Journal of Machine Learning Research 4 (2003) 633-648 Submitted 7/02; Published 9/03

Smooth Boosting and Learning with Malicious Noise

Rocco A. Servedio ROCCO@CS.COLUMBIA.EDU
Department of Computer Science
Columbia University
New York, NY 10027, USA

Editor: Manfred Warmuth

Abstract

We describe a new boosting algorithm which generates only smooth distributions which do not assign too much weight to any single example. We show that this new boosting algorithm can be used to construct efficient PAC learning algorithms which tolerate relatively high rates of malicious noise. In particular, we use the new smooth boosting algorithm to construct malicious noise tolerant versions of the PAC-model *p*-norm linear threshold learning algorithms described by Servedio (2002). The bounds on sample complexity and malicious noise tolerance of these new PAC algorithms closely correspond to known bounds for the online *p*-norm algorithms of Grove, Littlestone and Schuurmans (1997) and Gentile and Littlestone (1999). As special cases of our new algorithms we obtain linear threshold learning algorithms which match the sample complexity and malicious noise tolerance of the online Perceptron and Winnow algorithms. Our analysis reveals an interesting connection between boosting and noise tolerance in the PAC setting.

Keywords: Boosting, Learning with Noise, Linear Threshold Functions

1. Introduction

Any realistic model of learning from examples must address the issue of noisy data. In the Probably Approximately Correct learning framework, Valiant (1985) introduced the notion of PAC learning in the presence of *malicious noise*. This is a worst-case model of errors in which some fraction of the labeled examples given to a learning algorithm may be corrupted by an adversary who can modify both example points and labels in an arbitrary fashion (a detailed description of the model is given in Section 3). The frequency of such corrupted examples is known as the *malicious noise rate*.

Learning in the presence of malicious noise is in general quite difficult. Kearns and Li (1993) have shown that for many classes of Boolean functions (concept classes), it is impossible to learn to accuracy ε if the malicious noise rate exceeds $\frac{\varepsilon}{1+\varepsilon}$. In fact, for many interesting concept classes such as the class of linear threshold functions, the best efficient algorithms known can only tolerate malicious noise rates significantly lower than this general upper bound. Despite these difficulties, the importance of being able to cope with noisy data has led many researchers to study PAC learning in the presence of malicious noise (see Aslam and Decatur, 1998, Auer, 1997, Auer and Cesa-Bianchi, 1998, Cesa-Bianchi et al., 1999, Decatur, 1993, Mansour and Parnas, 1998).

In this paper we give a new *smooth boosting* algorithm which can be used to transform a malicious noise tolerant weak learning algorithm into a PAC algorithm which learns successfully in the presence of malicious noise. We use this smooth boosting algorithm to construct a family of PAC algorithms for learning linear threshold functions in the presence of malicious noise. These

new algorithms match the sample complexity and noise tolerance of the online p-norm algorithms of Grove, Littlestone, and Schuurmans (1997) and Gentile and Littlestone (1999), which include as special cases the well-known Perceptron and Winnow algorithms.

1.1 Smooth Boosting and Learning with Malicious Noise

Our basic approach is quite simple, as illustrated by the following example. Consider a learning scenario in which we have a weak learning algorithm L which takes as input a finite sample S of m labeled examples. Algorithm L is known to have some tolerance to malicious noise; specifically, L is guaranteed to generate a hypothesis with nonnegligible advantage provided that the frequency of noisy examples in its sample is at most 10%. We would like to learn to high accuracy in the presence of malicious noise at a rate of 1%.

The obvious approach in this setting is to use a boosting algorithm, which is an algorithm that can generate a high accuracy hypothesis given access to a weak learner; see the paper of Schapire (1999) for an overview of boosting. In the context of our learning scenario, a boosting algorithm will generate some sequence $D_1, D_2, \ldots$ of probability distributions over S and will run the weak learning algorithm L on each of these distributions. This approach can fail, though, if the boosting algorithm generates distributions which are very skewed from the uniform distribution on S; if distribution D_i assigns weights as large as $\frac{20}{m}$ to individual points in S, for instance, then the frequency of noisy examples for L in stage i could be as high as 20%. What we need instead is a *smooth* boosting algorithm which only constructs distributions D_i over S which never assign weight greater than $\frac{10}{m}$ to any single example. By using such a smooth booster we are assured that the weak learner will function successfully at each stage, so the overall boosting process will work correctly.

While the setting described above is artificial, we note that indirect empirical evidence has been given supporting the smooth boosting approach for noisy settings. It is well known (Dietterich, 2000, Schapire, 1999) that commonly used boosting algorithms such as `AdaBoost` (Freund and Schapire, 1997) can perform poorly on noisy data. Dietterich (2000) has suggested that this poor performance is due to `AdaBoost`'s tendency to generate very skewed distributions which put a great deal of weight on a few noisy examples. This overweighting of noisy examples cannot occur under a smooth boosting regimen.

In Section 2 we give a new boosting algorithm, `SmoothBoost`, which is guaranteed to generate only smooth distributions as described above. We show in Section 5 that the distributions generated by `SmoothBoost` are optimally smooth.

`SmoothBoost` is not the first boosting algorithm which attempts to avoid the skewed distributions of `AdaBoost`; algorithms with similar smoothness guarantees have been given by Domingo and Watanabe (2000) and Impagliazzo (1995). Freund (1999) has also described a boosting algorithm which uses a more moderate weighting scheme than `AdaBoost`. In Section 2.3 we show that our `SmoothBoost` algorithm has several other desirable properties, such as constructing a large margin final hypothesis, which are essential for the noisy linear threshold learning application of Section 3. We discuss the relationship between `SmoothBoost` and the algorithms of Domingo and Watanabe, Impagliazzo, and Freund in Section 2.4.

1.2 Learning Linear Threshold Functions with Malicious Noise

We use the `SmoothBoost` algorithm in Section 3 to construct a family of PAC-model malicious noise tolerant algorithms for learning linear threshold functions. A similar family was constructed

by Servedio (2000) using AdaBoost instead of SmoothBoost as the boosting component. It was shown by Servedio (2000) that for linearly separable data these PAC model algorithms have sample complexity bounds which are essentially identical to the bounds obtained from a standard PAC conversion of the online p-norm linear threshold learning algorithms of Grove et al. (1997). We note that the online p-norm algorithms include as special cases ($p = 2$ and $p = \infty$) the well-studied online Perceptron and Winnow algorithms.

Gentile and Littlestone (1999) have given mistake bounds for the online p-norm algorithms when run on examples which are not linearly separable, thus generalizing previous bounds on noise tolerance for Perceptron (Freund and Schapire, 1998) and Winnow (Littlestone, 1991). A significant drawback of the AdaBoost-based PAC-model p-norm algorithms of Servedio (2000) is that they do not appear to succeed in the presence of malicious noise. We show in Section 4 that for all values $2 \leq p \leq \infty$, our new PAC algorithms which use SmoothBoost match both the sample complexity and the malicious noise tolerance of the PAC conversions of the online p-norm algorithms. Our construction thus provides malicious noise tolerant PAC analogues of Perceptron and Winnow (and many other algorithms as well).

2. Smooth Boosting

In this section we describe a new boosting algorithm, SmoothBoost, which has several useful properties. SmoothBoost only constructs smooth distributions which do not put too much weight on any single example; it can be used to generate a large margin final hypothesis; and it can be used with a weak learning algorithm which outputs real-valued hypotheses. All of these properties are essential for the noisy linear threshold learning problem we address in Section 3.

