# SPARSE BOOSTING

Zhen James Xiang and Peter J. Ramadge

Dept. of Electrical Engineering, Princeton University, Princeton NJ

## ABSTRACT

We propose a boosting algorithm that seeks to minimize the AdaBoost exponential loss of a composite classifier using only a sparse set of base classifiers. The proposed algorithm is computationally efficient and in test examples produces composite classifiers that are sparser and generalize as well those produced by Adaboost. The algorithm can be viewed as a coordinate descent method for the  $l_1$ -regularized Adaboost exponential loss function.

*Index Terms*— Pattern classification, Algorithms, Signal representations, Optimization methods.

# 1. INTRODUCTION AND OUTLINE

AdaBoost [1] is used as a pattern classification algorithm in a wide variety of signal processing applications [2]. AdaBoost constructs a powerful composite classifier as a weighted linear combination of a large set of base classifiers. Although the discriminant power of a single base classifier is usually weak, the composite classifier can achieve an acceptable classification accuracy.

In addition to classification accuracy, sparsity of the composite classifier is a desirable attribute. By this we mean that relatively few base classifiers are assigned a nonzero weight in the linear combination. A sparse classifier is easier to store, process, and interpret and, most importantly, is less prone to over fitting. Through empirical studies, AdaBoost is known to be generally resistant to over fitting, even for a large number of iterations [3, 4]. However, we will show that there are additional gains to be made by adding a sparsity mechanism directly into the boosting algorithm.

Boosting algorithms can be interpreted as iterative gradient descent procedures that minimize a loss function on the training data [5, 6, 7]. In many cases, these schemes add at most one new weight per iteration. Hence early stopping of the boosting process is a simple method to ensure sparsity [8]. It has been shown that early stopping of AdaBoost approximately minimizes an exponential loss function subject to an  $l_1$  constraint on the coefficient vector [9]. This suggests that sparsity can be ensured by imposing  $l_1$ -regularization on the optimization of the loss function. This is in accord with results in compressed sensing, where  $l_1$ -regularization has proved an effective approximation to an  $l_0$  sparsity constraint [10]. The idea of  $l_1$ -regularized loss minimization has been explored in [9]. However, direct solution of the convex  $l_1$ -regularized loss problem is computationally too expensive in many real applications. This has lead to proposals for indirect methods of solution. To this end,  $\epsilon$ -boosting seeks to solve the regularized problem iteratively adding a small weight  $\epsilon$  to one base classifier each iteration [5]. However,  $\epsilon$ -boosting is too inefficient for practical application. Other work has examined combining Adaboost with smaller  $l_1$ -regularized loss optimization problems from the perspective of maximizing the margin on the training examples [11]. Despite this work, there remains a need for an efficient boosting algorithm that directly incorporates a sparsity mechanism.

This paper explores the idea of changing two coordinates in reverse directions to incorporate sparsity into coordinate descent boosting procedures. We demonstrate this idea by a new algorithm called RegBoost (short for *Regularized Boost*), which works in a similar fashion to AdaBoost except that it incorporates this simple, intuitive mechanism which is amenable to  $l_1$ -regularization. Moreover, with one simple change, RegBoost reverts to Adaboost. Our empirical studies show that RegBoost achieves similar generalization to AdaBoost but with sparser composite classifiers.

In Section 2, we introduce some basic notations and review the AdaBoost algorithm. We then present the RegBoost algorithm in Section 3 and examine its performance in Section 4, followed by conclusions in Section 5.

