
A Boosting Framework on Grounds of Online Learning

Tofigh Naghibi, Beat Pfister

Computer Engineering and Networks Laboratory

ETH Zurich, Switzerland

naghibi@tik.ee.ethz.ch, pfister@tik.ee.ethz.ch

Abstract

By exploiting the duality between boosting and online learning, we present a boosting framework which proves to be extremely powerful thanks to employing the vast knowledge available in the online learning area. Using this framework, we develop various algorithms to address multiple practically and theoretically interesting questions including sparse boosting, smooth-distribution boosting, agnostic learning and, as a by-product, some generalization to double-projection online learning algorithms.

1 Introduction

A boosting algorithm can be seen as a meta-algorithm that maintains a distribution over the sample space. At each iteration a weak hypothesis is learned and the distribution is updated, accordingly. The output (strong hypothesis) is a convex combination of the weak hypotheses. Two dominant views to describe and design boosting algorithms are “weak to strong learner” (WTSL), which is the original viewpoint presented in [1, 2], and boosting by “coordinate-wise gradient descent in the functional space” (CWGD) appearing in later works [3, 4, 5]. A boosting algorithm adhering to the first view guarantees that it only requires a finite number of iterations (equivalently, finite number of weak hypotheses) to learn a $(1 - \epsilon)$ -accurate hypothesis. In contrast, an algorithm resulting from the CWGD viewpoint (usually called potential booster) may not necessarily be a boosting algorithm in the probability approximately correct (PAC) learning sense. However, while it is rather difficult to construct a boosting algorithm based on the first view, the algorithmic frameworks, e.g., AnyBoost [4], resulting from the second viewpoint have proven to be particularly prolific when it comes to developing new boosting algorithms. Under the CWGD view, the choice of the convex loss function to be minimized is (arguably) the cornerstone of designing a boosting algorithm. This, however, is a severe disadvantage in some applications.

In CWGD, the weights are not directly controllable (designable) and are only viewed as the values of the gradient of the loss function. In many applications, some characteristics of the desired distribution are known or given as problem requirements while, finding a loss function that generates such a distribution is likely to be difficult. For instance, what loss functions can generate sparse distributions?¹ What family of loss functions results in a smooth distribution?² We even can go further and imagine the scenarios in which a loss function needs to put more weights on a given subset of examples than others, either because that subset has more reliable labels or it is a problem requirement to have a more accurate hypothesis for that part of the sample space. Then, what

¹In the boosting terminology, sparsity usually refers to the greedy hypothesis-selection strategy of boosting methods in the functional space. However, sparsity in this paper refers to the sparsity of the distribution (weights) over the sample space.

²A smooth distribution is a distribution that does not put too much weight on any single sample or in other words, a distribution emulated by the booster does not dramatically diverge from the target distribution [6, 7].

loss function can generate such a customized distribution? Moreover, does it result in a provable boosting algorithm? In general, how can we characterize the accuracy of the final hypothesis?

Although, to be fair, the so-called loss function hunting approach has given rise to useful boosting algorithms such as LogitBoost, FilterBoost, GiniBoost and MadaBoost [5, 8, 9, 10] which (to some extent) answer some of the above questions, it is an inflexible and relatively unsuccessful approach to addressing the boosting problems with distribution constraints.

Another approach to designing a boosting algorithm is to directly follow the WTSL viewpoint [11, 6, 12]. The immediate advantages of such an approach are, first, the resultant algorithms are provable boosting algorithms, i.e., they output a hypothesis of arbitrary accuracy. Second, the booster has direct control over the weights, making it more suitable for boosting problems subject to some distribution constraints. However, since the WTSL view does not offer any algorithmic framework (as opposed to the CWGD view), it is rather difficult to come up with a distribution update mechanism resulting in a provable boosting algorithm. There are, however, a few useful, and albeit fairly limited, algorithmic frameworks such as TotalBoost [13] that can be used to derive other provable boosting algorithms. The TotalBoost algorithm can maximize the margin by iteratively solving a convex problem with the totally corrective constraint. A more general family of boosting algorithms was later proposed by Shalev-Shwartz et. al. [14], where it was shown that weak learnability and linear separability are equivalent, a result following from von Neumann’s minmax theorem. Using this theorem, they constructed a family of algorithms that maintain smooth distributions over the sample space, and consequently are noise tolerant. Their proposed algorithms find an $(1 - \epsilon)$ -accurate solution after performing at most $O(\log(N)/\epsilon^2)$ iterations, where N is the number of training examples.

1.1 Our Results

We present a family of boosting algorithms that can be derived from well-known online learning algorithms, including projected gradient descent [15] and its generalization, mirror descent (both active and lazy updates, see [16]) and composite objective mirror descent (COMID) [17]. We prove the PAC learnability of the algorithms derived from this framework and we show that this framework in fact generates maximum margin algorithms. That is, given a desired accuracy level ν , it outputs a hypothesis of margin $\gamma_{\min} - \nu$ with γ_{\min} being the minimum edge that the weak classifier guarantees to return.

The duality between (linear) online learning and boosting is by no means new. This duality was first pointed out in [2] and was later elaborated and formalized by using the von Neumann’s minmax theorem [18]. Following this line, we provide several proof techniques required to show the PAC learnability of the derived boosting algorithms. These techniques are fairly versatile and can be used to translate many other online learning methods into our boosting framework. To motivate our boosting framework, we derive two practically and theoretically interesting algorithms: (I) SparseBoost algorithm which by maintaining a sparse distribution over the sample space tries to reduce the space and the computation complexity. In fact this problem, i.e., applying batch boosting on the successive subsets of data when there is not sufficient memory to store an entire dataset, was first discussed by Breiman in [19], though no algorithm with theoretical guarantee was suggested. SparseBoost is the first provable batch booster that can (partially) address this problem. By analyzing this algorithm, we show that the tuning parameter of the regularization term ℓ_1 at each round t should not exceed $\frac{\gamma_t}{2}\eta_t$ to still have a boosting algorithm, where η_t is the coefficient of the t^{th} weak hypothesis and γ_t is its edge. (II) A smooth boosting algorithm that requires only $O(\log 1/\epsilon)$ number of rounds to learn a $(1 - \epsilon)$ -accurate hypothesis. This algorithm can also be seen as an agnostic boosting algorithm³ due to the fact that smooth distributions provide a theoretical guarantee for noise tolerance in various noisy learning settings, such as agnostic boosting [21, 22].

Furthermore, we provide an interesting theoretical result about MadaBoost [10]. We give a proof (to the best of our knowledge the only available unconditional proof) for the boosting property of (a variant of) MadaBoost and show that, unlike the common presumption, its convergence rate is of $O(1/\epsilon^2)$ rather than $O(1/\epsilon)$.

³Unlike the PAC model, the agnostic learning model allows an arbitrary target function (labeling function) that may not belong to the class studied, and hence, can be viewed as a noise tolerant learning model [20].

Finally, we show our proof technique can be employed to generalize some of the known online learning algorithms. Specifically, consider the Lazy update variant of the online Mirror Descent (LMD) algorithm (see for instance [16]). The standard proof to show that the LMD update scheme achieves vanishing regret bound is through showing its equivalence to the FTRL algorithm [16] in the case that they are both linearized, i.e., the cost function is linear. However, this indirect proof is fairly restrictive when it comes to generalizing the LMD-type algorithms. Here, we present a direct proof for it, which can be easily adopted to generalize the LMD-type algorithms.

