

14 Boosting

Goal

Understand the ensemble method for combining classifiers. Bagging, Random Forest, and the celebrated Adaboost.

Alert 14.1: Convention

Gray boxes are not required hence can be omitted for unenthusiastic readers.

[This note is likely to be updated again soon.](#)

Remark 14.2: Together and stronger?

Often it is possible to train a variety of different classifiers for a particular problem at hand, and a lot of time, energy and discussion are spent on debating and choosing the most appropriate classifier. This makes sense when the classifiers are “expensive” to obtain (be it computationally or financially or resourcefully).

Putting operational costs aside, however, is it possible to **combine** a bunch of classifiers and get better performance, for instance when compared against the best classifier in the gang? Of course, one usually does not know which classifier in the gang is “best” (unless when we try all of them out).

Remark 14.3: Condorcet’s Jury Theorem (Condorcet 1785)

To motivate our development, let us consider an ideal scenario where we have a collection of classifiers $\{h_t : \mathcal{X} \rightarrow \{0, 1\}, t = 1, 2, \dots, 2T + 1\}$ for a binary classification problem (where we encode the labels as $\{0, 1\}$). **Conditioned** on some auxiliary information \mathcal{F} , we assume the classifiers h_t **independently** achieve the following accuracy on a test sample (X, Y) :

$$p := \Pr(h_t(X) = Y | \mathcal{F}) > \frac{1}{2}.$$

(For instance, if $\mathcal{F} = \{X, Y\}$ and the classifiers h_t are trained on an independent dataset \mathcal{D}_t .) We predict according to the majority:

$$h(X) = \begin{cases} 1, & \text{if } \#\{t : h_t(X) = 1\} \geq T + 1 \\ 0, & \text{if } \#\{t : h_t(X) = 0\} \geq T + 1 \end{cases}.$$

What is the accuracy of this “meta-classifier” h ? Simple calculation and Exercises 14.5 and 14.6 reveal:

$$\begin{aligned} \Pr(h(X) = Y | \mathcal{F}) &= \Pr\left(\sum_{t=1}^{2T+1} \mathbb{I}[h_t(X) = Y] \geq T + 1 \mid \mathcal{F}\right) = \sum_{k=T+1}^{2T+1} \binom{2T+1}{k} p^k (1-p)^{2T+1-k} \\ &= \frac{(2T+1)!}{(T!)^2} \int_0^p u^T (1-u)^T du, \end{aligned}$$

which is increasing w.r.t. p , and more importantly also increasing w.r.t. T (provided that $p > \frac{1}{2}$) and approaching 1 as $T \rightarrow \infty$!

Therefore, provided that we can combine a large number of classifiers with **conditionally independent** accuracy, each slightly better than random guessing, we can approach perfect accuracy! The caveat of course is the difficulty (or even impossibility) in obtaining these classifiers with conditionally independent accuracy in the first place. As pointed out by Vardeman and Morris (2013), independent classifiers do not necessarily yield independent accuracy (unless we condition on the test sample (X, Y) , for example).

See Grofman et al. (1983), Young (1988), and Austen-Smith and Banks (1996) for more interesting discussions around Condorcet’s Jury Theorem and voting.

[Condorcet, N. C. d. \(1785\). “Essai sur l’application de l’analyse á la probabilité des décisions”. AMS.](#)

- Vardeman, S. B. and M. D. Morris (2013). “Majority Voting by Independent Classifiers Can Increase Error Rates”. *The American Statistician*, vol. 67, no. 2, pp. 94–96.
- Grofman, B., G. Owen, and S. L. Feld (1983). “Thirteen theorems in search of the truth”. *Theory and Decision*, vol. 15, pp. 261–278.
- Young, H. P. (1988). “Condorcet’s Theory of Voting”. *The American Political Science Review*, vol. 82, no. 4, pp. 1231–1244.
- Austen-Smith, D. and J. S. Banks (1996). “Information Aggregation, Rationality, and the Condorcet Jury Theorem”. *The American Political Science Review*, vol. 90, no. 1, pp. 34–45.

Exercise 14.4: Odd vs. even

Perform a similar analysis as in Remark 14.3 for $2T$ classifiers. Is there any advantage in using the odd $2T + 1$ over the even $2T$?

Exercise 14.5: Binomial and Beta distributions

Let $U_1, \dots, U_n \stackrel{i.i.d.}{\sim} \text{uniform}(0, 1)$, and denote $U_{(1)} \leq U_{(2)} \leq \dots \leq U_{(n)}$ as the **ordered statistic**. Prove that

- $U_{(i)} \sim \text{Beta}(i, n - i + 1)$, where we recall that the density of the Beta distribution is

$$p(u) = \frac{1}{\text{B}(i, n - i + 1)} u^{i-1} (1 - u)^{(n-i+1)-1}, \quad \text{where } \text{B}(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}, \quad \text{and recall } \Gamma(n) = (n - 1)!.$$

- $\Pr(U_{(i)} \leq p) = \sum_{k=i}^n \binom{n}{k} p^k (1 - p)^{n-k}$, which is just the probability for a Binomial random variable to have at least i successes in n trials (with success probability p).
- The survival function of Binomial is just the cdf of Beta, i.e.,

$$\sum_{k=i}^n \binom{n}{k} p^k (1 - p)^{n-k} = \frac{1}{\text{B}(i, n - i + 1)} \int_0^p u^{i-1} (1 - u)^{(n-i+1)-1} du. \quad (14.1)$$

[One can also prove (14.1) directly by differentiating the left-hand side w.r.t. p , or by (repeatedly) applying integration by parts to the right-hand side.]

