Additive Logistic Regression: a Statistical View of Boosting

JEROME FRIEDMAN [*]
TREVOR HASTIE [*]
ROBERT TIBSHIRANI [†]

July 23, 1998

Abstract

Boosting (Freund & Schapire 1996, Schapire & Singer 1998) is one of the most important recent developments in classification methodology. The performance of many classification algorithms often can be dramatically improved by sequentially applying them to reweighted versions of the input data, and taking a weighted majority vote of the sequence of classifiers thereby produced. We show that this seemingly mysterious phenomenon can be understood in terms of well known statistical principles, namely additive modeling and maximum likelihood. For the two-class problem, boosting can be viewed as an approximation to additive modeling on the logistic scale using maximum Bernoulli likelihood as a criterion. We develop more direct approximations and show that they exhibit nearly identical results to that of boosting. Direct multi-class generalizations based on multinomial likelihood are derived that exhibit performance comparable to other recently proposed multi-class generalizations of boosting in most situations, and far superior in some. We suggest a minor modification to boosting that can reduce computation, often by factors of 10 to 50. Finally, we apply these insights to produce an alternative formulation of boosting decision trees. This approach, based on best-first truncated tree induction, often leads to better performance, and can provide interpretable descriptions of the aggregate decision rule. It is also much faster computationally making it more suitable to large scale data mining applications.

[*]Department of Statistics, Sequoia Hall, Stanford University, Stanford California 94305; {jhf,trevor}@stat.stanford.edu

[†]Department of Public Health Sciences, and Department of Statistics, University of Toronto; tibs@utstat.toronto.edu

1 Introduction

The starting point for this paper is an interesting procedure called "boosting", which is a way of combining or boosting the performance of many "weak" classifiers to produce a powerful "committee". Boosting was proposed in the machine learning literature (Freund & Schapire 1996) and has since received much attention.

While boosting has evolved somewhat over the years, we first describe the most commonly used version of the $AdaBoost$ procedure (Freund & Schapire 1996), which we call "Discrete" AdaBoost. Here is a concise description of AdaBoost in the two-class classification setting. We have training data $(x_1, y_1), \ldots (x_N, y_N)$ with x_i a vector valued feature and $y_i = -1$ or 1. We define $F(x) = \sum_1^M c_m f_m(x)$ where each $f_m(x)$ is a classifier producing values ± 1 and c_m are constants; the corresponding prediction is $\text{sign}(F(x))$. The AdaBoost procedure trains the classifiers $f_m(x)$ on weighted versions of the training sample, giving higher weight to cases that are currently misclassified. This is done for a sequence of weighted samples, and then the final classifier is defined to be a linear combination of the classifiers from each stage. We describe the procedure in more detail in Algorithm 1

Discrete AdaBoost(Freund & Schapire 1996)

1. Start with weights $w_i = 1/N$, $i = 1, \ldots, N$.

2. Repeat for $m = 1, 2, \ldots, M$:

    (a) Estimate the classifier $f_m(x)$ from the training data with weights w_i.

    (b) Compute $e_m = E_w[1_{(y \neq f_m(x))}]$, $c_m = \log((1 - e_m)/e_m)$.

    (c) Set $w_i \leftarrow w_i \exp[c_m \cdot 1_{(y_i \neq f_m(x_i))}]$, $i = 1, 2, \ldots N$, and renormalize so that $\sum_i w_i = 1$.

3. Output the classifier $\text{sign}[\sum_{m=1}^M c_m f_m(x)]$

Algorithm 1: *E_w is the expectation with respect to the weights $w = (w_1, w_2, \ldots w_n)$. At each iteration AdaBoost increases the weights of the observations misclassified by $f_m(x)$ by a factor that depends on the weighted training error.*

Much has been written about the success of AdaBoost in producing accurate classifiers. Many authors have explored the use of a tree-based classifier for $f_m(x)$ and have demonstrated that it consistently produces significantly

lower error rates than a single decision tree. In fact, Breiman (NIPS workshop, 1996) called AdaBoost with trees the "best off the shelf classifier in the world". Interestingly, the test error seems to consistently decrease and then level off as more classifiers are added, rather than ultimately increase. For some reason, it seems that AdaBoost is immune to overfitting.

Figure 1 shows the performance of Discrete AdaBoost on a synthetic classification task, using a adaptation of CART™(Breiman, Friedman, Olshen & Stone 1984) as the base classifier. This adaptation grows fixed-size trees in a "best-first" manner (see Section 7). Included in the figure is the *bagged* tree (Breiman 1996) which averages trees grown on bootstrap resampled versions of the training data. Bagging is purely a variance-reduction technique, and since trees tend to have high variance, bagging often produces good results.

Early versions of AdaBoost used a resampling scheme to implement step 2 of Algorithm 1, by weighted importance sampling from the training data. This suggested a connection with bagging, and that a major component of the success of boosting has to do with variance reduction.

However, boosting performs comparably well when:

- a weighted tree-growing algorithm is used in step 2 rather than weighted resampling, where each training observation is assigned its weight w_i. This removes the randomization component essential in bagging.

- "stumps" are used for the weak learners. Stumps are single-split trees with only two terminal nodes. These typically have low variance but high bias. Bagging performs very poorly with stumps (Fig. 1[top-right panel].)

These observations suggest that boosting is capable of both bias and variance reduction, and thus differs fundamentally from bagging.

The *base classifier* in Discrete AdaBoost produces a classification rule $f_m(x) : \mathcal{X} \mapsto \{-1, 1\}$, where $\mathcal{X}$ is the domain of the predictive features x. If the implementation of the base classifier cannot deal with observation weights, weighted resampling is used instead. Freund & Schapire (1996) and Schapire & Singer (1998) have suggested various modifications to improve the boosting algorithms; here we focus on a version due to Schapire & Singer (1998), which we call "Real AdaBoost", that uses real-valued "confidence-rated" predictions rather than the $\{-1, 1\}$ of Discrete AdaBoost. The base classifier for this generalized boosting produces a mapping $f_m(x) : \mathcal{X} \mapsto R$; the sign of $f_m(x)$ gives the classification, and $|f_m(x)|$ a measure of the "confidence" in the prediction. This real-valued boosting tends to perform the

Figure 1: *Test error for Bagging (BAG), Discrete AdaBoost (DAB) and Real AdaBoost (RAB) on a simulated two-class nested spheres problem (see Section 5.) There are 2000 training data points in 10 dimensions, and the Bayes error rate is zero. All trees are grown "best-first" without pruning. The left-most iteration corresponds to a single tree.*

best in our simulated examples in Fig. 1, especially with stumps, although we see with 100 node trees Discrete AdaBoost overtakes Real AdaBoost after 200 iterations.

Real AdaBoost(Schapire & Singer 1998)

1. Start with weights $w_i = 1/N$, $i = 1, 2, \ldots, N$.

2. Repeat for $m = 1, 2, \ldots, M$:

   (a) Estimate the "confidence rated" classifier $f_m(x) : \mathcal{X} \mapsto R$ and the constant c_m from the training data with weights w_i.

   (b) Set $w_i \leftarrow w_i \exp[-c_m \cdot y_i f_m(x_i)]$, $i = 1, 2, \ldots N$, and renormalize so that $\sum_i w_i = 1$.

3. Output the classifier sign$[\sum_{m=1}^{M} c_m f_m(x)]$

Algorithm 2: *The Real AdaBoost algorithm allows for the estimator $f_m(x)$ to range over R. In the special case that $f_m(x) \in \{-1, 1\}$ it reduces to AdaBoost, since $y_i f_m(x_i)$ is 1 for a correct and-1 for an incorrect classification. In the general case the constant c_m is absorbed into $f_m(x)$. We describe the Schapire-Singer estimate for $f_m(x)$ in Section 3.*

Freund & Schapire (1996) and Schapire & Singer (1998) provide some theory to support their algorithms, in the form of upper bounds on generalization error. This theory (Schapire 1990) has evolved in the machine learning community, initially based on the concepts of PAC learning (Kearns & Vazirani 1994), and later from game theory (Freund 1995, Breiman 1997). Early versions of boosting "weak learners" (Schapire 1990) are far simpler than those described here, and the theory is more precise. The bounds and the theory associated with the AdaBoost algorithms are interesting, but tend to be too loose to be of practical importance. In practice boosting achieves results far more impressive than the bounds would imply.

In this paper we analyze the AdaBoost procedures from a statistical perspective. We show that the underlying model they are fitting is additive logistic regression. The AdaBoost algorithms are Newton methods for optimizing a particular exponential loss function — a criterion which behaves much like the log-likelihood on the logistic scale. We also derive new boosting-like procedures for classification.

In Section 2 we briefly review additive modelling. Section 3 shows how boosting can be viewed as an additive model estimator, and proposes some new boosting methods for the two class case. The multiclass problem is

studied in Section 4. Simulated and real data experiments are discussed in Sections 5 and 6. Our tree-growing implementation, using truncated best-first trees, is described in Section 7. Weight trimming to speed up computation is discussed in Section 8, and we end with a discussion in Section 9.

2   Additive Models

AdaBoost produces an *additive* model $F(x) = \sum_{m=1}^{M} c_m f_m(x)$, although "weighted committee" or "ensemble" sound more glamorous. Additive models have a long history in statistics, and we give some examples here.