2.1 Preliminaries

We fix some terminology from Impagliazzo (1995) first. A *measure* on a finite set is a function $M : S \to [0,1]$. We write $|M|$ to denote $\sum_{x \in S} M(x)$. Given a measure M, there is a natural induced distribution D_M defined by $D_M(x) = M(x)/|M|$. This definition yields

Observation 1 $L_\infty(D_M) \leq \frac{1}{|M|}$.

Let D be a distribution over a set $S = \langle x^1, y_1 \rangle, \ldots, \langle x^m, y_m \rangle$ of labeled examples with each $y_j \in \{-1,1\}$ and let h be a real-valued function which maps $\{x^1, \ldots, x^m\}$ into $[-1,1]$. If $\frac{1}{2}\sum_{j=1}^m D(j)|h(x^j) - y_j| \leq \frac{1}{2} - \gamma$ then we say that the *advantage* of h under D is γ. We say that an algorithm which takes S and D as input and outputs an h which has advantage at least $\gamma > 0$ is a *weak learning algorithm* (this is somewhat less general than the notion of weak learning which was originally introduced by Kearns and Valiant (1994) but is sufficient for our purposes). Finally, let $f : X \to [-1,1]$ be a real-valued function. We say that the *margin* of f on a labeled example $\langle x, y \rangle \in X \times \{-1,1\}$ is $yf(x)$; intuitively, this is the amount by which f predicts y correctly. Note that the margin of f on $\langle x, y \rangle$ is nonnegative if and only if $\mathrm{sign}(f(x))$ predicts y correctly.

2.2 The SmoothBoost Algorithm

For our purposes, we can view a boosting algorithm as an algorithm which is given access to a weak learning algorithm and a data set of labeled examples. The boosting algorithm generates a sequence of probability distributions over the data set, runs the weak learning algorithm on each of

Input: parameters $0 < \kappa < 1$, $0 \leq \theta \leq \gamma < \frac{1}{2}$
sample $S = \langle x^1, y_1 \rangle, \ldots, \langle x^m, y_m \rangle$ where each $y_i \in \{-1, 1\}$
weak learner WL which takes input (S, D_t) and outputs
$h_t : \{x^1, \ldots, x^m\} \to [-1, 1]$

Output: hypothesis $h(x) = \text{sign}(f(x))$

1. **forall** $j = 1, \ldots, m$ **set** $M_1(j) = 1$
2. **forall** $j = 1, \ldots, m$ **set** $N_0(j) = 0$
3. **set** $t = 1$
4. **until** $|M_t|/m < \kappa$ **do**
5. **forall** $j = 1, \ldots, m$ **set** $D_t(j) = M_t(j)/|M_t|$
6. run $\mathtt{WL}(S, D_t)$ to get h_t such that $\frac{1}{2}\sum_{j=1}^{m} D_t(j)|h_t(x^j) - y_j| \leq \frac{1}{2} - \gamma$
7. **forall** $j = 1, \ldots, m$ **set** $N_t(j) = N_{t-1}(j) + y_j h_t(x^j) - \theta$
8. **forall** $j = 1, \ldots, m$ **set** $M_{t+1}(j) = \begin{cases} 1 & \text{if } N_t(j) < 0 \\ (1-\gamma)^{N_t(j)/2} & \text{if } N_t(j) \geq 0 \end{cases}$
9. **set** $t = t + 1$
10. **set** $T = t - 1$
11. **return** $h = \text{sign}(f(x))$ where $f(x) = \frac{1}{T}\sum_{i=1}^{T} h_i(x)$

Figure 1: The SmoothBoost algorithm.

these distributions, and combines the resulting hypotheses to obtain a final hypothesis which has high accuracy for the data set. (Boosting algorithms of this sort, which work with a fixed sample, are sometimes referred to as boosting-by-sampling algorithms). See Schapire (1999) for a detailed overview of boosting.

Our new boosting algorithm, SmoothBoost, is given in Figure 1. The parameter κ is the desired error rate of the final hypothesis, the parameter γ is the guaranteed advantage of the hypotheses returned by the weak learner, and θ is the desired margin of the final hypothesis. SmoothBoost runs the weak learning algorithm several times on a sequence of carefully constructed distributions and outputs a thresholded sum of the hypotheses thus generated. The quantity $N_t(j)$ in line 7 may be viewed as the cumulative amount by which the hypotheses $h_1, \ldots, h_t$ beat the desired margin θ on the labeled example $\langle x^j, y_j \rangle$. The measure M_{t+1} assigns more weight to examples where N_t is small and less weight to examples where N_t is large, thus forcing the weak learner to focus in stage $t+1$ on examples where previous hypotheses have done poorly. Note that since any measure maps into $[0, 1]$ there is a strict upper bound on the amount of weight which can be assigned to any example.

2.3 Proof of Correctness

Several useful properties of the SmoothBoost algorithm are easy to verify. The algorithm is called SmoothBoost because each distribution it constructs is guaranteed to be "smooth" in that no single point receives too much weight:

Lemma 1 *Each D_t defined in step 5 of* SmoothBoost *has $L_\infty(D_t) \leq \frac{1}{\kappa m}$.*

Proof Follows directly from Observation 1 and the condition in line 4. ■

Another useful property is that the final hypothesis h has margin at least θ on all but a κ fraction of the points in S:

Theorem 2 *If* SmoothBoost *terminates then f satisfies* $\frac{|\{j : y_j f(x^j) \leq \theta\}|}{m} < \kappa$.

Proof Since $N_T(j) = T(y_j f(x^j) - \theta)$, if $y_j f(x^j) \leq \theta$ then $N_T(j) \leq 0$ and hence $M_{T+1}(j) = 1$. Consequently we have

$$\frac{|\{j : y_j f(x^j) \leq \theta\}|}{m} \leq \frac{\sum_{j=1}^m M_{T+1}(j)}{m} = \frac{|M_{T+1}|}{m} < \kappa$$

by the condition in line 4. ■

Note that since $\theta \geq 0$ Theorem 2 implies that the final SmoothBoost hypothesis is correct on all but a κ fraction of S.

Finally we must show that the algorithm terminates in a reasonable amount of time. The following theorem bounds the number of times that SmoothBoost will execute its main loop:

Theorem 3 *If each hypothesis h_t returned by* WL *in line 6 has advantage at least γ under D_t (i.e., satisfies the condition of line 6) and θ is set to $\frac{\gamma}{2+\gamma}$, then* SmoothBoost *terminates with $T < \frac{2}{\kappa\gamma^2\sqrt{1-\gamma}}$.*

As will be evident from the proof, slightly different bounds on T can be established by choosing different values of θ in the range $[0,\gamma]$. We take $\theta = \frac{\gamma}{2+\gamma}$ in the theorem above both to obtain a margin of $\Omega(\gamma)$ and to obtain a clean bound in the theorem. Theorem 3 follows from the bounds established in the following two lemmas:

Lemma 4 *Under the conditions of Theorem 3, we have that*

$$\sum_{j=1}^m \sum_{t=1}^T M_t(j) y_j h_t(x^j) \geq 2\gamma \sum_{t=1}^T |M_t|.$$

Lemma 5 *Under the conditions of Theorem 3, we have that*

$$\sum_{j=1}^m \sum_{t=1}^T M_t(j) y_j h_t(x^j) < \frac{2m}{\gamma\sqrt{1-\gamma}} + \gamma \sum_{t=1}^T |M_t|.$$

Combining these bounds we obtain $\frac{2m}{\gamma\sqrt{1-\gamma}} > \gamma \sum_{t=1}^T |M_t| \geq \gamma\kappa m T$ where the last inequality is because $|M_t| \geq \kappa m$ for $t = 1, \ldots, T$.