2 Preliminaries

Let $\{(\mathbf{x}_i, a_i)\}, 1 \leq i \leq N$, be N training samples, where $\mathbf{x}_i \in \mathcal{X}$ and $a_i \in \{-1, +1\}$. Assume $h \in \mathcal{H}$ is a real-valued function mapping \mathcal{X} into $[-1, 1]$. Denote a distribution over the training data by $\mathbf{w} = [w_1, \dots, w_N]^\top$ and define a loss vector $\mathbf{d} = [-a_1 h(\mathbf{x}_1), \dots, -a_N h(\mathbf{x}_N)]^\top$. We define $\gamma = -\mathbf{w}^\top \mathbf{d}$ as the *edge* of the hypothesis h under the distribution w and it is assumed to be positive when h is returned by a weak learner. In this paper we do not consider the branching program based boosters and adhere to the typical boosting protocol (described in Section 1).

Since a central notion throughout this paper is that of Bregman divergences, we briefly revisit some of their properties. A Bregman divergence is defined with respect to a convex function \mathcal{R} as

$$B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) = \mathcal{R}(\mathbf{x}) - \mathcal{R}(\mathbf{y}) - \nabla \mathcal{R}(\mathbf{y})(\mathbf{x} - \mathbf{y})^\top \quad (1)$$

and can be interpreted as a distance measure between \mathbf{x} and \mathbf{y} . Due to the convexity of \mathcal{R} , a Bregman divergence is always non-negative, i.e., $B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq 0$. In this work we consider \mathcal{R} to be a β -strongly convex function⁴ with respect to a norm $\|\cdot\|$. With this choice of \mathcal{R} , the Bregman divergence $B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|^2$. As an example, if $\mathcal{R}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{x}$ (which is 1-strongly convex with respect to $\|\cdot\|_2$), then $B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$ is the Euclidean distance. Another example is the negative entropy function $\mathcal{R}(\mathbf{x}) = \sum_{i=1}^N x_i \log x_i$ (resulting in the KL-divergence) which is known to be 1-strongly convex over the probability simplex with respect to ℓ_1 norm.

The Bregman projection is another fundamental concept of our framework.

Definition 1 (Bregman Projection). *The Bregman projection of a vector \mathbf{y} onto a convex set S with respect to a Bregman divergence $B_{\mathcal{R}}$ is*

$$\Pi_S(\mathbf{y}) = \arg \min_{\mathbf{x} \in S} B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \quad (2)$$

Moreover, the following generalized Pythagorean theorem holds for Bregman projections.

Lemma 1 (Generalized Pythagorean) [23, Lemma 11.3]. *Given a point $\mathbf{y} \in \mathbb{R}^N$, a convex set S and $\hat{\mathbf{y}} = \Pi_S(\mathbf{y})$ as the Bregman projection of \mathbf{y} onto S , for all $\mathbf{x} \in S$ we have*

$$\text{Exact: } B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq B_{\mathcal{R}}(\mathbf{x}, \hat{\mathbf{y}}) + B_{\mathcal{R}}(\hat{\mathbf{y}}, \mathbf{y}) \quad (3)$$

$$\text{Relaxed: } B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq B_{\mathcal{R}}(\mathbf{x}, \hat{\mathbf{y}}) \quad (4)$$

The relaxed version follows from the fact that $B_{\mathcal{R}}(\hat{\mathbf{y}}, \mathbf{y}) \geq 0$ and thus can be ignored.

Lemma 2. *For any vectors $\mathbf{x}, \mathbf{y}, \mathbf{z}$, we have*

$$(\mathbf{x} - \mathbf{y})^\top (\nabla \mathcal{R}(\mathbf{z}) - \nabla \mathcal{R}(\mathbf{y})) = B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) - B_{\mathcal{R}}(\mathbf{x}, \mathbf{z}) + B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}) \quad (5)$$

The above lemma follows directly from the Bregman divergence definition in (1). Additionally, the following definitions from convex analysis are useful throughout the paper.

Definition 2 (Norm & dual norm). *Let $\|\cdot\|_A$ be a norm. Then its dual norm is defined as*

$$\|\mathbf{y}\|_{A^*} = \sup\{\mathbf{y}^\top \mathbf{x}, \|\mathbf{x}\|_A \leq 1\} \quad (6)$$

For instance, the dual norm of $\|\cdot\|_2 = \ell_2$ is $\|\cdot\|_{2^*} = \ell_2$ norm and the dual norm of ℓ_1 is ℓ_∞ norm. Further,

Lemma 3. *For any vectors \mathbf{x}, \mathbf{y} and any norm $\|\cdot\|_A$, the following inequality holds:*

$$\mathbf{x}^\top \mathbf{y} \leq \|\mathbf{x}\|_A \|\mathbf{y}\|_{A^*} \leq \frac{1}{2} \|\mathbf{x}\|_A^2 + \frac{1}{2} \|\mathbf{y}\|_{A^*}^2 \quad (7)$$

⁴That is, its second derivative (Hessian in higher dimensions) is bounded away from zero by at least β .

Throughout this paper, we use the shorthands $\|\cdot\|_A = \|\cdot\|$ and $\|\cdot\|_{A^*} = \|\cdot\|_*$ for the norm and its dual, respectively.

Finally, before continuing, we establish our notations. Vectors are lower case bold letters and their entries are non-bold letters with subscripts, such as x_i of \mathbf{x} , or non-bold letter with superscripts if the vector already has a subscript, such as x_t^i of \mathbf{x}_t . Moreover, an N -dimensional probability simplex is denoted by $\mathcal{S} = \{\mathbf{w} \mid \sum_{i=1}^N w_i = 1, w_i \geq 0\}$. The proofs of the theorems and the lemmas can be found in the Supplement.

3 Boosting Framework

Let $\mathcal{R}(\mathbf{x})$ be a 1-strongly convex function with respect to a norm $\|\cdot\|$ and denote its associated Bregman divergence $B_{\mathcal{R}}$. Moreover, let the dual norm of a loss vector \mathbf{d}_t be upper bounded, i.e., $\|\mathbf{d}_t\|_* \leq L$. It is easy to verify that for \mathbf{d}_t as defined in MABOOST, $L = 1$ when $\|\cdot\|_* = \ell_\infty$ and $L = N$ when $\|\cdot\|_* = \ell_2$. The following Mirror Ascent Boosting (MABOOST) algorithm is our boosting framework.

Algorithm 1: Mirror Ascent Boosting (MABOOST)

Input: $\mathcal{R}(\mathbf{x})$ 1-strongly convex function, $\mathbf{w}_1 = [\frac{1}{N}, \dots, \frac{1}{N}]^\top$ and $\mathbf{z}_1 = [\frac{1}{N}, \dots, \frac{1}{N}]^\top$

For $t = 1, \dots, T$ **do**

(a) Train classifier with \mathbf{w}_t and get h_t , let $\mathbf{d}_t = [-a_1 h_t(\mathbf{x}_1), \dots, -a_N h_t(\mathbf{x}_N)]$ and $\gamma_t = -\mathbf{w}_t^\top \mathbf{d}_t$.

(b) Set $\eta_t = \frac{\gamma_t}{L}$

(c) Update weights: $\nabla \mathcal{R}(\mathbf{z}_{t+1}) = \nabla \mathcal{R}(\mathbf{z}_t) + \eta_t \mathbf{d}_t$ (lazy update)
 $\nabla \mathcal{R}(\mathbf{z}_{t+1}) = \nabla \mathcal{R}(\mathbf{w}_t) + \eta_t \mathbf{d}_t$ (active update)

(d) Project onto \mathcal{S} : $\mathbf{w}_{t+1} = \underset{\mathbf{w} \in \mathcal{S}}{\operatorname{argmin}} B_{\mathcal{R}}(\mathbf{w}, \mathbf{z}_{t+1})$

End

Output: The final hypothesis $f(\mathbf{x}) = \operatorname{sign}\left(\sum_{t=1}^T \eta_t h_t(\mathbf{x})\right)$.