### 2. PRELIMINARIES

In a typical binary classification problem, one is given training data  $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$  (where  $\mathbf{x}_i \in \mathbb{R}^p$  are instances and  $y_i \in \{-1, +1\}$  are corresponding labels) and a set of base classifiers  $\mathcal{H} = \{h_j\}_{j=1}^{|\mathcal{H}|}$ . A base classifier  $h_j \in \mathcal{H}$ is a mapping from instances to possible labels,  $h_j : \mathbb{R}^p \mapsto$  $\{-1, +1\}$ . The performance of any classifier h under probability distribution  $\mathbf{w} = \{w_i\}_{i=1}^m$  is measured by its *edge*:  $edge(h, \mathbf{w}) = \sum_{i=1}^m w_i y_i h(x_i)$ , which is related to the error probability  $P_{err}(h)$  with respect to  $\mathbf{w}$  on the training set by  $edge(h, \mathbf{w}) = 1 - 2P_{err}(h)$ 

Let  $h_c(x) = \operatorname{sign}(\sum_{j=1}^{|\mathcal{H}|} \alpha_j h_j(x))$  denote the *composite* classifier. We assume that  $h \in \mathcal{H} \Rightarrow -h \in \mathcal{H}$  and require  $\alpha_j \ge 0, j = 1, \ldots, |\mathcal{H}|$ . Using this notation, AdaBoost can be specified as follows:

# Algorithm 1 AdaBoost1: $\vec{\alpha} \leftarrow \mathbf{0}$ 2: $w_i \leftarrow \frac{1}{m}, \quad i = 1, 2, \dots, m$ 3: for t = 1 to T do4: $k \leftarrow \arg\max_j edge(h_j, \mathbf{w})$ 5: $\varepsilon = \frac{1}{2} \ln \frac{1+edge(h_k, \mathbf{w})}{1-edge(h_k, \mathbf{w})} = \frac{1}{2} \ln \frac{1-P_{err}(h_k)}{P_{err}(h_k)}$ 6: $\alpha_k \leftarrow \alpha_k + \varepsilon$ 7: $w_i \leftarrow \begin{cases} w_i e^{\varepsilon} & \text{if } y_i h_k(x_i) = -1 \\ w_i e^{-\varepsilon} & \text{if } y_i h_k(x_i) = 1 \end{cases}$ 8: $\mathbf{w} \leftarrow \frac{\mathbf{w}}{||\mathbf{w}||_1}$ 9: end for10: Result: $h_c(x) = \operatorname{sign}(\sum_{j=1}^{|\mathcal{H}|} \alpha_j h_j(x))$

The margin of the classifier on example *i* is  $\mu_i = \sum_{j=1}^{|\mathcal{H}|} y_i \alpha_j h_j(x_i) / \|\vec{\alpha}\|_1$ . There is empirical and theoretical evidence that AdaBoost achieves good generalization by enlarging the margins on the training examples [12]. Maximizing the margins is closely related to minimizing the following exponential loss function:

$$\mathcal{L}(\vec{\alpha}) = \sum_{i=1}^{m} \exp(-y_i \sum_{j=1}^{|\mathcal{H}|} \alpha_j h_j(x_i))$$
(1)

In [6] and [7] it is shown that AdaBoost is equivalent to a coordinate descent procedure for minimizing (1). Moreover, in [9] it is shown that early stopping of AdaBoost approximately minimizes (1) subject to an  $l_1$  constraint on the coefficient vector  $\vec{\alpha}$ .

In summary, the basic mechanism of AdaBoost is currently understood in terms of minimizing the exponential loss (1) using coordinate descent and this in turn is closely related to enlarging the margins. Minimization of the loss function subject to  $l_1$ -regularization has also been considered [5, 11]. However, obtaining an efficient algorithm to solve the regularized problem remains a problem under study.