Proof of Lemma 4: Since $h_t(x^j) \in [-1,1]$ and $y_j \in \{-1,1\}$, we have $y_j h_t(x^j) = 1 - |h_t(x^j) - y_j|$, and thus

$$\sum_{j=1}^m D_t(j) y_j h_t(x^j) = \sum_{j=1}^m D_t(j)(1 - |h_t(x^j) - y_j|) \geq 2\gamma.$$

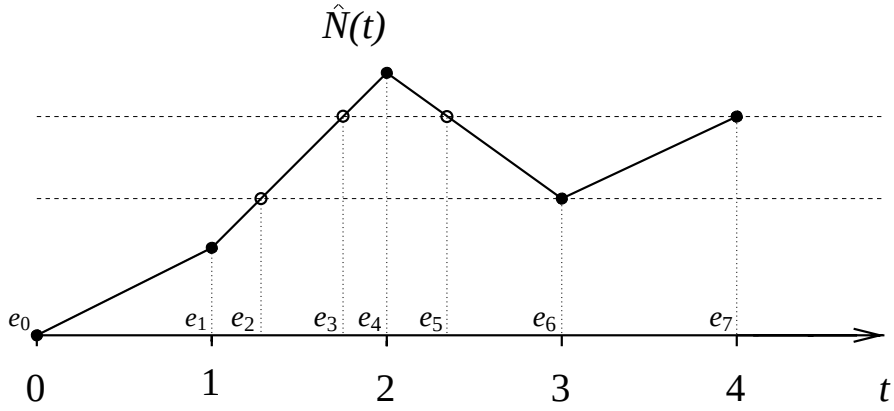

Figure 2: A plot of $\hat{N}$ with $T = 4$. Note that $\hat{N}$ is piecewise linear with joins at integer values of t. A possible pairing of segments matches $[e_2, e_3]$ with $[e_5, e_6]$ and $[e_3, e_4]$ with $[e_4, e_5]$, leaving $[e_0, e_1]$, $[e_1, e_2]$ and $[e_6, e_7]$ unpaired. In this example $\hat{N}$ is increasing on each unpaired segment.

This implies that

$$\sum_{j=1}^{m}\sum_{t=1}^{T} M_t(j)y_j h_t(x^j) = \sum_{t=1}^{T}|M_t|\sum_{j=1}^{m} D_t(j)y_j h_t(x^j) \geq \sum_{t=1}^{T} 2\gamma|M_t|.$$

■

Proof of Lemma 5: By the definition of $N_t(j)$, we have

$$\begin{aligned}\sum_{j=1}^{m}\sum_{t=1}^{T} M_t(j)y_j h_t(x^j) &= \sum_{j=1}^{m}\sum_{t=1}^{T} M_t(j)(N_t(j) - N_{t-1}(j) + \theta) \\ &= \theta\sum_{t=1}^{T}|M_t| + \sum_{t=1}^{T}\sum_{j=1}^{m} M_t(j)(N_t(j) - N_{t-1}(j)). \qquad (1)\end{aligned}$$

It thus suffices to show that if $\theta = \frac{\gamma}{2+\gamma}$, then for each $j = 1, \ldots, m$ we have

$$\sum_{t=1}^{T} M_t(j)(N_t(j) - N_{t-1}(j)) < \frac{2}{\gamma\sqrt{1-\gamma}} + (\gamma - \theta)\sum_{t=1}^{T} M_t(j) \qquad (2)$$

since summing this inequality over $j = 1, \ldots, m$ and substituting into (1) proves the lemma. Fix any $j \in \{1, \ldots, m\}$; for ease of notation we write N_t and M_t in place of $N_t(j)$ and $M_t(j)$ for the rest of the proof.

If $N_t = N_{t-1}$ for some integer t then the term $M_t(N_t - N_{t-1})$ contributes 0 to the sum in (2), so without loss of generality we assume that $N_t \neq N_{t-1}$ for all integers t. We extend the sequence $(N_0, N_1, \ldots, N_T)$ to a continuous piecewise linear function $\hat{N}$ on $[0, T]$ in the obvious way, that is, for t an integer and $\varepsilon \in [0,1]$ we have $\hat{N}(t+\varepsilon) = N_t + \varepsilon(N_{t+1} - N_t)$. Let

$$E = \{e \in [0,T] : \hat{N}(e) = N_t \text{ for some integer } t = 0, 1, \ldots, T\}.$$

The set E is finite so we have $0 = e_0 < e_1 \cdots < e_r = T$ with $E = \{e_0, \ldots, e_r\}$ (see Figure 2). Since for each integer $t \geq 1$ the interval $(t-1,t]$ must contain some e_i, we can reexpress the sum $\sum_{t=1}^{T} M_t(N_t - N_{t-1})$ as

$$\sum_{i=1}^{r} M_{\lceil e_i \rceil} \left(\hat{N}(e_i) - \hat{N}(e_{i-1}) \right). \tag{3}$$

We say that two segments $[e_{a-1}, e_a]$ and $[e_{b-1}, e_b]$ *match* if $\hat{N}(e_{a-1}) = \hat{N}(e_b)$ and $\hat{N}(e_{b-1}) = \hat{N}(e_a)$. For example, in Figure 2 the segment $[e_2, e_3]$ matches $[e_5, e_6]$ but does not match $[e_6, e_7]$. We pair up matching segments until no more pairs can be formed. Note that if any unpaired segments remain, it must be the case that either $\hat{N}$ is increasing on each unpaired segment (if $N_T > 0$) or $\hat{N}$ is decreasing on each unpaired segment (if $N_T < 0$). Now we separate the sum (3) into two pieces, i.e., $\sum_{i=1}^{r} M_{\lceil e_i \rceil}(\hat{N}(e_i) - \hat{N}(e_{i-1})) = P + U$, where P is the sum over all paired segments and U is the sum over all unpaired segments. We will show that $P < (\gamma - \theta) \sum_{t=1}^{T} M_t$ and $U < \frac{2}{\gamma\sqrt{1-\gamma}}$, thus proving the lemma.

First we bound P. Let $[e_{a-1}, e_a]$ and $[e_{b-1}, e_b]$ be a pair of matching segments where $\hat{N}$ is increasing on $[e_{a-1}, e_a]$ and decreasing on $[e_{b-1}, e_b]$. The contribution of these two segments to P is

$$\begin{aligned} &M_{\lceil e_a \rceil} \left(\hat{N}(e_a) - \hat{N}(e_{a-1}) \right) + M_{\lceil e_b \rceil} \left(\hat{N}(e_b) - \hat{N}(e_{b-1}) \right) \\ &\qquad = \left(M_{\lceil e_a \rceil} - M_{\lceil e_b \rceil} \right) \left(\hat{N}(e_a) - \hat{N}(e_{a-1}) \right). \end{aligned} \tag{4}$$

Since each segment $[e_{a-1}, e_a]$ is contained in $[t-1,t]$ for some integer t, we have that $\lceil e_a \rceil - 1 \leq e_{a-1} < e_a \leq \lceil e_a \rceil$. The linearity of $\hat{N}$ on $[\lceil e_a \rceil - 1, \lceil e_a \rceil]$ implies that

$$N_{\lceil e_a \rceil - 1} \leq \hat{N}(e_{a-1}) < \hat{N}(e_a) \leq N_{\lceil e_a \rceil} \leq N_{\lceil e_a \rceil - 1} + 1 - \theta \tag{5}$$

where the last inequality is because $y_j h_t(x^j) \leq 1$ in line 7 of `SmoothBoost`. Similarly, we have that $\lceil e_b \rceil - 1 \leq e_{b-1} < e_b \leq \lceil e_b \rceil$, and hence

$$N_{\lceil e_b \rceil - 1} \geq \hat{N}(e_{b-1}) > \hat{N}(e_b) \geq N_{\lceil e_b \rceil} \geq N_{\lceil e_b \rceil - 1} - 1 - \theta. \tag{6}$$