2.1   Additive Regression Models

We initially focus on the regression problem, where the response y is quantitative, and we are interested in modeling the mean $E(Y|x) = F(x)$. The additive model has the form

$$F(x) = \sum_{j=1}^{p} f_j(x_j).$$ (1)

Here there is a separate function $f_j(x_j)$ for each of the p input variables x_j. More generally, each component f_j is a function of a small, pre-specified subset of the input variables. The *backfitting algorithm* (Friedman & Stuetzle 1981, Buja, Hastie & Tibshirani 1989) is a convenient modular algorithm for fitting additive models. A backfitting update is

$$f_j(x_j) \leftarrow E\left[y - \sum_{k \neq j} f_k(x_k)|x_j\right].$$ (2)

Any method or algorithm for estimating a function of x_j can be used to obtain an estimate of the conditional expectation in (2). In particular, this can include nonparametric *smoothing* algorithms, such as local regression or smoothing splines. In the right hand side, all the latest versions of the functions f_k are used in forming the partial residuals. The backfitting cycles are repeated until convergence. Under fairly general conditions, backfitting can be shown to converge to the minimizer of $E(y - F(x))^2$ (Buja et al. 1989).

2.2   Extended Additive Models

More generally, one can consider additive models whose elements $\{f_m(x)\}_1^M$ are functions of potentially all of the input features x. Usually, in this

context, the $f_m(x)$ are taken to be simple functions characterized by a set of parameters γ and a multiplier β_m,

$$f_m(x) = \beta_m b(x\,;\gamma_m). \tag{3}$$

The additive model then becomes

$$F_M(x) = \sum_{m=1}^{M} \beta_m b(x\,;\gamma_m). \tag{4}$$

For example, in single hidden layer neural networks $b(x\,;\gamma) = \sigma(\gamma^t x)$ where $\sigma(\cdot)$ is a sigmoid function and γ parameterizes a linear combination of the input features. In signal processing, wavelets are a popular choice with γ parameterizing the location and scale of a "mother" wavelet $b(x\,;\gamma)$. In these applications $\{b(x\,;\gamma_m)\}_1^M$ are generally called "basis functions" since they span a function subspace.

If least–squares is used as a fitting criterion, one can solve for an optimal set of parameters through a generalized back–fitting algorithm with updates

$$\{\beta_m, \gamma_m\} \leftarrow \arg\min_{\beta,\gamma} E\left[y - \sum_{k\neq m}\beta_k b(x\,;\gamma_k) - \beta b(x\,;\gamma)\right]^2. \tag{5}$$

Alternatively, one can use a "greedy" forward stepwise approach

$$\{\beta_m, \gamma_m\} \leftarrow \arg\min_{\beta,\gamma} E\left[y - F_{m-1}(x) - \beta b(x\,;\gamma)\right]^2 \tag{6}$$

where $\{\beta_k, \gamma_k\}_1^{m-1}$ are fixed at their corresponding solution values at earlier iterations. This is the approach used by Mallat & Zhang (1993) in "matching pursuit". There the $b(x\,;\gamma)$ represent an over complete wavelet–like basis. In the language of boosting, $f(x) = \beta b(x\,;\gamma)$ would be called a "weak learner" and $F_M(x)$ (4) the "committee". If decision trees were used as the weak learner the parameters γ would represent the splitting variables, split points, the constants in each terminal node, and number of terminal nodes of each tree.

Note that the back–fitting procedure (5) or its greedy cousin (6) only require an algorithm for fitting a *single* weak learner (3) to data. This base algorithm is simply applied repeatedly to modified versions of the original data

$$y \leftarrow y - \sum_{k\neq m} f_k(x).$$

In the forward stepwise procedure (6) the modified output at the mth iteration y_m only depends on its value y_{m-1} and the solution $f_{m-1}(x)$ at the previous iteration

$$y_m = y_{m-1} - f_{m-1}(x). \tag{7}$$

At each step m, the previous output values y_{m-1} are modified (7) so that the previous model $f_{m-1}(x)$ has no explanatory power on the new outputs y_m. One can therefore view this as a procedure for boosting a weak learner $f_1(x) = \beta_1 b(x\,;\gamma_1)$ to form a powerful committee $F_M(x)$ (4).

2.3 Classification problems

For the classification problem, we learn from Bayes theorem that all we need is $P(y = j|x)$, the posterior or conditional class probabilities. One could transfer all the above regression machinery across to the classification domain by simply noting that $E(1_{(y=j)}|x) = P(y = j|x)$, where $1_{(y=j)}$ is the $0/1$ indicator variable representing class j. While this works fairly well in general, several problems have been noted (Hastie, Tibshirani & Buja 1994) for constrained regression methods. The estimates are typically not confined to $[0, 1]$, and severe masking problems can occur. A notable exception is when trees are used as the regression method, and in fact this is the approach used by Breiman et al. (1984).

Logistic regression is a popular approach used in statistics for overcoming these problems. For a two class problem, the model is

$$\log \frac{P(y = 1|x)}{P(y = 0|x)} = \sum_{m=1}^{M} f_m(x). \tag{8}$$

The monotone *logit* transformation on the left guarantees that for any values of $F(x) = \sum_{m=1}^{M} f_m(x) \in R$, the probability estimates lie in $[0, 1]$; inverting we get

$$p(x) = P(y = 1|x) = \frac{e^{F(x)}}{1 + e^{F(x)}}. \tag{9}$$

Here we have given a general additive form for $F(x)$; special cases exist that are well known in statistics. In particular, the linear logistic regression model (McCullagh & Nelder 1989, for example) and additive logistic regression model (Hastie & Tibshirani 1990) are popular. These models are usually fit by maximizing the binomial log-likelihood, and enjoy all the associated asymptotic optimality features of maximum likelihood estimation.

A generalized version of backfitting (2), called "Local Scoring" in (Hastie & Tibshirani 1990), is used to fit the additive logistic model. Starting with

guesses $f_1(x_1)\dots f_p(x_p)$, $F(x) = \sum f_k(x_k)$ and $p(x)$ defined in (9), we form the working response:

$$z = F(x) + \frac{y - p(x)}{p(x)(1 - p(x))}. \tag{10}$$

We then apply backfitting to the response z with observation weights $p(x)(1-p(x))$ to obtain new $f_k(x_k)$. This process is repeated until convergence. The forward stage-wise version (6) of this procedure bears a close similarity to the LogitBoost algorithm described later in the paper.

3 Boosting — an Additive Logistic Regression Model

In this Section we show that the boosting algorithms are stage-wise estimation procedures for fitting an additive logistic regression model. They optimize an exponential criterion which to second order is equivalent to the binomial log-likelihood criterion. We then propose a more standard likelihood-based procedure.

3.1 An Exponential Criterion

Consider minimizing the criterion

$$J(F) = E(e^{-yF(x)}) \tag{11}$$

for estimation of $F(x)$.[1]

Lemma 1 shows that the function $F(x)$ that minimizes the L_2 version of the exponential criterion is the symmetric logistic transform of $P(y = 1|x)$

Lemma 1 $E(e^{-yF(x)})$ *is minimized at*

$$F(x) = \frac{1}{2} \log \frac{P(y = 1|x)}{P(y = -1|x)}. \tag{12}$$

Hence

$$P(y = 1|x) = \frac{e^{F(x)}}{e^{-F(x)} + e^{F(x)}} \tag{13}$$

$$P(y = -1|x) = \frac{e^{-F(x)}}{e^{-F(x)} + e^{F(x)}}. \tag{14}$$

---

[1]E represents expectation; depending on the context, this may be an L_2 population expectation, or else a sample average. E_w means a weighted expectation.

Proof

While E entails expectation over the joint distribution of y and x, it is sufficient to minimize the criterion conditional on x.

$$E\left(e^{-yF(x)})|x\right) = P(y=1|x)e^{-F(x)} + P(y=-1|x)e^{F(x)}$$

$$\frac{\partial E\left(e^{-yF(x)})|x\right)}{\partial F(x)} = -P(y=1|x)e^{-F(x)} + P(y=-1|x)e^{F(x)}$$

Setting the derivative to zero the result follows.

$\square$

The usual logistic transform does not have the factor $\frac{1}{2}$ in (12); by multiplying the numerator and denominator in (13) by $e^{F(x)}$, we get the usual logistic model

$$p(x) = \frac{e^{2F(x)}}{1 + e^{2F(x)}} \tag{15}$$

Hence the two models are equivalent.

Corollary 1 *If E is replaced by averages over regions of x where $F(x)$ is constant (as in the terminal node of a decision tree), the same result applies to the sample proportions of $y = 1$ and $y = -1$.*

In proposition 1 we show that the Discrete AdaBoost increments in Algorithm 1 are Newton-style updates for minimizing the exponential criterion. This can be interpreted as a stage-wise estimation procedure for fitting an additive logistic regression model.

Proposition 1 *The Discrete AdaBoost algorithm produces adaptive Newton updates for minimizing $E(e^{-yF(x)})$, which are stage-wise contributions to an additive logistic regression model.*

Proof

Let $J(F) = E[e^{-yF(x)}]$. Suppose we have a current estimate $F(x)$ and seek an improved estimate $F(x) + cf(x)$. For fixed c (and x), we expand $J(F(x) + cf(x))$ to second order about $f(x) = 0$

$$J(F + cf) = E[e^{-y(F(x)+cf(x))}]$$
$$\approx E[e^{-yF(x)}(1 - ycf(x) + c^2 f(x)^2/2)]$$

Minimizing pointwise with respect to $f(x) \in \{-1, 1\}$, we find

$$\hat{f}(x) = \arg\min_{f} E_w(1 - ycf(x) + c^2 f(x)^2/2|x)$$

$$= \arg\min_f E_w[(y - cf(x))^2 | x] \qquad (16)$$

$$= \arg\min_f E_w[(y - f(x))^2 | x] \qquad (17)$$

where $w(y|x) = \exp(-yF(x))/E \exp(-yF(x))$, and (17) follows from (16) by considering the two possible choices for $f(x)$.