### 3. THE REGBOOST ALGORITHM

With the motivation provided above, we aim to efficiently solve the regularized loss minimization problem:

min: 
$$\mathcal{L}(\vec{\alpha}) = \sum_{i=1}^{m} \exp(-y_i \sum_{j=1}^{|\mathcal{H}|} \alpha_j h_j(x_i))$$
 (2)

s.t.: 
$$\sum_{j=1}^{|\mathcal{H}|} |\alpha_j| \le R \tag{3}$$

Our proposed algorithm, RegBoost, for obtaining a solution is defined as follows:

# Algorithm 2 RegBoost

1:  $\vec{\alpha} \leftarrow \vec{\alpha}_{0}$ 2:  $w_{i} = \exp(-y_{i} \sum_{j=1}^{|\mathcal{H}|} \alpha_{j}h_{j}(x_{i}))$ 3:  $\mathbf{w} \leftarrow \frac{\mathbf{w}}{||\mathbf{w}||_{1}}$ 4: **for** t = 1 to T **do** 5:  $h_{k} = \arg\max_{j} edge(h_{j}, \mathbf{w})$ 6:  $h_{l} = \arg\min_{j:\alpha_{j}>0} edge(h_{j}, \mathbf{w})$ 7:  $\varepsilon = \min(2\alpha_{l}, \frac{1}{2} \ln \frac{P(yh_{k}(x) > yh_{l}(x))}{P(yh_{k}(x) < yhl(x))})$ 8:  $\alpha_{k} \leftarrow \alpha_{k} + \frac{\varepsilon}{2}$   $\alpha_{l} \leftarrow \alpha_{l} - \frac{\varepsilon}{2}$ 9:  $w_{i} \leftarrow \begin{cases} w_{i}e^{\varepsilon} \text{ if } y_{i}h_{k}(x_{i}) < y_{i}h_{l}(x_{i}) \\ w_{i} e^{-\varepsilon} \text{ if } y_{i}h_{k}(x_{i}) > y_{i}h_{l}(x_{i}) \\ w_{i} \text{ if } y_{i}h_{k}(x_{i}) = y_{i}h_{l}(x_{i}) \end{cases}$ 10:  $\mathbf{w} \leftarrow \frac{\mathbf{w}}{||\mathbf{w}||_{1}}$ 11: **end for** 12: Result:  $h_{c}(x) = \operatorname{sign}(\sum_{j=1}^{|\mathcal{H}|} \alpha_{j}h_{j}(x))$ 

Because the loss function (1) decreases when we scale up  $\vec{\alpha}$ , (3) must be an equality at an optimum solution. Therefore we start with an initial point  $\vec{\alpha}_0$  on the boundary of the feasible set ( $\|\vec{\alpha}_0\|_1 = R$ ). Such initialization can be obtained by using AdaBoost until  $\|\vec{\alpha}\|_1$  reaches R or (more naively) by assigning the regularization budget R to a single classifier.

Unlike AdaBoost, which only "adds" base classifiers iteratively, a key component of RegBoost is that it can also "remove" classifiers at the same time. In each iteration, RegBoost not only chooses a best base classifier  $h_k$  (step 5), but also chooses a worst active (meaning with nonzero coefficient) classifier  $h_l$  (step 6). The  $l_1$  budget on the coefficients is then transferred from the worst classifier  $h_l$  to the best classifier  $h_k$  (step 8), while keeping  $\sum_{j=1}^{|\mathcal{H}|} |\alpha_j| = R$ .

This major difference ensures that RegBoost will produce a sparser composite classifier than AdaBoost. The reason for this is twofold: First, the new algorithm promptly removes poor classifiers and relocates resources to better classifiers. This results in fewer active classifiers, with each single classifier being more effective. Second, the algorithm seeks a solution of the  $l_1$ -regularized loss minimization problem (1) (We will prove this below.). The  $l_1$ -regularization will impose sparsity on the solution and minimizing the loss will ensure good generalization error [9, 12].