This algorithm is a variant of the mirror descent algorithm [16], modified to work as a boosting algorithm. The basic principle in this algorithm is quite clear. As in ADABOOST, the weight of a wrongly (correctly) classified sample increases (decreases). The weight vector is then projected onto the probability simplex in order to keep the weight sum equal to 1. The distinction between the active and lazy update versions and the fact that the algorithm may behave quite differently under different update strategies should be emphasized. In the lazy update version, the norm of the auxiliary variable \mathbf{z}_t is unbounded which makes the lazy update inappropriate in some situations. In the active update version, on the other hand, the algorithm always needs to access (compute) the previous projected weight \mathbf{w}_t to update the weight at round t and this may not be possible in some applications (such as boosting-by-filtering).

Due to the duality between online learning and boosting, it is not surprising that MABOOST (both the active and lazy versions) is a boosting algorithm. The proof of its boosting property, however, reveals some interesting properties which enables us to generalize the MABOOST framework. In the following, only the proof of the active update is given and the lazy update is left to Section 3.4.

Theorem 1. Suppose that MABOOST generates weak hypotheses h_1, \dots, h_T whose edges are $\gamma_1, \dots, \gamma_T$. Then the error ϵ of the combined hypothesis f on the training set is bounded and yields for the following \mathcal{R} functions:

$$\mathcal{R}(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_2^2 : \quad \epsilon \leq \frac{1}{\sum_{t=1}^T \frac{1}{2} \gamma_t^2 + 1} \quad (8)$$

$$\mathcal{R}(\mathbf{w}) = \sum_{i=1}^N w_i \log w_i : \quad \epsilon \leq e^{-\sum_{t=1}^T \frac{1}{2} \gamma_t^2} \quad (9)$$

In fact, the first bound (8) holds for any 1-strongly convex \mathcal{R} , though for some \mathcal{R} (e.g., negative entropy) a much tighter bound as in (9) can be achieved.

Proof: Assume $\mathbf{w}^* = [w_1^*, \dots, w_N^*]^\top$ is a distribution vector where $w_i^* = \frac{1}{N\epsilon}$ if $f(\mathbf{x}_i) \neq a_i$, and 0 otherwise. \mathbf{w}^* can be seen as a uniform distribution over the wrongly classified samples by the ensemble hypothesis f . Using this vector and following the approach in [16], we derive the upper bound of $\sum_{t=1}^T \eta_t (\mathbf{w}^{*\top} \mathbf{d}_t - \mathbf{w}_t^\top \mathbf{d}_t)$ where $\mathbf{d}_t = [d_t^1, \dots, d_t^N]$ is a loss vector as defined in Algorithm 1.

$$(\mathbf{w}^* - \mathbf{w}_t)^\top \eta_t \mathbf{d}_t = (\mathbf{w}^* - \mathbf{w}_t)^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{w}_t)) \quad (10a)$$

$$= B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_t) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1}) \quad (10b)$$

$$\leq B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_t) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1}) \quad (10c)$$

where the first equation follows Lemma 2 and inequality (10c) results from the relaxed version of Lemma 1. Note that Lemma 1 can be applied here because $\mathbf{w}^* \in \mathcal{S}$.

Further, the $B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1})$ term is bounded. By applying Lemma 3

$$B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{w}_t) = (\mathbf{z}_{t+1} - \mathbf{w}_t)^\top \eta_t \mathbf{d}_t \leq \frac{1}{2} \|\mathbf{z}_{t+1} - \mathbf{w}_t\|^2 + \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 \quad (11)$$

and since $B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{w}_t) \geq \frac{1}{2} \|\mathbf{z}_{t+1} - \mathbf{w}_t\|^2$ due to the 1-strongly convexity of \mathcal{R} , we have

$$B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1}) \leq \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 \quad (12)$$

Now, replacing (12) into (10c) and summing it up from $t = 1$ to T , yields

$$\sum_{t=1}^T \mathbf{w}^{*\top} \eta_t \mathbf{d}_t - \mathbf{w}_t^\top \eta_t \mathbf{d}_t \leq \sum_{t=1}^T \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_{T+1}) \quad (13)$$

Moreover, it is evident from the algorithm description that for mistakenly classified samples

$$-a_i f(\mathbf{x}_i) = -a_i \text{sign} \left(\sum_{t=1}^T \eta_t h_t(\mathbf{x}_i) \right) = \text{sign} \left(\sum_{t=1}^T \eta_t d_t^i \right) \geq 0 \quad \forall \mathbf{x}_i \in \{\mathbf{x} | f(\mathbf{x}_i) \neq a_i\} \quad (14)$$

Following (14), the first term in (13) will be $\mathbf{w}^{*\top} \sum_{t=1}^T \eta_t \mathbf{d}_t \geq 0$ and thus, can be ignored. Moreover, by the definition of γ , the second term is $\sum_{t=1}^T -\mathbf{w}_t^\top \eta_t \mathbf{d}_t = \sum_{t=1}^T \eta_t \gamma_t$. Putting all these together, ignoring the last term in (13) and replacing $\|\mathbf{d}_t\|_*^2$ with its upper bound L , yields

$$-B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) \leq L \sum_{t=1}^T \frac{1}{2} \eta_t^2 - \sum_{t=1}^T \eta_t \gamma_t \quad (15)$$

Replacing the left side with $-B_{\mathcal{R}} = -\|\mathbf{w}^* - \mathbf{w}_1\|^2 = \frac{\epsilon - 1}{N\epsilon}$ for the case of quadratic \mathcal{R} , and with $-B_{\mathcal{R}} = \log(\epsilon)$ when \mathcal{R} is a negative entropy function, taking the derivative w.r.t η_t and equating it to zero (which yields $\eta_t = \frac{\gamma_t}{L}$) we achieve the error bounds in (8) and (9). Note that in the case of \mathcal{R} being the negative entropy function, Algorithm 1 degenerates into ADABOOST with a different choice of η_t .

Before continuing our discussion, it is important to mention that the cornerstone concept of the proof is the choice of \mathbf{w}^* . For instance, a different choice of \mathbf{w}^* results in the following max-margin theorem.

Theorem 2. Setting $\eta_t = \frac{\gamma_t}{L\sqrt{t}}$, MABOOST outputs a hypothesis of margin at least $\gamma_{\min} - \nu$, where ν is a desired accuracy level and tends to zero in $O(\frac{\log T}{\sqrt{T}})$ rounds of boosting.

Observations: Two observations follow immediately from the proof of Theorem 1. First, the requirement of using Lemma 1 is $\mathbf{w}^* \in \mathcal{S}$, so in the case of projecting onto a smaller convex set $\mathcal{S}_k \subseteq \mathcal{S}$, as long as $\mathbf{w}^* \in \mathcal{S}_k$ holds, the proof is intact. Second, only the relaxed version of Lemma 1 is required in the proof (to obtain inequality (10c)). Hence, if there is an approximate projection operator $\hat{\Pi}_{\mathcal{S}}$ that satisfies the inequality $B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) \geq B_{\mathcal{R}}(\mathbf{w}^*, \hat{\Pi}_{\mathcal{S}}(\mathbf{z}_{t+1}))$, it can be substituted

for the exact projection operator Π_S and the active update version of the algorithm still works. A practical approximate operator of this type can be obtained by using the double-projection strategy as in Lemma 4.