Given $\hat{f} \in \{-1, 1\}$, we can directly minimize $J(F + c\hat{f})$ to determine c:

$$\hat{c} = \arg\min_c E_w e^{-cy\hat{f}(x)}$$
$$= \frac{1}{2} \log \frac{1 - e}{e}$$

where $e = E_w[1_{(y \neq \hat{f}(x))}]$. Combining these steps we get the update for $F(x)$

$$F(x) \leftarrow F(x) + \frac{1}{2} \log \frac{1 - e}{e} \hat{f}(x)$$

In the next iteration the new contribution $\hat{c}\hat{f}(x)$ to $F(x)$ augments the weights:

$$w(y|x) \leftarrow w(y|x) \cdot e^{-\hat{c}\hat{f}(x)y},$$

followed by a normalization. Since $y\hat{f}(x) = 2 \times 1_{(y \neq \hat{f}(x))} - 1$, we see that the update is equivalent to

$$w(y|x) \leftarrow w(y|x) \cdot \exp\left(\log\left(\frac{1 - e}{e}\right) 1_{(y \neq \hat{f}(x))}\right)$$

Thus the function and weight updates are identical to those used in Discrete AdaBoost.

$\square$

Parts of this derivation for AdaBoost can be found in Schapire & Singer (1998). Newton algorithms repeatedly optimize a quadratic approximation to a nonlinear criterion, which is the first part of the update in AdaBoost.

This L_2 version of AdaBoost translates naturally to a data version using trees. The weighted least squares criterion is used to grow the tree-based classifier $\hat{f}(x)$, and given $\hat{f}(x)$, the constant c is based on the weighted training error.

Note that after each Newton step, the weights change, and hence the tree configuration will change as well. This adds a nonlinear twist to the Newton algorithm.

Corollary 2 *After each update to the weights, the* weighted *misclassification error of the most recent weak learner is* 50%.

Proof

This follows by noting that the c that minimizes $J(F + cf)$ satisfies

$$\frac{\partial J(F + cf)}{\partial c} = -E[e^{-y(F(x)+cf(x))}yf(x)] = 0 \tag{18}$$

The result follows since $yf(x)$ is 1 for a correct classification, and -1 for a misclassification.

$\square$

Schapire & Singer (1998) give the interpretation that the weights are updated to make the new weighted problem maximally difficult for the next weak learner.

The Discrete AdaBoost algorithm expects the tree or other "weak learner" to deliver a classifier $f(x) \in \{-1, 1\}$. We now show that the Real AdaBoost algorithm uses the weighted probability estimates in the terminal nodes of the tree to update the additive logistic model, rather than the classifications themselves. Again we derive the population algorithm, and then apply it to data.

Proposition 2 *The Real AdaBoost algorithm fits an additive logistic regression model by stage-wise optimization of* $J(F) = E[e^{-yF(x)}]$

Proof

Suppose we have a current estimate $F(x)$ and seek an improved estimate $F(x) + f(x)$ by minimizing $J(F(x) + f(x))$.

$$
\begin{aligned}
\frac{\partial J(F(x) + f(x))}{\partial f(x)} &= -E(e^{-yF(x)}ye^{-yf(x)}|x) \\
&= -E[e^{-yF(x)}1_{(y=1)}e^{-f(x)}|x] + E[e^{-yF(x)}1_{(y=-1)}e^{f(x)}|x]
\end{aligned}
$$

Dividing through by $Ee^{-yF(x)}$ and setting the derivative to zero we get

$$\hat{f}(x) = \frac{1}{2}\log\frac{E_w[1_{(y=1)}|x]}{E_w[1_{(y=-1)}|x]} \tag{19}$$

$$= \frac{1}{2}\log\frac{P_w(y=1|x)}{P_w(y=-1|x)} \tag{20}$$

where $w(y|x) = \exp(-yF(x))/E(\exp(-yF(x))|x)$.

Careful examination of Schapire & Singer (1998) shows that this update matches theirs (and the c_m in Algorithm 2 are redundant.) The weights get updated by

$$w(y|x) \leftarrow w(y|x) \cdot e^{-y\hat{f}(x)}$$

□

Corollary 3 *At the optimal $F(x)$, the weighted conditional mean of y is 0.*

Proof
If $F(x)$ is optimal, we have

$$\frac{\partial J(F(x))}{F(x)} = -Ee^{-yF(x)}y = 0 \tag{21}$$

□

We can think of the weights as providing an alternative to residuals for the binary classification problem. At the optimal function F, there is no further information about F in the weighted conditional distribution of y. If there is, we use it to update F.

At iteration M in either the Discrete or Real AdaBoost algorithms, we have composed an additive function of the form

$$F(x) = \sum_{m=1}^{M} f_m(x) \tag{22}$$

where each of the components are found in a greedy forward stage-wise fashion, fixing the earlier components. Our term "stage-wise" refers to a similar approach in Statistics:

- Variables are included sequentially in a stepwise regression.

- The coefficients of variables already included receive no further adjustment.

3.2   Why $Ee^{-yF(x)}$?

So far the only justification for this exponential criterion is that it has a sensible population minimizer, and the algorithm described above performs well on real data. In addition

- Schapire & Singer (1998) motivate $e^{-yF(x)}$ as a differentiable upper-bound to misclassification error (see Fig. 2);

- the AdaBoost algorithm that it generates is extremely modular, requiring at each iteration the retraining of a classifier on a weighted training database.

Let $y^* = (y + 1)/2$, taking values $0, 1$, and parametrize the binomial probabilities by

$$p(x) = \frac{e^{F(x)}}{e^{F(x)} + e^{-F(x)}}$$

The expected binomial log-likelihood is

$$
\begin{align}
E\ell(y^*, p(x)) &= E[y^* \log(p(x)) + (1 - y^*) \log(1 - p(x))] \tag{23}\\
&= -E \log(1 + e^{-yF(x)}) \tag{24}
\end{align}
$$

- The population minimizers of $-E\ell(y^*, p(x))$ and $Ee^{-yF(x)}$ coincide. In fact, the exponential criterion and the (negative) log-likelihood are equivalent to second order in a Taylor series around $F = 0$:

$$-\ell(y^*, p) \approx \exp(-yF) + \log(2) - 1 \tag{25}$$

Graphs of $\exp(-yF)$ and $\log(1 + e^{-yF(x)})$ (suitably scaled) are shown in Fig. 2, as a function of yF — positive values of yF imply correct classification. Note that $-\exp(-yF)$ itself is not a proper log-likelihood, as it does not equal the log of any probability mass function on ± 1.

- Also shown in Fig. 2 is the indicator function $1_{(F>0)}$, which gives misclassification error.

- There is another way to view the criterion $J(F)$. It is easy to show that

$$e^{-yF(x)} = \frac{|y^* - p(x)|}{\sqrt{p(x)(1 - p(x))}}, \tag{26}$$

with $F(x) = \log(p(x)/(1 - p(x)))$, The right-hand side is known as the *Chi* statistic in the statistical literature.

One feature of both the exponential and log-likelihood criteria is that they are monotone and smooth. Even if the training error is zero, the criteria will drive the estimates towards purer solutions (in terms of probability estimates).

Why not estimate the f_m by minimizing the squared error $E(y - F(x))^2$? If $F_{m-1}(x) = \sum_1^{m-1} f_j(x)$ is the current prediction, this leads to a forward

Figure 2: *A variety of loss functions for estimating a function $F(x)$ for classi-fication. The horizontal axis is yF, which is negative for errors and positive for correct classifications. All the loss functions are monotone in yF. The curve la-beled "Squared Error(p)" is $(y^* - p)^2$, and gives a uniformly better approximation to misclassification loss than the exponential criterion (Schapire-Singer). The curve labeled "Squared Error(F)" is $(y - F)^2$, and increases once yF exceeds 1, thereby increasingly penalizing classifications that are "too correct".*

stage-wise procedure that does an unweighted fit to the response $y - F_{m-1}(x)$ at step m (6). Empirically we have found that this approach works quite well, but is dominated by those that use monotone loss criteria. We believe that the non-monotonicity of squared error loss (Fig. 2) is the reason. Correct classifications ($yF(x) > 1$) incur increasing loss for increasing values of $|F(x)|$. This makes squared-error loss an especially poor approximation to misclassification error rate. Classifications that are "too correct" are penalized as much as misclassification errors.

3.3 Using the log-likelihood criterion

In this Section we explore algorithms for fitting additive logistic regression models by stage-wise optimization of the Bernoulli log-likelihood. Here we focus again on the two-class case, and will use a 0/1 response y^* to represent the outcome. We represent the probability of $y^* = 1$ by $p(x)$, where

$$p(x) = \frac{e^{F(x)}}{e^{F(x)} + e^{-F(x)}} \tag{27}$$

Algorithm 3 gives the details.

LogitBoost (2 classes)

1. Start with weights $w_i = 1/N$ $i = 1, 2, \ldots, N$, $F(x) = 0$ and probability estimates $p_i = \frac{1}{2}$.

2. Repeat for $m = 1, 2, \ldots, M$:

   (a) Compute the working response and weights

   $$
   \begin{aligned}
   z_i &= \frac{y_i^* - p_i}{p_i(1 - p_i)} \\
   w_i &= p_i(1 - p_i)
   \end{aligned}
   $$

   (b) Estimate $f_m(x)$ by weighted least-squares fitting of z to x.
   (c) Update $F(x) \leftarrow F(x) + \frac{1}{2}f_m(x)$ and $p(x)$ via (27).