Other than this key difference, and the changes it necessitates, RegBoost is highly analogous to AdaBoost. This minimal change ensures that RegBoost inherits the simplicity and efficiency of AdaBoost while directly incorporating a sparsity mechanism. In this sense it is a natural extension of Adaboost. Indeed, each iteration step of AdaBoost can be viewed as a degenerate case of an iteration step of RegBoost . To see this, first note that the AdaBoost step length is  $\varepsilon = \frac{1}{2} \ln \frac{1-P_{err}(h_k)}{P_{err}(h_k)}$ . This reflects the value of the newly chosen classifier. Similarly, in RegBoost , the step length  $\frac{1}{2} \ln \frac{P(yh_k(x) > yh_l(x))}{P(yh_k(x) < yh_l(x))}$  can be viewed as a measure of how

better  $h_k$  is compared to  $h_l$ . The important point is that these two step sizes agree when  $h_l = -h_k$ , in which case we have:

$$\frac{1}{2} \ln \frac{P(yh_k(x) > yh_l(x))}{P(yh_k(x) < yh_l(x))} = \frac{1}{2} \ln \frac{P(yh_k(x) = 1)}{P(yh_k(x) = -1)} = \frac{1}{2} \ln \frac{1 - P_{err}(h_k)}{P_{err}(h_k)} \quad (4)$$

Now, if we remove the requirement  $\alpha_j > 0$  (meaning  $h_j$  does not have to be an *active* classifier) in step 6 of RegBoost, then the worst classifier is just the negation of the best classifier:  $h_l = -h_k$ . In this case,  $\varepsilon$  is the same as that of AdaBoost and all the adjustments in steps 8 and 9 reduce to a normal AdaBoost iteration.

It should also be noted that RegBoost has the same theoretical foundation as AdaBoost: it is a "coordinate descent" procedure for solving (2,3). To illustrate this, consider all possible adjustments along the coordinate directions from point  $\vec{\alpha}$ . For any two indices  $j_1, j_2$ , We can subtract a/2 from coefficient  $\alpha_{j_1}$  and add it to  $\alpha_{j_2}$  to keep  $\sum_{j=1}^{|\mathcal{H}|} |\alpha_j| = R$ . Denote the value of the loss function after such adjustment as:

$$\mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, a) = \mathcal{L}(\alpha_1, \dots, \alpha_{j_1} - \frac{a}{2}, \dots, \alpha_{j_2} + \frac{a}{2}, \dots, \alpha_{|\mathcal{H}|})$$

Notice that  $\mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, 0) = \mathcal{L}(\vec{\alpha}).$ 

The following theorem shows that RegBoost operates on a diagonal of the two-coordinate plane yielding the largest gradient descent. Moreover, the step size is chosen to achieve the minimum along that direction.

**Theorem 1.** In each iteration of Algorithm 2, the choice of k, l and  $\varepsilon$  satisfy the following properties:

$$(l,k) = \arg\max_{(j_1,j_2):\alpha_{j_1}>0} \left( -\frac{\partial \mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, a)}{\partial a} \Big|_{a=0} \right)$$
(5)

$$\varepsilon = \underset{a:a/2 \le \alpha_l}{\arg \min} \mathcal{L}_{l \to k}(\vec{\alpha}, a) \tag{6}$$

*Proof.* First, the weight distribution on each example is always proportional to the exponential loss of the example:

$$w_i = \frac{1}{Z} \exp\left(-y_i \sum_{j=1}^{|\mathcal{H}|} \alpha_j h_j(x_i)\right) \tag{7}$$

where Z is a normalization factor. Equation (7) is clearly true for the initialization step (step 2). Every time we remove  $\frac{\varepsilon}{2}$ from  $\alpha_l$  to  $\alpha_k$ , the weight adjustments in step 9 guarantee that equation (7) continues to hold.