Since $\hat{N}(e_a) = \hat{N}(e_{b-1})$ inequalities (5) and (6) imply that $N_{\lceil e_a \rceil - 1} \geq N_{\lceil e_b \rceil - 1} - 2$. The definition of M now implies that $M_{\lceil e_b \rceil} \geq (1-\gamma) M_{\lceil e_a \rceil}$. Since $\hat{N}(e_a) - \hat{N}(e_{a-1}) > 0$, we thus have that (4) is at most

$$\gamma M_{\lceil e_a \rceil} \left(\hat{N}(e_a) - \hat{N}(e_{a-1}) \right) \leq \gamma(1-\theta) M_{\lceil e_a \rceil}(e_a - e_{a-1}) \tag{7}$$

where the inequality follows from (5) and the linearity of $\hat{N}$ on $[e_{a-1}, e_a]$. Since $\hat{N}(e_a) - \hat{N}(e_{a-1}) = \hat{N}(e_{b-1}) - \hat{N}(e_b)$, we similarly have that (4) is at most

$$\begin{aligned} \gamma M_{\lceil e_a \rceil} \left(\hat{N}(e_{b-1}) - \hat{N}(e_b) \right) &\leq \frac{\gamma}{1-\gamma} M_{\lceil e_b \rceil} \left(\hat{N}(e_{b-1}) - \hat{N}(e_b) \right) \\ &\leq \frac{\gamma}{1-\gamma}(1+\theta) M_{\lceil e_b \rceil}(e_{b-1} - e_b). \end{aligned} \tag{8}$$

Since $\theta = \frac{\gamma}{2+\gamma}$ we have that the right side of (7) equals $\frac{2\gamma}{2+\gamma} M_{\lceil e_a \rceil}(e_a - e_{a-1})$ and the right side of (8) equals $\frac{2\gamma(1+\gamma)}{(1-\gamma)(2+\gamma)} M_{\lceil e_b \rceil}(e_{b-1} - e_b)$ Since (4) is upper bounded by each of these quantities, taking a

convex combination of $\frac{1+\gamma}{2}$ times the first plus $\frac{1-\gamma}{2}$ times the second and simplifying we find that (4) is at most

$$\frac{\gamma(1+\gamma)}{2+\gamma}\left(M_{\lceil e_a \rceil}(e_a - e_{a-1}) + M_{\lceil e_b \rceil}(e_{b-1} - e_b)\right). \tag{9}$$

If we sum (9) over all pairs of matching segments the resulting quantity is an upper bound on P. In this sum, for each value of $t = 1, \ldots, T$, the coefficient of M_t will be at most $\frac{\gamma(1+\gamma)}{2+\gamma} = \gamma - \theta$. (This bound on the coefficient of M_t holds because for each t, the total length of all paired segments in $[t-1,t]$ is at most 1). Consequently we have $P < (\gamma - \theta)\sum_{t=1}^{T} M_t$ as desired.

Now we show that U, the sum over unpaired segments, is at most $\frac{2}{\gamma\sqrt{1-\gamma}}$. If $\hat{N}$ is decreasing on each unpaired segment then clearly $U < 0$, so we suppose that $\hat{N}$ is increasing on each unpaired segment. Let $[e_{c_1-1}, e_{c_1}], \ldots, [e_{c_d-1}, e_{c_d}]$ be all the unpaired segments. As in Figure 2 it must be the case that the intervals $[\hat{N}(e_{c_i-1}), \hat{N}(e_{c_i}))$ are all disjoint and their union is $[0, N_T)$. By the definition of M, we have $U = \sum_{i=1}^{d}(1-\gamma)^{(N_{\lceil e_{c_i} \rceil - 1})/2}\left(\hat{N}(e_{c_i}) - \hat{N}(e_{c_i-1})\right)$. As in the bound for P, we have

$$N_{\lceil e_{c_i} \rceil - 1} \leq \hat{N}(e_{c_i-1}) < \hat{N}(e_{c_i}) \leq N_{\lceil e_{c_i} \rceil} \leq N_{\lceil e_{c_i} \rceil - 1} + 1 - \theta < N_{\lceil e_{c_i} \rceil - 1} + 1$$

and hence

$$\begin{aligned} U &\leq \sum_{i=1}^{d}(1-\gamma)^{(\hat{N}(e_{c_i})-1)/2}\left(\hat{N}(e_{c_i}) - \hat{N}(e_{c_i-1})\right) \\ &= (1-\gamma)^{-1/2}\sum_{i=1}^{d}(1-\gamma)^{\hat{N}(e_{c_i})/2}\left(\hat{N}(e_{c_i}) - \hat{N}(e_{c_i-1})\right). \end{aligned}$$

Since $\hat{N}$ is increasing, for each i we have

$$(1-\gamma)^{\hat{N}(e_{c_i})/2}\left(\hat{N}(e_{c_i}) - \hat{N}(e_{c_i-1})\right) < \int_{z=\hat{N}(e_{c_i-1})}^{\hat{N}(e_{c_i})}(1-\gamma)^{z/2}dz.$$

Since the disjoint intervals $[\hat{N}(e_{c_i-1}), \hat{N}(e_{c_i}))$ cover $[0, N_T)$ we thus have

$$\begin{aligned} U &< (1-\gamma)^{-1/2}\int_{z=0}^{N_T}(1-\gamma)^{z/2}dz \\ &< (1-\gamma)^{-1/2}\int_{z=0}^{\infty}(1-\gamma)^{z/2}dz \\ &= \frac{-2}{\sqrt{1-\gamma}\ln(1-\gamma)} < \frac{2}{\gamma\sqrt{1-\gamma}} \quad \text{for } 0 < \gamma < 1/2. \end{aligned}$$

(Lemma 5) ■

2.4 Comparison with Other Boosting Algorithms

The `SmoothBoost` algorithm was inspired by an algorithm given by Impagliazzo (1995) in the context of hard-core set constructions in complexity theory. Klivans and Servedio (1999) observed that Impagliazzo's algorithm can be reinterpreted as a boosting algorithm which generates distributions D_t which, like the distributions generated by `SmoothBoost`, satisfy $L_\infty(D_t) \leq \frac{1}{\kappa m}$. However, our

SmoothBoost algorithm differs from Impagliazzo's algorithm in several important ways. The algorithm of Impagliazzo (1995) uses additive rather than multiplicative updates for $M_t(j)$, and the bound on T which is given for the algorithm by Impagliazzo (1995) is $O(\frac{1}{\kappa^2\gamma^2})$ which is worse than our bound by essentially a factor of $\frac{1}{\kappa}$. Another important difference is that Impagliazzo's algorithm has no θ parameter and does not appear to output a large margin final hypothesis. Finally, the analysis given by Impagliazzo only covers the case where the weak hypotheses are binary-valued rather than real-valued.

The well-known boosting algorithm AdaBoost of Freund and Schapire (1997) is somewhat faster than SmoothBoost, requiring only $T = O(\frac{\log(1/\kappa)}{\gamma^2})$ stages. Like SmoothBoost, AdaBoost can be used with real-valued weak hypotheses and can be used to output a large margin final hypothesis (Schapire et al., 1998). However, AdaBoost is not guaranteed to generate only smooth distributions, and thus does not appear to be useful in a malicious noise context.