3. Output the classifier $\text{sign}[F(x)] = \text{sign}[\sum_{m=1}^{M} f_m(x)]$

Algorithm 3: *An adaptive Newton algorithm for fitting an additive logistic regression model.*

16

Proposition 3 *The LogitBoost algorithm uses adaptive Newton steps for fitting an additive symmetric logistic model by maximum likelihood.*

Proof

Consider the update $F(x) + f(x)$ and the expected log-likelihood

$$\ell(F + f) = E[2y^*(F(x) + f(x)) - \log[1 + \exp(2F(x)]. \qquad (28)$$

Conditioning on x, we compute the first and second derivative at $f(x) = 0$:

$$
\begin{aligned}
s(x) &= \left.\frac{\partial \ell(F(x) + f(x)}{\partial f(x)}\right|_{f(x)=0} \\
&= 2E(y^* - p(x)|x) \qquad (29) \\
H(x) &= \left.\frac{\partial^2 \ell(F(x) + f(x)}{\partial f(x)^2}\right|_{f(x)=0} \\
&= -4E(p(x)(1 - p(x))|x) \qquad (30)
\end{aligned}
$$

where $p(x)$ is defined in terms of $F(x)$. The Newton update is then

$$
\begin{aligned}
F(x) &\leftarrow F(x) - H(x)^{-1}s(x) \\
&= F(x) + \frac{1}{2}\frac{E(y^* - p(x)|x)}{Ep(x)(1 - p(x))|x)} \qquad (31) \\
&= F(x) + \frac{1}{2}E_w\left(\frac{y^* - p(x)}{p(x)(1 - p(x))}|x\right) \qquad (32)
\end{aligned}
$$

where $w(x) = p(x)(1 - p(x))$. Equivalently, the Newton update solves the weighted least squares criterion

$$\min_{f(x)} E_{w(x)}\left(\frac{1}{2}\frac{y^* - p(x)}{p(x)(1 - p(x))} - f(x)\right)^2 \qquad (33)$$

$\square$

The population algorithm described here translates immediately to an implementation on data when $E(\cdot|x)$ is replaced by a regression method, such as regression trees (Breiman et al. 1984). While the role of the weights are somewhat artificial in the L_2 case, they are not in any implementation; $w(x)$ is constant when conditioned on x, but the $w(x_i)$ in a terminal node of a tree, for example, depend on the current values $F(x_i)$, and will typically not be constant.

Sometimes the $w(x)$ get very small in regions of (x) perceived (by $F(x)$) to be *pure*—that is, when $p(x)$ is close to 0 or 1. This can cause numerical problems in the construction of z, and led to the following crucial implementation protections:

- If $y^* = 1$, then compute $z = \frac{y^* - p}{p(1-p)}$ as $\frac{1}{p}$. Since this number can get large if p is small, threshold this ratio at $zmax$. The particular value chosen for $zmax$ is not crucial; we have found empirically that $zmax \in [2, 4]$ works well. Likewise, if $y^* = 0$, compute $z = \frac{-1}{(1-p)}$ with a lower threshold of $-zmax$.

- Enforce a lower threshold on the weights: $w = \max(w, 2 \times machine\text{-}zero)$.

3.4 Optimizing $Ee^{-yF(x)}$ by Newton stepping

The L_2 Real Adaboost procedure (Algorithm 2) optimizes $Ee^{-y(F(x)+f(x))}$ exactly with respect to f at each iteration. Here we explore a "gentler" version that instead takes adaptive Newton steps much like the LogitBoost algorithm just described.

Gentle AdaBoost

1. Start with weights $w_i = 1/N$ $i = 1, 2, \ldots, N$, $F(x) = 0$.

2. Repeat for $m = 1, 2, \ldots, M$:

    (a) Estimate $f_m(x)$ by weighted a fit of y to x.

    (b) Update $F(x) \leftarrow F(x) + f_m(x)$

    (c) Update $w_i \leftarrow w_i e^{-y_i f_m(x_i)}$ and renormalize.

3. Output the classifier $\text{sign}[F(x)] = \text{sign}[\sum_{m=1}^{M} f_m(x)]$

Algorithm 4: *A modified version of the Real AdaBoost algorithm, using Newton stepping rather than exact optimization at each step*

Proposition 4 *The Gentle AdaBoost algorithm uses adaptive Newton steps for minimizing $Ee^{-yF(x)}$.*

Proof

$$\frac{\partial J(F(x) + f(x))}{\partial f(x)}\Big|_{f(x)=0} = -E(e^{-yF(x)}y|x)$$

$$\frac{\partial^2 J(F(x) + f(x))}{\partial f(x)^2}\Big|_{f(x)=0} = E(e^{-yF(x)}|x) \text{ since } y^2 = 1$$

18

Hence the Newton update is

$$
\begin{aligned}
F(x) &\leftarrow F(x) + \frac{E(e^{-yF(x)}y}{E(e^{-yF(x)}|x)} \\
&= F(x) + E_w y
\end{aligned}
$$

where

$$
w(y|x) = \frac{e^{-yF(x)}}{E(e^{-yF(x)}|x)}
$$

$\square$

The main difference between this and the Real AdaBoost algorithm is how it uses its estimates of the weighted class probabilities to update the functions. Here the update is $f_m(x) = P_w(y = 1|x) - P_w(y = -1|x)$, rather than half the log-ratio as in (20): $f_m(x) = \frac{1}{2} \log \frac{P_w(y=1|x)}{P_w(y=-1|x)}$. Log-ratios can be numerically unstable, leading to very large updates in pure regions, while the update here lies in the range $[-1, 1]$. Empirical evidence suggests (see Section 6) that this more conservative algorithm has similar performance to both the Real AdaBoost and LogitBoost algorithms, and often outperforms them both, especially when stability is an issue.

Freund & Schapire (1996) also propose an AdaBoost algorithm similar to Gentle Adaboost, except the function $f_m(x)$ is multiplied by a constant $(c_m f_m(x))$ which is estimated in a global fashion much like in Discrete AdaBoost. We do not pursue this particular variant further here.

There is a strong similarity between the updates for the Gentle AdaBoost algorithm and those for the LogitBoost algorithm. Let $p = P(y = 1|x)$, and $p_m = \frac{e^{F(x)}}{e^{F(x)} + e^{-F(x)}}$. Then

$$
\begin{aligned}
\frac{E(e^{-yF(x)}y|x)}{E(e^{-yF(x)}|x)} &= \frac{e^{-F(x)}p - e^{F(x)}(1 - p)}{e^{-F(x)}p + e^{F(x)}(1 - p)} \\
&= \frac{p - p_m}{(1 - p_m)p + p_m(1 - p)} \tag{34}
\end{aligned}
$$

The analogous expression for LogitBoost from (31) is

$$
\frac{1}{2} \frac{p - p_m}{p_m(1 - p_m)} \tag{35}
$$

At $p_m \approx \frac{1}{2}$ these are nearly the same, but they differ as the p_m become extreme. For example, if $p \approx 1$ and $p_m \approx 0$, (35) blows up, while (34) is about 1 (and always falls in $[-1, 1]$.)

4 Multiclass procedures

Here we explore extensions of boosting to classification with multiple classes. We start off by proposing a natural generalization of the two-class symmetric logistic transformation, and then consider specific algorithms. In this context Schapire & Singer (1998) define J responses y_j for a J class problem, each taking values in $\{-1, 1\}$. Similarly the *indicator response vector* with elements y_j^* is more standard in the statistics literature. Assume the classes are mutually exclusive.

Definition 1 *For a J class problem let $P_j(x) = P(y_j = 1|x)$. We define the symmetric multiple logistic transformation*

$$F_j(x) = \log P_j(x) - \frac{1}{J} \sum_{k=1}^{J} \log P_k(x) \tag{36}$$

Equivalently,

$$P_j(x) = \frac{e^{F_j(x)}}{\sum_{k=1}^{J} e^{F_k(x)}}, \quad \sum_{k=1}^{J} F_k(x) = 0 \tag{37}$$

The centering condition in (37) is for numerical stability only; it simply pins the F_j down, else we could add an arbitrary constant to each F_j and the probabilities remain the same. The equivalence of these two definitions is easily established, as well as the equivalence with the two-class case.

Schapire & Singer (1998) provide several generalizations of AdaBoost for the multiclass case; we describe their *AdaBoost.MH* algorithm, since it seemed to dominate the others in their empirical studies. We then connect it to the models presented here. We will refer to the augmented variable in Algorithm 5 as the "class" variable C. We make a few observations:

- The L_2 version of this algorithm minimizes $\sum_{j=1}^{J} Ee^{-y_j F_j(x)}$, which is equivalent to running separate L_2 boosting algorithms on each of the J problems of size N obtained by partitioning the $N \times J$ samples in the obvious fashion. This is seen trivially by first conditioning on $C = j$, and then $x|C = j$, when computing conditional expectations.

- The same is almost true for their tree-based algorithm. We see this because

  1. If the first split is on C — either a J-nary split if permitted, or else $J - 1$ binary splits — then the sub-trees are identical to separate trees grown to each of the J groups. This will always be the case for the first tree.

20

AdaBoost.MH (Schapire & Singer 1998)

The original N observations are expanded into $N \times J$ pairs $((x_i, 1), y_{i1}), ((x_i, 2), y_{i2}), \ldots, ((x_i, J), y_{iJ}), \; i = 1, \ldots, N$.

1. Start with weights $w_{ij} = 1/NJ, \; i = 1, \ldots, N, \; j = 1, \ldots, J$.

2. Repeat for $m = 1, 2, \ldots, M$:

   (a) Estimate the "confidence rated" classifier $f_m(x, j) \; : \; (\mathcal{X} \times (1, \ldots, J)) \mapsto R$ from the training data with weights w_{ij}.

   (b) Set $w_{ij} \leftarrow w_{ij} \exp[-y_{ij} f_m(x_i, j)], \; i = 1, 2, \ldots, N, \; j = 1, \ldots, J$, and renormalize so that $\sum_{i,j} w_{ij} = 1$.