The loss function can now be written as:  $\mathcal{L}(\vec{\alpha}) = \sum_{i=1}^{m} Zw_i$ . We then divide the sum  $\sum_{i=1}^{m}$  into four groups, according to whether  $(y_i h_{j_1}(x_i), y_i h_{j_2}(x_i)) = (1, 1), (-1, 1), (1, -1)$  or (-1, -1). These cases correspond to whether  $h_{j_1}$  and  $h_{j_2}$  are right or wrong on an example. After substituting

 $\alpha_{j_1}$  with  $\alpha_{j_1} - a/2$ , and  $\alpha_{j_2}$  with  $\alpha_{j_2} + a/2$ , we obtain

$$\mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, a) = \sum_{\substack{y_i h_{j_1}(x_i) = y_i h_{j_2}(x_i) \\ y_i h_{j_1}(x_i) = 1 \\ y_i h_{j_2}(x_i) = -1 }} Ze^a w_i + \sum_{\substack{i: y_i h_{j_1}(x_i) = -1 \\ y_i h_{j_2}(x_i) = 1 }} Ze^{-a} w_i$$
(8)

So

$$-\frac{\partial \mathcal{L}_{j_1 \to j_2}(\vec{a}, a)}{\partial a} = Ze^{-a}P(yh_{j_1} = -1, yh_{j_2} = 1)$$
(9)  
$$-Ze^aP(yh_{j_1} = 1, yh_{j_2} = -1)$$

Letting a = 0 yields:

$$\begin{aligned} & -\frac{\partial \mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, a)}{\partial a} \Big|_{a=0} \\ &= ZP(yh_{j_1} = -1, yh_{j_2} = 1) - ZP(yh_{j_1} = 1, yh_{j_2} = -1) \\ &= ZP(yh_{j_1} = -1) - ZP(yh_{j_2} = -1) \\ &= \frac{Z}{2}edge(h_{j_2}, \mathbf{w}) - \frac{Z}{2}edge(h_{j_1}, \mathbf{w}) \end{aligned}$$

To maximize this quantity under the constraint of  $\alpha_{j_1} > 0$ , we must choose  $j_2 = \arg \max_j edge(h_j, \mathbf{w}), j_1 = \arg \min_{j:\alpha_j>0} edge(h_j, \mathbf{w})$ . This is exactly the way we choose l and k in the algorithm. So equation (5) is proved.

From equation (9), with  $j_1 = l, j_2 = k$ ,  $\mathcal{L}_{j_1 \to j_2}(\vec{\alpha}, a)$  achieves minimum when  $Ze^{-a}P(yh_{j_1} = -1, yh_{j_2} = 1) - Ze^aP(yh_{j_1} = 1, yh_{j_2} = -1) = 0$ . Therefore

$$a = \frac{1}{2} \ln \frac{ZP(yh_l(x) = -1, yh_k(x) = 1)}{ZP(yh_l(x) = 1, yh_k(x) = -1)}$$
  
=  $\frac{1}{2} \ln \frac{P(yh_k(x) > yh_l(x))}{P(yh_k(x) < yh_l(x))}$  (10)

That is how  $\varepsilon$  is selected in the algorithm, with the additional constraint  $\varepsilon/2 \le \alpha_l$  to ensure all coefficients remain positive. This proves (6).

### 4. EXPERIMENTS

In all experiments, we used 100 training examples and 500 testing examples. We used single stumps as base classifiers. We first compare RegBoost with AdaBoost using a benchmark data set *ringnorm*. The ringnorm examples are generated from two 20-dimensional Gaussian distributions, one  $N(\mathbf{0}, 4I)$ , the other  $N(\mu, I)$  with  $\mu = (a, a, \dots, a)$  and  $a = 1/\sqrt{20}$ . We performed two experiments; each with 20 trials.

In experiment 1, we used AdaBoost to initialize RegBoost by first running AdaBoost until  $\sum_k \alpha_k = R$ , then running RegBoost from that initial condition. The results, Fig1(a) and 1(c), show that RegBoost sparsifies the result of AdaBoost while continuing to lower the test error. For fair comparison, we also included the generalization results for *clipped* 



Fig. 1. Comparison of RegBoost and AdaBoost on ringnorm

AdaBoost, which simply drops the base classifiers with the smallest coefficients in the AdaBoost ensemble in order to match the number of active classifiers in RegBoost . Fig1(a) shows that this naive method of obtaining sparsity sacrifices classification accuracy. In experiment 2, we used an even simpler initialization of RegBoost (assign all weight R to a single classifier) and set R = 100. The results, Fig1(b) and 1(d), show that RegBoost recovers from this naive initialization and eventually surpasses AdaBoost in both sparsity and accuracy.