Lemma 4. *Consider the convex sets \mathcal{K} and \mathcal{S} , where $\mathcal{S} \subseteq \mathcal{K}$. Then for any $\mathbf{x} \in \mathcal{S}$ and $\mathbf{y} \in \mathbb{R}^N$, $\hat{\Pi}_S(\mathbf{y}) = \Pi_S(\Pi_{\mathcal{K}}(\mathbf{y}))$ is an approximate projection operator that satisfies $B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq B_{\mathcal{R}}(\mathbf{x}, \hat{\Pi}_S(\mathbf{y}))$.*

These observations are employed to generalize Algorithm 1. However, we want to emphasize that the approximate Bregman projection is only valid for the active update version of MABOOST.

3.1 Smooth Boosting

Let $k > 0$ be a smoothness parameter. A distribution \mathbf{w} is smooth w.r.t a given distribution \mathbf{D} if $w_i \leq kD_i$ for all $1 \leq i \leq N$. Here, we consider the smoothness w.r.t to the uniform distribution, i.e., $D_i = \frac{1}{N}$. Then, given a desired smoothness parameter k , we require a boosting algorithm that only constructs distributions \mathbf{w} such that $w_i \leq \frac{k}{N}$, while guaranteeing to output a $(1 - \frac{1}{k})$ -accurate hypothesis. To this end, we only need to replace the probability simplex \mathcal{S} with $\mathcal{S}_k = \{\mathbf{w} | \sum_{i=1}^N w_i = 1, 0 \leq w_i \leq \frac{k}{N}\}$ in MABOOST to obtain a smooth distribution boosting algorithm, called smooth-MABOOST. That is, the update rule is: $\mathbf{w}_{t+1} = \underset{\mathbf{w} \in \mathcal{S}_k}{\operatorname{argmin}} B_{\mathcal{R}}(\mathbf{w}, \mathbf{z}_{t+1})$.

Note that the proof of Theorem 1 holds for smooth-MABOOST, as well. As long as $\epsilon \geq \frac{1}{k}$, the error distribution \mathbf{w}^* ($w_i^* = \frac{1}{N\epsilon}$ if $f(\mathbf{x}_i) \neq a_i$, and 0 otherwise) is in \mathcal{S}_k because $\frac{1}{N\epsilon} \leq \frac{k}{N}$. Thus, based on the first observation, the error bounds achieved in Theorem 1 hold for $\epsilon \geq \frac{1}{k}$. In particular, $\epsilon = \frac{1}{k}$ is reached after a finite number of iterations. This projection problem has already appeared in the literature. An entropic projection algorithm (\mathcal{R} is negative entropy), for instance, was proposed in [14]. Using negative entropy and their suggested projection algorithm results in a fast smooth boosting algorithm with the following convergence rate.

Theorem 3. *Given $\mathcal{R}(\mathbf{w}) = \sum_{i=1}^N w_i \log w_i$ and a desired ϵ , smooth-MABOOST finds a $(1 - \epsilon)$ -accurate hypothesis in $O(\log(\frac{1}{\epsilon})/\gamma^2)$ of iterations.*

3.2 Combining Datasets

Let's assume we have two sets of data. A primary dataset \mathcal{A} and a secondary dataset \mathcal{B} . The goal is to train a classifier that achieves $(1 - \epsilon)$ accuracy on \mathcal{A} while limiting the error on dataset \mathcal{B} to $\epsilon_{\mathcal{B}} \leq \frac{1}{k}$. This scenario has many potential applications including transfer learning [24], weighted combination of datasets based on their noise level and emphasizing on a particular region of a sample space as a problem requirement (e.g., a medical diagnostic test that should not make a wrong diagnosis when the sample is a pregnant woman). To address this problem, we only need to replace \mathcal{S} in MABOOST with $\mathcal{S}_c = \{\mathbf{w} | \sum_{i=1}^N w_i = 1, 0 \leq w_i \forall i \in \mathcal{A} \wedge 0 \leq w_i \leq \frac{k}{N} \forall i \in \mathcal{B}\}$ where $i \in \mathcal{A}$ shorthands the indices of samples in \mathcal{A} . By generating smooth distributions on \mathcal{B} , this algorithm limits the weight of the secondary dataset, which intuitively results in limiting its effect on the final hypothesis. The proof of its boosting property is quite similar to Theorem 1 and can be found in the Supplement.

3.3 Sparse Boosting

Let $\mathcal{R}(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_2^2$. Since in this case the projection onto the simplex is in fact an ℓ_1 -constrained optimization problem, it is plausible that some of the weights are zero (sparse distribution), which is already a useful observation. To promote the sparsity of the weight vector, we want to directly regularize the projection with the ℓ_1 norm, i.e., adding $\|\mathbf{w}\|_1$ to the objective function in the projection step. It is, however, not possible in MABOOST, since $\|\mathbf{w}\|_1$ is trivially constant on the simplex. Therefore, following the second observation, we split the projection step into two consecutive projections. The first projection is onto \mathcal{K} , an N -dimensional unit hypercube $\mathcal{K} = \{\mathbf{y} | 0 \leq y_i \leq 1\}$. This projection is regularized with the ℓ_1 norm and the solution is then projected onto a simplex. Note

that the second projection can only make the solution sparser (look at the projection onto simplex algorithm in [25]).

Algorithm 2: SparseBoost

Let \mathcal{K} be a hypercube and \mathcal{S} a probability simplex; Set $\mathbf{w}_1 = [\frac{1}{N}, \dots, \frac{1}{N}]^\top$;
At $t = 1, \dots, T$, train h_t with \mathbf{w}_t , set $\eta_t = \frac{\gamma_t}{N}$ and $0 \leq \alpha_t < \frac{\gamma_t}{2}$, and update

$$\begin{aligned}\mathbf{z}_{t+1} &= \mathbf{w}_t + \eta_t \mathbf{d}_t \\ \mathbf{y}_{t+1} &= \arg \min_{\mathbf{y} \in \mathcal{K}} \|\mathbf{y} - \mathbf{z}_{t+1}\|^2 + \alpha_t \eta_t \|\mathbf{y}\|_1 \\ \mathbf{w}_{t+1} &= \arg \min_{\mathbf{w} \in \mathcal{S}} \|\mathbf{w} - \mathbf{y}_{t+1}\|^2\end{aligned}$$

Output the final hypothesis $f(\mathbf{x}) = \text{sign}\left(\sum_{t=1}^T \eta_t h_t(\mathbf{x})\right)$.

α_t is the regularization factor at round t . Since $\alpha_t \eta_t$ controls the sparsity of the solution, it is natural to investigate the maximum value that α_t can take, provided that the boosting property still holds. This bound is implicit in the following theorem.

Theorem 4. *Suppose that SparseBoost generates weak hypotheses h_1, \dots, h_T whose edges are $\gamma_1, \dots, \gamma_T$. Then, as long as $\alpha_t \leq \frac{\gamma_t}{2}$, the error ϵ of the combined hypothesis f on the training set is bounded as follows:*

$$\epsilon \leq \frac{1}{\sum_{t=1}^T \frac{1}{2} \gamma_t (\gamma_t - 2\alpha_t) + 1} \quad (16)$$

See the Supplement for the proof. It is noteworthy that SparseBoost can be seen as a variant of the COMID algorithm [17] with the difference that SparseBoost uses a double-projection or as called in Lemma 4, approximate projection strategy.