3. Output the classifier $\mathrm{argmax}_j F(x, j)$ where $F(x, j) = \sum_m f_m(x, j)$.

Algorithm 5: *The AdaBoost.MH algorithm converts the J class problem into that of estimating a 2 class classifier on a training set J times as large, with an additional "feature" defined by the set of class labels.*

2. If a tree does not split on C anywhere on the path to a terminal node, then that node returns a function $f_m(x, j) = g_m(x)$ that contributes nothing to the classification decision. However, as long as a tree includes a split on C at least once on every path to a terminal node, it will make a contribution to the classifier for all input feature values.

The advantage/disadvantage of building one large tree using class label as an additional input feature is not clear. No motivation is provided. We therefore implement AdaBoost.MH using the more traditional direct approach of building J separate trees to minimize $\sum_{j=1}^{J} E e^{-y_j F_j(x)}$

We have thus shown

Proposition 5 *The AdaBoost.MH algorithm for a J-class problem fits J uncoupled additive logistic models, $G_j(x) = \frac{1}{2} \log P_j(x)/(1 - P_j(x))$, each class against the rest.*

In principal this parametrization is fine, since $G_j(x)$ is monotone in $P_j(x)$. However, we are estimating the $G_j(x)$ in an uncoupled fashion, and there is no guarantee that the implied probabilities sum to 1. We give some examples where this makes a difference, and AdaBoost.MH performs more poorly than an alternative coupled likelihood procedure.

21

Schapire and Singer's AdaBoost.MH was also intended to cover situations when observations can belong to more than one class. The "MH" represents "Multi-Label Hamming", Hamming loss being used to measure the errors in the space of 2^J possible class labels. In this context fitting a separate classifier for each label is a reasonable strategy. Schapire and Singer also propose using AdaBoost.MH when the class labels are mutually exclusive, which is the focus in this paper.

Algorithm 6 is a natural generalization of algorithm 3 for fitting the J-class logistic regression model (37).

LogitBoost (J classes)

1. Start with weights $w_{ij} = 1/N$, $i = 1, \ldots, N$, $j = 1, \ldots, J$, $F_j(x) = 0$ and $P_j(x) = 1/J \; \forall j$.

2. Repeat for $m = 1, 2, \ldots, M$:

   (a) Repeat for $j = 1, \ldots, J$:

      i. Compute working responses and weights in the jth class

      $$
      \begin{aligned}
      z_{ij} &= \frac{y_{ij}^* - p_{ij}}{p_{ij}(1 - p_{ij})} \\
      w_{ij} &= p_{ij}(1 - p_{ij})
      \end{aligned}
      $$

      ii. Estimate $f_{mj}(x)$ by a weighted least-squares fit of z_{ij} to x_i

   (b) Set $f_{mj}(x) \leftarrow \frac{J-1}{J}(f_{mj}(x) - \frac{1}{J}\sum_{k=1}^{J} f_{mk}(x))$, and $F_j(x) \leftarrow F_j(x) + f_{mj}(x)$

   (c) Update $P_j(x)$ via (37).

3. Output the classifier $\text{argmax}_j F_j(x)$

Algorithm 6: *An adaptive Newton algorithm for fitting an additive multiple logistic regression model.*

Proposition 6 *The LogitBoost algorithm 6 uses adaptive quasi-Newton steps for fitting an additive symmetric logistic model by maximum-likelihood*

We sketch an informal proof.

Proof

- We first give the L_2 score and Hessian for the Newton algorithm corresponding to a standard multi-logit parametrization $G_j(x)$, where say

22

$G_J(x) = 0$. The expected log-likelihood is

$$\ell(G+g) = \sum_{j=1}^{J-1} E(y_j^*|x)(G_j(x) + g_j(x)) - \log(1 + \sum_{k=1}^{J-1} e^{G_k(x)+g_k(x)})$$

$$\begin{aligned} s_j(x) &= E(y_j^* - p_j(x)|x), \; j = 1, \ldots, J-1 \\ H_{j,k}(x) &= -p_j(x)(\delta_{jk} - p_k(x)), \; j, k = 1, \ldots, J-1 \end{aligned}$$

- Our quasi-Newton update amounts to using a diagonal approximation to the Hessian:

$$g_j(x) = \frac{E(y_j^* - p_j(x)|x)}{p_j(x)(1 - p_j(x))}, \; j = 1, \ldots, J-1$$

- To convert to the symmetric parametrization, we would note that $g_J = 0$, and set $f_j(x) = g_j(x) - \frac{1}{J}\sum_{k=1}^{J} g_k(x)$. However, this procedure could be applied using any class as the base, not just the Jth. By averaging over all choices for the base class, we get the update

$$f_j(x) = \left(\frac{J-1}{J}\right)\left(\frac{E(y_j^* - p_j(x)|x)}{p_j(x)(1 - p_j(x))} - \frac{1}{J}\sum_{k=1}^{J}\frac{E(y_k^* - p_k(x)|x)}{p_k(x)(1 - p_k(x))}\right)$$

□

5 Simulation studies

In this Section the four flavors of boosting outlined above are applied to several artificially constructed problems. Comparisons based on real data are presented in Section 6.

An advantage of comparisons made in a simulation setting is that all aspects of each example are known, including the Bayes error rate and the complexity of the decision boundary. In addition, the population expected error rates achieved by each of the respective methods can be estimated to arbitrary accuracy by averaging over a large number of different training and test data sets drawn from the population. The four boosting methods compared here are

DAB: Discrete AdaBoost — Algorithm 1

RAB: Real AdaBoost — Algorithm 2 for two classes, and Algorithm 5 (AdaBoost.MH) for more than two classes.

LB: LogitBoost — Algorithms 3 and 6.

GAB: Gentle AdaBoost — Algorithm 4

In an attempt to differentiate performance, all of the simulated examples involve fairly complex decision boundaries. The ten predictive features for all examples are randomly drawn from a ten-dimensional standard normal distribution $x \sim N^{10}(0, I)$. For the first three examples the decision boundaries separating successive classes are nested concentric ten-dimensional spheres constructed by thresholding the squared-radius from the origin

$$r^2 = \sum_{j=1}^{10} x_j^2.$$ (38)

Each class C_k $(1 \leq k \leq K)$ is defined as the subset of observations

$$C_k = \{x_i \, | \, t_{k-1} \leq r_i^2 < t_k\}$$ (39)

with $t_0 = 0$ and $t_K = \infty$. The $\{t_k\}_1^{K-1}$ for each example were chosen so as to put approximately equal numbers of observations in each class. The training sample size is $N = K \cdot 1000$ so that approximately 1000 training observations are in each class. An independently drawn test set of 10000 observations was used to estimate error rates for each training set. Averaged results over ten such independently drawn training/test set combinations were used for the final error rate estimates. The corresponding statistical uncertainties (standard errors) of these final estimates (averages) are approximately a line width on each plot.

Figure 3 [top-left] compares the four algorithms in the two-class ($K = 2$) case using a two–terminal node decision tree ("stump") as the base classifier. Shown is error rate as a function of number of boosting iterations. The upper (black) line represents DAB and the other three nearly coincident lines are the other three methods (dotted red = RAB, short-dashed green = LB, and long-dashed blue=GAB). Note that the somewhat erratic behavior of DAB, especially for less that 200 iterations, is not due to statistical uncertainty. For less than 400 iterations LB has a minuscule edge, after that it is a dead heat with RAB and GAB. DAB shows substantially inferior performance here with roughly twice the error rate at all iterations.

Figure 3 [lower-left] shows the corresponding results for three classes ($K = 3$) again with two–terminal node trees. Here the problem is more

Figure 3: *Additive Decision Boundary. In all panels except the the top right, the solid curve (representing discrete AdaBoost) lies alone above the other three curves.*

difficult as represented by increased error rates for all four methods, but their relationship is roughly the same: the upper (black) line represents DAB and the other three nearly coincident lines are the other three methods. The situation is somewhat different for larger number of classes ($K \geq 4$). Figure 3 [lower-right] shows results for $K = 5$ which are typical for $K \geq 4$. As before, DAB incurs much higher error rates than all the others, and RAB and GAB have nearly identical performance. However, the performance of LB relative to RAB and GAB has changed. Up to about 40 iterations it has the same error rate. From 40 to about 100 iterations LB's error rates are somewhat higher than the other two. After 100 iterations the error rate for LB continues to improve whereas that for RAB and GAB level off, decreasing much more slowly. By 800 iterations the error rate for LB is 0.19 whereas that for RAB and GAB is 0.32. Speculation as to the reason for LB's performance gain in these situations is presented below.

In the above examples a two-terminal-node tree (stump) was used as the base classifier. One might expect the use of larger trees would do better for these rather complex problems. Figure 3 [top-right] shows results for the two–class problem, here boosting trees with eight terminal nodes. These results can be compared to those for stumps in Fig. 3 [top-left]. Initially, error rates for boosting eight node trees decrease much more rapidly than for stumps, with each successive iteration, for all methods. However, the error rates quickly level off and improvement is very slow after about 100 iterations. The overall performance of DAB is much improved with the bigger trees, coming close to that of the other three methods. As before RAB, GAB, and LB exhibit nearly identical performance. Note that at each iteration the eight–node tree model consists of four–times the number of additive terms as does the corresponding stump model. This is why the error rates decrease so much more rapidly in the early iterations. In terms of model complexity (and training time), a 100 iteration model using eight-terminal node trees is equivalent to a 400 iteration stump model .

