FRIEDMAN, J. (1996). Another approach to polychotomous classification. Technical report, Stanford Univ.

FRIEDMAN, J. (1999). Greedy function approximation: the gradient boosting machine. Technical report, Stanford Univ.

FRIEDMAN, J. and STUETZLE, W. (1981). Projection pursuit regression. *J. Amer. Statist. Assoc.* **76** 817–823.

HASTIE, T. and TIBSHIRANI, R. (1990). *Generalized Additive Models*. Chapman and Hall, London.

HASTIE, T. and TIBSHIRANI, R. (1998). Classification by pairwise coupling. *Ann. Statist.* **26** 451–471.

HASTIE, T., TIBSHIRANI, R. and BUJA, A. (1994). Flexible discriminant analysis by optimal scoring. *J. Amer. Statist. Assoc.* **89** 1255–1270.

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

HUBER, P. (1964). Robust estimation of a location parameter. *Ann. Math. Statist.* **53** 73–101.

KEARNS, M. and VAZIRANI, U. (1994). *An Introduction to Computational Learning Theory*. MIT Press.

MALLAT, S. and ZHANG, Z. (1993). Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Processing* **41** 3397–3415.

MCCULLAGH, P. and NELDER, J. (1989). *Generalized Linear Models*. Chapman and Hall, London.

SCHAPIRE, R. (1997). Using output codes to boost multiclass learning problems. In *Proceedings of the Fourteenth International Conference on Machine Learning* 313–321. Morgan Kaufman, San Francisco.

SCHAPIRE, R. E. (1990). The strength of weak learnability. *Machine Learning* **5** 197–227.

SCHAPIRE, R. E. and SINGER, Y. (1998). Improved boosting algorithms using confidence-rated predictions. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*.

SCHAPIRE, R., FREUND, Y., BARTLETT, P. and LEE, W. (1998). Boosting the margin: a new explanation for the effectiveness of voting methods. *Ann. Statist.* **26** 1651–1686.

VALIANT, L. G. (1984). A theory of the learnable. *Comm. ACM* **27** 1134–1142.

DEPARTMENT OF STATISTICS
SEQUOIA HALL
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305
E-MAIL: {jnf, hastie, tibs}@stat.stanford.edu

DISCUSSION

LEO BREIMAN

University of California, Berkeley

The authors have presented a fascinating view of boosting and I congratulate them for the steps they have taken to clear up the mystery of why the AdaBoost algorithm works so well. Yet I have misgivings about their explanation (which I expressed earlier in my comments to them on receiving the first draft of their paper).

A crucial property of AdaBoost is that it almost never overfits the data no matter how many iterations it is run. The authors counterexample is the first convincing one I have seen but stands in contrast to the hundreds of data sets on which AdaBoost has not produced overfitting. Any complete explanation of AdaBoost has to explain this empirical fact. Ordinarily, in logistic regression,

if one adds more and more variables, increasing the likelihood at each step, then at some point overfitting sets in and the test-set error blossoms. The fact that this does not happen in AdaBoost or in LogitBoost is extraordinary. And, unless I am missing something, there is no explanation in the paper.

I approach the problem differently. In Breiman (1999a), I assume an ensemble of classifiers $\{h(x, \theta), \theta \in \Theta\}$. Let the sequence $\{\theta_n\}$ be iid selections from Θ according to some probability Q and let the classifiers $h(\mathbf{x}, \theta_1), h(\mathbf{x}, \theta_2), \ldots, h(\mathbf{x}, \theta_N)$ each cast a unit vote for the predicted class at $\mathbf{x}$. Then, under weak conditions, as N goes to infinity, the selected class converges a.s. to $\arg\max_j Q(h(\mathbf{x}, \theta) = j)$. This holds, for example, if the classifiers are trees, in which case I refer to the ensemble as a "random forest." The theorem is a simple consequence of the law of large numbers.

For example, in bagging, each classifier is grown on a bootstrap sample from the original training set. If the training set has K examples, let θ be a K-vector of nonnegative integers adding to K. Under Q, let θ have the distribution of sampling K times with replacement from K objects. Thus, using the results cited, bagging converges a.s. as the number of iterations becomes large.

AdaBoost is probably a random forest. Consider the following version of AdaBoost: given a probability P on the training set, use a fixed algorithm to construct a classifier $h(\mathbf{x}, P)$ using the weighted training set. The next probability P' on the training set depends only on P, the indices of the examples misclassified and the weighted misclassification error. Then there is a deterministic function Φ defined on the space of all probabilities on the training set such that $P' = \Phi(P)$. Assume that Φ is ergodic with invariant measure Q and that the voting is unweighted. Then the class selected converges a.s. to $\arg\max_j Q(h(\mathbf{x}, P) = j)$.

The implication is that the random forest constructed by letting the sequence $\{P_n\}$ be iid from Q is equivalent to AdaBoost. AdaBoost uses weighted voting, but a small alteration in the argument gives the same result. In Breiman (1999a), empirical evidence is given to show that AdaBoost is a random forest. This would settle the question of why AdaBoost does not overfit, but not the question of why AdaBoost gives such low test-set error.

There are many different random forests, and some are more accurate than others. For instance, Dieterrich (1998) constructs a random forest by choosing at random among the ten best splits at any node and shows empirically that this gives generally lower test-set error than bagging. In (1999a) we show that the test-set misclassification error is bounded above by an expression that depends on two parameters: the first is the expected correlation between pairs of classifiers, and the second is the expected accuracy of the individual classifiers.

One very effective way to reduce correlation is by random choice of features at each node. That is, suppose there are ten predictor variables. At each node, choose a small subgroup of them at random to split on. This idea first appeared in a paper by Amit and Geman (1997). Using this approach and 100 iterations gives the following test-set errors as compared to the best corresponding values for LogitBoost.

Test-Set Errors (%)

Dataset	LogitBoost	Random Forest
Breast	3.8	2.9
Ionosphere	6.3	7.1
Glass	23.8	20.6
Sonar	15.4	15.9
Waveform	19.1	17.2
Sat-images	11.2	8.6

The comparison is on runs I have already done. Since I have not done a run on the vowel dataset that is similar to that of the paper, it does not appear. Similarly, I have not run the letter dataset referred to. The major point of this comparison is to show that correlation and strength are the essential ingredients of an accurate ensemble, and that if boosting succeeds, it is because it gives low correlation between classifiers of reasonable strength. The fact that a similar algorithm maximizes the likelihood in a logistic setting is not the primary reason for its success.

So what is known about the factors leading to AdaBoost's success? One clear factor is that at each step it decouples itself from the previous step. This would intuitively seem to lead to low correlation between pairs of classifiers in AdaBoost considered as a random forest. However, this has not yet been analytically or empirically established.

A recently discovered empirical property of AdaBoost is mysterious but may lead to better understanding. At each iteration, the examples currently misclassified by AdaBoost are weighted more heavily for the next iteration, to encourage their correct classification. The commonly accepted interpretation, therefore, is that AdaBoost tries to classify correctly the hard-to-classify examples, and ignores the easy-to-classify examples. This interpretation is incorrect, and I have to admit some culpability.

Suppose AdaBoost is run for a while. Then for each case in the training set, look at the weighted proportion of times that it has been misclassified (the weights are just those that AdaBoost assigns to each tree in the ensemble). For all data sets I've looked at so far, these proportions are nearly the same over the entire training set. That is, AdaBoost is an *equalizer*. When it puts more emphasis on points that have been recently misclassified, it is trying to get them correctly classified next time around [see Wheway (1999)].

Let the points in the training set be $\{(y_n, \mathbf{x}_n)\}$. Let $\text{pr}_k(n)$ be the unweighted proportion of times that the nth case in the training set is misclassified by the first k classifiers. To check on my hypothesis that AdaBoost worked so well because it is an equalizer, using the methods in Breiman (1999b), I devised an algorithm called arc-equal which sequentially minimizes $\text{Av}_n(\text{pr}_k(n) - t)^2$, where t is an adjustable parameter.

Then I did runs on a number of datasets, proceeding as follows: first AdaBoost was run 100 times. The t was taken equal to the average weighted misclassification rate over the training set produced by AdaBoost and arc-

equal run 100 times. When this was repeated 50 times, each time leaving out 10% of the data as a test set, the test-set error rate of arc-equal was always nearly the same as that of AdaBoost.

The experiments performed and the empirical observation that our datasets, AdaBoost tends to equalize the misclassification error over the training set is suggestive. But it does not yet explain why equalization leads to low test-set error. This mystery should be most interesting to some of our more theoretically motivated statisticians.

REFERENCES

AMIT, Y. and GEMAN, D. (1997). Shape quantization and recognition with randomized trees. *Neural Comput.* **9** 1545–1588.

BREIMAN, L. (1999a). Random forests. Technical report, available at www.stat.berkeley.edu.

BREIMAN, L. (1999b). Prediction games and arcing algorithms. *Neural Comput.* **11** 1493–1517.

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

WHEWAY, V. (1999). Variance reduction trends on 'boosted' classifiers. Available from virg@cse.unsw.edu.au

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA
367 EVANS HALL, #3860
BERKELEY, CALIFORNIA 94720-3860
E-MAIL: leo@stat.berkeley.edu

DISCUSSION

PETER BÜHLMANN AND BIN YU

ETH Zürich and Bell Laboratories, Murray Hill, Lucent Technologies and University of California, Berkeley

We congratulate Friedman, Hastie and Tibshirani (FHT) for their significant work connecting boosting and (kinds of) additive models. FHT provide a much-needed bridge between an important and effective machine learning procedure and traditional statistical modeling ideas. With this bridge, ideas can now flow easily in both directions so that a deeper and more thorough understanding of boosting will eventually emerge.

In this discussion, we would like to share our thoughts and reflections on boosting and related statistical ideas, inspired by FHT.

1. What makes boosting work in classification? Let us consider the two-class problem in this section. In an elegant and convincing way, FHT bring Freund and Schapire's Discrete AdaBoosting to our home domain by rederiving its population version as a familiar Newton-like step in the optimization of a not-so-familiar exponential loss function serving as a surrogate for the zero–one loss function. This surrogate is sensible since it is uniquely minimized by

(half of) the log odds ratio of the conditional probability of $Y = 1$ given $X = x$. Under this light, it is not hard to find instances of statistical procedures sharing traces of similarities with boosting. For example, in parametric estimation, we indeed have been "boosting" a consistent estimator to an efficient estimator by taking a Newton step in the optimization of the likelihood function and indeed have been using estimating equations as sensible surrogates because the likelihood equations are intractable. Despite these similarities, there are fundamental differences between boosting and these statistical procedures. As recounted by FHT, boosting as a conjecture was proposed as a theoretical learning question in the machine learning community and is now being viewed as a nonparametric stepwise fitting technique in an additive style. Its superb empirical performance in high-dimensional classification problems has very few, if any, rivals. Even though the concept of a "true" model does get used in the evaluation stage of boosting, the starting point of boosting is not a "true model" as is commonly done in statistics. Its starting point is a "weak learner" and the question posed was how to make it better or "boost" it. In hindsight, this is a natural and realistic approach in modern data analysis where the size of datasets could not be imagined in the time of our forebears such as Fisher or Neyman. Because of the complexity and scale of these problems, it is impossible to come up with an effective likelihood approach in one step. Often, if not always, a sensible procedure or a "weak learner" is available, either ad hoc or based on a simple approximation. Moreover, an evaluation criterion or a loss function is always available. The boosting question now becomes how to improve the weak learner in terms of this loss function. With a starting point and an objective function, a greedy approach for improvement is very natural. Numerical considerations of the greedy approach explain why boosting in terms of the *evaluating* zero–one loss function might not be a good idea. Any derivative-based greedy algorithm such as Newton's method is not appropriate for this discontinuous loss function whose minimizer is not unique but a whole class. A surrogate, the exponential loss function, is used in AdaBoosting as the *implementing* objective function although the evaluation function remains the zero–one loss. From a numerical point of view, the exponential function is an excellent function to apply Newton's method, because of its convexity and gradually changing derivatives. Furthermore, this is a relevant loss function to optimize since its minimizer is sensible for the classification problem. A second surrogate, proposed by FHT, is the binomial likelihood that as a function is very close to the exponential loss and has the same minimizer. FHT devise the Gentle Boosting algorithm based on the binomial surrogate and show a similar performance to those based on the exponential loss. A third surrogate, the squared loss function, is also discussed in FHT. The squared function is the best possible for Newton's method: convex and having constant second derivatives. FHT report good performance but think that it is dominated by schemes based on either the exponential or the binomial likelihood loss functions and hint that the reason might be that the squared loss function loses its monotonicity when $yF(x) > 1$. It is well known that the unique minimizer of the squared loss is the conditional expectation of Y given $X = x$

which is naturally bounded between -1 and 1 since $|Y| = 1$. If this constraint is taken into account in the greedy search for the optimizer of the squared loss, then one should not wander out to the region where $yF(x) > 1$ but stay in the region where the squared loss is monotone. It is curious to know whether or not taking this constraint into account (obviously numerically more complicated) would bring the performance of squared loss boosting on a par with those based on the exponential or binomial likelihood loss.