3.4 Lazy Update Boosting

In this section, we present the proof for the lazy update version of MABoost (LAMABoost) in Theorem 1. The proof technique is novel and can be used to generalize several known online learning algorithms such as OMDA in [26] and Meta algorithm in [27]. Moreover, we show that MadaBoost [10] can be presented in the LAMABoost setting. This gives a simple proof for MadaBoost without making the assumption that the edge sequence is monotonically decreasing (as in [10]).

Proof: Assume $\mathbf{w}^* = [w_1^*, \dots, w_N^*]^\top$ is a distribution vector where $w_i^* = \frac{1}{N\epsilon}$ if $f(\mathbf{x}_i) \neq a_i$, and 0 otherwise. Then,

$$\begin{aligned}(\mathbf{w}^* - \mathbf{w}_t)^\top \eta_t \mathbf{d}_t &= (\mathbf{w}_{t+1} - \mathbf{w}_t)^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) \\ &\quad + (\mathbf{z}_{t+1} - \mathbf{w}_{t+1})^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) + (\mathbf{w}^* - \mathbf{z}_{t+1})^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) \\ &\leq \frac{1}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2 + \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{z}_{t+1}) - B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{z}_t) + B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{z}_t) \\ &\quad - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_t) - B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{z}_t) \\ &\leq \frac{1}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2 + \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 - B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{w}_t) \\ &\quad + B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{z}_{t+1}) - B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_t) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_t)\end{aligned} \quad (17)$$

where the first inequality follows applying Lemma 3 to the first term and Lemma 2 to the rest of the terms and the second inequality is the result of applying the exact version of Lemma 1 to $B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{z}_t)$. Moreover, since $B_{\mathcal{R}}(\mathbf{w}_{t+1}, \mathbf{w}_t) - \frac{1}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|^2 \geq 0$, they can be ignored in (17). Summing up the inequality (17) from $t = 1$ to T , yields

$$-B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_1) \leq L \sum_{t=1}^T \frac{1}{2} \eta_t^2 - \sum_{t=1}^T \eta_t \gamma_t \quad (18)$$

where we used the facts that $\mathbf{w}^{*\top} \sum_{t=1}^T \eta_t \mathbf{d}_t \geq 0$ and $\sum_{t=1}^T -\mathbf{w}_t^\top \eta_t \mathbf{d}_t = \sum_{t=1}^T \eta_t \gamma_t$. The above inequality is exactly the same as (15), and replacing $-B_{\mathcal{R}}$ with $\frac{\epsilon-1}{N\epsilon}$ or $\log(\epsilon)$ yields the same

error bounds in Theorem 1. Note that, since the exact version of Lemma 1 is required to obtain (17), this proof does not reveal whether LAMABOOST can be generalized to employ the double-projection strategy. In some particular cases, however, we may show that a double-projection variant of LAMABOOST is still a provable boosting algorithm.

In the following, we briefly show that MadaBoost can be seen as a double-projection LAMABOOST.

Algorithm 3: Variant of MadaBoost

Let $\mathcal{R}(\mathbf{w})$ be the negative entropy and \mathcal{K} a unit hypercube; Set $\mathbf{z}_1 = [1, \dots, 1]^\top$;

At $t = 1, \dots, T$, train h_t with \mathbf{w}_t , set $f_t(\mathbf{x}) = \text{sign}\left(\sum_{t'=1}^t \eta_{t'} h_{t'}(\mathbf{x})\right)$ and calculate

$$\epsilon_t = \frac{\sum_{i=1}^N \frac{1}{2} |f_t(\mathbf{x}_i) - a_i|}{N}, \text{ set } \eta_t = \epsilon_t \gamma_t \text{ and update}$$

$$\begin{aligned} \nabla \mathcal{R}(\mathbf{z}_{t+1}) &= \nabla \mathcal{R}(\mathbf{z}_t) + \eta_t \mathbf{d}_t & \rightarrow z_{t+1}^i &= z_t^i e^{\eta_t d_t^i} \\ \mathbf{y}_{t+1} &= \arg \min_{\mathbf{y} \in \mathcal{K}} B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}_{t+1}) & \rightarrow y_{t+1}^i &= \min(1, z_{t+1}^i) \\ \mathbf{w}_{t+1} &= \arg \min_{\mathbf{w} \in \mathcal{S}} B_{\mathcal{R}}(\mathbf{w}, \mathbf{y}_{t+1}) & \rightarrow w_{t+1}^i &= \frac{y_{t+1}^i}{\|\mathbf{y}_{t+1}\|_1} \end{aligned}$$

Output the final hypothesis $f(\mathbf{x}) = \text{sign}\left(\sum_{t=1}^T \eta_t h_t(\mathbf{x})\right)$.

Algorithm 3 is essentially MadaBoost, only with a different choice of η_t . It is well-known that the entropy projection onto the probability simplex results in the normalization and thus, the second projection of Algorithm 3. The entropy projection onto the unit hypercube, however, maybe less known and thus, its proof is given in the Supplement.

Theorem 5. *Algorithm 3 yields a $(1 - \epsilon)$ -accurate hypothesis after at most $T = O(\frac{1}{\epsilon^2 \gamma^2})$.*

This is an important result since it shows that MadaBoost seems, at least in theory, to be slower than what we hoped, namely $O(\frac{1}{\epsilon \gamma^2})$.

4 Conclusion and Discussion

In this work, we provided a boosting framework that can produce provable boosting algorithms. This framework is mainly suitable for designing boosting algorithms with distribution constraints. A sparse boosting algorithm that samples only a fraction of examples at each round was derived from this framework. However, since our proposed algorithm cannot control the exact number of zeros in the weight vector, a natural extension to this algorithm is to develop a boosting algorithm that receives the sparsity level as an input. However, this immediately raises the question: what is the maximum number of examples that can be removed at each round from the dataset, while still achieving a $(1 - \epsilon)$ -accurate hypothesis?

The boosting framework derived in this work is essentially the dual of the online mirror descent algorithm. This framework can be generalized in different ways. Here, we showed that replacing the Bregman projection step with the double-projection strategy, or as we call it approximate Bregman projection, still results in a boosting algorithm in the active version of MABOOST, though this may not hold for the lazy version. In some special cases (MadaBoost for instance), however, it can be shown that this double-projection strategy works for the lazy version as well. Our conjecture is that under some conditions on the first convex set, the lazy version can also be generalized to work with the approximate projection operator. Finally, we provided a new error bound for the MadaBoost algorithm that does not depend on any assumption. Unlike the common conjecture, the convergence rate of MadaBoost (at least with our choice of η) is of $O(1/\epsilon^2)$.

Acknowledgments

This work was partially supported by SNSF. We would like to thank Professor Rocco Servedio for an inspiring email conversation and our colleague Hui Liang for his helpful comments.