Comparing the top-two panels in Fig. 3 one sees that for RAB, GAB, and LB the error rate using the bigger trees (.072) is in fact 33% higher than that for stumps (.054) at 800 iterations, even though the former is four times more complex. This seemingly mysterious behavior is easily understood by examining the nature of the decision boundary separating the classes. In general, the decision boundary between any two classes can be defined by an equation of the form $B(x) = c$; using $sign[B(x) - c]$ to discriminate between the two classes results in the optimal Bayes rule. It is the goal of any classifier to approximate $B(x)$ (which is not necessarily unique) as closely as possible under whatever constraints (concept bias) are associated

with the classifier. As discussed above, boosting produces an additive model whose components (basis functions) are represented by the base classifier. If a boundary function $B(x)$ for a particular problem happens to be well approximated by such an additive model then performance is likely to be good; if not, high error rates are a likely result.

With stumps as the base classifier each basis function involves only a single (splitting) predictor variable

$$f_k(x) = c_k 1_{(s_k \cdot (x_{j(k)} - t_k) > 0)}. \tag{40}$$

Here $j(k)$ is the split variable, t_k the split point , and $s_k = \pm 1$ is the direction ($+ =$ right son, $- =$ left son) associated with the kth term. The boosted model is a linear combination of such single variable functions. Let

$$g_j(x_j) = \sum_{j(k)=j} f_k(x_j). \tag{41}$$

That is, each $g_j(x_j)$ is the sum of all (stump) functions involving the same (jth) predictor, with $g_j(x_j) = 0$ if none exist. Then the model produced by boosting stumps is additive in the *original* features

$$F(x) = \sum_{j=1}^{p} g_j(x_j). \tag{42}$$

Examination of (38) and (39) reveals that an optimal decision boundary for the above examples is also additive in the original features; $f_j(x_j) = x_j^2 + const$. Thus, in the context of decision trees, stumps are ideally matched to these problems; larger trees are not needed. However boosting larger trees need not be counter productive in this case if all of the splits in each individual tree are made on the same predictor variable. This would also produce an additive model in the *original* features (42). However, due to the forward greedy stage-wise strategy used by boosting, this is not likely to happen if the decision boundary function involves more than one predictor; each individual tree will try to do its best to involve all of the important predictors. Owing to the nature of decision trees, this will produce models with *interaction effects*; most terms in the model will involve products in more than one variable. Such non-additive models are not as well suited for approximating truly additive decision boundaries such as (38) and (39). This is reflected in increased error rate as observed in Fig. 3.

It should be noted that if $B(x) = c$ describes a decision boundary, any monotone transformation of both sides describes the same boundary. This

fact makes boosting stumps more general than one might at first think. For example, suppose $B(x) = \prod_{j=1}^{p} f(x_j)$ with all quantities being strictly positive. Such a boundary is very complex involving interactions to the highest order. However, $\tilde{B}(x) = \log B(x) = \sum_{j=1}^{p} \log f_j(x_j)$ is an additive function (in the original variables) and $\tilde{B}(x) = \log c$ describes the same boundary. Thus, boosting stumps would be optimal in spite of the apparent complexity of this $B(x)$. More generally if any monotone transformation of the decision boundary definition renders it approximately additive, boosting stumps should be competitive. Also note that the "naive" Bayes procedure, which is often highly competitive, produces decision boundary estimates that, after logarithmic transformation, are additive in the original predictors. Boosting decision stumps is considerably more flexible than most implementations of "naive" Bayes.

The above discussion suggests that if the decision boundary separating pairs of classes were inherently *non-additive* in the original predictor variables, then boosting stumps would be less advantageous than using larger trees. A tree with m terminal nodes can produce basis functions with a maximum interaction order of $\min(m-1, p)$ where p is the number of predictor features. These higher order basis functions provide the possibility to more accurately estimate those $B(x)$ with high order interactions. The purpose of the next example is to verify this intuition. There are two classes ($K = 2$) and 5000 training observations with the $\{x_i\}_1^{5000}$ drawn for a ten-dimensional normal distribution as in the previous examples. Class labels were randomly assigned to each observation with log-odds

$$\log\left(\frac{\Pr[y = 1 \mid x]}{\Pr[y = -1 \mid x]}\right) = 10 \sum_{j=1}^{6} x_j \left(1 + \sum_{l=1}^{6} (-1)^l x_l\right).$$

Approximately equal numbers of observations are assigned to each of the two classes, and the Bayes error rate is 0.046. The decision boundary for this problem is a complicated function of the first six predictor variables involving all of them in second order interactions of equal strength. As in the above examples, test sets of 10000 observations was used to estimate error rates for each training set, and final estimates were averages over ten replications.

Figure 4 [top-left] shows test-error rate as a function of iteration number for each of the four boosting methods using stumps. As in the previous examples, RAB and GAB track each other very closely. DAB begins very slowly, being dominated by all of the others until around 180 iterations, where it passes below RAB and GAB. LB mostly dominates, having the

Figure 4: *Interactive Decision Boundary*

lowest error rate until about 650 iterations. At that point DAB catches up and by 800 iterations it may have a very slight edge. However, none of these boosting methods perform well with stumps on this problem, the best error rate being 0.35.

Figure 4 [top-right] shows the corresponding plot when four terminal node trees are boosted. Here there is a dramatic improvement with all of the four methods. For the first time there is some small differentiation between RAB and GAB. At nearly all iterations the performance ranking is LB best, followed by GAB, RAB, and DAB in order. At 800 iterations LB achieves an error rate of 0.134. Figure 4 [lower-left] shows results when eight terminal node trees are boosted. Here, error rates are generally further reduced with LB improving the least (0.130), but still dominating. The performance ranking among the other three methods changes with increasing iterations; DAB overtakes RAB at around 150 iterations and GAB at about 230 becoming fairly close to LB by 800 iterations with an error rate of 0.138.

Although limited in scope, these simulation studies suggest several trends. They explain why boosting stumps can sometimes be superior to using larger trees, and suggest situations where this is likely to be the case; that is when decision boundaries $B(x)$ can be closely approximated by functions that are additive in the original predictor features. When higher order interactions are required stumps exhibit poor performance. These examples illustrate the close similarity between RAB and GAB. In all cases the difference in performance between DAB and the others decreases when larger trees and more iterations are used, sometimes overtaking the others. More generally, relative performance of these four methods depends on the problem at hand in terms of the nature of the decision boundaries, the complexity of the base classifier, and the number of boosting iterations.

The superior performance of LB in Fig. 3 [lower-right] appears to be a consequence of the multi–class logistic model (Algorithm 6). All of the other methods use the asymmetric AdaBoost.MH strategy (Algorithm 5) of building separate two–class models for each individual class against the pooled complement classes. Even if the decision boundaries separating all class pairs are relatively simple, pooling classes can produce complex decision boundaries that are difficult to approximate (Friedman 1996). By considering all of the classes simultaneously, the symmetric multi–class model is better able to take advantage of simple pairwise boundaries when they exist (Hastie & Tibshirani 1998). As noted above, the pairwise boundaries induced by (38) and (39) are simple when viewed in the context of additive modeling, whereas the pooled boundaries are more complex; they cannot be well approximated by functions that are additive in the original predictor

30

variables.

The decision boundaries associated with these examples were deliberately chosen to be geometrically complex in an attempt to illicit performance differences among the methods being tested. Such complicated boundaries are not likely to often occur in practice. Many practical problems involve comparatively simple boundaries (Holte 1993); in such cases performance differences will still be situation dependent, but correspondingly less pronounced.

6 Some experiments with data

In this section we show the results of running the four fitting methods: LogitBoost, Discrete AdaBoost, Real AdaBoost, and Gentle AdaBoost on a collection of datasets from the UC-Irvine machine learning archive, plus a popular simulated dataset. The base learner is a tree in each case, with either 2 terminal nodes ("stumps") or 8 terminal nodes. For comparison, a single CART decision tree was also fit, with the tree size determined by 5-fold cross-validation.

The datasets are summarized in Table 1. The test error rates are shown in Table 2 for the smaller datasets, and in Table 3 for the larger ones. The `vowel`, `sonar`, `satimage` and `letter` datasets come with a pre-specified test set. The `waveform` data is simulated, as described in (Breiman et al. 1984). For the others, 5-fold cross-validation was used to estimate the test error.

It is difficult to discern trends on the small data sets (Table 2) because all but quite large observed differences in performance could be attributed to sampling fluctuations. On the `vowel`, `breast cancer`, `ionosphere`, `sonar`, and `waveform` data, purely additive (two-terminal node tree) models seem to perform comparably to the larger (eight-node) trees. The `glass` data seems to benefit a little from larger trees. There is no clear differentiation in performance among the boosting methods.

On the larger data sets (Table 3) clearer trends are discernible. For the `satimage` data the eight-node tree models are only slightly, but significantly, more accurate than the purely additive models. For the `letter` data there is no contest. Boosting stumps is clearly inadequate. There is no clear differentiation among the boosting methods for eight-node trees. For the stumps, LogitBoost, Real AdaBoost, and Gentle AdaBoost have comparable performance, distinctly superior to Discrete Adaboost. This is consistent with the results of the simulation study (Section 5).

Except perhaps for Discrete AdaBoost, the real data examples fail to

Table 1: *Datasets used in the experiments*

Data set	# Train	# Test	# Inputs	# Classes
vowel	528	462	10	11
breast cancer	699	5-fold CV	9	2
ionosphere	351	5-fold CV	34	2
glass	214	5-fold CV	10	7
sonar	210	5-fold CV	60	2
waveform	300	5000	21	3
satimage	4435	2000	36	6
letter	16000	4000	16	26

demonstrate performance differences between the various boosting methods. This is in contrast to the simulated data sets of Section 5. There LogitBoost generally dominated, although often by a small margin. The inability of the real data examples to discriminate may reflect statistical difficulties in estimating subtle differences with small samples. Alternatively, it may be that the their underlying decision boundaries are all relatively simple (Holte 1993) so that all reasonable methods exhibit similar performance.