Freund (1999) has recently introduced and studied a sophisticated boosting algorithm called BrownBoost which uses a gentler weighting scheme than AdaBoost. Freund suggests that BrownBoost should be well suited for dealing with noisy data; however it is not clear from the analysis of Freund (1999) whether BrownBoost-generated distributions satisfy a smoothness property such as the $L_\infty(D_t) \leq \frac{1}{\kappa m}$ property of SmoothBoost, or whether BrownBoost can be used to generate a large margin final hypothesis. We note that the BrownBoost algorithm is much more complicated to run than SmoothBoost, as it involves solving a differential equation at each stage of boosting.

SmoothBoost is perhaps most similar to the modified AdaBoost algorithm MadaBoost which was defined and analyzed by Domingo and Watanabe (2000). Like SmoothBoost, MadaBoost uses multiplicative updates on weights and never allows weights to exceed 1 in value. Domingo and Watanabe proved that MadaBoost takes at most $T \leq \frac{2}{\kappa\gamma^2}$ stages, which is quite similar to our bound in Theorem 3. (If we set $\theta = 0$ in SmoothBoost, a slight modification of the proof of Theorem 3 gives a bound of roughly $\frac{4}{3\kappa\gamma^2}$, which improves the Madaboost bound by a constant factor.) However, the analysis for MadaBoost given by Domingo and Watanabe (2000) only covers the case of binary-valued weak hypotheses, and does not establish that MadaBoost generates a large margin final hypothesis. We also note that our proof technique of simultaneously upper and lower bounding $\sum_{j=1}^{m}\sum_{t=1}^{T} M_t(j) y_j h_t(x^j)$ is different from the approach used by Domingo and Watanabe (2000).

3. Learning Linear Threshold Functions with Malicious Noise

In this section we show how the SmoothBoost algorithm can be used in conjunction with a simple noise tolerant weak learning algorithm to obtain a PAC learning algorithm for learning linear threshold functions with malicious noise.

3.1 Geometric Preliminaries

For $\overline{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$ and $p \geq 1$ we write $\|\overline{x}\|_p$ to denote the p-norm of $\overline{x}$, namely $\|\overline{x}\|_p = (\sum_{i=1}^{n} |x_i|^p)^{1/p}$. The ∞-norm of $\overline{x}$ is $\|\overline{x}\|_\infty = \max_{i=1,\ldots,n} |x_i|$. We write $B_p(R)$ to denote the p-norm ball of radius R, i.e., $B_p(R) = \{\overline{x} \in \mathbb{R}^n : \|\overline{x}\|_p \leq R\}$.

For $p, q \geq 1$ the q-norm is *dual* to the p-norm if $\frac{1}{p} + \frac{1}{q} = 1$; so the 1-norm and the ∞-norm are dual to each other and the 2-norm is dual to itself. For the rest of the paper p and q always denote dual norms. The following facts (see Taylor and Mann, 1972, pp. 203-204) will be useful:

Hölder Inequality: $|\overline{u} \cdot \overline{v}| \leq \|\overline{u}\|_p \|\overline{v}\|_q$ for all $\overline{u}, \overline{v} \in \mathbb{R}^n$ and $1 \leq p \leq \infty$.

Minkowski Inequality: $\|\overline{u}+\overline{v}\|_p \leq \|\overline{u}\|_p + \|\overline{v}\|_p$ for all $\overline{u},\overline{v} \in \mathbb{R}^n$ and $1 \leq p \leq \infty$.

Finally, recall that a *linear threshold function* is a function $f : \mathbb{R}^n \to \{-1,1\}$ such that $f(\overline{x}) = \text{sign}(\overline{u} \cdot \overline{x})$ for some $\overline{u} \in \mathbb{R}^n$.

3.2 PAC Learning with Malicious Noise

Let $EX^{\eta}_{MAL}(\overline{u},D)$ be a *malicious example oracle with noise rate* η that behaves as follows when invoked: with probability $1-\eta$ the oracle returns a *clean* example $\langle \overline{x}, \text{sign}(\overline{u} \cdot \overline{x}) \rangle$ where $\overline{x}$ is drawn from the probability distribution D over $B_p(R)$. With probability η, though, $EX^{\eta}_{MAL}(\overline{u},D)$ returns a *dirty* example $\langle \overline{x}, y \rangle \in B_p(R) \times \{-1,1\}$ about which nothing can be assumed. Such a malicious example $\langle \overline{x}, y \rangle$ may be chosen by a computationally unbounded adversary which has complete knowledge of $\overline{u}$, D, and the state of the learning algorithm when the oracle is invoked.

The goal of a learning algorithm in this model is to construct an approximation to the target concept $\text{sign}(\overline{u} \cdot \overline{x})$. More formally, we say that a Boolean function $h : \mathbb{R}^n \to \{-1,1\}$ is an ε-*approximator for* $\overline{u}$ *under* D if $\Pr_{\overline{x} \in D}[h(\overline{x}) \neq \text{sign}(\overline{u} \cdot \overline{x})] \leq \varepsilon$. The learning algorithm is given an accuracy parameter ε and a confidence parameter δ, has access to $EX^{\eta}_{MAL}(\overline{u},D)$, and must output a hypothesis h which, with probability at least $1-\delta$, is an ε-approximator for $\overline{u}$ under D. The *sample complexity* of a learning algorithm in this model is the number of times it queries the malicious example oracle.[1]

A final note: like the Perceptron algorithm, the learning algorithms which we consider will require that the quantity $\overline{u} \cdot \overline{x}$ be bounded away from zero (at least most of the time). We thus say that a distribution D is ξ-*good for* $\overline{u}$ if $|\overline{u} \cdot \overline{x}| \geq \xi$ for all $\overline{x}$ which have nonzero probability under D, and we restrict our attention to learning under ξ-good distributions. (Of course, dirty examples drawn from $EX^{\eta}_{MAL}(\overline{u},D)$ need not satisfy $|\overline{u} \cdot \overline{x}| \geq \xi$.)

3.3 A Noise Tolerant Weak Learning Algorithm

As shown in Figure 3, our weak learning algorithm for linear threshold functions, called WLA, takes as input a data set S and a distribution D over S. The algorithm computes the vector $\overline{z}$ which is the average location of the (label-normalized) points in S under D, transforms $\overline{z}$ to obtain a vector $\overline{w}$, and predicts using the linear functional defined by $\overline{w}$. As motivation for the algorithm, note that if every example pair $\langle \overline{x}, y \rangle$ satisfies $y = \text{sign}(\overline{u} \cdot \overline{x})$ for some $\overline{u}$, then each point $y\overline{x}$ would lie on the same side of the hyperplane defined by $\overline{u}$ as $\overline{u}$ itself, and hence the average vector $\overline{z}$ defined in Step 1 of the algorithm intuitively should point in roughly the same direction as $\overline{u}$.

Servedio (2000) showed that the WLA algorithm is a weak learning algorithm for linear threshold functions in a noise-free setting. The following theorem shows that if a small fraction of the examples in S are affected by malicious noise, WLA will still generates a hypothesis with nonnegligible advantage provided that the input distribution D is sufficiently smooth.

Theorem 6 *Fix* $2 \leq p \leq \infty$ *and let* $S = \langle \overline{x}^1, y_1 \rangle, \ldots, \langle \overline{x}^m, y_m \rangle$ *be a set of labeled examples with each* $\overline{x}^j \in B_p(R)$. *Let* D *be a distribution over* S *such that* $L_\infty(D) \leq \frac{1}{\kappa m}$. *Suppose that* $\xi > 0$ *and* $\overline{u} \in \mathbb{R}^n$ *are*

1. A slightly stronger model of PAC learning with malicious noise has also been proposed by Aslam and Decatur (1998) and Cesa-Bianchi et al. (1999). In this model first a clean sample of the desired size is drawn from a noise-free oracle; then each point in the sample is independently selected with probability η; then an adversary replaces each selected point with a dirty example of its choice; and finally the corrupted sample is provided to the learning algorithm. This model is stronger than the original malicious noise model since each dirty example is chosen by the adversary with full knowledge of the entire sample. All of our results also hold in this stronger model.

