

## Boosting Neural Networks

**Holger Schwenk**

*LIMSI-CNRS, 91403 Orsay cedex, FRANCE*

**Yoshua Bengio**

*DIRO, University of Montréal, Montréal, Quebec, H3C 3J7, Canada*

Boosting is a general method for improving the performance of learning algorithms. A recently proposed boosting algorithm, AdaBoost, has been applied with great success to several benchmark machine learning problems using mainly decision trees as base classifiers. In this article we investigate whether AdaBoost also works as well with neural networks, and we discuss the advantages and drawbacks of different versions of the AdaBoost algorithm. In particular, we compare training methods based on sampling the training set and weighting the cost function. The results suggest that random resampling of the training data is not the main explanation of the success of the improvements brought by AdaBoost. This is in contrast to bagging, which directly aims at reducing variance and for which random resampling is essential to obtain the reduction in generalization error. Our system achieves about 1.4% error on a data set of on-line handwritten digits from more than 200 writers. A boosted multi-layer network achieved 1.5% error on the UCI letters and 8.1% error on the UCI satellite data set, which is significantly better than boosted decision trees.

### 1 Introduction ---

Boosting, a general method for improving the performance of a learning algorithm, is a method for finding a highly accurate classifier on the training set by combining “weak hypotheses” (Schapire, 1990), each of which needs only to be moderately accurate on the training set. (Perrone, 1994, is an earlier overview of different ways to combine neural networks.) A recently proposed boosting algorithm is AdaBoost (Freund, 1995), which stands for “adaptive boosting.” During the past two years, many empirical studies have been published that use decision trees as base classifiers for AdaBoost (Breiman, 1996; Drucker & Cortes, 1996; Freund & Schapire, 1996a; Quinlan, 1996; Maclin & Opitz, 1997; Grove & Schuurmans, 1998; Bauer & Kohavi, 1999; Dietterich, forthcoming). All of these experiments have shown impressive improvements in the generalization behavior and

suggest that AdaBoost tends to be robust to overfitting. In fact, in many experiments it has been observed that the generalization error continues to decrease toward an apparent asymptote after the training error has reached zero. (Schapire, Freund, Bartlett, & Lee, 1997), suggest a possible explanation for this unusual behavior based on the definition of *margin of classification*. Other attempts to understand boosting theoretically can be found in Schapire et al., 1997; Breiman 1997a, 1998; Friedman, Hastie, and Tibshirani, 1998; and Schapire, 1999. AdaBoost has also been linked with game theory (Breiman, 1997b; Grove & Schuurmans, 1998; Freund & Schapire, 1996b, 1999) in order to understand the behavior of AdaBoost and to propose alternative algorithms. Mason and Baxter (1999) propose a new variant of boosting based on the direct optimization of margins.

There is recent evidence that AdaBoost may very well overfit if we combine several hundred thousand classifiers (Grove & Schuurmans, 1998). It also seems that the performance of AdaBoost degrades a lot in the presence of significant amounts of noise (Dietterich, forthcoming; Rätsch, Onoda, & Müller, 1998). Although much useful theoretical and experimental work has been done, there is still a lot that is not well understood about AdaBoost's impressive generalization behavior. To the best of our knowledge, all applications of AdaBoost have been to decision trees and no applications to multilayer artificial neural networks have been reported in the literature. This article extends and provides a deeper experimental analysis of our first experiments with the application of AdaBoost to neural networks (Schwenk & Bengio, 1997, 1998).

In this article we consider the following questions: Does AdaBoost work as well for neural networks as for decision trees? The short answer is yes, and sometimes even better. Does it behave in a similar way (as was observed previously in the literature)? The short answer is yes. Furthermore, are there particulars in the way neural networks are trained with gradient backpropagation that should be taken into account when choosing a particular version of AdaBoost? The short answer is yes, because it is possible to weight the cost function of neural networks directly. Is overfitting of the individual neural networks a concern? The short answer is, not as much as when not using boosting. Is the random resampling used in previous implementations of AdaBoost critical, or can we get similar performances by weighing the training criterion (which can easily be done with neural networks)? The short answer is that it is not critical for generalization but helps to obtain faster convergence of individual networks when coupled with stochastic gradient descent.

In the next section, we describe the AdaBoost algorithm and discuss several implementation issues when using neural networks as base classifiers. In section 3, we present results that we have obtained on three medium-sized tasks: a data set of handwritten on-line digits and the letter and satellite data sets of the UCI repository. We finish with a conclusion and perspectives for future research.

## 2 AdaBoost

---

It is well known that it is often possible to increase the accuracy of a classifier by averaging the decisions of an ensemble of classifiers (Perrone, 1993; Krogh & Vedelsby, 1995). In general, more improvement can be expected when the individual classifiers are diverse and yet accurate. One can try to obtain this result by taking a base learning algorithm and invoking it several times on different training sets. Two popular techniques exist that differ in the way they construct these training sets: bagging (Breiman, 1994) and boosting (Schapire, 1990; Freund, 1995). In bagging, each classifier is trained on a bootstrap replicate of the original training set. Given a training set  $S$  of  $N$  examples, the new training set is created by resampling  $N$  examples uniformly with replacement. Some examples may occur several times, and others may not occur in the sample at all. One can show that, on average, only about two-thirds of the examples occur in each bootstrap replicate. The individual training sets are independent, and the classifiers could be trained in parallel. Bagging is known to be particularly effective when the classifiers are unstable, that is, when perturbing the learning set can cause significant changes in the classification behavior. Formulated in the context of the bias-variance decomposition (Geman, Bienenstok, & Doursat, 1992), bagging improves generalization performance due to a reduction in variance while maintaining or only slightly increasing bias. Note, however, there is no unique bias-variance decomposition for classification tasks (Kong & Dietterich, 1995; Breiman, 1996; Kohavi & Wolpert, 1996; Tibshirani, 1996).