We then tested RegBoost on four data sets from the UCI repository: *German, Heart, Sonar* and *Spam.* For each data set we are interested in the best error rate achieved before over fitting occurs. In each experiment, we averaged the results of 10 trials. We ran 500 iterations for each trial. We used the naive initialization strategy for RegBoost with R = 40 (larger R yield the same result). The results are shown in Table 1. In all data sets, RegBoost learns a sparser classifier with comparable generalization performance.

As expected, we found little difference in running times between two algorithms. The only additional computation in a RegBoost iteration is searching for a worst active classifier among 100 or so classifiers (Table 1). Since both algorithms have to search over 10000 or so weak classifiers in the first place, this increment is negligible.

Table 1. Comparison of RegBoost and AdaBoost on real data

	RegBoost		AdaBoost	
	error rate	number	error rate	number
German	24.9%	47	25.3%	60
Heart	18.1%	11	19.2%	16
Sonar	12.1%	52	13.8%	157
Spam	10.7%	26	10.7%	59

## 5. CONCLUSION

The RegBoost boosting algorithm has been introduced to address the need for an efficient solution method for the  $l_1$ regularized loss minimization problem. We provided both theoretical and experimental evidence that RegBoost produces sparse classifiers without compromising classification accuracy or computational efficiency. The algorithm is robust to the selection of the regularization parameter R and outperforms AdaBoost even when started with very naive initializations. When started from an initial condition produced by Adaboost, it actually sparsifies the Adaboost solution *and* improves generalization. The sparsity mechanism behind RegBoost can potentially be applied to other coordinate descent algorithms.

### 6. REFERENCES

- Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *European Conference on Computational Learning Theory*, 1995, pp. 23–37.
- [2] Y. Freund and R. E. Schapire, "Experiments with a new boosting algorithm," in *International Conference* on Machine Learning, 1996, pp. 148–156.
- [3] H. Drucker and C. Cortes, "Boosting decision trees," in Advances in Neural Information Processing Systems 8: Proceedings of the 1995 Conference. MIT Press, 1996.
- [4] L. Breiman, "Arcing classifiers," Annals of Statistics, vol. 26, pp. 801–823, 1998.
- [5] J.H. Friedman, "Greedy function approximation: A gradient boosting machine," *Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
- [6] G. Rätsch, T. Onoda, and K.R. Müller, "Soft margins for adaboost," *Machine Learning*, vol. 42, no. 3, pp. 287–320, 2001.
- [7] L. Breiman, "Prediction games and arcing algorithms," *Neural Comp.*, vol. 11, no. 7, pp. 1493–1517, 1999.
- [8] W. Jiang, "Process consistency for adaboost," Annals of Statistics, vol. 32, no. 1, pp. 13–29, 2004.
- [9] S. Rosset, J. Zhu, and T. Hastie, "Boosting as a regularized path to a maximum margin classifier," *Journal of Machine Learning Research*, vol. 5, pp. 941–973, 2004.
- [10] E.J. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [11] Y. T. Xi, Z. J. Xiang, P. J. Ramadge, and R. E. Schapire, "Speed and sparsity of regularized boosting," in *12th International Conference on Artificial Intelligence and Statistics*, April 2009.
- [12] R.E. Schapire, Y. Freund, P. Bartlett, and W.S. Lee, "Boosting the margin: A new explanation for the effectiveness of voting methods," *Annals of Statistics*, vol. 26, pp. 1651–1686, 1998.