7   Additive Logistic Trees

In most applications of boosting the base classifier is considered to be a primitive, repeatedly called by the boosting procedure as iterations proceed. The operations performed by the base classifier are the same as they would be in any other context given the same data and weights. The fact that the final model is going to be a linear combination of a large number of such classifiers is not taken into account. In particular, when using decision trees, the same tree growing and pruning algorithms are generally employed. Sometimes alterations are made (such as no pruning) for programming convenience and speed.

When boosting is viewed in the light of additive modeling, however, this greedy approach can be seen to be far from optimal in many situations. As discussed in Section 5 the goal of the final classifier is to produce an accurate approximation to the decision boundary function $B(x)$. In the context of boosting, this goal applies to the final additive model, not to the individual terms (base classifiers) at the time they were constructed. For example, it was seen in Section 5 that if $B(x)$ was close to being additive in the original

Table 2: *Test error rates on small real examples*

Method		2 Terminal Nodes			8 Terminal Nodes		
	Iterations	50	100	200	50	100	200
Vowel	CART error= .642						
LogitBoost		.532	.524	.511	.517	.517	.517
Real AdaBoost		.565	.561	.548	.496	.496	.496
Gentle AdaBoost		.556	.571	.584	.515	.496	.496
Discrete AdaBoost		.563	.535	.563	.511	.500	.500
Breast	CART error= .045						
LogitBoost		.028	.031	.029	.034	.038	.038
Real AdaBoost		.038	.038	.040	.032	.034	.034
Gentle AdaBoost		.037	.037	.041	.032	.031	.031
Discrete AdaBoost		.042	.040	.040	.032	.035	.037
Ion	CART error= .076						
LogitBoost		.074	.077	.071	.068	.063	.063
Real AdaBoost		.068	.066	.068	.054	.054	.054
Gentle AdaBoost		.085	.074	.077	.066	.063	.063
Discrete AdaBoost		.088	.080	.080	.068	.063	.063
Glass	CART error= .400						
LogitBoost		.266	.257	.266	.243	.238	.238
Real AdaBoost		.276	.247	.257	.234	.234	.234
Gentle AdaBoost		.276	.261	.252	.219	.233	.238
Discrete AdaBoost		.285	.285	.271	.238	.234	.243
Sonar	CART error= .596						
LogitBoost		.231	.231	.202	.163	.154	.154
GA		.154	.163	.202	.173	.173	.173
SA		.183	.183	.173	.154	.154	.154
Discrete AdaBoost		.154	.144	.183	.163	.144	.144
Waveform	CART error= .364						
LogitBoost		.196	.195	.206	.192	.191	.191
Real AdaBoost		.193	.197	.195	.185	.182	.182
Gentle AdaBoost		.190	.188	.193	.185	.185	.186
Discrete AdaBoost		.188	.185	.191	.186	.183	.183

Table 3: *Test error rates on larger data examples.*

Method	Terminal Nodes	Iterations				Fraction
		20	50	100	200	
Satimage	CART error = .148					
LogitBoost	2	.140	.120	.112	.102	
Real AdaBoost	2	.148	.126	.117	.119	
Gentle AdaBoost	2	.148	.129	.119	.119	
Discrete AdaBoost	2	.174	.156	.140	.128	
LogitBoost	8	.096	.095	.092	.088	
Real AdaBoost	8	.105	.102	.092	.091	
Gentle AdaBoost	8	.106	.103	.095	.089	
Discrete AdaBoost	8	.122	.107	.100	.099	
Letter	CART error = .124					
LogitBoost	2	.250	.182	.159	.145	.06
Real AdaBoost	2	.244	.181	.160	.150	.12
Gentle AdaBoost	2	.246	.187	.157	.145	.14
Discrete AdaBoost	2	.310	.226	.196	.185	.18
LogitBoost	8	.075	.047	.036	.033	.03
Real AdaBoost	8	.068	.041	.033	.032	.03
Gentle AdaBoost	8	.068	.040	.030	.028	.03
Discrete AdaBoost	8	.080	.045	.035	.029	.03

predictive features, then boosting stumps was optimal since it produced an approximation with the same structure. Building larger trees increased the error rate of the final model because the resulting approximation involved high order interactions among the features. The larger trees optimized error rates of the individual base classifiers, given the weights at that step, and even produced lower unweighted error rates in the early stages. But, after a sufficient number of boosts, the stump based model achieved superior performance.

More generally, one can consider an expansion of the of the decision boundary function in a functional ANOVA decomposition (Friedman 1991)

$$B(x) = \sum_j f_j(x_j) + \sum_{j,k} f_{jk}(x_j, x_k) + \sum_{j,k,l} f_{jkl}(x_j, x_k, x_l) + ... \tag{43}$$

The first sum represents the closest function to $B(x)$ that is additive in the original features, the first two represent the closest approximation involving at most two–feature interactions, the first three represent three–feature interactions, and so on. If $B(x)$ can be accurately approximated by such an expansion, truncated at low interaction order, then allowing the base classifier to produce higher order interactions can reduce the accuracy of the final boosted model. In the context of decision trees, higher order interactions are produced by deeper trees.

In situations where the true underlying decision boundary function admits a low order ANOVA decomposition, one can take advantage of this structure to improve accuracy by restricting the depth of the base decision trees to be not much larger than the actual interaction order of $B(x)$. Since this is not likely to be known in advance for any particular problem, this maximum depth becomes a "meta-parameter" of the procedure to be estimated by some model selection technique, such as cross-validation.

One can restrict the depth of an induced decision tree by using its standard pruning procedure, starting from the largest possible tree, but requiring it to delete enough splits to achieve the desired maximum depth. This can be computationally wasteful when this depth is small. The time required to build the tree is proportional to the depth of the largest possible tree before pruning. Therefore, dramatic computational savings can be achieved by simply stopping the growing process at the maximum depth, or alternatively at a maximum number of terminal nodes. The standard heuristic arguments in favor of growing large trees and then pruning do not apply in the context of boosting. Shortcomings in any individual tree can be compensated by trees grown later in the boosting sequence.

If a truncation strategy based on number of terminal nodes is to be employed, it is necessary to define an order in which splitting takes place. We adopt a "best– first" strategy. An optimal split is computed for each currently terminal node. The node whose split would achieve the greatest reduction in the tree building criterion is then actually split. This increases the number of terminal nodes by one. This continues until a maximum number M of terminal notes is induced. Standard computational tricks can be employed so that inducing trees in this order requires no more computation than other orderings commonly used in decision tree induction.

The truncation limit M is applied to all trees in the boosting sequence. It is thus a meta–parameter of the entire boosting procedure. An optimal value can be estimated through standard model selection techniques such as minimizing cross-validated error rate of the final boosted model. We refer


Figure 5: *Coordinate functions for the additive logistic tree obtained by boosting with stumps, for the two-class nested sphere example from Section 5.*

to this combination of truncated best–first trees, with boosting, as "additive logistic trees" (ALT). This is the procedure used in all of the simulated and real examples. One can compare results on the latter (Tables 2 and 3) to corresponding results reported by Dietterich (1998, Table 1) on common data sets. Error rates achieved by ALT with very small truncation values are seen to compare quite favorably with other committee approaches using much larger trees at each boosting step. Even when error rates are the same, the computational savings associated with ALT can be quite important in data mining contexts where large data sets cause computation time to become an issue.

Another advantage of low order approximations is model visualization. In particular, for models additive in the original features (42) , the contri-

36

bution of each feature x_j can be viewed as a graph of $g_j(x_j)$ (41) plotted against x_j. Figure 5 shows such plots for the ten features of the two–class nested spheres example of Fig. 3. The functions are shown for the first class concentrated near the origin; the corresponding functions for the other class are the negatives of these functions.

The plots in Fig. 5 clearly show that the contribution to the log–odds of each individual feature, for given values of the other features, is approximately quadratic, except perhaps for the extreme values where there is little data. They indicate that observations with feature values closer to the origin have increased odds of being from the first class. Their sum is clearly indicative of the additive quadratic nature of the decision boundary (38) and (39).

When there are more than two classes plots similar to Fig. 5 can be made for each class, and analogously interpreted. Higher order interactions models are more difficult to visualize. If there are at most two–feature interactions, the two–variable contributions can be visualized using contour or perspective mesh plots. Beyond two-feature interactions, visualization techniques are even less effective. Even when non-interaction (stump) models do not achieve the highest accuracy, they can be very useful as descriptive statistics owing to the interpretability of the resulting model.

8   Weight trimming

In this section we propose a simple idea and show that it can dramatically reduce computation for boosted models without sacrificing accuracy. Despite its apparent simplicity this approach does not appear to be in common use. At each boosting iteration there is a distribution of weights over the training sample. As iterations proceed this distribution tends to become highly skewed towards smaller weight values. A larger fraction of the training sample becomes correctly classified with increasing confidence, thereby receiving smaller weights. Observations with very low relative weight have little impact on training of the base classifier; only those that carry the dominant proportion of the weight mass are influential. The fraction of such high weight observations can become very small in later iterations. This suggests that at any iteration one can simply delete from the training sample the large fraction of observations with very low weight without having much effect on the resulting induced classifier. However, computation is reduced since it tends to be proportional to the size of the training sample, regardless of weights.

At each boosting iteration, training observations i whose weight w_i is less than a threshold $w_i < t(\beta)$ are not used to train the classifier. We take the value of $t(\beta)$ to be the βth quantile of the weight distribution over the training data at the corresponding iteration. That is, only those observations that carry the fraction $1-\beta$ of the total weight mass are used for training. Typically $\beta \in [0.01, 0.1]$ so that the data used for training carries from 90 to 99 percent of the total weight mass. Note that the weights for *all* training observations are recomputed at each iteration. Observations deleted at a particular iteration may therefore re-enter at later iterations if their weights subsequently increase relative to other observations.