Although the three surrogates mentioned so far have sensible minimizers, they are only qualitative approximations to the zero–one loss function. They are actually quite bad quantitative approximations (cf. Figure 2 of FHT). But the population minimizers of the expected zero–one loss function include those of exponential, binomial likelihood and the squared loss. The first question then simply becomes:

1. Which surrogate (or implementing loss function for boosting) to use so that the boosting *estimate* best approximates its corresponding population minimizer? This (difficult) question might involve numerical as well as statistical–stochastic efficiency issues.

Now let us step back from these successful (implementing) surrogates to look at the original zero–one loss function. From the quantitative approximation point of view, an obvious surrogate is a smoothed version of the zero–one loss function $1 - \Phi(yF(x)/\sigma)$ where Φ is the cdf of $\mathcal{N}(0, 1)$ and σ is a tuning parameter to control the approximation of such a smoothed version to the original zero–one loss function. This makes the numerical greedy approach possible. If one follows the exact steps (when c is fixed) to this objective function as in the derivation of the population version of Discrete AdaBoost (Result 1 in FHT), one gets a boosting algorithm with a *different reweighting* scheme having weights.

$$\varphi(yF(x)/\sigma) = \varphi(F(x)/\sigma),$$

because $y^2 = 1$ and $\varphi(\cdot)$, the standard normal density, is symmetric.

These weights ignore the values of y or avoid reality check and concentrate more and more on the cases where the previous classifier F is unsure about itself. Since this is the opposite of the boosting reweighting philosophy, one might expect that it would not work well. The smoothed zero–one loss function is not convex (or concave) and its derivatives change rapidly, especially for a small tuning parameter σ or a closer approximation to the zero–one loss function. For such functions, the quadratic approximation is not accurate and Newton's method easily overshoots to miss the true optimizer. Hence this recipe should not work from a numerical point of view.

2. The question is, however, for this smoothed zero–one loss, will a more suitable numerical optimization method such as trust region [cf. Gill, Murray and Wright (1981)] lead to a sensible boosting procedure having a different weighting philosophy, presumably with higher computational cost?

As argued above, boosting algorithms benefit greatly from a smart choice of a surrogate implementing objective function from both statistical and numerical points of view. FHT (Section 4.4) also point out that Gentle Boosting has an edge because of the numerical stability of the derivative calculation of the binomial likelihood function. However, boosting's most powerful advocate is its stunning success on real data sets and its "mysterious" resistance to overfitting in most cases in classification. The summary in Table 2 of FHT on real dataset results suggests that the performance of a particular variant of boosting depends on the interaction among the choice of the weak learner, the underlying problem and the effectiveness of the Newton method as a numerical approximation to the surrogate loss function. Figuring out how these factors interact should bring us a deeper understanding of boosting and might shed light on its resistance to overfitting.

We explore this interaction in the rest of this discussion. In the context of boosting L_2 regression, we compare boosting with another ensemble scheme, bagging [Breiman (1996)], for which we have gained some understanding recently [Bühlmann and Yu (2000)]. Different weak learners are looked at, the overfitting issue is touched upon, and a bag-boosting scheme is proposed. In the last section, we return to classification to emphasize the importance of the choice of the weak learner and to make the point that the resistance to overfitting of boosting is probably due to the zero–one evaluation loss function.

2. Boosting and bagging with L_2 loss for regression. Boosting as explained and analyzed by FHT is a general method: a greedy forward stagewise technique to fit a model of additive style by minimizing a loss function. This view opens the door for boosting in other contexts than classification (although L_2 loss is also an appropriate surrogate as described in the previous section). For L_2 regression, the boosting algorithm works as follows:

(a) Set $F_0 \equiv 0$ (or another more sensible starting value).

(b) For $m = 1, 2, \ldots, M$, fit the function estimator $f_m(\cdot)$ which is parameterized as $f_m(x) = \beta b(x, \gamma)$ (as in FHT):

$$\beta_m, \gamma_m = \arg\min_{\beta,\gamma} \sum_{i=1}^{n} \big(Y_i - F_{m-1}(X_i) - \beta b(X_i; \gamma)\big)^2.$$

Set $F_m(\cdot) = F_{m-1}(\cdot) + f_m(\cdot)$.

(c) Output the function estimator

$$F_M(\cdot) = \sum_{m=1}^{M} f_m(\cdot).$$

This algorithm is indicated by FHT [formula (6)] and was also given by Friedman (1999). It does not involve any "reweighting". The weak learner $b(\cdot; \gamma)$ is fitted in the mth step to the current residuals $Y_i - F_{m-1}(X_i)$. There is no need for a surrogate loss function in this case since the evaluating L_2 loss is the best possible for Newton's method and the quadratic approximation is

exact. We now discuss the issue of choosing the weak learner. This discussion is continued in Section 3 for the case of classification.

Linear weak learners. It is well known that boosting's cousin bagging does not change any linear procedures. It can be shown easily that it is the same with boosting in L_2 regression. When the weak learner is linear $b(x, \gamma) = x^T\gamma$ and $f_m(x) = x^T\gamma_m$, the following two statements hold:

(a) For L_2 boosting in regression as described earlier,

$$f_m(\cdot) \equiv 0 \text{ for all } m = 2, 3, \ldots, M,$$

$$F_M(\cdot) = \text{least-squares linear regression predictor for all } M \geq 1.$$

(b) For LogitBoost as presented by FHT, with

$$p_M(x) = \exp(F_M(x))/\big(\exp(F_M(x)) + \exp(-F_M(x))\big),$$

$p_M(x)$ converges to the corresponding MLE in logistic linear regression as $M \to \infty$, provided that standard conditions for convergence of the MLE hold.

Stumps as the weak learner. We now investigate the effect of boosting in a concrete simple linear model,

$$Y_i = 2 + 3X_i + \varepsilon_i, \qquad i = 1, \ldots, n, \tag{1}$$

with ε_i's i.i.d. $\sim \mathcal{N}(0, 1)$, X_i's i.i.d. $\sim$ Uniform($[0, 1]$) and all X_i's independent from all ε_i's. The target is the true function $F_{\text{true}}(x) = 2 + 3x$, $x \in [0, 1]$.

We use stumps (with one-dimensional covariate space) as weak learners and run the L_2 boosting algorithm for regression from above. For a typical sample, the boosting stumps estimate is displayed in Figure 1; it is still a rather crude and erratic approximation for the target. Figure 1 uses six boosting iterations which is optimal on average with respect to MSE as indicated in Figure 2. Based on 50 independent simulations of model (1), we have estimated bias, variance and hence also mean-squared error as a function of the number of boosting iterations. The result is displayed in Figure 2.

3. Interestingly, there is a clear indication for overfitting, starting after six iterations already. We see here a simple example in L_2 boosting for regression where overfitting occurs easily in contrast to classification; similar phenomena are reported in Friedman (1999).

From our own work [Bühlmann and Yu (2000)] we know that stumps evaluated at x have high variances for x in a whole region of the covariate space. From an asymptotic point of view, this region is "centered around" the true optimal split point for a stump and has "substantial" size $O(n^{-1/3})$. That is, stumps do have high variances even in low dimensions as in this simple case (with only three parameters) as long as one is looking at the "right scale" $O(n^{-1/3})$; such a high variance presumably propagates when combining stumps in boosting. This observation is the starting point for another boosting machine to be described next.

The bag-boosting machine with bagged weak learners. Bagging is known to smooth the weak learner such as stumps and thus achieves a variance reduction. The precise (asymptotic) smoothing function is given in Bühlmann and Yu (2000), which also characterizes its variance reduction effect in simple yet canonical cases. The high variance of a stump is reduced up to a factor of size 2 to 3, while leaving the bias approximately unchanged. This implies that bagging is also effective for a low-dimensional predictor such as stumps. Hence, we combine bagging with boosting and we call it bag-boosting. For the weak learner in boosting, just replace the stump (or a more general tree) by the bagged stump (or bagged tree). This is a very natural idea and has been thought about by a couple of researchers although we have not seen anything in writing. The resulted fit from bag-boosting is shown in Figures 1 and 2. Note that performance for bagging alone is given by bag-boosting with one boost. Instead of a classical bagging step, we actually use "sub-bagging" or bagging on subsamples of size $[n/2]$ (resampling without replacement) which is computationally cheaper while still being as accurate as bagging [Bühlmann and Yu (2000)]. The visual improvement of bag-boosting in Figure 1 can be

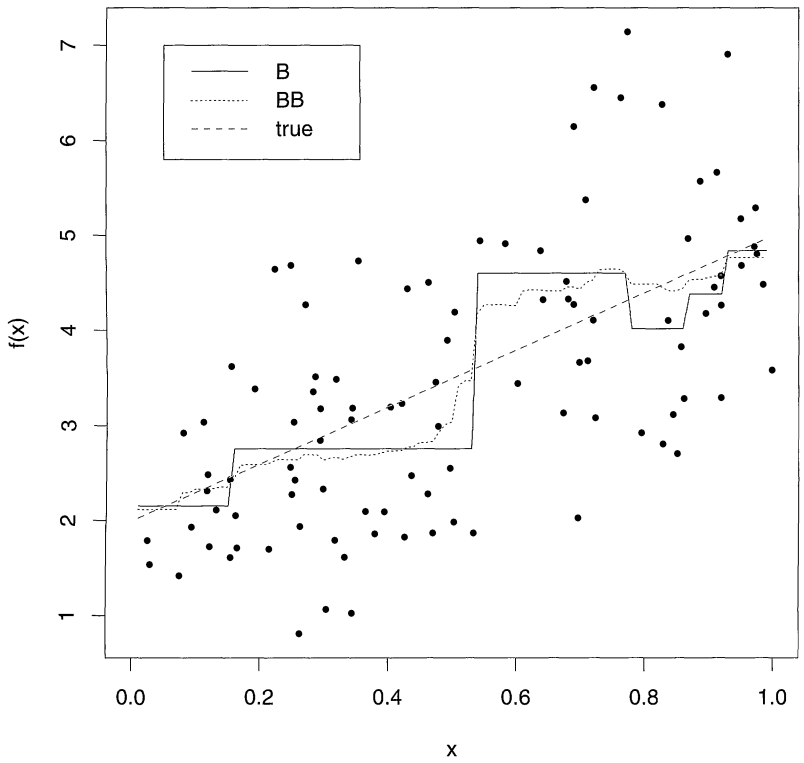


FIG. 1. *Boosting stumps* (B, *black*) *and bag-boosting stumps* (BB, *red*) *estimate based on a typical sample* (*dots*) *from* (1). *The target function is indicated in green. Boosting is done with* 6 *and bag-boosting with* 4 *iterations.*