References

- [1] R. E. Schapire. The strength of weak learnability. *Journal of Machine Learning Research*, 1990.
- [2] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 1997.
- [3] L. Breiman. Prediction games and arcing algorithms. *Neural Computation*, 1999.
- [4] L. Mason, J. Baxter, P. Bartlett, and M. Frean. Boosting algorithms as gradient descent. In *NIPS*, 1999.
- [5] J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 1998.
- [6] R. A. Servedio. Smooth boosting and learning with malicious noise. *Journal of Machine Learning Research*, 2003.
- [7] D. Gavinsky. Optimally-smooth adaptive boosting and application to agnostic learning. *Journal of Machine Learning Research*, 2003.
- [8] J. K. Bradley and R. E. Schapire. Filterboost: Regression and classification on large datasets. In *NIPS*. 2008.
- [9] K. Hatano. Smooth boosting using an information-based criterion. In *Algorithmic Learning Theory*. 2006.
- [10] C. Domingo and O. Watanabe. Madaboost: A modification of AdaBoost. In *COLT*, 2000.
- [11] Y. Freund. Boosting a weak learning algorithm by majority. *Journal of Information and Computation*, 1995.
- [12] N. H. Bshouty, D. Gavinsky, and M. Long. On boosting with polynomially bounded distributions. *Journal of Machine Learning Research*, 2002.
- [13] M. K. Warmuth, J. Liao, and G. Rätsch. Totally corrective boosting algorithms that maximize the margin. In *ICML*, 2006.
- [14] S. Shalev-Shwartz and Y. Singer. On the equivalence of weak learnability and linear separability: new relaxations and efficient boosting algorithms. In *COLT*, 2008.
- [15] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, 2003.
- [16] E. Hazan. A survey: The convex optimization approach to regret minimization. Working draft, 2009.
- [17] J. C. Duchi, S. Shalev-shwartz, Y. Singer, and A. Tewari. Composite objective mirror descent. In *COLT*, 2010.
- [18] Y. Freund and R. E. Schapire. Game theory, on-line prediction and boosting. In *COLT*, 1996.
- [19] L. Breiman. Pasting bites together for prediction in large data sets and on-line. Technical report, Dept. Statistics, Univ. California, Berkeley, 1997.
- [20] M. J. Kearns, R. E. Schapire, and L. M. Sellie. Toward efficient agnostic learning. In *COLT*, 1992.
- [21] A. Kalai and V. Kanade. Potential-based agnostic boosting. In *NIPS*. 2009.
- [22] S. Ben-David, P. Long, and Y. Mansour. Agnostic boosting. In *COLT*. 2001.
- [23] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [24] W. Dai, Q. Yang, G. Xue, and Y. Yong. Boosting for transfer learning. In *ICML*, 2007.
- [25] W. Wang and M. A. Carreira-Perpiñán. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. arXiv:1309.1541, 2013.
- [26] A. Rakhlin and K. Sridharan. Online learning with predictable sequences. In *COLT*, 2013.
- [27] C. Chiang, T. Yang, C. Lee, M. Mahdavi, C. Lu, R. Jin, and S. Zhu. Online optimization with gradual variations. In *COLT*, 2012.
- [28] Y. Freund and R. E. Schapire. Experiments with a New Boosting Algorithm. In *ICML*, 1996.
- [29] G. Rätsch and M. Warmuth. Efficient margin maximization with boosting. *Journal of Machine Learning Research*, 2005.
- [30] Shai Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Machine Learning*, 2012.

Supplement

Before proceeding with the proofs, some definitions and facts need to be reminded.

Definition: Margin

Given a final hypothesis $f(\mathbf{x}) = \sum_{t=1}^T \eta_t h_t(\mathbf{x})$, the margin of a sample (\mathbf{x}_j, a_j) is defined as $m(\mathbf{x}_j) = a_j f(\mathbf{x}_j) / \sum_{t=1}^T \eta_t$. Moreover, the margin of a set of examples denoted by $m_{\mathcal{D}}$ is the minimum of margins over the examples, i.e., $m_{\mathcal{D}} = \min_{\mathbf{x}} m(\mathbf{x}_j)$.

Fact: Duality between max-margin and min-edge

The minimum edge γ_{\min} that can be achieved over all possible distributions of the training set is equal to the maximum margin ($m^* = \max_{\eta} m_{\mathcal{D}}$) of any linear combination of hypotheses from the hypotheses space.

This fact is discussed in details in [28] and [29]. It is the direct result of von Neumann's minmax theorem and simply means that the maximum achievable margin is γ_{\min} .

Proof of Theorem 2

The proof for the maximum margin property of MABOOST, is almost the same as the proof of Theorem 1.

Let's assume the i^{th} sample has the worst margin, i.e., $m_{\mathcal{D}} = m(\mathbf{x}_i)$. Let all entries of the error vector \mathbf{w}^* to be zero except its i^{th} entry which is set to be 1. Following the same approach as in Theorem 1, (see equation (13)), we get

$$\sum_{t=1}^T \mathbf{w}^{*\top} \eta_t \mathbf{d}_t - \mathbf{w}_t^\top \eta_t \mathbf{d}_t \leq \sum_{t=1}^T \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_{T+1}) \quad (19)$$

With our choice of \mathbf{w}^* it is easy to verify that the first term on the left side of the inequality is $m_{\mathcal{D}} \sum_{t=1}^T \eta_t = -\sum_{t=1}^T \mathbf{w}^{*\top} \eta_t \mathbf{d}_t$. By setting $C = B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1)$, ignoring the last term $B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_{T+1})$, replacing $\|\mathbf{d}_t\|_*^2$ with its upper bound L and using the identity $\sum_{t=1}^T \mathbf{w}_t^\top \eta_t \mathbf{d}_t = -\sum_{t=1}^T \eta_t \gamma_t$ the above inequality is simplified to

$$-m_{\mathcal{D}} \sum_{t=1}^T \eta_t \leq L \sum_{t=1}^T \frac{1}{2} \eta_t^2 - \sum_{t=1}^T \eta_t \gamma_t + C \quad (20)$$

Replacing η_t with the value suggested in Theorem 2, i.e., $\eta_t = \frac{\gamma_t}{L\sqrt{t}}$ and dividing both sides by $\sum_{t=1}^T \eta_t$, gives

$$\frac{\sum_{t=1}^T (\frac{1}{\sqrt{t}} - \frac{1}{t}) \gamma_t^2}{\sum_{t=1}^T \frac{1}{\sqrt{t}} \gamma_t} - \frac{LC}{\sum_{t=1}^T \frac{1}{\sqrt{t}} \gamma_t} \leq m_{\mathcal{D}} \quad (21)$$

The first term is minimized when $\gamma_t = \gamma_{\min}$. Similarly to the first term, the second term is maximized when γ_t is replaced by its minimum value. This gives the following lower bound for $m_{\mathcal{D}}$:

$$\gamma_{\min} \frac{\sum_{t=1}^T \frac{1}{\sqrt{t}} - \frac{1}{t}}{\sum_{t=1}^T \frac{1}{\sqrt{t}}} - \frac{LC}{\gamma_{\min} \sum_{t=1}^T \frac{1}{\sqrt{t}}} \leq m_{\mathcal{D}} \quad (22)$$

Considering the facts that $\int_1^{T+1} \frac{dx}{\sqrt{x}} \leq \sum_{t=1}^T \frac{1}{\sqrt{t}}$ and $1 + \int_1^T \frac{dx}{x} \geq \sum_{t=1}^T \frac{1}{t}$, we get

$$\gamma_{\min} - \frac{1 + \log T}{2\sqrt{T+1} - 2} \gamma_{\min} - \frac{LC}{\gamma_{\min}(\sqrt{T+1} - 1)} \leq m_{\mathcal{D}} \quad (23)$$

Now by taking $\nu = \frac{1 + \log T}{2\sqrt{T+1} - 2} \gamma_{\min} + \frac{LC}{\gamma_{\min}(\sqrt{T+1} - 1)}$, we have $\gamma_{\min} - \nu \leq m_{\mathcal{D}}$. It is clear from (23) that ν approaches zero as T tends to infinity with a convergence rate proportional to $\frac{\log T}{\sqrt{T}}$. It is noteworthy that this convergence rate is slightly worse than of TotalBoost which is $O(\frac{1}{\sqrt{T}})$.