Figure 6: *The left panel shows the test error for the letter recognition problem as a function of iteration number. The black solid curve uses all the training data, the blue dashed curve uses a subset based on weight thresholding. The right panel shows the percent of training data used for both approaches.*

Figure 6 [left panel] shows test-error rate as a function of iteration number for the `letter` recognition problem described in Section 6, here using Gentle AdaBoost and eight node trees as the base classifier. Two error rate curves are shown. The black solid one represents using the full training sample at each iteration ($\beta = 0$), whereas the blue dashed curve represents the corresponding error rate for $\beta = 0.1$. The two curves track each other

very closely especially at the later iterations. Figure 6 [right panel] shows the corresponding fraction of observations used to train the base classifier as a function of iteration number. Here the two curves are not similar. With $\beta = 0.1$ the number of observations used for training drops very rapidly reaching roughly 5% of the total at 20 iterations. By 50 iterations it is down to about 3% where it stays throughout the rest of the boosting procedure. Thus, computation is reduced by over a factor of 30 with no apparent loss in classification accuracy. The reason why sample size in this case decreases for $\beta = 0$ after 150 iterations, is that if all of the observations in a particular class are classified correctly with very high confidence ($F_k > 15 + \log(N)$) training for that class stops, and continues only for the remaining classes. At 200 iterations, 10 classes remained of the original 26 classes.

The last column labeled *fraction* in Table 3 shows the average fraction of observations used in training the base classifiers over the 200 iterations, for all boosting methods and tree sizes. For eight-node trees, all methods behave as shown in Fig. 6. With stumps, LogitBoost uses considerably less data than the others and is thereby correspondingly faster.

This is a genuine property of LogitBoost that sometimes gives it an advantage with weight trimming. Unlike the other methods, the LogitBoost weights $w_i = p_i(1 - p_i)$ do not in any way involve the class outputs y_i; they simply measure nearness to the currently estimated decision boundary $F_M(x) = 0$. Discarding small weights thus retains only those training observations that are estimated to be close to the boundary. For the other three procedures the weight is monotone in $-y_i F_M(x_i)$. This gives highest weight to currently misclassified training observations, especially those far from the boundary. If after trimming the fraction of observations remaining is less than the error rate, the subsample passed to the base learner will be highly unbalanced containing very few correctly classified observations. This imbalance seems to inhibit learning. No such imbalance occurs with LogitBoost since near the decision boundary, correctly and misclassified observations appear in roughly equal numbers.

As this example illustrates, very large reductions in computation for boosting can be achieved by this simple trick. A variety of other examples (not shown) exhibit similar behavior with all boosting methods. Note that other committee approaches to classification such as bagging (Breiman 1996) and randomized trees (Dietterich 1998), while admitting parallel implementations, cannot take advantage of this approach to reduce computation.

9 Concluding remarks

In order to understand a learning procedure statistically it is necessary to identify two important aspects: its structural model and its error model. The former is most important since it determines the function (concept) space of the approximator, thereby characterizing the class of concepts that can be accurately approximated (learned) with it. The error model specifies the distribution of random departures of sampled data from the structural model. It thereby defines the criterion to be optimized in the estimation of the structural model.

We have shown that the structural model for boosting is additive on the logistic scale with the base (weak) learner as basis elements. This understanding alone explains many of the properties of boosting. It is no surprise that a large number of such (jointly optimized) basis elements defines a much richer class of learners than one of them alone. It reveals that in the context of boosting all base (weak) learners are not equivalent, and there is no universally best choice over all situations. As illustrated in Section 5 the base learners need to be chosen so that the resulting additive expansion matches the particular decision boundary encountered. Even in the limited context of boosting decision trees the interaction order, as characterized by the number of terminal nodes, needs to be chosen with care. Purely additive models induced by decision stumps are sometimes, but not always, the best. However, we conjecture that boundaries involving very high order interactions are likely to be encountered rarely in practice. This motivates our additive logistic trees (ALT) procedure described in Section 7.

The error model for two-class boosting is the obvious one for binary variables, namely the Bernoulli distribution. We show that the AdaBoost procedures maximize a criterion that is closely related to expected log–Bernoulli likelihood, having the identical solution in the distributional (L_2) limit of infinite data. We derived a more direct procedure for maximizing this log-likelihood (LogitBoost) and show that it exhibits properties nearly identical to those of Real AdaBoost.

In the multi-class case, the AdaBoost procedures maximize a separate Bernoulli likelihood for each class versus the others. This is a natural choice and is especially appropriate when observations can belong to more than one class (Schapire & Singer 1998). In the more usual setting of a unique class label for each observation, the symmetric multinomial distribution is a more appropriate error model. We develop a multi-class LogitBoost procedure that maximizes the corresponding log-likelihood by quasi-Newton stepping. We show through simulated examples that there exist settings

where this approach leads to superior performance, although none of these situations seems to have been encountered in the set of real data examples used for illustration; the performance of both approaches had quite similar performance over these examples.

The concepts developed in this paper suggest that there is very little, if any, connection between (deterministic) weighted boosting and other (randomized) ensemble methods such as bagging (Breiman 1996) and randomized trees (Dietterich 1998). In the language least squares regression, the latter are purely "variance" reducing procedures intended to mitigate instability, especially that associated with (large) decision trees. Boosting on the other hand seems fundamentally different. It appears to be a purely "bias" reducing procedure, intended to increase the flexibility of stable (highly biased) weak learners by incorporating them in a jointly fitted additive expansion.

The distinction becomes less clear when boosting is implemented by finite random importance sampling instead of weights. The advantages/disadvantages of introducing randomization into boosting by drawing finite samples is not clear. If there turns out to be an advantage with randomization in some situations, then the degree of randomization, as reflected by the sample size, is an open question. It is not obvious that the common choice of using the size of the original training sample is optimal in all (or any) situations.

One fascinating issue not covered in this paper is the fact that boosting, whatever flavor, seldom seems to overfit, no matter how many terms are included in the additive expansion. Some possible explanations are:

- As the LogitBoost iterations proceed, the overall impact of changes introduced by $f_m(x)$ reduces. Only observations with appreciable weight determine the new functions — those near the decision boundary. By definition these observations have $F(x)$ near zero and can be affected by changes, while those in pure regions have large values of $|F(x)|$ are are less likely to be modified.

- The stage-wise nature of the boosting algorithms do not allow the full collection of parameters to be jointly fit, and thus have far lower variance than the full parameterization might suggest. In the Machine-Learning literature this is explained in terms of VC dimension of the ensemble compared to that of each weak learner.

- Classifiers are hurt less by overfitting than other function estimators (e.g. the famous risk bound of the 1-nearest-neighbor classifier (Cover & Hart 1967)).

Whatever the explanation, the empirical evidence is strong; the introduction of boosting by Schapire, Freund and colleagues has brought an exciting and important set of new ideas to the table.

Acknowledgements

References

Breiman, L. (1996), 'Bagging predictors', *Machine Learning 26* .

Breiman, L. (1997), Prediction games and arcing algorithms, Technical Report Technical Report 504, Statistics Department, University of California, Berkeley. Submitted to Neural Computing.

Breiman, L., Friedman, J., Olshen, R. & Stone, C. (1984), *Classification and Regression Trees*, Wadsworth, Belmont, California.

Buja, A., Hastie, T. & Tibshirani, R. (1989), 'Linear smoothers and additive models (with discussion)', *Annals of Statistics* **17**, 453–555.

Cover, T. & Hart, P. (1967), 'Nearest neighbor pattern classification', *Proc. IEEE Trans. Inform. Theory* pp. 21–27.

Dietterich, T. (1998), 'An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization', *Machine Learning* **?**, 1–22.

Freund, Y. (1995), 'Boosting a weak learning algorithm by majority', *Information and Computation* **121**(2), 256–285.

Freund, Y. & Schapire, R. (1996), Experiments with a new boosting algorithm, *in* 'Machine Learning: Proceedings of the Thirteenth International Conference', pp. 148–156.

Friedman, J. (1991), 'Multivariate adaptive regression splines (with discussion)', *Annals of Statistics* **19**(1), 1–141.

Friedman, J. (1996), Another approach to polychotomous classification, Technical report, Stanford University.

Friedman, J. & Stuetzle, W. (1981), 'Projection pursuit regression', *Journal of the American Statistical Association* **76**, 817–823.

Hastie, T. & Tibshirani, R. (1990), *Generalized Additive Models*, Chapman and Hall.

Hastie, T. & Tibshirani, R. (1998), 'Classification by pairwise coupling', *Annals of Statistics* . (to appear).

Hastie, T., Tibshirani, R. & Buja, A. (1994), 'Flexible discriminant analysis by optimal scoring', *Journal of the American Statistical Association* **89**, 1255–1270.

Holte, R. (1993), 'Very simple classification rules perform well on most commonly used datasets', *Machine Learning* **11**, 63–90.

Kearns, M. & Vazirani, U. (1994), *An Introduction to Computational Learning Theory*, MIT Press.

Mallat, S. & Zhang, Z. (1993), 'Matching pursuits with time-frequency dictionaries', *IEEE Transactions on Signal Processing* **41**, 3397–3415.

McCullagh, P. & Nelder, J. (1989), *Generalized Linear Models*, Chapman and Hall.

Schapire, R. (1990), 'The strength of weak learnability', *Machine Learning* **5**(2), 197–227.

Schapire, R. & Singer, Y. (1998), Improved boosting algorithms using confidence-rated predictions, *in* 'Proceedings of the Eleventh Annual Conference on Computational Learning Theory'.