-4435.
- Martin Jaggi, Virginia Smith, Martin Takác, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I Jordan. Communication-efficient distributed dual coordinate ascent. In *Proceedings of NIPS*, pages 3068–3076, 2014.
- MI Jordan and TM Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.
- Adam Tauman Kalai, Yishay Mansour, and Elad Verbin. On agnostic boosting and parity learning. In *Proceedings of STOC*, pages 629–638. ACM, 2008.
- Satyen Kale. Boosting and hard-core set constructions: a simplified approach. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(131), 2007.
- Varun Kanade and Adam Kalai. Potential-based agnostic boosting. In *Proceedings of NIPS*, pages 880–888, 2009.
- Michael J Kearns, Robert E Schapire, and Linda M Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141, 1994.
- Mu Li, Dave Andersen, Alex Smola, and Kai Yu. Parameter server for distributed machine learning. In *Proceedings of NIPS*, 2014.
- Phil Long and Rocco A Servedio. Random classification noise defeats all convex potential boosters. In *Proceedings of ICML*, pages 608–615. ACM, 2008.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT press, 2012.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning*. Cambridge University Press., 2014.
- Yuchen Zhang, John C. Duchi, and Martin Wainwright. Communication-efficient algorithms for statistical optimization. In *Proceedings of NIPS*, 2012.
- Yuchen Zhang, John Duchi, Michael Jordan, and Martin Wainwright. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. In *Proceedings of NIPS*, 2013.
- Yong Zhuang, Wei-Sheng Chin, Yu-Chin Juan, and Chih-Jen Lin. Distributed newton methods for regularized logistic regression. In *Proceedings of the PAKDD*, pages 690–703, 2015.