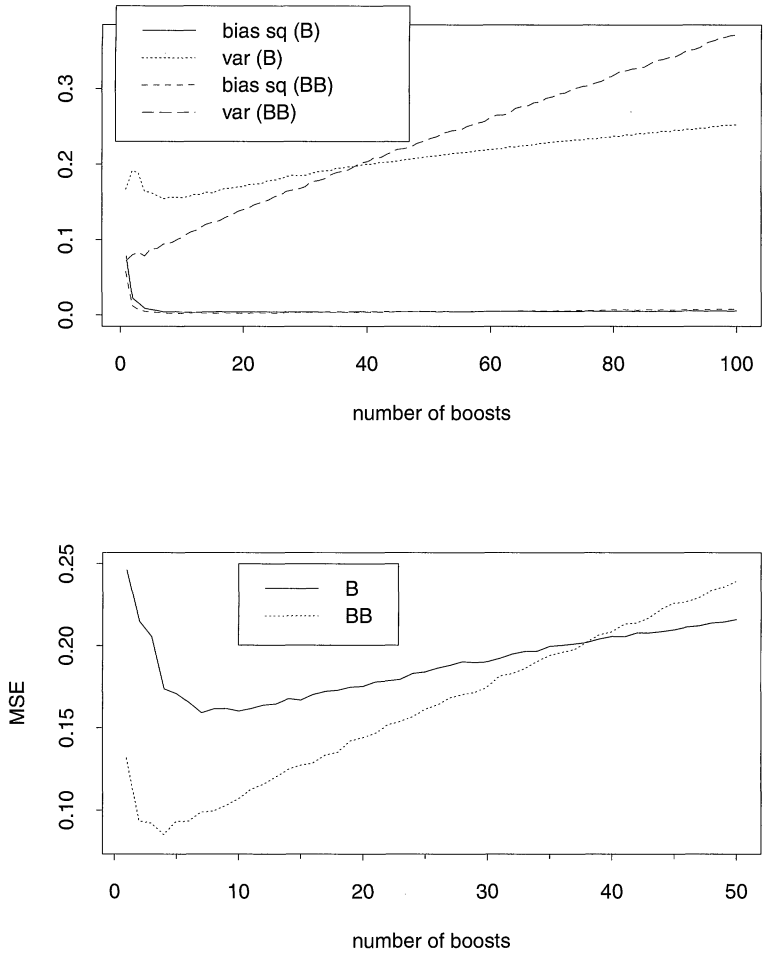


FIG. 2. *Boosting stumps and bag-boosting stumps for model* (1). *Top: squared bias and variance of boosting* (B, *black and red*) *and bag-boosting* (BB, *purple and green*). *Bottom*: *mean-squared error for boosting* (B, *black*) *and bag-boosting* (BB, *red*).

explained by the fact that a bagged stump is a smooth rather than a step function [Bühlmann and Yu (2000)]. Its variance and MSE improvements are impressively described by Figure 2, provided that we know roughly where to stop with boosting iterations. Figure 2 illustrates that the more efficient base learner (namely the bagged stump) has a faster increase in variance with the number of boosting iterations: It would be interesting to know whether this is a more general fact.

4. Thus bag-boosting has a potential to improve upon boosting with stumps or larger decision trees. Its drawback is the higher, although still feasible, computational cost.

3. Back to classification. Many issues remain. For example, the question about convergence is not touched by FHT. In particular, for a fixed sample size, there are two issues: the (non)convergence of the boosting algorithm and whether it is better to stop iterations before having approximately converged, if the algorithm actually would converge. The latter issue is known in fitting complex parametric models such as neural networks where a possible regularization is given by stopping before convergence. There is another convergence issue about the next level when the sample size tends to infinity.

5. A more important question is how much FHT's population story tells about the data driven algorithms. The population quantities are always approximated: the accuracy depends on the choice of weak learners (see also our remarks about linear weak learners, stumps and bagged stumps in Section 2) and the data-generating distributions. Finally, the hope is that the errors do not propagate with the boosting iterations!

Particularly, even with the connection made by FHT, it is hard to know whether one should use a simple learner like stumps or a more complex one such as trees with more terminal nodes.

In the boosting literature, which is almost exclusively on classification, decision trees are the most popular weak learners. We do agree that trees are generally attractive, but it is worth pointing out that other (nonlinear) weak learners may exhibit good performance as well. Even when using trees as weak learners, the choice of the number of terminal nodes (e.g., with best-first induction) in the learners does matter. Figure 1 in FHT can be misleading for a superficial reader. It wrongly suggests that the number of terminal nodes in the learner is irrelevant when there is sufficient computing power for performing many boosting steps; and that with stumps, Discrete AdaBoost is outperformed by Real AdaBoost consistently. The careful reader will exploit a more subtle sensitivity with respect to choosing the learner in Section 6 of FHT [e.g., Figure 4 (top left and top right)] and results on real data sets (Table 2).

FHT distinguish between additive and nonadditive decision boundaries. They argue convincingly that in the first case, choosing stumps as weak learners is better than using a larger tree which is expected to overestimate nonexisting interaction terms. If we knew that the decision boundary could be well approximated by an additive function, we would feel more comfortable by fitting an additive model (with backfitting) as in Hastie and Tibshirani (1990), rather than boosting stumps. Or in other words, if the conjecture by FHT [Section 11] "that boundaries involving very high-order interactions will rarely be encountered in practice" (whatever "high" means here) is really relevant, there would not be any good reason to prefer boosting over the established and popular GAM backfitting estimate.

Boosting's resistance to overfitting has been one of its major virtues in comparison with other nonparametric techniques. Overfitting in boosting algorithms in the sense of fitting a large model can happen in two ways. The first is having a large model fitted as the weak learner and the other having to run

TABLE 1
Misclassification rates (in percentages) for the breast cancer data original tree (unpruned; the program from R): 5.643

Subbagging	4.812
Boosting stumps	4.029 (with optimal 97 boosts)
Bag-boosting stumps	3.652 (with optimal 47 boosts)
Boosting large tree	2.929 (with optimal 117 boosts)
Bag-boosting large tree	2.768 (with optimal 11 boosts)

the algorithm for many iterations. The most overfitted model comes from a large or complex weak learner in combination with a long iteration in boosting. FHT touch on this overfitting issue in the concluding section by listing three plausible explanations. As emphasized in Section 1 of this Discussion, the use of the (implementing) surrogate exponential loss function is crucial from the numerical point of view and is thus partially responsible for the success of boosting. However, the evaluation in a classification problem has always been the zero–one loss.

6. This divorce of the implementation and the evaluation loss function does not exist in L_2 boosting where we saw overfitting easily in Section 2. We concur with the last point of FHT in Section 11 that the zero–one loss is very robust against overfitting and we further argue that an evaluation based on the exponential implementing loss does show strong overfitting. This conjecture is supported by the breast cancer dataset below. Moreover, we believe that "parts of the real world" are nonadditive. Under this belief and knowing boosting's resistance to overfitting, it can be very valuable to use best-first inducted trees with *more than two* terminal nodes. This is confirmed also by our analysis of the breast cancer data in Table 1.

Analysis of breast cancer data. We partially redo the analysis of FHT for the breast cancer data and thereby add insights on bag-boosting and overfitting. Table 1 and Figure 3 give comparative results to the ones in FHT. We use "sub-bagging" for computational efficiency and it has a performance similar to bagging [Bühlmann and Yu (2000)]. The misclassification rates (called test error in FHT) are estimated by averaging over 100 random partitions of the data set into a 90% training and a 10% testing, while FHT use a 5-fold CV. So our results differ slightly from FHT (also because presumably we are using different tree algorithms), but we believe ours have better accuracies. Boosting and its variants are done with Real AdaBoost, and the trees are run with the program from R. Bag-boosting large trees gives the best result, followed by boosting large trees, bag-boosting stumps, boosting stumps and then bagging. Bag-boosting only needs few iterations and brings nontrivial improvement over boosting alone, at the expense of a more computationally costly weak learner. Note that Figure 3 (bottom panel) shows strong overfitting with respect to $\mathbf{E}[\exp(-YF(X))]$!

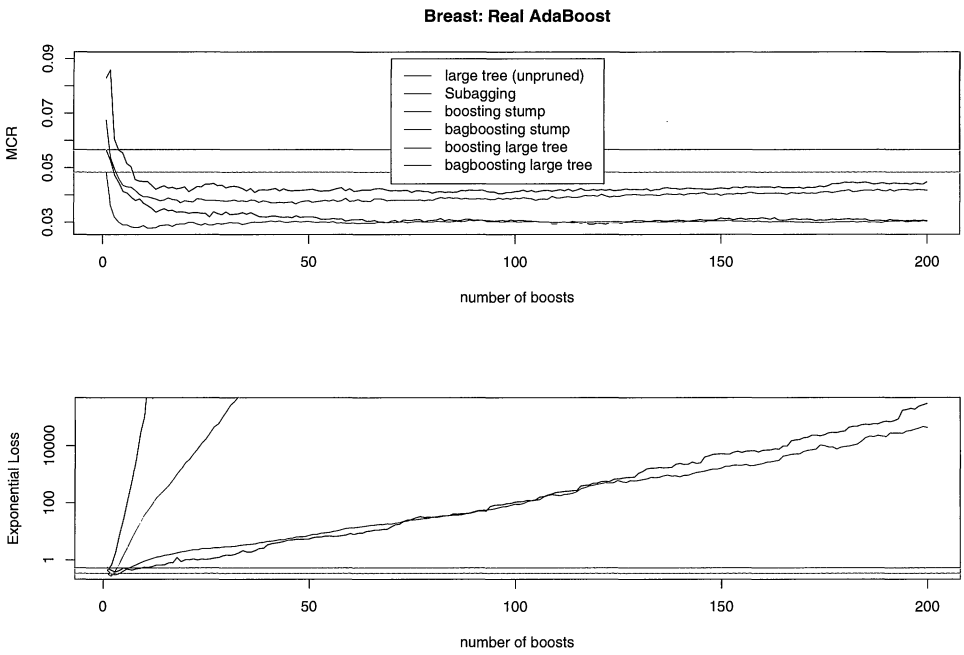


FIG. 3. *Different weak learners in Real AdaBoost for breast cancer data (Section 7 in FHT). Top: misclassification rates. Bottom: expected exponential losses* $\mathbf{E}[\exp(-YF(X))]$ *on the logarithmic scale.*

7. For this particular dataset and with respect to the misclassification rate, there is hardly any overfitting. It suggests that there is no loss in using large trees, if there is never any substantial overfitting. Boosting stumps on the other hand restricts to additive models. With respect to the misclassification rate, should we always boost with large trees as weak learners?

Acknowledgments. Bin Yu thanks Mark Hansen for many stimulating discussions and helpful comments on the draft and Margaret Wright for helpful discussions on optimization. Special thanks are due to Josef Vogt for the results on the breast cancer data.

REFERENCES

BREIMAN, L. (1996). Bagging predictors. *Machine Learning* **24** 123–140.

BÜHLMANN, P. and YU, B. (2000). Explaining bagging. Preprint.