Input: parameter $p \geq 2$
sample $S = \langle \overline{x}^1, y_1 \rangle, \ldots, \langle \overline{x}^m, y_m \rangle$ where each $y_i \in \{-1,1\}$
distribution D over S
upper bound R on $\|\overline{x}\|_p$

Output: hypothesis $h(\overline{x})$

1. **set** $\overline{z} = \sum_{j=1}^{m} D(j) y_j \overline{x}^j$
2. **for all** $i = 1, \ldots, n$ **set** $w_i = \text{sign}(z_i)|z_i|^{p-1}$
3. **return** hypothesis $h(\overline{x}) \equiv \overline{v} \cdot \overline{x}$ where $\overline{v} = \frac{\overline{w}}{\|\overline{w}\|_q R}$

Figure 3: The p-norm weak learning algorithm WLA.

such that $\xi \leq R\|\overline{u}\|_q$ *and at most* $\eta' m$ *examples in* S *do not satisfy* $y_j(\overline{u} \cdot \overline{x}^j) \geq \xi$, *where* $\eta' \leq \frac{\kappa\xi}{4R\|\overline{u}\|_q}$. *Then* $\mathtt{WLA}(p, S, D)$ *returns a hypothesis* $h : B_p(R) \to [-1,1]$ *which has advantage at least* $\frac{\xi}{4R\|\overline{u}\|_q}$ *under* D.

Proof: By Hölder's inequality, for any $\overline{x} \in B_p(R)$ we have

$$|h(\overline{x})| = \frac{|\overline{w} \cdot \overline{x}|}{\|\overline{w}\|_q R} \leq \frac{\|\overline{w}\|_q \|\overline{x}\|_p}{\|\overline{w}\|_q R} \leq 1,$$

and thus h indeed maps $B_p(R)$ into $[-1,1]$.

Now we show that h has the desired advantage. Since $h_t(\overline{x}^j) \in [-1,1]$ and $y_j \in \{-1,1\}$, we have $|h(\overline{x}^j) - y_j| = 1 - y_j h(\overline{x}^j)$, so

$$\frac{1}{2}\sum_{j=1}^{m} D(j)|h(\overline{x}^j) - y_j| = \frac{1}{2}\sum_{j=1}^{m} D(j)(1 - y_j h(\overline{x}^j)) = \frac{1}{2} - \left(\frac{\sum_{j=1}^{m} D(j) y_j (\overline{w} \cdot \overline{x}^j)}{2\|\overline{w}\|_q R}\right).$$

To prove the theorem it thus suffices to show that $\frac{\sum_{j=1}^{m} D(j) y_j (\overline{w} \cdot \overline{x}^j)}{\|\overline{w}\|_q} \geq \frac{\xi}{2\|\overline{u}\|_q}$. The numerator of the left side is $\overline{w} \cdot \left(\sum_{j=1}^{m} D(j) y_j \overline{x}^j\right) = \overline{w} \cdot \overline{z} = \sum_{i=1}^{n} |z_i|^p = \|\overline{z}\|_p^p$. Using the fact that $(p-1)q = p$, the denominator is

$$\|\overline{w}\|_q = \left(\sum_{i=1}^{n} \left(|z_i|^{p-1}\right)^q\right)^{1/q} = \left(\sum_{i=1}^{n} |z_i|^p\right)^{1/q} = \|\overline{z}\|_p^{p/q}.$$

We can therefore rewrite the left side as $\|\overline{z}\|_p^p / \|\overline{z}\|_p^{p/q} = \|\overline{z}\|_p$, and thus our goal is to show that $\|\overline{z}\|_p \geq \frac{\xi}{2\|\overline{u}\|_q}$. By Hölder's inequality it suffices to show that $\overline{z} \cdot \overline{u} \geq \frac{\xi}{2}$, which we now prove.

Let $S_1 = \{\langle \overline{x}^j, y_j \rangle \in S : y_j(\overline{u} \cdot \overline{x}^j) \geq \xi\}$ and let $S_2 = S \setminus S_1$. The definition of S_1 immediately yields $\sum_{j \in S_1} D(j) y_j (\overline{u} \cdot \overline{x}^j) \geq D(S_1)\xi$. Moreover, since each $\|\overline{x}^j\|_p \leq R$, by Hölder's inequality we have $y_j(\overline{u} \cdot \overline{x}^j) \geq -\|\overline{x}^j\|_p \cdot \|\overline{u}\|_q \geq -R\|\overline{u}\|_q$ for each $\langle \overline{x}^j, y_j \rangle \in S_2$. Since each example in S_2 has weight at most $\frac{1}{\kappa m}$ under D, we have $D(S_2) \leq \frac{\eta'}{\kappa}$, and hence

$$\overline{z} \cdot \overline{u} \;=\; \sum_{j=1}^{m} D(j) y_j (\overline{u} \cdot \overline{x}^j) = \sum_{j \in S_1} D(j) y_j (\overline{u} \cdot \overline{x}^j) + \sum_{j \in S_2} D(j) y_j (\overline{u} \cdot \overline{x}^j)$$

$$\begin{aligned} &\geq D(S_1)\xi - D(S_2)R\|\overline{u}\|_q \geq \left(1-\frac{\eta'}{\kappa}\right)\xi - \frac{\eta' R\|\overline{u}\|_q}{\kappa} \\ &\geq \frac{3\xi}{4} - \frac{\xi}{4} = \frac{\xi}{2}, \end{aligned}$$

where the inequality $(1-\frac{\eta'}{\kappa}) \geq \frac{3}{4}$ follows from the bound on η' and the fact that $\xi \leq R\|\overline{u}\|_q$. ■

3.4 Putting it All Together

The algorithm for learning $\text{sign}(\overline{u}\cdot\overline{x})$ with respect to a ξ-good distribution D over $B_p(R)$ is as follows:

- Draw from $EX^{\eta}_{MAL}(\overline{u},D)$ a sample $S = \langle\overline{x}^1,y_1\rangle,\ldots,\langle\overline{x}^m,y_m\rangle$ of m labeled examples.
- Run `SmoothBoost` on S with parameters $\kappa = \frac{\varepsilon}{4}$, $\gamma = \frac{\xi}{4R\|\overline{u}\|_q}$, $\theta = \frac{\gamma}{2+\gamma}$ using `WLA` as the weak learning algorithm.

We now determine constraints on the sample size m and the malicious noise rate η under which this is a successful and efficient learning algorithm.