Proof of Lemma 4

Remember that $\hat{\Pi}_{\mathcal{S}}(\mathbf{y}) = \Pi_{\mathcal{S}}(\Pi_{\mathcal{K}}(\mathbf{y}))$. Our goal is to show that $B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq B_{\mathcal{R}}(\mathbf{x}, \hat{\Pi}_{\mathcal{S}}(\mathbf{y}))$.

To this end, we only need to repeatedly apply Lemma 1, as follows

$$B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq B_{\mathcal{R}}(\mathbf{x}, \Pi_{\mathcal{K}}(\mathbf{y})) \quad (24)$$

$$B_{\mathcal{R}}(\mathbf{x}, \Pi_{\mathcal{K}}(\mathbf{y})) \geq B_{\mathcal{R}}(\mathbf{x}, \hat{\Pi}_{\mathcal{S}}(\mathbf{y})) \quad (25)$$

which completes the proof.

Proof of combining datasets boosting algorithm

We have to show that when the convex set is defined as

$$\mathcal{S}_c = \{\mathbf{w} \mid \sum_{i=1}^N w_i = 1, 0 \leq w_i \forall i \in \mathcal{A} \wedge 0 \leq w_i \leq \frac{k}{N} \forall i \in \mathcal{B}\} \quad (26)$$

the error of the final hypothesis on \mathcal{A} , i.e., $\epsilon_{\mathcal{A}}$, converges to zero while the error on \mathcal{B} is guaranteed to be $\epsilon_{\mathcal{B}} \leq \frac{1}{k}$.

First, we show the convergence of $\epsilon_{\mathcal{A}}$ to zero. This is easily obtained by setting \mathbf{w}^* to be an error vector with zero weights over the training samples from \mathcal{B} and $\frac{1}{\epsilon_{\mathcal{A}} N_{\mathcal{A}}}$ weights over the training set \mathcal{A} . One can verify that $\mathbf{w}^* \in \mathcal{S}_c$, thus the proof of Theorem 1 holds and subsequently, the error bounds in (8) stating that $\epsilon_{\mathcal{A}} \rightarrow 0$ as the number of iterations increases.

To show the second part of the theorem that is $\epsilon_{\mathcal{B}} \leq \frac{1}{k}$, vector \mathbf{w}^* is selected to be an error vector with zero weights over the training samples from \mathcal{A} and $\frac{1}{\epsilon_{\mathcal{B}} N_{\mathcal{B}}}$ weights over the training set \mathcal{B} . Note that, as long as $\epsilon_{\mathcal{B}}$ is greater than $\frac{1}{k}$, this $\mathbf{w}^* \in \mathcal{S}_c$. Thus, for all $\frac{1}{k} \leq \epsilon_{\mathcal{B}}$ the proof of Theorem 1 holds and as the bounds in (8) show, the error decreases as the number of iterations increases. In particular in a finite number of rounds, the classification error on \mathcal{B} reduces to $\frac{1}{k}$ which completes the proof.

Proof of Theorem 4

The proof is almost identical to the proof given in [17], with a slight change to take the double-projection strategy into account. Let \mathbf{w}^* to be the same error vector as defined in Theorem 1. We start this proof by again bounding the

$$(\mathbf{w}^* - \mathbf{w}_t)^\top \eta_t \mathbf{d}_t = B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_t) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1}) \quad (27)$$

Remember that \mathbf{y}_{t+1} is the solution of an optimization problem (Bregman projection). By writing the optimality condition at \mathbf{y}_{t+1} we have

$$(\mathbf{w}^* - \mathbf{y}_{t+1})^\top (\nabla \mathcal{R}(\mathbf{y}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_{t+1})) + \alpha_t \eta_t \dot{r}(\mathbf{y}_{t+1})^\top (\mathbf{w}^* - \mathbf{y}_{t+1}) \geq 0 \quad (28)$$

where $\dot{r}(\mathbf{y})$ is a sub-gradient vector of the ℓ_1 norm function $r(\mathbf{y}) = \sum_{i=1}^N y_i$. By applying the three point identity in Lemma 2, we get

$$B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) \geq B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{y}_{t+1}) + B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_{t+1}) - \alpha_t \eta_t \dot{r}(\mathbf{y}_{t+1})^\top (\mathbf{w}^* - \mathbf{y}_{t+1}) \quad (29)$$

To bound the $\dot{r}(\mathbf{y}_{t+1})^\top (\mathbf{w}^* - \mathbf{y}_{t+1})$ in the above expression, we use the fact that $r(\mathbf{y})$ is a convex function and thus satisfies

$$\alpha_t \eta_t \dot{r}(\mathbf{y}_{t+1})^\top (\mathbf{w}^* - \mathbf{y}_{t+1}) \leq \alpha_t \eta_t (r(\mathbf{w}^*) - r(\mathbf{y}_{t+1})) \quad (30)$$

Hence,

$$-B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) \leq -B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{y}_{t+1}) - B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_{t+1}) + \alpha_t \eta_t (r(\mathbf{w}^*) - r(\mathbf{y}_{t+1})) \quad (31)$$

Be applying Lemma 1 (generalized Pythagorean theorem) to $B_{\mathcal{R}}(\mathbf{w}_t, \mathbf{z}_{t+1})$ in (27), and replacing $B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1})$ with its upper bound from (31), we get the following bound:

$$\sum_{t=1}^T \mathbf{w}^{*\top} \eta_t \mathbf{d}_t \leq \sum_{t=1}^T \mathbf{w}_t^\top \eta_t \mathbf{d}_t + \sum_{t=1}^T \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) + \sum_{t=1}^T \alpha_t \eta_t (1 - r(\mathbf{y}_{t+1})) \quad (32)$$

Now, replacing $r(\mathbf{y}_{t+1})$ with its lower bound, i.e, 0 and using the fact that $\sum_{t=1}^T \mathbf{w}^{*\top} \eta_t \mathbf{d}_t \geq 0$ (as shown in (14)) and $\sum_{t=1}^T \mathbf{w}_t^\top \eta_t \mathbf{d}_t = -\sum_{t=1}^T \eta_t \gamma_t$, yields

$$0 \leq -\sum_{t=1}^T \eta_t \gamma_t + \sum_{t=1}^T \frac{1}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) + \sum_{t=1}^T \alpha_t \eta_t \quad (33)$$

Since in SparseBoost $\mathcal{R}(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_2^2$, the Bregman divergence $B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{w}_1) = \frac{1-\epsilon}{N\epsilon}$ and $\|\mathbf{d}_t\|_*^2 \leq N$. Replacing them in (33) and setting $\eta_t = \frac{\gamma_t}{N}$, gives the bound in Theorem 4.

Proof of Entropy Projection onto Hypercube (Second Update Step in MadaBoost)

Lemma 5. Let $\mathcal{R}(\mathbf{w}) = \sum_{i=1}^N w_i \log w_i - w_i$. Then the Bregman projection of a positive vector $\mathbf{z} \in \mathbb{R}_+^N$ onto the unit hypercube $\mathcal{K} = [0, 1]^N$ is $y_i = \min(1, z_i), i = 1, \dots, N$.