FRIEDMAN, J. (1999). Greedy function approximation: the gradient boosting machine. Technical report, Stanford Univ.

GILL, E. P., MURRAY, W. and WRIGHT, M. H. (1981). *Practical Optimization.* Academic Press, New York.

HASTIE, T. and TIBSHIRANI, R. (1990). *Generalized Additive Models.* Chapman and Hall, London.

SEMINAR FÜR STATISTIK
ETH-ZENTRUM, LEO D72
CH-8092 ZURICH
SWITZERLAND
E-MAIL buhlmann@stat.math.ethz.ch

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720-3860

DISCUSSION

ANDREAS BUJA

AT&T Labs

1. The cultural background of boosting: machine learning versus statistics. I have worked under the same roof with the inventors of boosting, Freund and Schapire, for several years, but I might as well have been on another planet. The physical vicinity was of no help in coming to grips with boosting. I listened to their talks, but I heard the words and missed the point. Later I read Leo Breiman's debates with Freund and Schapire, and it really got my attention when I heard his praise of boosting. His penetrating analysis was an effort to understand the performance of boosting, in particular its mysterious immunity to overfitting. What he did not provide was a translation of boosting into terms that are familiar to statisticians. He did not bridge the culture gap that separates the stodgy main stream of statistics from its dynamic younger sibling in computer science, machine learning. This is one of the achievements of the paper by Friedman, Hastie and Tibshirani (FHT): It gives statisticians a language with which to think about boosting.

If Freund and Schapire feel misinterpreted or if they think of FHT's statistical modeling interpretation as misguided, as I expect they will, it is beside the point. There is no single "true" interpretation of anything; interpretation is a vehicle in the service of human comprehension. The value of an interpretation is in enabling others to fruitfully think about an idea, and this is what the FHT paper certainly does for statisticians and the idea of boosting.

An interpretation of new ideas in terms of old ones is not a mere reduction of the ideas to an "old hat." On the contrary, such a reduction can clarify what is novel. The emphasis of the FHT interpretations may be too long on the reduction aspect and too short on working out the novelty in boosting; I will try to explain what I mean below. The novelty that we see in light of FHT's interpretations, though, may not agree with what Freund and Schapire think is novel. Thus, while the FHT paper may provide a gate for statisticians to a new area, the same gate may not be open in reverse for the folks in machine learning (although I'd love to be proved wrong).

2. Boosting arbitrary generalized linear and additive models. According to FHT, boosting can be seen as a two-class classification method that fits additive models on a logistic scale in a forward stagewise manner with regard to a loss function that is second-order equivalent to the negative log likelihood for the conditional Bernoulli model. This interpretation leads FHT naturally to propose their "LogitBoost" method by substituting the exact negative log likelihood for the original loss function.

The test of every interpretation is in its capacity to spawn new ideas and generalizations. Differing interpretations are worthwhile because they lead to differing generalizations. The FHT interpretation resulted immediately in a

very competitive new version of boosting, as well as an improvement in boosting multiclass classification.

It seems to me, though, that the FHT interpretation calls, even screams, for a much more global extension of boosting: Now that we know how to boost additive logistic regression, we also know how to boost any generalized additive model. If Algorithm 3 of FHT for LogitBoost is the familiar Newton–Raphson algorithm modified for the acquisition of nonlinear terms in a stagewise logistic regression, then the same can be used for the acquisition of nonlinear terms in any stagewise generalized regression.

In particular, we now know how to boost ordinary least-squares regression: by stagewise LS fitting of nonlinear terms to the previous residuals. This may seem trite, but it isn't if one has ever tried to come up with a version of boosting for regression with continuous response. Freund and Schapire have tried, and their tentative proposal was to reduce regression to some form of classification, in a classical attempt at reducing a new problem to an old one. This seems like a circuitous approach, however, once the FHT interpretation of boosting is at hand. With the extension of the FHT interpretation to generalized additive regression, it seems most natural to interpret stagewise LS regression as the regression version of boosting.

3. The role of greedy stagewise fitting in the protection against overfit. FHT's interpretation of boosting is lucid, but their presentation of stagewise fitting needs to be put in perspective. I have no quibbles with the facts, but some with the positioning of the facts.

FHT do state that the fitting à la boosting is "forward stagewise," but the impression we (and above all the readers in machine learning) may get is that this term describes a common mode of model fitting in statistics. It doesn't. We have the term, but it is not very well known, and it does not enjoy high standing among most statisticians who know about it. It is a computational device, or more accurately, a capitulation, in those situations where one is unable to fit a large number of model terms jointly. We even teach our students that fitting one term at a time and not adjusting the others is *not* the way to fit a model; we teach them that changing one term in a model affects all others in a joint model fit, assuming that the joint model fit is the norm. Interestingly, one popular method for joint fitting, the backfitting algorithm, was devised by Friedman and Stuetzle (1981) to more closely achieve joint optimality in the otherwise greedy fitting procedure of projection pursuit regression (PPR). Had they omitted the backfitting part in their algorithm and allowed the acquisition of a greater number of ridge function terms instead, PPR would have been a regression version of boosting. It appears that even the unconventional minds of Friedman and Stuetzle were involuntarily constrained by statisticians' prejudice in favor of parsimonious models obtained by joint model fits. Doing away with this prejudice is certainly one of the novelties in the boosting endeavor. Ockham's razor is being replaced by Freund and Schapire's sledgehammer.

Omitting joint fitting is crucial to the success of boosting: Its purported advantage, relative protection against overfit, disappears instantaneously if termwise (stagewise) fitting is replaced with joint fitting of all terms. The 500 or so stumps and twigs that I have seen fitted in a single boosting fit would come crashing down if they were fitted jointly. Therefore, the essence of the protection against overfit is in the suboptimality of the fitting method.

This also explains why boosting could not have been invented by statisticians: Our preoccupations are such that the statistical benefits of a suboptimal computational device would not have appeared on our radar. There are exceptions to the rule, though: Statisticians have, for example, developed a notion of one-step estimates in robustness, where it was shown that one Newton step in the computation of an M-estimate provides essentially the robustness benefits of full iteration. Similarly, in order to understand the performance of boosting, we will need a theory that explains why this particular type of suboptimality often avoids overfit.

A curious observation in this context concerns the cavalier attitude in the boosting community, including FHT, toward the stopping problem. They seem to consider it sufficient to provide plots or tables that show misclassification rates for 100, 200, 300, 400 terms. So we stop around 300 or 400, and this is part of the definition of the method. The assumption, empirically substantiated, is that it really doesn't matter much as long as we stop after the acquisition of a sufficiently large number of terms. The situation is similar to that in artificial neural networks (NNs), where the stopping time of online optimization acts as a smoothing parameter. In NNs, however, one tries to detect the onset of overfit with cross-validation. In boosting, by contrast, the testing error remains flat for a long time, overfit does not occur before the human runs out of patience and a precise stopping rule seems therefore unnecessary.

A satisfactory theory of protection against overfit would have to explain the relative flatness of the test error in boosting. The flatness is only relative, however, and even boosting has to submit to overfit at some point. This is seen from the fact that usually the class of fitted terms $b(x, \gamma)$ is rich in the sense that for a given finite sample size N the N-dimensional design vectors $\mathbf{b}(\gamma) = (b(x_n, \gamma))_{n=1\cdots N}$ span all of R^N as γ runs across its domain.

This is obviously not a proof, but a slightly tighter argument can be given. First, we note that, unlike in joint fitting, any iterative termwise fitting procedure depends not only on the space spanned by the terms, but the actual basis they represent and the order in which they are fitted. (I thank Werner Stuetzle who spelled this out most clearly to me.)

In the simple case where the domain of the γ's is finite and consists of K elements $\gamma_1, \ldots, \gamma_K$ only, boosting is a greedy form of backfitting the *linear* model $\sum_{k=1\cdots K} \beta_k b(x, \gamma_k)$. Instead of cycling through the terms in an orderly fashion $k = 1, 2, \ldots, K, 1, 2, \ldots$ as in conventional backfitting, the next k is picked greedily so as to lower the loss function maximally. We can conceive of at least three types of backfitting according to the order in which the terms are fitted: (1) cycle over the terms in a fixed order (traditional backfitting), (2) pick the next term randomly, (3) pick the next term greedily. It would

seem plausible to examine combinations of type 2 and type 3, whereby the picking probability of a term is a function of its loss reduction; as is, type 2 uses a uniform distribution on the terms and type 3 puts probability 1 on the term with maximal reduction.

It is clear from the known convergence of type 1 backfitting to the joint least-squares solution that overfit must occur if the joint solution overfits. If similar convergence properties for type 2 and type 3 backfitting hold, overfit will occur there also. In light of these considerations, it may be just a matter of time till boosting follows the lead of NNs and cross-validation is introduced for the stopping problem in boosting as well.

An issue that a theory of protection against overfit would have to tackle is that there are two ways for a learner or base term $b(x, \gamma)$ to be "weak": weakness can arise from bias *or* from variance. In either case the weak learner usually contains signal, but the signal is weak for different reasons. Weakness due to bias means underfit, weakness due to variance means overfit. Boosting is designed to reduce bias but not variance. The latter is addressed by Breiman's (1996) "bagging" proposal and by Bayesian model averaging. As things stand, one either approaches good fits from poor fits with boosting, or from overly flexible fits with model averaging. It would be desirable if both bias and variance could be tackled in one unified approach.

4. The role of the functional form of the weak learner. Additive models with small trees as base terms are an innovation we owe the machine learning literature. Trees are not essential to boosting, but the idea of building additive models from multiple trees is a welcome side benefit.

Functional form hasn't been a topic of great debate in the boosting literature, probably because of the original justifications in terms of the "weak learner" metaphor. These justifications do not carry us far, however. The question is not *whether* but *how much* a weak learner can be boosted, or in statistical terminology: how well an additive model can approximate an underlying response surface. A major determinant for answering this second question is the functional form of the weak learner $b(x, \gamma)$. The fact that the boosting folks haven't hit on this the hard way is probably due to the Holte (1993) effect, cited a couple of times by FHT: "Very simple classification rules perform well on most commonly used datasets." That is, nature rarely chooses complex decision boundaries that couldn't be described by thresholding additive functions or at most two-variable or three-variable interactions.

Inspired by ANOVA practices, FHT propose to measure and control complexity of a functional form in terms of interaction order, that is, the maximum number of variables admitted in the weak learner at any given time. For trees, the interaction order is identical with the maximum admitted tree depth. But the idea of interaction order can be extended to other functional forms as well. The most obvious are Friedman's MARS fits, in which interaction order can be controlled as easily as in trees. Another possibility are projection pursuit regression and classification (PPR and PPC, respectively), although a major modification is necessary: the linear combinations that enter the ridge

functions would have to be constrained to, say, two or three predictors at a time.

In order to more fully understand boosting, it will be necessary to apply it to functional forms other than trees. This would not be difficult in the case of MARS and PPR/C. In fact, for a boosting version of either one would simply replace the joint fitting with unadjusted greedy fitting. Ironically, stagewise modifications of MARS and PPR/C require that the technically most creative ideas of their inventors be thrown out: Friedman used sophisticated stable Cholesky updates for the fast computation of joint fits in MARS, and Friedman and Stuetzle devised backfitting for the same purpose in PPR/C. Now we are led to expect that giving up these technical marvels will be beneficial! Having experienced first hand the problems of overfit in the original MARS, I would be thrilled if the anticipated benefits of boosting showed up.