We first note that since D is ξ-good for $\overline{u}$, we have that $\xi \leq R\|\overline{u}\|_q$. Furthermore, since $\kappa = \frac{\varepsilon}{4}$, Lemma 1 implies that each distribution D_t which is given to `WLA` by `SmoothBoost` has $L_\infty(D_t) \leq \frac{4}{\varepsilon m}$. Let $S_C \subseteq S$ be the clean examples and $S_D = S \setminus S_C$ the dirty examples in S. If $\eta \leq \frac{\varepsilon\xi}{32R\|\overline{u}\|_q}$ and $m \geq \frac{96R\|\overline{u}\|_q}{\varepsilon\xi}\log\frac{2}{\delta}$, then a simple Chernoff bound implies that with probability at least $1-\frac{\delta}{2}$ we have $|S_D| \leq \frac{\varepsilon\xi}{16R\|\overline{u}\|_q}m$. Thus, we can apply Theorem 6 with $\eta' = \frac{\varepsilon\xi}{16R\|\overline{u}\|_q}$; so each weak hypothesis $h_t(\overline{x}) = \overline{v}^t\cdot\overline{x}$ generated by `WLA` has advantage $\frac{\xi}{4R\|\overline{u}\|_q}$ under D_t. Consequently, by Theorems 2 and 3, `SmoothBoost` efficiently outputs a final hypothesis $h(\overline{x}) = \text{sign}(f(\overline{x}))$ which has margin less than θ on at most an $\frac{\varepsilon}{4}$ fraction of S. Since $|S_C|$ is easily seen to be at least $\frac{m}{2}$, we have that the margin of h is less than θ on at most an $\frac{\varepsilon}{2}$ fraction of S_C. This means that we can apply powerful methods from the theory of data-dependent structural risk minimization (Bartlett and Shawe-Taylor, 1999, Shawe-Taylor et al., 1998) to bound the error of h under D.

Recall that the final `SmoothBoost` hypothesis is $h(\overline{x}) = \text{sign}(f(\overline{x}))$ where $f(\overline{x}) = \overline{v}\cdot\overline{x}$ is a convex combination of hypotheses $h_t(\overline{x}) = \overline{v}^t\cdot\overline{x}$. Since each vector $\overline{v}^t$ satisfies $\|\overline{v}^t\|_q \leq \frac{1}{R}$, by Minkowski's inequality we have that $\|\overline{v}\|_q \leq \frac{1}{R}$ as well. The following theorem was proved by Servedio (2000):

Theorem 7 *Fix any value $2 \leq p \leq \infty$ and let F be the class of functions $\{\overline{x}\mapsto\overline{v}\cdot\overline{x} : \|\overline{v}\|_q \leq \frac{1}{R}, \overline{x}\in B_p(R)\}$. Then $fat_F(\mu) \leq \frac{2\log 4n}{\mu^2}$, where $fat_F(\mu)$ is the* fat-shattering dimension of F at scale μ *as defined in Bartlett et al. (1996), Bartlett and Shawe-Taylor (1999), Shawe-Taylor et al. (1998).*

The following theorem is due to Bartlett and Shawe-Taylor (1999):

Theorem 8 *Let F be a collection of real-valued functions over some domain X, let D be a distribution over $X\times\{-1,1\}$, let $S = \langle\overline{x}^1,y_1\rangle,\ldots,\langle\overline{x}^m,y_m\rangle$ be a sequence of labeled examples drawn from D, and let $h(\overline{x}) = sign(f(\overline{x}))$ for some $f\in F$. If h has margin less than θ on at most k examples in S, then with probability at least $1-\delta$ we have that $\Pr_{\langle\overline{x},y\rangle\in D}[h(\overline{x})\neq y]$ is at most*

$$\frac{k}{m} + \sqrt{\frac{2}{m}(d\ln(34e/m)\log(578m) + \ln(4/\delta))}, \tag{10}$$

where $d = fat_F(\theta/16)$.

We have that h has margin less than θ on at most an $\frac{\varepsilon}{2}$ fraction of the clean examples S_C, so we may take k/m to be $\frac{\varepsilon}{2}$ in the above theorem. Now if we apply Theorem 7 and solve for m the inequality obtained by setting (10) to be at most ε, we obtain

Theorem 9 *Fix* $2 \leq p \leq \infty$ *and let* D *be a distribution over* $B_p(R)$ *which is* ξ*-good for* $\overline{u}$*. The algorithm described above uses* $m = \tilde{O}\left(\left(\frac{R\|\overline{u}\|_q}{\xi\varepsilon}\right)^2\right)$ *examples and outputs an* ε*-approximator for* $\overline{u}$ *under* D *with probability* $1-\delta$ *in the presence of malicious noise at a rate* $\eta = \Omega\left(\varepsilon \cdot \frac{\xi}{R\|\overline{u}\|_q}\right)$.

4. Comparison with Online Algorithms

The bounds given by Theorem 9 on sample complexity and malicious noise tolerance of our algorithms based on `SmoothBoost` are remarkably similar to the bounds which can be obtained through a natural PAC conversion of the online p-norm algorithms introduced by Grove et al. (1997) and studied by Gentile and Littlestone (1999). Grove, Littlestone and Schuurmans (Theorem 6.1) proved that the online p-norm algorithm makes at most $O\left(\left(\frac{R\|\overline{u}\|_q}{\xi}\right)^2\right)$ mistakes on linearly separable data. Subsequently Gentile and Littlestone (1999) extended the earlier analysis of Grove, Littlestone and Schuurmans and considered a setting in which the examples are not linearly separable. Their analysis (Theorem 6) shows that if an example sequence containing K malicious errors is provided to the online p-norm algorithm, then the algorithm will make at most

$$O\left(\left(\frac{R\|\overline{u}\|_q}{\xi}\right)^2 + K \cdot \frac{R\|\overline{u}\|_q}{\xi}\right)$$

mistakes. To obtain PAC-model bounds on the online p-norm algorithms in the presence of malicious noise, we use the following theorem due to Auer and Cesa-Bianchi (1998, Theorem 6.2):

Theorem 10 *Fix a hypothesis class* H *of Vapnik-Chervonenkis dimension* d*. Let* A *be an online learning algorithm with the following properties: (1)* A *only uses hypotheses which belong to* H*, (2) if* A *is given a noise-free example sequence then* A *makes at most* m_0 *mistakes, and (3) if* A *is given an example sequence with* K *malicious errors then* A *makes at most* $m_0 + BK$ *mistakes. Then there is a PAC algorithm* A' *which learns to accuracy* ε *and confidence* δ*, uses* $\tilde{O}(\frac{B^2}{\varepsilon^2} + \frac{m_0}{\varepsilon} + \frac{d}{\varepsilon})$ *examples, and can tolerate a malicious noise rate* $\eta = \frac{\varepsilon}{2B}$.

Applying this theorem, we find that these PAC conversions of the online p-norm algorithms have sample complexity and malicious noise tolerance bounds which are essentially identical to the bounds given for our `SmoothBoost`-based algorithm.

5. `SmoothBoost` is Optimally Smooth

It is evident from the proof of Theorem 9 that the smoothness of the distributions generated by `SmoothBoost` relates directly to the level of malicious noise which our linear threshold learning algorithm can tolerate. On the other hand, as mentioned in Section 1, Kearns and Li have shown

that for a broad range of concept classes no algorithm can learn to accuracy ε in the presence of malicious noise at a rate $\eta > \frac{\varepsilon}{1+\varepsilon}$. Using the Kearns-Li upper bound on malicious noise tolerance, we prove in this section that SmoothBoost is optimal up to constant factors in terms of the smoothness of the distributions which it generates. This demonstrates an interesting connection between bounds on noise-tolerant learning and bounds on boosting algorithms.

Recall that if SmoothBoost is run on a set of m examples with input parameters κ, γ, θ, then each distribution D_t which SmoothBoost constructs will satisfy $L_\infty(D_t) \leq \frac{1}{\kappa m}$. The proof is by contradiction; so suppose that there exists a boosting algorithm called SuperSmoothBoost which is similar to SmoothBoost but which has an even stronger guarantee on its distributions. More precisely we suppose that SuperSmoothBoost takes as input parameters κ, γ and a labeled sample S of size m, has access to a weak learning algorithm WL, generates a sequence $D_1, D_2, \ldots$ of distributions over S, and outputs a Boolean-valued final hypothesis h. As in Section 2.3, we suppose that if the weak learning algorithm WL always returns a hypothesis h_t which has advantage γ under D_t, then SuperSmoothBoost will eventually halt and the final hypothesis h will agree with at least a $1-\kappa$ fraction of the labeled examples in S. Finally, we suppose that each distribution D_t is guaranteed to satisfy $L_\infty(D_t) \leq \frac{1}{64\kappa m}$.