AdaBoost, on the other hand, constructs a composite classifier by sequentially training classifiers while putting more and more emphasis on certain patterns. For this, AdaBoost maintains a probability distribution  $D_t(i)$  over the original training set. In each round  $t$ , the classifier is trained with respect to this distribution. Some learning algorithms do not allow training with respect to a weighted cost function. In this case, sampling with replacement (using the probability distribution  $D_t$ ) can be used to approximate a weighted cost function. Examples with high probability would then occur more often than those with low probability, and some examples may not occur in the sample at all, although their probability is not zero.

After each AdaBoost round, the probability of incorrectly labeled examples is increased, and the probability of correctly labeled examples is decreased. The result of training the  $t$ th classifier is a hypothesis  $h_t: X \rightarrow Y$  where  $Y = \{1, \dots, k\}$  is the space of labels, and  $X$  is the space of input features. After the  $t$ th round, the weighted error  $\epsilon_t = \sum_i h_t(x_i) \neq y_i D_t(i)$  of the resulting classifier is calculated, and the distribution  $D_{t+1}$  is computed from  $D_t$  by increasing the probability of incorrectly labeled examples. The probabilities are changed so that the error of the  $t$ th classifier using these new “weights”  $D_{t+1}$  would be 0.5. In this way, the classifiers are optimally decoupled. The global decision  $f$  is obtained by weighted voting. This ba-



Table 1: Pseudo-Loss AdaBoost (AdaBoost.M2).

<b>Input:</b>	Sequence of $N$ examples $(x_1, y_1), \dots, (x_N, y_N)$ with labels $y_j \in Y = \{1, \dots, k\}$ .
<b>Init:</b>	Let $B = \{(i, y) : i \in \{1, \dots, N\}, y \neq y_i\}$ $D_1(i, y) = 1/ B $ for all $(i, y) \in B$ .
<b>Repeat:</b>	<ol style="list-style-type: none"> <li>1. Train neural network with respect to distribution <math>D_t</math> and obtain hypothesis <math>h_t : X \times Y \rightarrow [0, 1]</math>.</li> <li>2. Calculate the pseudo-loss of <math>h_t</math>:  <math display="block">\epsilon_t = \frac{1}{2} \sum_{(i,y) \in B} D_t(i, y)(1 - h_t(x_i, y_i) + h_t(x_i, y)).</math> </li> <li>3. Set <math>\beta_t = \epsilon_t / (1 - \epsilon_t)</math>.</li> <li>4. Update distribution <math>D_t</math>:  <math display="block">D_{t+1}(i, y) = \frac{D_t(i, y)}{Z_t} \beta_t^{\frac{1}{2}(1 + h_t(x_i, y_i) - h_t(x_i, y))},</math>                     where <math>Z_t</math> is a normalization constant.                 </li> </ol>
<b>Output:</b>	Final hypothesis: $f(x) = \arg \max_{y \in Y} \sum_t \left( \log \frac{1}{\beta_t} \right) h_t(x, y).$

the margin distribution for the generalization behavior of ensemble techniques can be found in Freund & Schapire (1996b), Schapire et al. (1997), Breiman, 1997a,1997b), Grove and Schuurmans (1998), and Rätsch et al. (1998).

In this article, an important focus is on whether the good generalization performance of AdaBoost is partially explained by the random resampling of the training sets generally used in its implementation. This issue will be addressed by comparing three versions of AdaBoost in which randomization is used (or not used) in three different ways.

**2.1 Applying AdaBoost to Neural Networks.** In this article we investigate different techniques of using neural networks as base classifiers for AdaBoost. In all cases, we have trained the neural networks by minimizing a quadratic criterion that is a weighted sum of the squared differences  $(z_{ij} - \hat{z}_{ij})^2$ , where  $z_i = (z_{i1}, z_{i2}, \dots, z_{ik})$  is the desired output vector (with a low target value everywhere except at the position corresponding to the target class) and  $\hat{z}_i$  is the output vector of the network. A score for class  $j$  for pattern  $i$  can be directly obtained from the  $j$ th element  $\hat{z}_{ij}$  of the output vector  $\hat{z}_i$ . When a class must be chosen, the one with the highest score is selected. Let  $V_t(i, j) = D_t(i, j) / \max_{k \neq y_i} D_t(i, k)$  for  $j \neq y_i$  and  $V_t(i, y_i) = 1$ . These weights are used to give more emphasis to certain incorrect labels according to the pseudo-loss AdaBoost.