5. Conclusion. The FHT interpretation of boosting has implications for machine learning as well as for statistics. For machine learning, the generalization of FHT's ideas lead us straight to a regression version of boosting. For statistics, FHT's ideas raise the expectation that the benefits of boosting might carry over to other stagewise modifications of nonparametric fitting procedures. The latter is just a guess so far, and in the absence of a convincing theory of boosting's protection against overfit, we wouldn't even know why it held if it did. Clearly, there is a lot of work ahead of us. FHT have given us a means for generating new ideas that can guide our future research.

REFERENCES

BREIMANN, L. (1996). Bagging predictors. *Machine Learning* **24** 123–140.

FRIEDMAN, J. and STUETZLE, W. (1981). Projection pursuit regression. *J. Amer. Statist. Assoc.* **76** 817–823.

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

AT&T LABS
SHANNON LABORATORY
180 PARK AVENUE
FLORHAM PARK, NEW JERSEY 07932-0971
E-MAIL: andreas@research.att.com

DISCUSSION

YOAV FREUND AND ROBERT E. SCHAPIRE

AT&T Labs

The main and important contribution of this paper is in establishing a connection between boosting, a newcomer to the statistics scene, and additive models.

One of the main properties of boosting that has made it interesting to statisticians and others is its relative (but not complete) immunity to overfitting. As pointed out by the authors, the current paper does not address this issue. Leo Breiman (1998) tried to explain this behaviour in terms of bias and variance. In our paper with Bartlett and Lee (1997), we gave an explanation in terms of the "margins" of the training examples and the VC-dimension of the base class. Breiman, as well as the current paper, point out that our bounds are very rough and are thus not useful in practice. While this is clearly true at this time, it is also true that the analysis given by Breiman and by this paper yield no provable bounds whatsoever. It is completely unclear whether this analysis can be used to predict the performance of classification rules outside of the training sample.

At the root of this argument about boosting is a much more fundamental argument about the type of prior assumptions that one should make when embarking on the task of inducing a classification rule from data. The assumption that seems to underlie the use of maximum likelihood in the current paper is that data are generated by a distribution from a *prespecified class*. In this case, this is the class of distributions in which the relationship between the features and the labels is described by a log-linear function. In comparison, the assumption that we make in our analysis is that the data are generated from some *arbitrary* distribution in an i.i.d. fashion. Clearly, our assumption is the weaker one and this leads to a theory that is more generally applicable.

From a related but more practical point of view, one main issue when applying boosting or boosting-like techniques in practice is how to choose the base class. The approach taken in this paper is that this choice is made based on our prior beliefs regarding the type of log-linear dependencies that might exist between the features and the label. On the other hand, in the boosting approach, we make an assumption about what kind of rules might have slight but significant correlations with the label. This is the essence of the "weak learning" assumption upon which the theory of boosting is founded.

In the current paper, boosting is analyzed mostly in the context of decision stumps and decision trees. The argument seems to be that while in most real-world cases decision stumps are powerful enough, in some less common cases the type of dependencies that exist in the data require a more powerful base class, such as two- or three-level decision trees. A rather different approach to the combination of decision trees and boosting was recently proposed by Freund and Mason (1999). They represent decision trees as sums of very simple functions and use boosting to simultaneously learn both the decision rules and the way to average them.

Another important issue discussed in this paper is the performance of boosting methods on data which are generated by classes that have a significant overlap, in other words, classification problems in which even the Bayes optimal prediction rule has a significant error. It has been observed by several authors, including those of the current paper, that AdaBoost is not an optimal method in this case. The problem seems to be that AdaBoost overemphasizes the atypical examples which eventually results in inferior rules. In the current

paper, the authors suggest "GentleBoost" as a better method than AdaBoost for this case. The reason that this might be a better method is that it gives less emphasis to misclassified examples. The increase in the weight of the example is quadratic in the negative margin, rather than exponential.

However, one can argue that this alteration of AdaBoost, while being a step in the right direction, is not large enough. In fact, one can argue that once an example has a very large negative margin it is best to assume that it is an outlier that should be completely removed from the training set. A new boosting algorithm based on this radical approach was recently proposed by Freund (1999).

REFERENCES

BREIMAN, L. (1998). Arcing classifiers. *Ann. Statist.* **26** 801–849.

FREUND, Y. (1999). An adaptive version of the boost by majority algorithm. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory*.

FREUND, Y. AND MASON, L. (1999). The alternating decision tree learning algorithm. In *Machine Learning: Proceedings of the Sixteenth International Conference* 124–133.

SCHAPIRE, R. E., FREUND, Y., BARTLETT, P. and LEE, W. S. (1998). Boosting the margin: A new explanation for the effectiveness of voting methods. *Ann. Statist.* **26** 1651–1686.

AT&T LABS–RESEARCH
180 PARK AVENUE
FLORHAM PARK, NEW JERSEY 07932-0971
E-MAIL: yoav@research.att.com

DISCUSSION

GREG RIDGEWAY

University of Washington

The authors have done a great service to the field of statistics by branching out to the problems of interests to the computational learning theorist. The connection between boosting and maximum likelihood not only helps to explain why boosting is an effective classification method but it also opens the door to the application of boosting to many other statistical prediction problems.

It is difficult not to understate the attention that boosting has received in the machine learning and data mining communities. Soon after Freund and Schapire's algorithm made its debut, a flurry of empirical work showcased its uncanny ability to improve prediction accuracy [Quinlan (1996), Bauer and Kohavi (1999), Drucker and Cortes (1996)]. Charles Elkan's application of the boosted naïve Bayes won first place out of 45 entries in the 1997 data mining competition [Elkan (1997)], demonstrating the strength of boosting against research and industrial prediction tools. Elkan also noted in this work that his model had the form of nonlinear logistic regression but did not completely

connect boosting with optimization and likelihood. Boosting is even making its way into popular science as noted in the *New York Times* [Lee (1999)].

Low variance, bias reducing steps. Likelihood based inference is in the business of obtaining good parameter estimates, which is not necessarily the same as producing a good predictor. In the classification game we win if we produce a classifier with the lowest misclassification rate. A good estimator of the underlying class distribution conditional on x is of secondary importance. Consider the following example.

EXAMPLE 1. Generate $N = 1000$ observations as

$$x \sim N(0,1), \qquad F(x) = -\frac{1}{2}x^2 - 1,$$

$$p(x) = \frac{\exp(2F(x))}{1+\exp(2F(x))}, \qquad y|x \sim \text{Bern}(p(x))$$

The Bayes decision rule is to classify a new observation as a 1 if $|x| > \sqrt{2}$ and 0 otherwise. The misclassification error for this rule is about 20.2%. Figure 1 shows the $p(x)$ that generated the data as a smooth curve. After 100

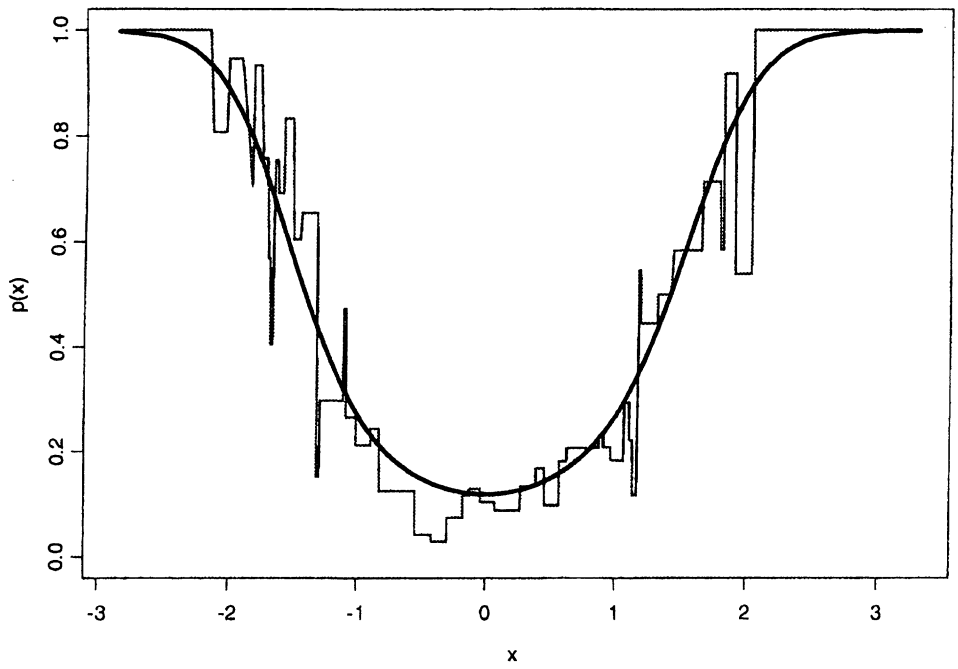


FIG. 1. *The smooth curve is* $P(Y=1|x)$ *and the jagged curve is the estimate that the Real AdaBoost algorithm produces after* 100 *iterations using a stump classifier on the Example* 1 *data.*

boosting iterations using a stump classifier (a piecewise constant with only one knot fit using CART) we see that boosting produces an estimate of $p(x)$ that is overfit in terms of providing an accurate picture of the underlying data generating mechanism. However, the misclassification rate on a test dataset with 10,000 observations is only 21.5%, not far off from the Bayes error rate. More iterations have little effect. This counts as a win for the analyst interested solely in predictive performance but a loss for those seeking to uncover the data generating mechanism. Even though the authors' Lemma 1 shows that the population minimizer of $\mathrm{E}e^{-yF(x)}$ is the true conditional class probability, with a finite sample it is easy for AdaBoost to overfit from a function estimate or calibration point of view. The debate between accurate prediction and the discovery of the contents of the black box is certainly not new but it is also has not gone away.

Some of the amazement surrounding boosting was that it could take a simple classifier like a stump and turn it into a near-Bayes optimal classifier. The choice of using a stump at each iteration in the example is critical. After two boosting iterations the classifier will pretty much nail the correct decision boundary since the two knots will likely be in the neighborhood of $-\sqrt{2}$ and $+\sqrt{2}$. Due to the constraints on the base classifier, subsequent iterations can do little to modify the decision boundary. Therefore, local jaggedness develops as successive stumps try to fit individual observations but it cannot be substantial enough to have an effect on the misclassification rate. This leads to the question of the importance of selecting the base classifier. In this example, fitting slightly more complex trees at each iteration has disastrous effects on the performance. Figure 2 tracks the change in misclassification rate with each boosting iteration when the base classifier is a tree with four terminal nodes. A four-terminal node tree is likely to nearly capture the decision boundary on the first iteration. Subsequent iterations have enough flexibility to fit individual observations and generalization performance declines. From Figure 2 we can see that the performance degrades to slightly worse than predicting 0, the majority class, for all observations.

The results of Example 1 yield the following insights into boosting.

1. AdaBoost (in its present form) is not necessarily useful in function estimation. Calibration of the boosted probability estimates is questionable.
2. The idea that AdaBoost is resistant to overfitting needs some qualification. The choice of base classifier can make or break the success of utilizing boosting.
3. Boosting seems to work well when we use a sequence of simple, low variance classifiers. Each stage makes a small, low variance step, sequentially chipping away at the bias.