Consider the following severely restricted linear threshold learning problem: the domain is $\{-1,1\}^2 \subset \Re^2$, so any distribution D can assign weight only to these four points. Moreover, we only allow two possibilities for the target concept $\text{sign}(\overline{u} \cdot \overline{x})$: the vector $\overline{u}$ is either $(1,0)$ or $(0,1)$. The point $(1,1)$ is classified positive by both of these concepts; the point $(1,-1)$ is classified positive only by the first of these concepts; the point $(-1,1)$ is classified positive only by the second of these concepts; and the point $(-1,-1)$ is classified positive by neither of these concepts. Hence the concept class consisting of these two concepts over these four points is a *distinct* concept class as defined by Kearns and Li (1993). It is clear that every example belongs to $B_\infty(1)$ (that is, $R = 1$), that $\|\overline{u}\|_1 = 1$, and that any distribution D over $\{-1,1\}^2$ is 1-good for $\overline{u}$ (i.e., $\xi = 1$).

Consider the following algorithm for this restricted learning problem:

- Draw from $EX^{\eta}_{MAL}(\overline{u}, D)$ a sample $S = \langle \overline{x}^1, y_1 \rangle, \ldots, \langle \overline{x}^m, y_m \rangle$ of m labeled examples.
- Run SuperSmoothBoost on S with parameters $\kappa = \frac{\varepsilon}{4}$, $\gamma = \frac{\xi}{4R\|\overline{u}\|_q} = \frac{1}{4}$ using WLA with $p = \infty$ as the weak learning algorithm.

Suppose that the malicious noise rate η is 2ε. As in Section 3.4, a Chernoff bound shows that for $m = O(\frac{1}{\varepsilon} \log \frac{1}{\delta})$, with probability at least $1 - \frac{\delta}{2}$ we have that the sample S contains at most $4\varepsilon m$ dirty examples. By the SuperSmoothBoost smoothness property and our choice of κ, we have that $L_\infty(D_t) \leq \frac{1}{16\varepsilon m}$. Theorem 6 now implies that each WLA hypothesis h_t has advantage at least $\frac{\xi}{4R\|\overline{u}\|_q} = \frac{1}{4}$ with respect to D_t. As in Section 3.4, we have that with probability at least $1 - \frac{\delta}{2}$ the final hypothesis h output by SuperSmoothBoost disagrees with at most an $\frac{\varepsilon}{2}$ fraction of the clean examples S_C.

Since the domain is finite (in fact of size four) we can bound generalization error directly. A simple Chernoff bound argument shows that if m is sufficiently large, then with probability at least $1 - \delta$ the hypothesis h will be an ε-approximator for $\text{sign}(\overline{u} \cdot \overline{x})$ under D. However, Kearns and Li (1993, Theorem 1) have shown that no learning algorithm for a distinct concept class can learn to accuracy ε with probability $1 - \delta$ in the presence of malicious noise at rate $\eta \geq \frac{\varepsilon}{1+\varepsilon}$. This contradiction proves that the SuperSmoothBoost algorithm cannot exist, and hence the distributions generated by SmoothBoost are optimal up to constant factors.

Acknowledgments

We thank Avrim Blum for a helpful discussion concerning the malicious noise tolerance of the Perceptron algorithm. We also thank Les Valiant for suggesting that the techniques in this paper could be used to prove a lower bound on the smoothness of an arbitrary boosting algorithm. This work was done while the author was at the Division of Engineering and Applied Sciences, Harvard University, and was supported by NSF Grant CCR-98-77049 and by an NSF Mathematical Sciences Postdoctoral Research Fellowship.

References

J. Aslam and S. Decatur. Specification and simulation of statistical query algorithms for efficiency and noise tolerance. *Journal of Computer and System Sciences*, 56:191–208, 1998.

P. Auer. Learning nested differences in the presence of malicious noise. *Theoretical Computer Science*, 185(1):159–175, 1997.

P. Auer and N. Cesa-Bianchi. On-line learning with malicious noise and the closure algorithm. *Annals of Mathematics and Artificial Intelligence*, 23:83–99, 1998.

P. Bartlett, P. Long, and R. Williamson. Fat-shattering and the learnability of real-valued functions. *Journal of Computer and System Sciences*, 52(3):434–452, 1996.

P. Bartlett and J. Shawe-Taylor. *Generalization performance of support vector machines and other pattern classifiers*, pages 43–54. MIT Press, 1999.

N. Cesa-Bianchi, E. Dichterman, P. Fischer, E. Shamir, and H.U. Simon. Sample-efficient strategies for learning in the presence of noise. *Journal of the ACM*, 46(5):684–719, 1999.

S. Decatur. Statistical queries and faulty PAC oracles. In *Proceedings of the Sixth Workshop on Computational Learning Theory*, pages 262–268, 1993.

T.G. Dietterich. An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning*, 40(2):139–158, 2000.

C. Domingo and O. Watanabe. Madaboost: a modified version of adaboost. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, pages 180–189, 2000.

Y. Freund. An adaptive version of the boost-by-majority algorithm. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory*, pages 102–113, 1999.

Y. Freund and R. 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.

Y. Freund and R. Schapire. Large margin classification using the perceptron algorithm. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, pages 209–217., 1998.

C. Gentile and N. Littlestone. The robustness of the p-norm algorithms. In *Proceedings of the 12th Annual Conference on Computational Learning Theory*, pages 1–11, 1999.

A. Grove, N. Littlestone, and D. Schuurmans. General convergence results for linear discriminant updates. In *Proceedings of the Tenth Annual Conference on Computational Learning Theory*, pages 171–183, 1997.

R. Impagliazzo. Hard-core distributions for somewhat hard problems. In *Proceedings of the Thirty-Sixth Annual Symposium on Foundations of Computer Science*, pages 538–545, 1995.

M. Kearns and M. Li. Learning in the presence of malicious errors. *SIAM Journal on Computing*, 22(4):807–837, 1993.

M. Kearns and L. Valiant. Cryptographic limitations on learning boolean formulae and finite automata. *Journal of the ACM*, 41(1):67–95, 1994.

A. Klivans and R. Servedio. Boosting and hard-core sets. In *Proceedings of the Fortieth Annual Symposium on Foundations of Computer Science*, pages 624–633, 1999.

N. Littlestone. Redundant noisy attributes, attribute errors, and linear-threshold learning using winnow. In *Proceedings of the Fourth Annual Workshop on Computational Learning Theory*, pages 147–156, 1991.

Y. Mansour and M. Parnas. Learning conjunctions with noise under product distributions. *Information Processing Letters*, 68(4):189–196, 1998.

R. Schapire. Theoretical views of boosting and applications. In *Proceedings of the Tenth International Conference on Algorithmic Learning Theory*, pages 12–24, 1999.

R. Schapire, Y. Freund, P. Bartlett, and W. Lee. Boosting the margin: a new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26(5):1651–1686, 1998.

R. Servedio. PAC analogues of perceptron and winnow via boosting the margin. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, pages 148–157, 2000.

J. Shawe-Taylor, P. Bartlett, R. Williamson, and M. Anthony. Structural risk minimization over data-dependent hierarchies. *IEEE Transactions on Information Theory*, 44(5):1926–1940, 1998.

A. Taylor and W. Mann. *Advanced Calculus*. Wiley & Sons, 1972.

L. Valiant. Learning disjunctions of conjunctions. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 560–566, 1985.