What we call *epoch* is a pass of the training algorithm through all the examples in a training set. We compare three different versions of AdaBoost:

- R: Training the  $t$ th classifier with a fixed training set obtained by resampling with replacement once from the original training set. Before starting training the  $t$ th network, we sample  $N$  patterns from the original training set, each time with a probability  $P_t(i)$  of picking pattern  $i$ . Training is performed for a fixed number of iterations always using this same resampled training set. This is basically the scheme that has been used in the past when applying AdaBoost to decision trees, except that we used the pseudo-loss AdaBoost. To approximate the pseudo-loss, the training cost that is minimized for a pattern that is the  $i$ th one from the original training set is  $\frac{1}{2} \sum_j V_t(i, j)(z_{ij} - \hat{z}_{ij})^2$ .
- E: Training the  $t$ th classifier using a different training set at each epoch, by resampling with replacement after each training epoch. After each epoch, a new training set is obtained by sampling from the original training set with probabilities  $P_t(i)$ . Since we used an on-line (stochastic) gradient in this case, this is equivalent to sampling a new pattern from the original training set with probability  $P_t(i)$  before each forward or backward pass through the neural network. Training continues until a fixed number of pattern presentations has been performed. As for version R, the training cost that is minimized for a pattern that is the  $i$ th one from the original training set is  $\frac{1}{2} \sum_j V_t(i, j)(z_{ij} - \hat{z}_{ij})^2$ .
- W: Training the  $t$ th classifier by directly weighting the cost function (here the squared error) of the  $t$ th neural network, that is, all the original training patterns are in the training set, but the cost is weighted by the probability of each example:  $\frac{1}{2} \sum_j D_t(i, j)(z_{ij} - \hat{z}_{ij})^2$ . If we used directly this formula, the gradients would be very small, even when all probabilities  $D_t(i, j)$  are identical. To avoid having to scale learning rates differently depending on the number of examples, the following “normalized” error function was used:

$$\frac{1}{2} \frac{P_t(i)}{\max_k P_t(k)} \sum_j V_t(i, j)(z_{ij} - \hat{z}_{ij})^2. \quad (2.2)$$

In versions E and W, what makes the combined networks essentially different from each other is the fact that they are trained with respect to different weightings  $D_t$  of the original training set. In R, an additional element of diversity is built in because the criterion used for the  $t$ th network is not exactly the errors weighted by  $P_t(i)$ . Instead, more emphasis is put on certain patterns while completely ignoring others (because of the initial random sampling of the training set). The E version can be seen as a stochastic version of the W version, as the number of iterations through the data increases and the learning rate decreases, E becomes a

very good approximation of  $W$ .  $W$  itself is closest to the recipe mandated by the AdaBoost algorithm (but, as we will see below, it suffers from numerical problems). Note that  $E$  is a better approximation of the weighted cost function than  $R$ , in particular when many epochs are performed. If random resampling of the training data explained a good part of the generalization performance of AdaBoost, then the weighted training version ( $W$ ) should perform worse than the resampling versions, and the fixed sample version ( $R$ ) should perform better than the continuously resampled version ( $E$ ). Note that for bagging, which directly aims at reducing variance, random resampling is essential to obtain the reduction in generalization error.

### 3 Results

---

Experiments have been performed on three data sets: a data set of on-line handwritten digits, the UCI letters data set of off-line machine-printed alphabetical characters, and the UCI satellite data set generated from Landsat multispectral scanner image data. All data sets have a predefined training and test set. All the  $p$ -values that are given in this section concern a pair  $(\hat{p}_1, \hat{p}_2)$  of test performance results (on  $n$  test points) for two classification systems with unknown true error rates  $p_1$  and  $p_2$ . The null hypothesis is that the true expected performance for the two systems is not different, that is,  $p_1 = p_2$ . Let  $\hat{p} = 0.5(\hat{p}_1 + \hat{p}_2)$  be the estimator of the common error rate under the null hypothesis. The alternative hypothesis is that  $p_1 < p_2$ , so the  $p$ -value is obtained as the probability of observing such a large difference under the null hypothesis, that is,  $P(Z > z)$  for a normal  $Z$ , with  $z = \frac{\sqrt{n}(\hat{p}_1 - \hat{p}_2)}{\sqrt{2\hat{p}(1-\hat{p})}}$ . This is based on the normal approximation of the binomial, which is appropriate for large  $n$  (however, see Dietterich, 1998, for a discussion of this and other tests to compare algorithms).

**3.1 Results on the On-line Data Set.** The on-line data set was collected at Paris 6 University (Schwenk & Milgram, 1996). A WACOM A5 tablet with a cordless pen was used in order to allow natural writing. Since we wanted to build a writer-independent recognition system, we tried to use many writers and impose as few constraints as possible on the writing style. In total, 203 students wrote down isolated numbers that have been divided into a learning set (1200 examples) and a test set (830 examples). The writers of both sets are completely distinct. A particular property of this data set is the notable variety of writing styles that are not equally frequent at all. There are, for instance, 100 zeros written counterclockwise, but only 3 written clockwise. Figure 1 gives an idea of the great variety of writing styles of this data set. We applied only a simple preprocessing: the characters were resampled to 11 points, centered and size normalized to an  $(x, y)$ -coordinate sequence in  $[-1, 1]^{22}$ .

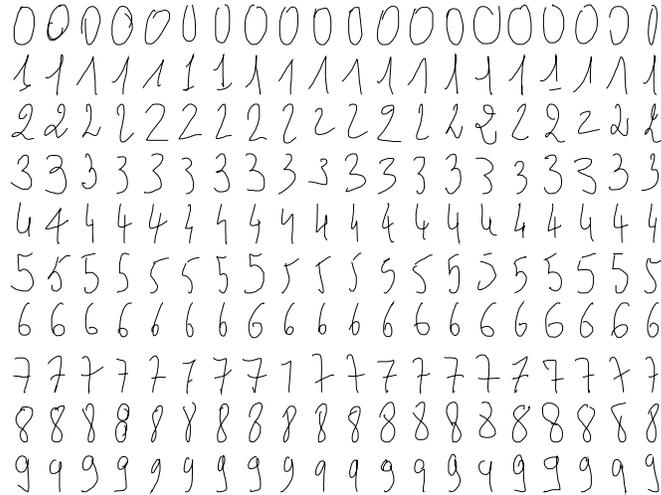


Figure 1: Some examples of the on-line handwritten digits data set (test set).