Other looks at functional optimization. Once boosting is cast as functional optimization we can approach the problem from other avenues. In the AdaBoost setup we wish to find an $F(x)$ that minimizes $J(F) = \mathrm{E}e^{-yF(x)}$. An additive improvement to $F(x)$ would find an α and an $f(x)$ such that

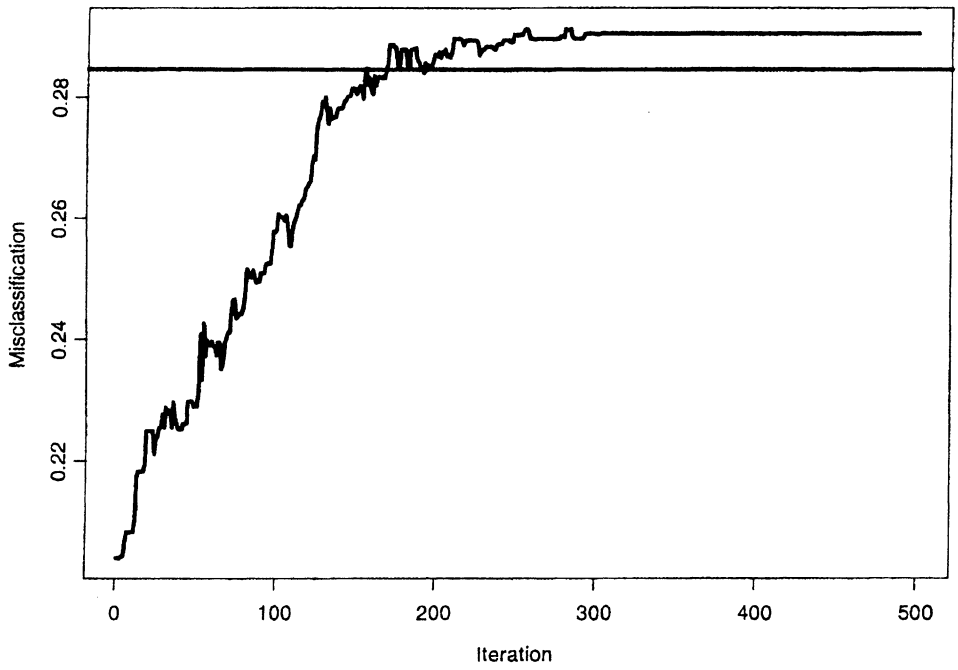


FIG. 2. *Misclassification rate when fitting trees with four terminal nodes to the Example* 1 *test data on each Real AdaBoost iteration. The horizontal line marks the error rate if we predicted all observations to be zeros.*

$J(F+\alpha f)$ decreases. The authors describe how AdaBoost selects α and $f(x)$. The Gentle AdaBoost algorithm that the authors propose performs a pointwise Newton step to suggest $f(x)$ with α fixed at 1. We can also view this as a gradient ascent in function space with a squared penalty on the magnitude of $f(x)$. Consider the following analysis.

The directional derivative, or Gâteaux variation, of $J(F)$ in the direction of some function $f(x)$ computes the local change in $J(F)$ if we make the modification $F(x) \leftarrow F(x)+\alpha f(x)$, moving $F(x)$ in the direction of $f(x)$. The directional derivative of $J(F)$ is

$$\begin{aligned}\delta J(F,f) &= \frac{d}{d\alpha}\mathrm{E}\,\exp(-y(F(x)+\alpha f(x)))\big|_{\alpha=0} \\ &= \mathrm{E} - yf(x)\exp(-yF(x)).\end{aligned} \tag{1}$$

A gradient ascent strategy would have us select from among some class of functions the $f(x)$ that minimizes $\delta J(F, f)$, or equivalently offers the greatest local decrease in $J(F)$. Without further restrictions this method is rather unstable since we can make $\delta J(F, f)$ arbitrarily small by setting

$$f(x) = \begin{cases} -\infty, & \text{if } \mathrm{E}y|x < 0, \\ \infty, & \text{if } \mathrm{E}y|x > 0. \end{cases} \tag{2}$$

On the other hand, if we penalize $f(x)$ in (1) for being too big then we might select $f(x)$ as

$$
\begin{aligned}
f(x) &= \arg\min_f \mathrm{E} - yf(x)e^{-yF(x)} + \tfrac{1}{2}f(x)^2 \\
&= \arg\min_f \mathrm{E} - yf(x)e^{-yF(x)} + \tfrac{1}{2}f(x)^2 + \tfrac{1}{2}(ye^{-yF(x)})^2 \\
&= \arg\min_f \mathrm{E}(f(x) - ye^{-yF(x)})^2.
\end{aligned} \tag{3}
$$

Note that (3) is equivalent to choosing the $f(x)$ that is similar to a least-squares fit to the gradient that the authors compute in their derivation of the Gentle AdaBoost algorithm. Obviously, in practice we only have a sample with which to estimate the expectation in (3). Therefore, it becomes critical that we use a low variance least-squares regression method to minimize the empirical estimate of (3). I will return to this topic in a moment. After estimating $f(x)$ we can use a line search to obtain the α that minimizes $\mathrm{E}\exp(-y(F(x)+\alpha f(x)))$.

Applying this same technique within the LogitBoost scenario, that is, maximizing the directional derivative of the Bernoulli log-likelihood with a squared penalty on the direction $f(x)$, the gradient step is

$$
f(x) = \arg\min_f \mathrm{E}\left(f(x) - \left(y^* - \frac{e^{F(x)}}{e^{-F(x)} + e^{F(x)}}\right)\right)^2. \tag{4}
$$

After estimating the optimal step direction the choice of the optimal step size, α, takes on an interesting form when we have a sample.

$$
\alpha = \arg\min_\alpha \sum_{i=1}^{N} 2y_i^*(F(x_i) + \alpha f(x_i)) - \log\bigl(1 + \exp\bigl(2(F(x_i) + \alpha f(x_i))\bigr)\bigr). \tag{5}
$$

Equation (5) has the form of a likelihood for the linear logistic regression model with $f(x_i)$ as a covariate and $F(x_i)$ as an offset term. The optimal α then is just the maximum likelihood estimate of the coefficient of $f(x_i)$ for which efficient software exists. Similar results hold for the class of exponential family regression models and the Cox model for survival analysis [Ridgeway (1999)].

Both the authors' results and this section arrive at the situation where we need to propose a direction in which to move our current guess for the regression function. When handed only finite training datasets, one is faced with estimating a direction that seems likely to offer a decrease in the loss function. Example 1 showed that high variance estimators of descent directions (e.g., four-terminal node trees) can cause problems. In the past statisticians have controlled the variance of predictors in other manners. As the authors point out, the LogitBoost algorithm has a familiar feel to statisticians. Traditional statistical practice uses linear regression instead of trees at each stage to estimate the conditional expectation. This yields the iteratively reweighted

least-squares algorithm for fitting linear logistic regression models. We can also use smoothers to estimate the conditional expectation as in the local scoring algorithm. These types of methods controlled variance at the expense of bias reduction by limiting the functional form, restricting variable interactions, or constraining the jaggedness of the curve.

When we introduce regression trees to estimate these gradient steps instability becomes an issue. However, various forms of averaging may help to keep variance in check. If we use bagged trees to estimate the gradient we can often reduce the variance introduced on the iteration. This is the basis of Breiman's adaptive bagging [Breiman (1999)]. Denison (2000) proposes averaging over all possible stumps at each iteration with each weighted by their posterior distribution. The amount of needed computation is nearly equivalent to CART with the benefit of more stable gradient steps. It seems that the future of nonparametric regression might consist of choosing an ideal loss function and a sequence of low variance estimates of the gradient direction in which each step decreases bias.

Bayesian prediction and function estimation. When I first learned of the problem of boosting I attempted to understand it from a Bayesian perspective. The redistribution of weights on observations, the model mixing and the improved performance seemed to have a Bayesian flavor to it. In fact, Freund and Schapire (1997) even offered a naïve Bayes justification for the AdaBoost algorithm. With the current results in hand it now seems clear that the Bayesian interpretation of boosting is that it is simply maximum likelihood! Boosting suggests moving the current function in some optimal direction and assembles a path through function space toward a maximum likelihood estimate. The Bayesian wants to move the current function in all directions at each iteration, eventually considering all paths through the function space with each weighted by some posterior probability. As is often the case, computation can impede a fully Bayesian solution. However, the lessons learned from boosting could be emminently useful for Bayesian prediction and function estimation problems.

In Bayesian prediction problems we want to find the predictive distribution of a y given its covariates, x, and a dataset containing previous observations,

$$\Pr(y|x, D) = \int \Pr(y|x, F)\Pr(F|D)\, dF. \tag{6}$$

The predictive distribution integrates over all members, F, in some class of functions. The crudest approximation simply assumes that $\Pr(F|D)$ has most of its mass in a small neighborhood of the maximum likelihood estimate of F so that $\Pr(y|x, D) \approx \Pr(y|x, \hat{F})$. To make a point prediction one would compute some functional of $\Pr(y|x, \hat{F})$ (e.g., mean, class with the largest probability). Recent efforts have proposed methods for sampling from $\Pr(F|D)$ in order to compute a Monte Carlo estimate of (6). Some of these ideas include MCMC on the space of trees [Denison, Mallick and Smith (1996), Chipman, George and McCulloch (1998)], on the space of additive models with a partic-

ular smoothness [Hastie and Tibshirani (1998)] and reversible jump MCMC on the space of piecewise constants [Heikkinen (1998)] or splines [DiMatteo, Genovese and Kass (1999)]. I think boosting may have a lot to offer the latter of these methods.

Heikkinen (1998) assumes that the regression function is a piecewise constant where the number of pieces and the constants associated with those pieces are random quantities. The simulation from $\Pr(F|D)$ proceeds by either randomly combining two neighboring pieces, splitting a random piece into two, or randomly perturbing the constant associated with one of the pieces. Among the curses of MCMC methods is the problem of wasting iterations exploring regions of the parameter space that are not particularly of interest or not exploring the entire interesting region. Each boosting iteration computes some kind of gradient in the function space indicating with respect to the current state which neighboring functions are most interesting in some sense. From the insight gained from boosting one could potentially produce more efficient proposal mechanisms for exploring $\Pr(F|D)$; that is in the works.

Closing comments. This work melds some exciting work in both computational learning theory and statistics. The authors have elegantly connected the AdaBoost algorithm with a novel interpretation of fitting additive logistic regression. Researchers from a diverse set of fields are in search of better methods for modeling datasets that are massive, noisy, complex and high dimensional. Boosting, as a modular procedure for constructing such models, offers an exciting paradigm for future prediction methods.

REFERENCES

BAUER, E. and KOHAVI, R. (1999). An empirical comparison of voting classification algorithms: bagging, boosting, and variants. *Machine Learning* **36** 105–139.

BREIMAN, L. (1999). Using adaptive bagging to debias regressions. Technical Report 547, Dept. Statistics, Univ. California, Berkeley.

CHIPMAN, H., GEORGE, E. and MCCULLOCH, R. (1998). Bayesian CART model search (with discussion). *J. Amer. Statist. Assoc.* **93** 935–960.

DENISON, D. (2000). Boosting with Bayesian stumps. Technical report, Dept. Mathematics, Imperial College.

DENISON, D., MALLICK, B. and SMITH, A. (1996). Bayesian CART. Technical report, Dept. Mathematics, Imperial College.

DIMATTEO, I., GENOVESE, C. and KASS, R. (1999). Bayesian curve fitting with free-knot splines. Technical report, Carnegie Mellon Univ.

DRUCKER, H. and CORTES, C. (1996). Boosting decision trees. In *Proceedings of Neural Information Processing 8* 479–485. MIT Press.

ELKAN, C. (1997). Boosting and naïve Bayes learning. Technical Report CS97-557, Univ. California, San Diego.

FREUND, Y. and SCHAPIRE, R. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* **55** 119–139.

HASTIE, T. and TIBSHIRANI, R. (1998). Bayesian backfitting. Technical report, Stanford Univ.

HEIKKINEN, J. (1998). Curve and surface estimation using dynamic step functions. In *Practical Nonparametric and Semiparametric Bayesian Statistics* (D. Dey, P. Müller and D. Sinha, eds.) 255–272. Springer, New York.