To show the correctness of the above lemma, i.e., that the solution of the Bregman projection

$$\mathbf{y} = \arg \min_{\mathbf{y} \in \mathcal{K}} B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}) \quad (34)$$

is $y_i = \min(1, z_i)$, we only need to show that \mathbf{y} satisfies the optimality condition

$$(\mathbf{v} - \mathbf{y})^\top \nabla B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}) \geq 0 \quad \forall \mathbf{v} \in \mathcal{K} \quad (35)$$

Given $\mathcal{R}(\mathbf{w}) = \sum_{i=1}^N w_i \log w_i - w_i$, the gradient of $B_{\mathcal{R}}$ is

$$\nabla B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}) = \sum_{i=1}^T \log \frac{y_i}{z_i} \quad (36)$$

Hence,

$$(\mathbf{v} - \mathbf{y})^\top \nabla B_{\mathcal{R}}(\mathbf{y}, \mathbf{z}) = \sum_{i \in \{i: z_i \geq 1\}} (v_i - y_i) \log \frac{y_i}{z_i} + \sum_{i \in \{i: z_i < 1\}} (v_i - y_i) \log \frac{y_i}{z_i} \quad (37)$$

For $z_i \geq 1$, y_i is equal to 1. That is, $\log \frac{y_i}{z_i} = \log \frac{1}{z_i} < 0$. On the other hand, since $v_i \leq 1$, $(v_i - y_i) = (v_i - 1) \leq 0$. Thus, the first sum in (37) is always non-negative. The second sum is always zero since $\log \frac{y_i}{z_i} = \log 1 = 0$. That is, the optimality condition (37) is non-negative for all \mathbf{v} which completes the proof.

Proof of Theorem 5

Its proof is essentially the same as the proof of the lazy version of MABOOST with a few differences. Before proceeding further, some definitions and facts should be re-emphasized.

First of all, since $\mathcal{R}(\mathbf{w}) = \sum_{i=1}^N w_i \log w_i - w_i$ is $\frac{1}{N}$ -strongly convex (see [30, p. 136]) with respect to ℓ_1 norm (and not 1-strongly as in Theorem 1), the following inequality holds for the Bregman divergence:

$$B_{\mathcal{R}}(\mathbf{x}, \mathbf{y}) \geq \frac{1}{2N} \|\mathbf{x} - \mathbf{y}\|_1^2 \quad (38)$$

Moreover, the following lemma which bounds $\|\mathbf{y}_t\|_1$ is essential for our proof.

Lemma 6. For all t , $\|\mathbf{y}_t\|_1 \geq N\epsilon_t$ where ϵ_t is the error of the ensemble hypothesis $H_t(\mathbf{x}) = \sum_{l=1}^t \eta_l h_l(\mathbf{x})$ at round t .

This lemma holds due to the fact that

$$y_t^i = \min(1, z_t^i) = \min(1, e^{\sum_{l=1}^t \eta_l d_l^i}) = \min(1, e^{-a_i H_t(\mathbf{x}_i)}) \quad (39)$$

where $H_t(\mathbf{x}) = \sum_{l=1}^t \eta_l h_l(\mathbf{x})$ is the output of the algorithm at round t . If $H_t(\mathbf{x}_i)$ makes a mistake on classifying \mathbf{x}_i , $-a_i H_t(\mathbf{x}_i)$ will be greater than zero and thus, $y_t^i = 1$. For the samples that are classified correctly, $-a_i H_t(\mathbf{x}_i) \leq 0$ and thus, $0 \leq y_t^i \leq 1$. That is, $N\epsilon_t =$ number of wrongly classified samples at round $t \leq \sum_{i=1}^N y_t^i = \|\mathbf{y}_t\|_1$.

We are now ready to proceed with the proof of Theorem 5. Let $\mathbf{w}^* = [w_1^*, \dots, w_N^*]^\top$ to be a vector where $w_i^* = 1$ if $f(\mathbf{x}_i) \neq a_i$, and 0 otherwise. Similar to the proof of the lazy update, we are going to bound the $\sum_{t=1}^T (\mathbf{w}^* - \mathbf{y}_t)^\top \eta_t \mathbf{d}_t$.

$$\begin{aligned} (\mathbf{w}^* - \mathbf{y}_t)^\top \eta_t \mathbf{d}_t &= (\mathbf{y}_{t+1} - \mathbf{y}_t)^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) \\ &\quad + (\mathbf{z}_{t+1} - \mathbf{y}_{t+1})^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) + (\mathbf{w}^* - \mathbf{z}_{t+1})^\top (\nabla \mathcal{R}(\mathbf{z}_{t+1}) - \nabla \mathcal{R}(\mathbf{z}_t)) \\ &\leq \frac{1}{2N} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|^2 + \frac{N}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 + B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_{t+1}) - B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_t) + B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{z}_t) \\ &\quad - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_t) - B_{\mathcal{R}}(\mathbf{z}_{t+1}, \mathbf{z}_t) \\ &\leq \frac{1}{2N} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|^2 + \frac{N}{2} \eta_t^2 \|\mathbf{d}_t\|_*^2 - B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{y}_t) \\ &\quad + B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_{t+1}) - B_{\mathcal{R}}(\mathbf{y}_t, \mathbf{z}_t) - B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_{t+1}) + B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_t) \end{aligned} \quad (40)$$

where the first inequality follows from applying Lemma 3 to the first term and Lemma 2 to the rest of the terms and the second inequality is the result of applying the exact version of Lemma 1 to $B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{z}_t)$. Moreover, since according to inequality (38) $B_{\mathcal{R}}(\mathbf{y}_{t+1}, \mathbf{y}_t) - \frac{1}{2N} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|^2 \geq 0$ and hence these terms can be ignored in (40). Summing up the inequality (40) from $t = 1$ to T , yields:

$$-B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_1) \leq \sum_{t=1}^T \frac{N}{2} \eta_t^2 - \sum_{t=1}^T \eta_t \gamma_t \|\mathbf{y}_t\|_1 \quad (41)$$

It is important to remark that $\|\mathbf{y}_t\|_1$ appearing in the last term is due to the fact that $\mathbf{w}_t = \frac{\mathbf{y}_t}{\|\mathbf{y}_t\|_1}$ and thus, $\mathbf{y}_t^\top \eta_t \mathbf{d}_t = \mathbf{w}_t^\top \eta_t \mathbf{d}_t \|\mathbf{y}_t\|_1 = \eta_t \gamma_t \|\mathbf{y}_t\|_1$.

Now, by replacing $\eta_t = \epsilon_t \gamma_t$ in the above equation and noting that $B_{\mathcal{R}}(\mathbf{w}^*, \mathbf{z}_1) = N - N\epsilon$, we get:

$$-N(1 - \epsilon) \leq \sum_{t=1}^T \frac{N}{2} \epsilon_t^2 \gamma_t^2 - \sum_{t=1}^T \epsilon_t \gamma_t^2 \|\mathbf{y}_t\|_1 \quad (42)$$

From Lemma 6, it is evident that $\|\mathbf{y}_t\|_1 \geq N\epsilon_t$. Moreover, since $\epsilon \leq \epsilon_t$, it can be replaced by ϵ , as well (though very pessimistic). As usual, γ_t is also replaced with the min edge, denoted by γ . Applying these lower bounds in (42), yields

$$\epsilon^2 \leq \frac{2(1 - \epsilon)\gamma^2}{T} \leq \frac{\gamma^2}{T} \quad (43)$$

which indicates that the proposed version of MadaBoost needs at most $O(\frac{1}{\epsilon^2 \gamma^2})$ iterations to converge.