Table 2 summarizes the results on the test set before using AdaBoost. Note that the differences among the test results on the last three networks are not statistically significant ( $p$ -value  $> 23\%$ ), whereas the difference with the first network is significant ( $p$ -value  $< 10^{-5}$ ). Tenfold cross-validation within the training set was used to find the optimal number of training epochs (typically about 200). Note that if training is continued until 1000 epochs, the test error increases by up to 1%.

Table 3 shows the results of bagged and boosted multilayer perceptrons with 10, 30, or 50 hidden units, trained for either 100, 200, 500, or 1000 epochs, and using the ordinary resampling scheme (R), resampling with different random selections at each epoch (E), or training with weights  $D_t$  on the squared error criterion for each pattern (W). In all cases, 100 neural networks were combined.

Table 2: On-line Digits Data Set Error Rates for Fully Connected MLPs (Not Boosted).

Architecture <sup>a</sup>	22-10-10	22-30-10	22-50-10	22-80-10
Train:	5.7%	0.8%	0.4%	0.08%
Test:	8.8	3.3	2.8	2.7

<sup>a</sup>The notation 22- $h$ -10 designates a fully connected neural network with 22 input nodes, one hidden layer with  $h$  neurons, and a 10-dimensional output layer.

Table 3: On-line Digits Test Error Rates for Boosted MLPs.

Architecture Version	22-10-10			22-30-10			22-50-10		
	R	E	W	R	E	W	R	E	W
<b>Bagging</b>	5.4%			2.8%			1.8%		
500 epochs	5.4%			2.8%			1.8%		
<b>AdaBoost</b>									
100 epochs	2.9%	3.2%	6.0%	1.7%	1.8%	5.1%	2.1%	1.8%	4.9%
200 epochs	3.0	2.8	5.6	1.8	1.8	4.2	1.8	1.7	3.5
500 epochs	2.5	2.7	3.3	1.7	1.5	3.0	1.7	1.7	2.8
1000 epochs	2.8	2.7	3.2	1.8	1.6	2.6	1.6	1.5	2.2
5000 epochs	—	—	2.9	—	—	1.6	—	—	1.6

AdaBoost improved in all cases the generalization error of the MLPs—for instance, from 8.8% to about 2.7% for the 22-10-10 architecture. The improvement with 50 hidden units from 2.8% (without AdaBoost) to 1.6% (with AdaBoost) is significant ( $p$ -value of 0.38%), despite the small number of examples. Boosting was also always superior to bagging, although the differences are not always very significant because of the small number of examples. Furthermore, it seems that the number of training epochs of each individual classifier has no significant impact on the results of the combined classifier, at least on this data set. AdaBoost with weighted training of MLPs (W version), however, does not work as well if the learning of each individual MLP is stopped too early (1000 epochs). The networks did not learn well enough the weighted examples, and  $\epsilon_t$  rapidly approached 0.5. When training each MLP for 5000 epochs, however, the weighted training (W) version achieved the same low test error rate.

AdaBoost is less useful for very big networks (50 or more hidden units for this data) since an individual classifier can achieve zero error on the original training set (using the E or W method). Such large networks probably have a very low bias but high variance. This may explain why bagging, a pure variance-reduction method, can do as well as AdaBoost, which is believed to reduce bias and variance. However, AdaBoost can achieve the same low error rates with the smaller 22-30-10 networks. Figure 2 shows the error rates of some of the boosted classifiers as the number of networks is increased. AdaBoost brings the training error to zero after only a few steps, even with an MLP with only 10 hidden units. The generalization error is also considerably improved, and it continues to decrease to an apparent asymptote after zero training error has been reached.

The surprising effect of continuously decreasing generalization error even after training error reaches zero has already been observed by others (Breiman, 1996; Drucker & Cortes, 1996; Freund & Schapire, 1996a; Quinlan, 1996). This seems to contradict Occam's razor, but a recent theorem (Schapire et al., 1997) suggests that the margin distribution may be relevant



to the generalization error. Although previous empirical results (Schapire et al., 1997) indicate that pushing the margin cumulative distribution to the right may improve generalization, other recent results (Breiman, 1997a, 1997b; Grove & Schuurmans, 1998) show that “improving” the whole margin distribution can also yield worse generalization. Figures 3 and 4 show several margin cumulative distributions, that is, the fraction of examples whose margin is at most  $x$  as a function of  $x \in [-1, 1]$ . The networks have been trained for 1000 epochs (5000 for the W version).

It is clear in Figures 3 and 4 that the number of examples with high margin increases when more classifiers are combined by boosting. When boosting neural networks with 10 hidden units, for instance, there are some examples with a margin smaller than  $-0.5$  when only two networks are combined. However, all examples have a positive margin when 10 nets are combined, and all examples have a margin higher than  $0.5$  for 100 networks. Bagging, on the other hand, has no significant influence on the margin distributions. There is almost no difference in the margin distributions of the R, E, or W version of AdaBoost either.<sup>2</sup> However, there is a difference between the margin distributions and the test set errors when the complexity of the neural networks is varied (hidden layer size). Finally, it seems that sometimes AdaBoost must allow some examples with very high margins in order to improve the minimal margin. This can best be seen for the 22-10-10 architecture.