LEE, J. (1999). A computer program that plays a hunch. *New York Times* August 17.

QUINLAN, J. (1996). Bagging, boosting, and C4.5. In *Proceedings Thirteenth American Association for Artificial Intelligence National Conference on Artificial Intelligence* 725–730. AAAI Press, Menlo Park, CA.

RIDGEWAY, G. (1999). The state of boosting. In *Proceedings of the Thirty-first Symposium on the Interface* 172–181.

RIDGEWAY, G. (1999). The state of boosting. In *Computing Science and Statistics* **31** (K. Berk, M. Pourahmadi, eds.) Interface Foundation of North America, 172–181. Fairfax, VA.

DEPARTMENT OF STATISTICS
UNIVERSITY OF WASHINGTON
SEATTLE, WASHINGTON 98195-4322
E-MAIL: greg@stat.washington.edu

REJOINDER

JEROME FRIEDMAN, TREVOR HASTIE AND ROBERT TIBSHIRANI

Stanford University

We thank the discussants for their generous and thoughtful comments, and the editors for arranging this discussion. Since the paper and discussions are long, we can only reply to a few of the many important points raised by the discussants.

1. Overfitting. All of the discussions raise the issue of overfitting in the context of boosting. Breiman, and to a lesser extent, Freund and Schapire suggest that our explanation is inadequate since it suggests that enough boosting steps can lead to overfitting. Breiman cites empirical evidence suggesting that "AdaBoost almost never overfits the data no matter how many iterations it is run." It would be a nice world if that were the case, but unfortunately it isn't quite. Besides the example in our paper, Ridgeway provides a simple example that exhibits dramatic overfitting in Figure 2 of his discussion. Quoting Ratsch, Onoda and Muller (2000), "Recent studies with highly noisy patterns [Quinlan (1996), Grove and Schuurmans (1998), Ratsch (1998)] showed that it is clearly a myth that Boosting methods do not overfit." Still, it does appear that boosting is more resistant to overfitting that one might expect based on general modeling experience. Buja suggests that the suboptimal greedy "stagewise" nature of the fitting procedure may provide some resistance. This could be the case, and merits further investigation. However, we agree with Bühlmann and Yu that it is largely the nature of zero–one loss that provides general resistance to overfitting as judged in classification.

We illustrate this with a simple type of "nearest-neighbor" procedure. Consider a fixed set of data points $\{\mathbf{x}_i\}_1^N$. A training data set consists of $\{y_i, \mathbf{x}_i\}_1^N$, where each $y_i \in \{-1, 1\}$ is a random variable with $p_i = \Pr(y_i = 1)$. Given a training set, our estimate for each p_i is $\hat{p}_i = (1 + \alpha y_i)/2$, $0 \leq \alpha \leq 1$. Here α is a parameter that controls the degree of fit to the training data. Larger values provide a closer fit; the value $\alpha = 1$ interpolates the data, whereas $\alpha = 0$

TABLE 1

Criterion	$L(y, F)$	$\hat{R}$	R	a^*
exponential	$\exp(-yF)$	$(\frac{1-a}{1+a})^{1/2}$	$(1-\bar{e})(\frac{1-a}{1+a})^{1/2}$ $+\bar{e}(\frac{1+a}{1-a})^{1/2}$	$1-2\bar{e}$
likelihood	$\log(1+\exp(-2yF))$	$\log\frac{2}{1+a}$	$(1-\bar{e})\log\frac{2}{1+a}$ $+\bar{e}\log\frac{2}{1-a}$	$1-2\bar{e}$
squared error	$(y-F)^2$	$(1-a)^2$	$(1-a)^2+4a\bar{e}$	$1-2\bar{e}$
absolute error	$\lvert y-F\rvert$	$1-a$	$1-(1-2\bar{e})a$	1
zero–one loss	$1_{[y\neq\text{sign}(F)]}$	0	$\bar{e}$	>0

estimates $\hat{p}_i = 1/2$ irrespective of the data. We consider five loss functions $L(y, F)$—exponential, likelihood, squared error, absolute error, and zero–one loss—all defined in Table 1.

For each loss function, F is defined to be the population minimizer

$$\arg\min_F EL(y, F) = \arg\min_F pL(1, F) + (1-p)L(-1, F).$$

For the exponential and likelihood, this is $F = \frac{1}{2}\log(p/(1-p))$, for squared and absolute error $F = 2p-1$, and for zero–one loss $F = \text{sign}(2p-1)$.

Table 1 shows the average loss on the training data

$$\hat{R} = \frac{1}{N}\sum_{i=1}^{N} L(y_i, \hat{F}_i)$$

and the corresponding average population risk

$$R = \frac{1}{N}\sum_{i=1}^{N} EL(y_i^*, \hat{F}_i).$$

The later expectation is with respect to the training y_i (which produces $\hat{F}_i$), and the test observation y_i^*, both drawn independently with probability p_i. The $\hat{F}_i$ are obtained by plugging $\hat{p}_i$ into the expressions for F. For the exponential and likelihood criteria,

$$\hat{F}_i = \frac{1}{2}\log\frac{\hat{p}_i}{1-\hat{p}_i} = \frac{1}{2}\log\frac{1+ay_i}{1-ay_i},$$

whereas for the squared and absolute error criteria,

$$\hat{F}_i = ay_i.$$

The risk values are reported in terms of the parameter a and

$$\bar{e} = \frac{1}{N}\sum_{i=1}^{N} E1_{[y_i^*\neq\text{sign}(F_i^*)]} = \frac{2}{N}\sum_{i=1}^{N} p_i(1-p_i),$$

which is just the average risk based on zero–one loss (and is twice the Bayes risk). This error rate does not depend on the value of a since it is independent of the magnitude of $\hat{F}_i$, depending only on its sign. Thus, overfitting as measured by zero–one loss does not exist for this procedure. Interpolating the data by setting $a = 1$ does not degrade zero–one loss. Also shown in Table 1 is the parameter value a^* that minimizes the population risk for each respective criterion.

As seen from Table 1, both the average training set loss and population risk for the four criteria, other than zero–one loss, do depend on the magnitude of $\hat{F}_i$, and thus depend strongly on how close the training data are fit, as regulated by the value of a. The average training set loss $\hat{R}$ is monotone decreasing with increasing values of a. Starting at $a = 0$, the population risk R for all four criteria initially decreases. For the first three a minimum is reached at the value $a^* = 1 - 2\bar{e}$. Increasing the value of a beyond that point monotonically increases population risk, symptomatic of classic overfitting. The degree of overfitting depends on the criterion. Criteria listed higher in Table 1 report more overfitting. For absolute error there is no overfitting; interpolating the training data here minimizes population risk for this criterion.

Figure 1 of this Rejoinder shows both training error (lowest green curve), population risk for $\bar{e} = 0.1$ (middle red curve), and $\bar{e} = 0.2$ (highest blue curve), as a function of $1/(1 - a)$ for the loss criteria shown in Table 1. All curves are normalized to have unit risk at $a = 0$. Note the difference in vertical scale between the lower and upper plots. The first three criteria exhibit classic overfitting. The degree of this overfit increases with increasing intrinsic noise as reflected by the value of $\bar{e}$. For exponential and likelihood loss the expected test error approaches infinity as the value of a approaches one (data interpolation), the latter much more slowly. Both squared and absolute loss approach finite values of $4\bar{e}$ and $2\bar{e}$, respectively. Except for the discontinuity at $a = 0$ (where the risk is that of a coin flip, $R = \frac{1}{2}$, normalized to be 1), the population zero–one risk curves drop immediately to $\bar{e}$. When the abscissa is plotted on a linear rather than logarithmic scale, the plot for exponential loss (upper left panel) closely resembles that displayed in Figure 3 (lower panel) of the Bühlmann and Yu discussion.

It is important to note that the extent to which overfitting is observed here depends only on the loss criterion used to measure degree-of-fit. For a given training data set, equal values of the parameter a produce identical models for the probabilities $\hat{p}_i$. The judgment as to whether, or how much, a particular model overfits the training data is determined solely by the criterion used to define "fit," and not by the training procedure used to obtain the model. Exponential loss reports dramatic overfitting, likelihood substantial overfitting, squared error loss moderate overfitting and absolute and zero–one loss no overfitting at all, on exactly the same sequence of models.

Of course, this example is highly idealized so as to permit exact calculations of risk. In other situations absolute and zero–one loss may exhibit some degree of overfitting. The assumption that the x-values are fixed prevents zero–one loss from reporting overfitting for this procedure. In actual practice,

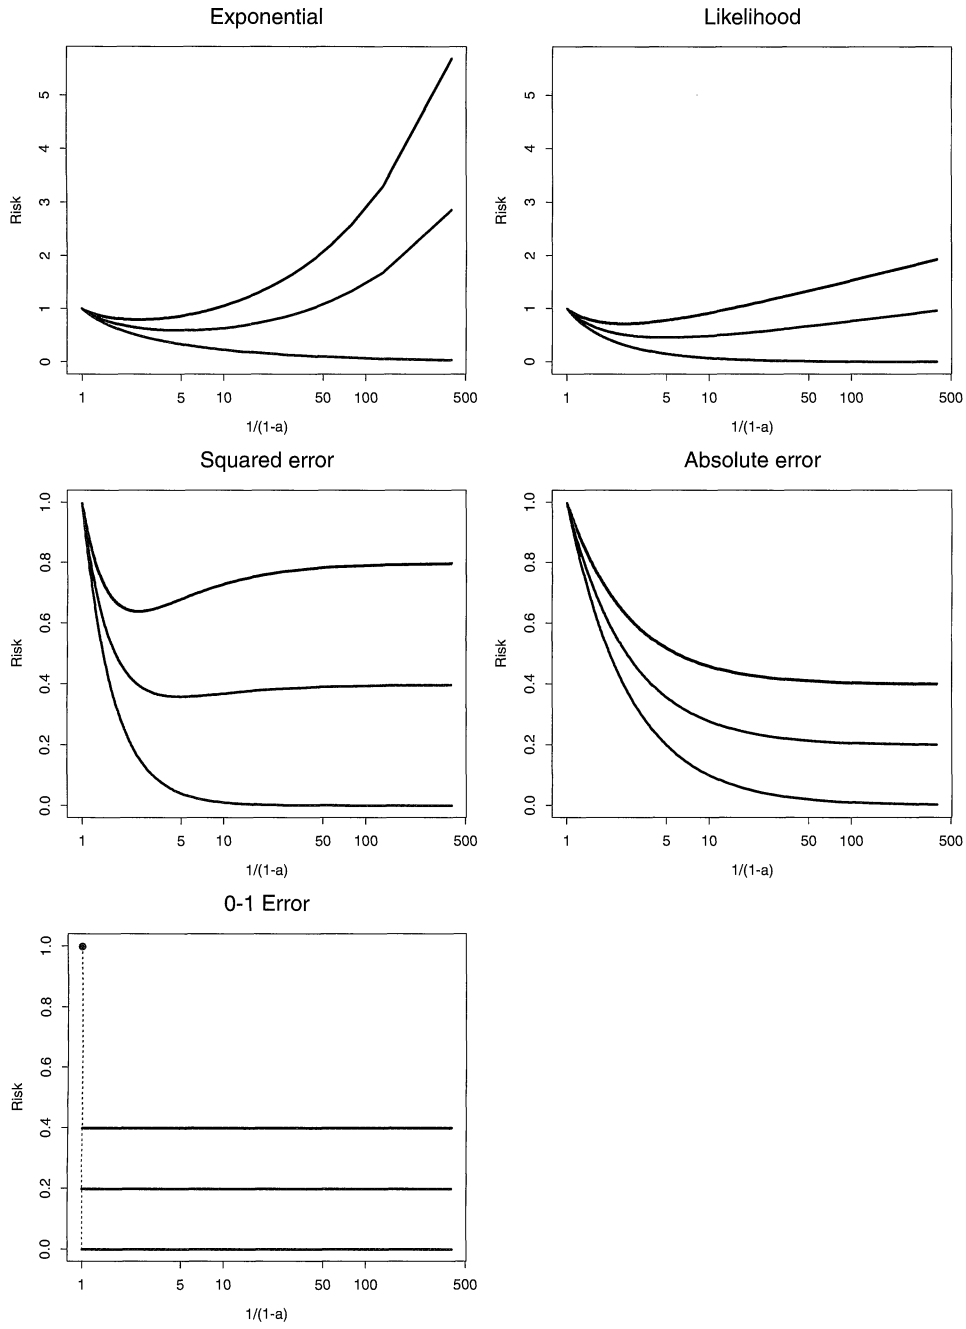


FIG. 1. *Training and population risk curves for the five loss functions in Table* 1, *as a function of* $1/(1-a)$. *The green curves are training risk. The red curves are population risk for* $\bar{e} = 0.1$, *and the blue curves are for* $\bar{e} = 0.2$. *All curves are normalized to have unit risk at* $a = 0$.