One should keep in mind that this data set contains only small amounts of noise. In application domains with high amounts of noise, it may be less advantageous to improve the minimal margin at any price (Grove & Schuurmans, 1998; Rätsch et al., 1998), since this would mean putting too much weight to noisy or wrongly labeled examples.

**3.2 Results on the UCI Letters and Satellite Data Sets.** Similar experiments were performed with MLPs on the letters data set from the UCI machine learning data set. It has 16,000 training and 4000 test patterns, 16 input features, and 26 classes (A–Z) of distorted machine-printed characters from 20 fonts. A few preliminary experiments on the training set only were used to choose a 16-70-50-26 architecture. Each input feature was normalized according to its mean and variance on the training set.

Two types of experiments were performed: (1) doing resampling after each epoch (E) and using stochastic gradient descent and (2) without resampling but using reweighting of the squared error (W) and conjugate gradient descent. In both cases, a fixed number of training epochs (500) was used. The plain, bagged, and boosted networks are compared to decision trees in Table 4.

In both cases (E and W) the same final generalization error results were obtained (1.50% for E and 1.47% for W), but the training time using the

---

<sup>2</sup> The W and E versions achieve slightly higher margins than R.

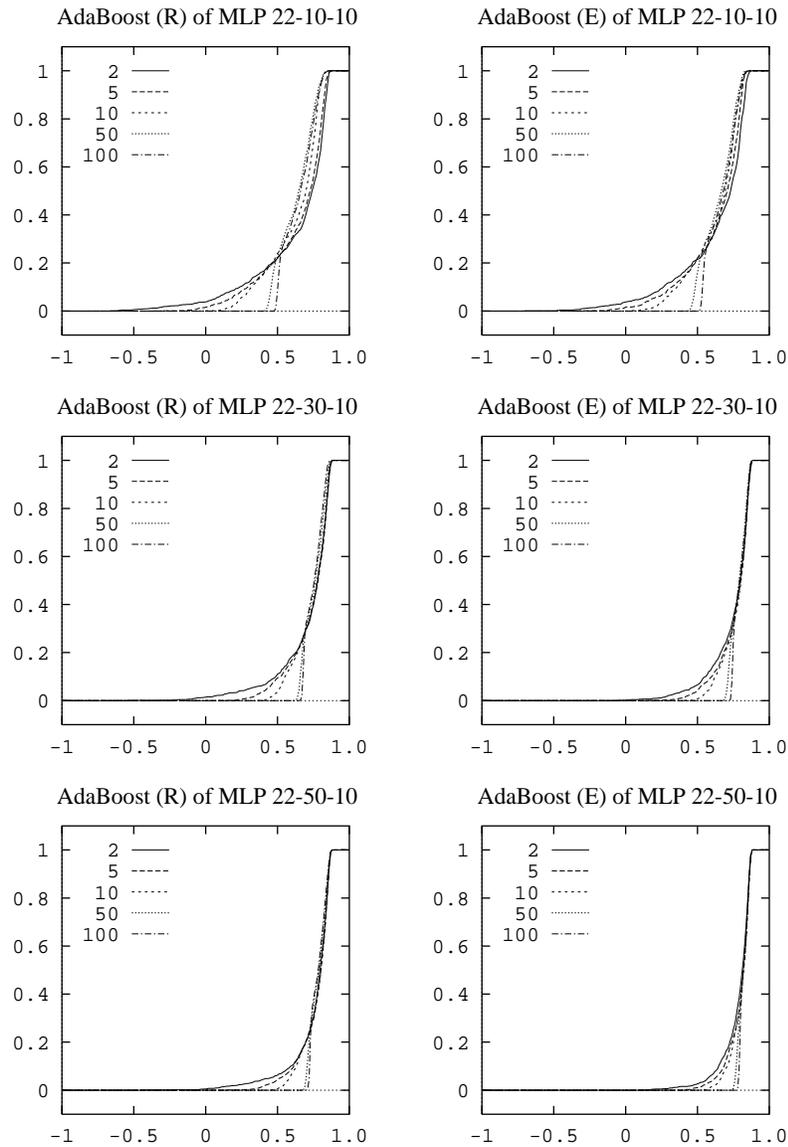


Figure 3: Margin distributions using 2, 5, 10, 50, and 100 networks, respectively.

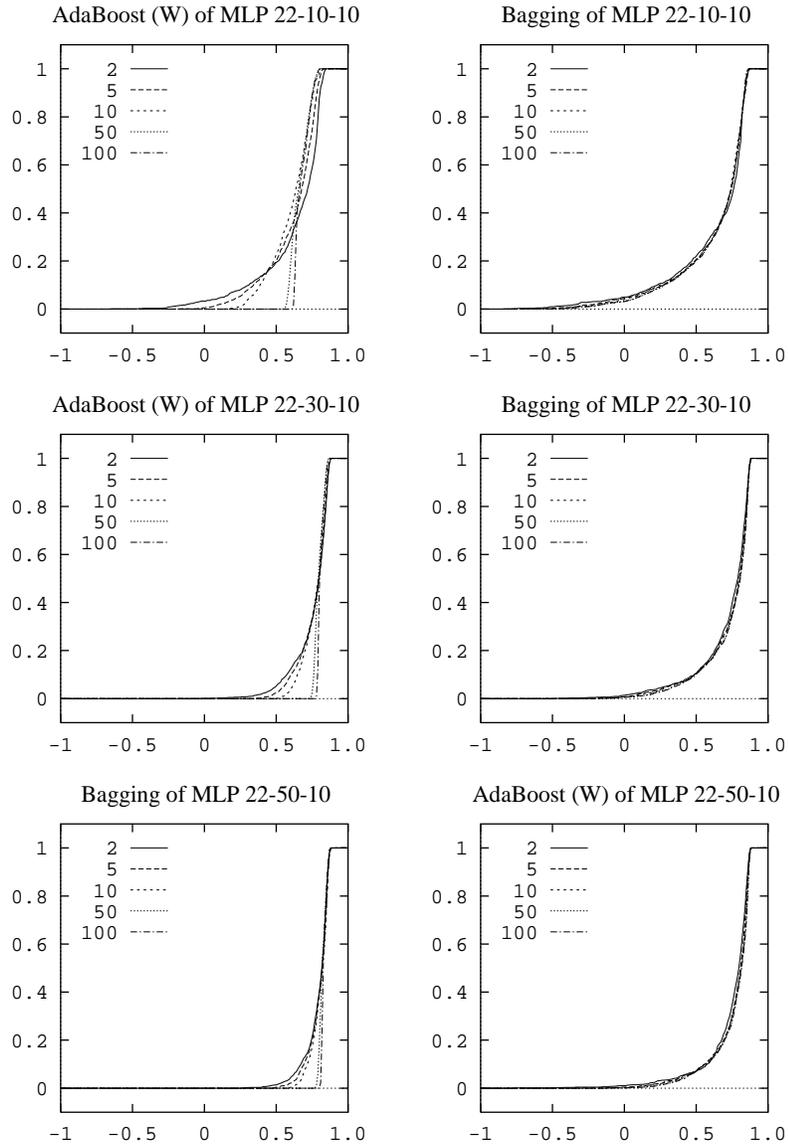


Figure 4: Margin distributions using 2, 5, 10, 50, and 100 networks, respectively.

Table 4: Test Error Rates on the UCI Data Sets.

Data Set	CART <sup>a</sup>			C4.5 <sup>b</sup>			MLP		
	Alone	Bagged	Boosted	Alone	Bagged	Boosted	Alone	Bagged	Boosted
Letter	12.4%	6.4%	3.4%	13.8%	6.8%	3.3%	6.1%	4.3%	1.5%
Satellite	14.8	10.3	8.8	14.8	10.6	8.9	12.8	8.7	<b>8.1</b>

<sup>a</sup> Results from Breiman (1996).

<sup>b</sup> Results from Freund & Schapire (1996a).

weighted squared error (W) was about 10 times greater. This shows that using random resampling (as in E or R) is not necessary to obtain good generalization (whereas it is clearly necessary for bagging). However, the experiments show that it is still preferable to use a random sampling method such as R or E for numerical reasons: convergence of each network is faster.

For this reason, for the E experiments with stochastic gradient descent, 100 networks were boosted, whereas we stopped training on the W network after 25 networks (when the generalization error seemed to have flattened out), which took more than a week on a fast processor (SGI Origin-2000). We believe that the main reason for this difference in training time is that the conjugate gradient method is a batch method and is therefore slower than stochastic gradient descent on redundant data sets with many thousands of examples, such as this one. (See comparisons between batch and on-line methods in Bourely, 1989, and conjugate gradients for classification tasks in particular Moller, 1992, 1993.) For the W version with stochastic gradient descent, the weighted training error of individual networks does not decrease as much as when using conjugate gradient descent, so AdaBoost itself did not work as well. We believe that this is due to the fact that it is difficult for stochastic gradient descent to approach a minimum when the output error is weighted with very different weights for different patterns (the patterns with small weights make almost no progress). On the other hand, the conjugate gradient descent method can approach a minimum of the weighted cost function more precisely, but inefficiently, when there are thousands of training examples.

The results obtained with the boosted network are extremely good—1.5% error, whether using the W version with conjugate gradients or the E version with stochastic gradient—and are the best ever published to date, as far as we know, for this data set. In a comparison with the boosted trees (3.3% error), the *p*-value of the null hypothesis is less than  $10^{-7}$ . The best performance reported in STATLOG (Feng, Sutherland, King, Muggleton, & Henery, 1993) is 6.4%. Note also that we need to combine only a few neural networks to get immediate important improvements: with the E version, 20 neural networks suffice for the error to fall under 2%, whereas boosted decision trees typically “converge” later. The W version of AdaBoost actually converged faster in terms of number of networks—after about 7 networks,

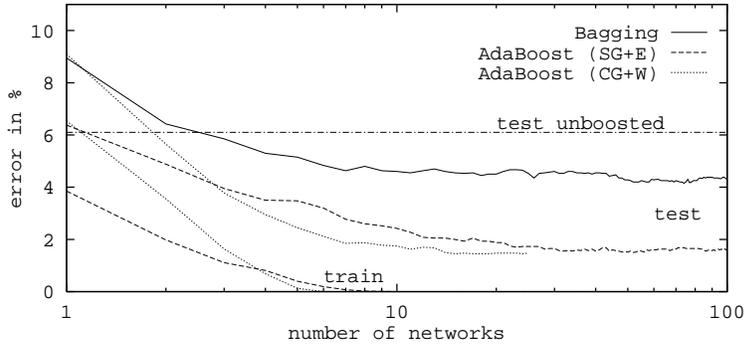


Figure 5: Error rates of the bagged and boosted neural networks for the UCI letter data set (log-scale). SG+E denotes stochastic gradient descent and resampling after each epoch. CG+W means conjugate gradient descent and weighting of the squared error. For clarity, the training error of bagging is not shown (it flattens out to about 1.8%). The dotted constant horizontal line corresponds to the test error of the unboosted classifier.

the 2% mark was reached, and after 14 networks the 1.5% apparent asymptote was reached (see Figure 5)—but converged much slower in terms of training time. Figure 6 shows the margin distributions for bagging and AdaBoost applied to this data set. Again, bagging has no effect on the margin distribution, whereas AdaBoost clearly increases the number of examples with large margins.