the x-values themselves are random variables and fitting the data too closely can result in some incorrect estimates of the sign of F, thereby degrading population zero–one risk. However, the analysis above suggests that overfitting as judged by absolute and especially zero–one loss should be less severe than when judged by exponential, likelihood or squared error loss. Also, there is no reason to expect that the degree-of-fit to the training data that results in minimum population risk will be the same for the former two measures as with the latter three. This is illustrated in Figure 3 of Bühlmann and Yu's discussion.

This analysis has no bearing on the relative merits of any of these criteria as implementing loss functions for boosting, or any other procedure. As pointed out by Bühlmann and Yu, it is often the case that the best criterion to drive a search strategy is not the criterion used to judge the value of the resulting solution. However, the phenomenon of "overfitting" is largely dictated by the latter.

In the field of statistics most of our experience has been derived from likelihood and squared error loss. As Breiman points out "adding more and more variables to a logistic regression, increasing the training likelihood at each step, usually results in overfitting at some point." This is surely the case when overfitting is judged by the likelihood criterion. It can sometimes be the case when judged by zero–one loss. However, the simple example presented here, as well as considerable empirical evidence elsewhere, suggest that zero–one loss, being less sensitive to variations in the estimates of the underlying probabilities (they depend only on the sign of $2\hat{p}_i - 1$), exhibit considerably less overfitting than other loss criteria that depend more heavily on the detailed probability estimates. This will be the case whether those estimates are obtained by AdaBoost, LogitBoost or any other procedure.

2. Regression. Most of the discussions consider the extension of boosting to regression. Buja and Bühlmann and Yu, as well as Ridgeway (1999), observe that our view of boosting leads to natural extensions to regression based on squared error loss. This is discussed in Section 10 of the paper. It is explored in more detail in Friedman (1999a), where gradient based boosting is extended to arbitrary differentiable loss criteria. As noted in Friedman (1999a), shrinkage is an important ingredient to the success of boosting in the regression context. At each step m of the iterative procedure, the current approximation $F_{m-1}(\mathbf{x})$ is updated by

$$F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \nu \cdot f_m(\mathbf{x}). \tag{7}$$

Here $f_m(\mathbf{x})$ is the update computed in the usual manner by the boosting algorithm at the mth step, and $\nu \in (0, 1]$ is a parameter whose value controls the amount of shrinkage. Empirical evidence was presented that such shrinkage dramatically improves the accuracy of the target function estimate when measured by likelihood, squared error or absolute error. Less dramatic but significant gains were also observed for zero–one loss, where the quality of the function estimate (as judged by the other criteria) is less important. In

all cases, accuracy improved monotonically with decreasing value of ν (more shrinkage), provided that the number of iterations M was correspondingly increased so as to reach minimum test error.

Breiman and Bühlmann and Yu suggest that incorporating a stochastic component into boosting algorithms can substantially improve performance. Bühlmann and Yu propose "bag-boosting" which, as Ridgeway points out, is the essence of Breiman's (1999) "adaptive bagging" procedure. The base learner is replaced by the corresponding bagged learner at each boosting step. This can be viewed as an alternative to shrinkage for reducing the learning rate. The individual bagging substeps can be regarded as randomized boosting steps without updating the approximation. Friedman (1999b) investigated combining randomization with shrinkage in an integrated manner. The approximation is slowly but continuously updated at each randomized boost using a slow learning rate, rather than alternating between a large number of steps without updating (bagging), followed by a single step with a full update (boost). Combining randomization with shrinkage in this manner provides a modest but significant additional improvement over the already large gain produced by shrinkage alone. However, there is a substantial computational advantage by employing what Bühlmann and Yu call "subbagging." Friedman and Hall (1999) discuss the equivalence of bagging based on boostrapping, and half-sampling without replacement and show that different subsample sizes can often lead to better performance.

3. Interpretations of boosting. For their invention of boosting, Freund and Schapire deserve a great deal of praise. It has led to both practical learning algorithms and much intellectual study and debate. As Buja so poetically states, "There is no single 'true' interpretation of anything; interpretation is a vehicle in the service of human comprehension. The proof of an interpretation is in enabling others to fruitfully think about an idea." We hope that our work has helped serve that goal. Freund and Schapire and Breiman present two additional interpretations, different from ours and each other. Ridgeway considers yet another view along Bayesian lines. Having multiple interpretations can only increase the chance of obtaining new insights.

We don't see a great deal of contradiction between our view and that of Freund and Schapire. "Provable bounds" can be derived for LogitBoost (Nigel Duffy and David Helmbold, private communication). Also, we don't see how our "assumptions" are necessarily stronger (or weaker) than those motivating their interpretation. It is hard to imagine a binary random variable that is not Bernoulli distributed. We do not assume "a log-linear relationship between the features and the labels." The function class describing this relationship is determined by the weak learner employed. AdaBoost, LogitBoost, or any other boosting method using the same learner is making similar assumptions concerning the label–feature relationship. As was pointed out in the paper, and amplified by Bühlmann and Yu and Buja, "The question is not *whether* but *how much* a weak learner can be boosted, or, in statistical terminology, how well an additive model can approximate an underlying response surface."

The underlying response surface is the unknown label–feature relationship. The weak learner provides the additive components. Different weak learners will perform differently with different response surfaces.

Freund and Mason's (1999) "alternating decision tree" approach is quite interesting. It shares many common elements with MARS [Friedman (1991)], substituting exponential loss for squared error loss, and using zero-order rather than first-order splines. Both of these modifications may be beneficial in the classification context.

Freund's (1999) suggestion of removing outliers is a good one, but not "radical." This approach is at the center of the field of robust statistics. The source of AdaBoost's lack of robustness is its use of exponential loss as the implementing loss function, which increases exponentially with large negative margin. Negative log-likelihood increases only linearly with large negative margin, thereby being much more robust. Still, sometimes performance can be further improved a little by completely removing the influence of observations with very large negative margins. This is accomplished as a beneficial side effect of the weight trimming strategy described in Section 9 of the paper, when used with LogitBoost. The LogitBoost weights are $w_i = p(x_i)(1 - p(x_i))$. Observations with the smallest weights at the current iteration are removed from consideration for that iteration. Observations with large negative margin have either $p(x_i) \simeq 0$ (for $y_i = 1$) or $p(x_i) \simeq 1$ (for $y_i = -1$), thereby being deleted.

Breiman views boosting as an "ensemble" method for combining classifiers. He seeks to incorporate boosting and bagging into a common framework based on "random forests." This is an intriguing concept. However, it seems to us that boosting operates differently than random forests. From the description, a random forests appears to be an ensembles of learners that are i.i.d. given the data. Hence the bias of the ensemble is the same as that of each individual learner. By contrast, in boosting the individual members are not independent given the data: the construction of each depends explicitly on those members that have been previously entered. The bias of the ensemble is typically much less than that of each constituent learner. Boosting turns a weak learner to strong ensemble, whereas random forests combine strong learners in an i.i.d. fashion in order to reduce variance.

Viewing boosting as stagewise additive regression does however suggest Breiman's conclusion that "if boosting succeeds, it is because it gives low correlation between classifiers of reasonable strength." In stagewise regression "each step decouples itself from the previous step" leading to "low correlation between pairs of classifiers." In general, a greedy search strategy will not successively add highly correlated terms, since that would not produce very much local reduction in the criterion being minimized. Also, there will be a tendency to add terms of "reasonable strength," since that will tend to reduce the criterion. Thus, one could conclude that AdaBoost owes its success to these particular aspects of stagewise additive regression. This is consistent with Breiman's explanation, although motivated somewhat differently. Unfortunately, this conclusion appears to be in conflict with one empirical fact,

namely, the success of shrinking (1) in increasing accuracy with all boosting methods, including AdaBoost [Friedman (1999a)].

Shrinking strongly inhibits the natural tendency of greedy search strategies to add low correlation terms of reasonable strength. For smaller values of the shrinkage parameter ν, the strength of each added term $\nu \cdot f_m(\mathbf{x})$ becomes smaller, and more highly correlated with recently entered terms. Yet, the evidence so far indicates that the smaller the value of ν, the higher the overall accuracy, as long as there are enough iterations.

The source of this mystery in under investigation. It is but one of the many aspect of this general class of procedures that are not yet well understood. Whether these procedures are referred to as "boosting," or more prosaically as "stagewise additive regression," there are still many problems left to solve that can challenge us for some time to come.

REFERENCES

BREIMAN, L. (1999). Using adaptive bagging to debias regressions. Technical Report 547, Dept. Statistics, Univ. California, Berkeley.

FREUND, Y. (1999). An adaptive version of boost by majority algorithm. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory*.

FREUND, Y. and MASON, L. (1999). The alternating decision tree learning algorithm. In *Machine Learning: Proceedings of the Sixteenth International Conference* 124–133.

FRIEDMAN, J. H. (1991). Multivariate adaptive regression splines (with discussion). *Ann. Statist.* **19** 1–141.

FRIEDMAN, J. H. (1999a). Greedy function approximation: a gradient boosting machine. *Ann. Statist.* To appear.

FRIEDMAN, J. H. (1999b). Stochastic gradient boosting. Technical report, Dept. Statistics, Stanford Univ.

FRIEDMAN, J. H. and HALL, P. (1999). On bagging and nonlinear estimation. *J. Comput. Graph. Statist.* To appear.

GROVE, A. and SCHUURMANS, D. (1998). Boosting in the limit: maximizing the margin of learned ensembles. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*.

QUINLAN, J. (1996). Boosting first order learning. In *Proceedings of the Seventh International Workshop on Algorithmic Learning Theory* (S. Arikawa and A. Sharma, eds.) *Lecture Notes in Artificial Intelligence* **1160** 143–155. Springer, Berlin.

RATSCH, G. (1998). Ensemble learning methods for classification. Masters thesis, Dept. Computer Science, Univ. Potsdam.

RATSCH, G., ONODA, T. and MULLER, K. R. (2000). Soft margins for AdaBoost. *Machine Learning* 1–35.

RIDGEWAY, G. (1999). The state of boosting. In *Proceedings of the Thirty-first Symposium on the Interface* 172–181.

DEPARTMENT OF STATISTICS
SEQUOIA HALL
370 SERRA MALL
STANFORD UNIVERSITY
STANFORD CALIFORNIA 94305-4065
E-MAIL: hastie@stat.stanford.edu