Similar conclusions hold for the UCI satellite data set (see Table 4), although the improvements are not as dramatic as in the case of the letter data set. The improvement due to AdaBoost is statistically significant ( $p$ -

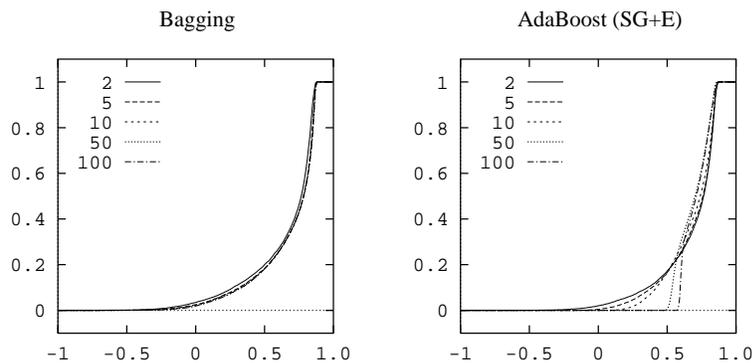


Figure 6: Margin distributions for the UCI letter data set.

value  $< 10^{-6}$ ), but the difference in performance between boosted MLPs and boosted decision trees is not ( $p$ -value  $> 21\%$ ). This data set has 6435 examples, with the first 4435 used for training and the last 2000 used for testing generalization. There are 36 inputs and 6 classes, and a 36-30-15-6 network was used. The two best training methods are the epoch resampling (E) with stochastic gradient and the weighted squared error (W) with conjugate gradient descent.

#### 4 Conclusion

---

As demonstrated here in three real-world applications, AdaBoost can significantly improve neural classifiers. In particular, the results obtained on the UCI letters data set (1.5% test error) are significantly better than the best published results to date, as far as we know. The behavior of AdaBoost for neural networks confirms previous observations on other learning algorithms (Breiman, 1996; Drucker & Cortes, 1996; Freund & Schapire, 1996a; Quinlan, 1996; Schapire et al., 1997), such as the continued generalization improvement after zero training error has been reached and the associated improvement in the margin distribution. It seems also that AdaBoost is not very sensitive to overtraining of the individual classifiers, so that the neural networks can be trained for a fixed (preferably high) number of training epochs. A similar observation was recently made with decision trees (Breiman, 1997b). This apparent insensitivity to overtraining of individual classifiers simplifies the choice of neural network design parameters.

Another interesting finding is that the weighted training version (W) of AdaBoost gives good generalization results for MLPs but requires many more training epochs or the use of a second-order (and, unfortunately, batch) method, such as conjugate gradients. We conjecture that this happens because of the weights on the cost function terms (especially when the weights are small), which could worsen the conditioning of the Hessian matrix. So in terms of generalization error, all three methods (R, E, W) gave similar results, but training time was lowest with the E method (with stochastic gradient descent), which samples each new training pattern from the original data with the AdaBoost weights. Although our experiments are insufficient to conclude, it is possible that the weighted training method (W) with conjugate gradients might be faster than the others for small training sets (a few hundred examples).

There are various ways to define *variance* for classifiers (e.g. Kong & Dietterich, 1995; Breiman, 1996; Kohavi & Wolpert, 1996; Tibshirani, 1996). It basically represents how the resulting classifier will vary when a different training set is sampled from the true generating distribution of the data. Our comparative results on the R, E, and W versions add credence to the view that randomness induced by resampling the training data is not the main reason for AdaBoost's reduction of the generalization error. This is in contrast to bagging, which is a pure variance-reduction method. For

bagging, random resampling is essential to obtain the observed variance reduction.

Another interesting issue is whether the boosted neural networks could be trained with a criterion other than the mean squared error criterion, one that would better approximate the goal of the AdaBoost criterion: minimizing a weighted classification error. Schapire & Singer (1998) addresses this issue.

### Acknowledgments

---

Most of the work reported here was done while H. S. was doing a post-doctorate at the University of Montreal. We thank the National Science and Engineering Research Council of Canada and the government of Quebec for financial support.

### References

---

- Bauer, E., & Kohavi, R. (1999). An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine Learning*, 36, 105–139.
- Bourrely, J. (1989). Parallelization of a neural learning algorithm on a hypercube. In *Hypercube and distributed computers* (pp. 219–229). Amsterdam: Elsevier Science, North-Holland.
- Breiman, L. (1994). Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Breiman, L. (1996). *Bias, variance, and arcing classifiers* (Tech. Rep. No. 460). Berkeley: Statistics Department, University of California at Berkeley.
- Breiman, L. (1997a). *Arcing the edge* (Tech. Rep. No. 486). Berkeley: Statistics Department, University of California at Berkeley.
- Breiman, L. (1997b). *Prediction games and arcing classifiers* (Tech. Rep. No. 504). Berkeley: Statistics Department, University of California at Berkeley.
- Breiman, L. (1998). Arcing classifiers. *Annals of Statistics*, 26(3), 801–849.
- Dietterich, T. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7), 1895–1924.
- Dietterich, T. G. (forthcoming). An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Machine Learning*. Available online at: <ftp://ftp.cs.orst.edu/pub/tgd/papers/tr-randomized-c4.ps.gz>.
- Drucker, H., & Cortes, C. (1996). Boosting decision trees. In D. S. Touretzky, M. C. Mozer, & M. E. Hasselmo (Eds.), *Advances in neural information processing systems* (pp. 479–485). Cambridge, MA: MIT Press.
- Feng, C., Sutherland, A., King, R., Muggleton, S., & Henery, R. (1993). Comparison of machine learning classifiers to statistics and neural networks. In *Proceedings of the Fourth International Workshop on Artificial Intelligence and Statistics* (pp. 41–52).
- Freund, Y. (1995). Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2), 256–285.
- Freund, Y., & Schapire, R. E. (1996a). Experiments with a new boosting algorithm.

- In *Machine Learning: Proceedings of Thirteenth International Conference* (pp. 148–156).
- Freund, Y., & Schapire, R. E. (1996b). Game theory, on-line prediction and boosting. In *Proceedings of the Ninth Annual Conference on Computational Learning Theory* (pp. 325–332).
- Freund, Y., & Schapire, R. E. (1997). A decision theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Science*, 55(1), 119–139.
- Freund, Y., & Schapire, R. E. (1999). Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29, 79–103.
- Friedman, J., Hastie, T., & Tibshirani, R. (1998). *Additive logistic regression: A statistical view of boosting* (Tech. Rep.). Stanford: Department of Statistics, Stanford University.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1), 1–58.
- Grove, A. J., & Schuurmans, D. (1998). Boosting in the limit: Maximizing the margin of learned ensembles. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*.
- Kohavi, R., & Wolpert, D. H. (1996). Bias plus variance decomposition for zero-one loss functions. In *Machine Learning: Proceedings of Thirteenth International Conference* (pp. 275–283).
- Kong, E. B., & Dietterich, T. G. (1995). Error-correcting output coding corrects bias and variance. In *Machine Learning: Proceedings of Twelfth International Conference* (pp. 313–321).
- Krogh, A., & Vedelsby, J. (1995). Neural network ensembles, cross validation and active learning. In G. Tesauro, D. S. Touretzky, & T. K. Leen, (Eds.), *Advances in neural information processing systems*, 7, (pp. 231–238). Cambridge, MA: MIT Press.
- Maclin, R., & Opitz, D. (1997). An empirical evaluation of bagging and boosting. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence* (pp. 546–551).
- Mason, L., & Baxter, P. B. J. (1999). Direct optimization of margins improves generalization in combined classifiers. In M. S. Kearns, S. Solla, & D. Cohn (Eds.), *Advances in neural information processing systems*, 11, Cambridge, MA: MIT Press.
- Moller, M. (1992). Supervised learning on large redundant training sets. In *Neural networks for signal processing 2*. New York: IEEE Press.
- Moller, M. (1993). *Efficient training of feed-forward neural networks*. Unpublished doctoral dissertation, Aarhus University, Aarhus, Denmark.
- Perrone, M. P. (1993). *Improving regression estimation: Averaging methods for variance reduction with extensions to general convex measure optimization*. Unpublished doctoral dissertation, Brown University, Providence, RI.
- Perrone, M. P. (1994). Putting it all together: Methods for combining neural networks. In J. D. Cowan, G. Tesauro, & J. Alspector, (Eds.), *Advances in neural information processing systems*, 6, (pp. 1188–1189). San Mateo, CA: Morgan Kaufmann.

- Quinlan, J. R. (1996). Bagging, boosting and C4.5. In *Machine Learning: Proceedings of the Fourteenth International Conference* (pp. 725–730).
- Rätsch, G., Onoda, T., & Müller, K.-R. (1998). *Soft margins for AdaBoost* (Tech. Rep. No. NC-TR-1998-021). Egham, Surrey: Royal Holloway College, University of London.
- Schapire, R. E. (1990). The strength of weak learnability. *Machine Learning*, 5(2), 197–227.
- Schapire, R. E. (1999). Theoretical views of boosting. In *Computational Learning Theory: Fourth European Conference, EuroCOLT* (pp. 1–10).
- Schapire, R. E., Freund, Y., Bartlett, P., & Lee, W. S. (1997). Boosting the margin: A new explanation for the effectiveness of voting methods. In *Machine Learning: Proceedings of Fourteenth International Conference* (pp. 322–330).
- Schapire, R. E., & Singer, Y. (1998). Improved boosting algorithms using confidence rated predictions. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*.
- Schwenk, H., & Bengio, Y. (1997). Adaboosting neural networks: Application to on-line character recognition. In *International Conference on Artificial Neural Networks* (pp. 967–972). Berlin: Springer-Verlag.
- Schwenk, H., & Bengio, Y. (1998). Training methods for adaptive boosting of neural networks. In M. I. Jordan, M. J. Kearns, & S. A. Solla, (Eds.), *Advances in neural information processing systems*, 10 (pp. 647–653). Cambridge, MA: MIT Press.
- Schwenk, H., & Milgram, M. (1996). Constraint tangent distance for online character recognition. In *International Conference on Pattern Recognition* (pp. D:520–524). New York: IEEE Computer Society Press.
- Tibshirani, R. (1996). *Bias, variance and prediction error for classification rules* (Tech. Rep.). Toronto: Department of Statistics, University of Toronto.

---

Received May 13, 1998; accepted March 24, 1999.