Exercise 14.6: Monotonicity and asymptotics

Let $p > \frac{1}{2}$ and define

$$s(T, p) = \sum_{k=T+1}^{2T+1} \binom{2T+1}{k} p^k (1 - p)^{2T+1-k} = \frac{(2T+1)!}{(T!)^2} \int_0^p u^T (1 - u)^T du.$$

Prove that

- $s(T, p)$ is increasing w.r.t. T . In particular, $s(T, p) \geq s(0, p) = p$.
- $s(T, p) \rightarrow 1$ as $T \rightarrow \infty$.

More general results can be found in Boland (1989), Berend and Paroush (1998), Ben-Yashar and Paroush (2000), and Berend and Sapir (2005), among many others.

Boland, P. J. (1989). “Majority Systems and the Condorcet Jury Theorem”. *Journal of the Royal Statistical Society. Series D (The Statistician)*, vol. 38, no. 3, pp. 181–189.

Berend, D. and J. Paroush (1998). “When is Condorcet’s Jury Theorem valid?” *Social Choice and Welfare*, vol. 15, no. 4, pp. 481–488.

Ben-Yashar, R. and J. Paroush (2000). “A nonasymptotic Condorcet jury theorem”. *Social Choice and Welfare*, vol. 17, no. 2, pp. 189–199.

Berend, D. and L. Sapir (2005). “Monotonicity in Condorcet Jury Theorem”. *Social Choice and Welfare*, vol. 24, no. 1, pp. 83–92.

Algorithm 14.7: Bootstrap Aggregating (Bagging, Breiman 1996)

One way to achieve (approximately) independent classifiers is to simply train them on independent datasets, which in most (if not all) applications is simply a luxury. However, we can use the same **bootstrapping** idea as in cross-validation (see Line 7):

Algorithm: Bagging predictors.

Input: training set \mathcal{D} , number of repetitions T

Output: meta-predictor h

```

1 for  $t = 1, 2, \dots, T$  do // in parallel
2   sample a new training set  $\mathcal{D}_t \subseteq \mathcal{D}$  // with or without replacement
3   train predictor  $h_t$  on  $\mathcal{D}_t$ 
4  $h \leftarrow \text{aggregate}(h_1, \dots, h_T)$  // majority vote for classification; average for regression
```

There are two ways to sample a new training set:

- Sample with replacement: **copy** a (uniformly) **random** element from \mathcal{D} to \mathcal{D}_t and repeat $|\mathcal{D}|$ times. Usually there will be repetitions of the same element in \mathcal{D}_t (which has no effect on most machine learning algorithms). *This is the common choice.*
- Sample without replacement: **“cut”** a (uniformly) **random** element from (a copy of) \mathcal{D} to \mathcal{D}_t and repeat say $70\% \times |\mathcal{D}|$ times. There will be no repetitions in each \mathcal{D}_t . Note that had we repeated $|\mathcal{D}|$ times (just as in sample with replacement) we would have $\mathcal{D}_t \equiv \mathcal{D}$, which is not very useful.

Of course the training sets $\{\mathcal{D}_t : t = 1, \dots, T\}$ are not really independent of each other, but aggregating predictors trained on them **usually (but not always)** improves performance.

Breiman, L. (1996). “Bagging Predictors”. *Machine Learning*, vol. 24, no. 2, pp. 123–140.

Algorithm 14.8: Randomizing output (e.g., Breiman 2000)

Bagging perturbs the training set by taking a random subset. We can also perturb the training set by simply adding noise:

- for regression tasks: replace line 2 in the bagging algorithm by “adding small Gaussian noise to each response y_i ”
- for classification tasks: replace line 2 in the bagging algorithm by “randomly flip a small portion of the labels y_i ”

We will come back to the idea of perturbation later. Intuitively, adding noise can prevent overfitting, and aggregation reduces variance.

Breiman, L. (2000). “Randomizing outputs to increase prediction accuracy”. *Machine Learning*, vol. 40, no. 3, pp. 229–242.

Algorithm 14.9: Random forest (Breiman 2001)

One of the most popular machine learning algorithms is random forest, where we combine bagging and decision trees. Instead of training a usual decision tree, we introduce yet another layer of randomization:

- during training, at each internal node where we need to split the training samples, we select a **random subset of features** and perform the splitting. A different subset of features is used at a different internal node.

Then we apply bagging and aggregate a bunch of the above “randomized” decision trees.

Breiman, L. (2001). “Random Forest”. *Machine Learning*, vol. 45, no. 1, pp. 5–32.

Algorithm 14.10: Hedging (Freund and Schapire 1997)

Algorithm: Hedging.

Input: initial weight vector $\mathbf{w}_1 \in \mathbb{R}_{++}^n$, discount factor $\beta \in [0, 1]$

Output: last weight vector \mathbf{w}_{T+1}

```

1 for  $t = 1, 2, \dots, T$  do
2   learner chooses probability vector  $\mathbf{p}_t = \frac{\mathbf{w}_t}{\mathbf{1}^T \mathbf{w}_t}$  // normalization
3   environment chooses loss vector  $\ell_t \in [0, 1]^n$  //  $\ell_t$  may depend on  $\mathbf{p}_t!$ 
4   learner suffers (expected) loss  $\langle \mathbf{p}_t, \ell_t \rangle$ 
5   learner updates weights  $\mathbf{w}_{t+1} = \mathbf{w}_t \odot \beta^{\ell_t}$  // element-wise product  $\odot$  and power
6   optional scaling:  $\mathbf{w}_{t+1} \leftarrow c_{t+1} \mathbf{w}_{t+1}$  //  $c_{t+1} > 0$  can be arbitrary

```

Imagine the following **horse racing** game^a: There are n horses in a race, which we repeat for T rounds. On each round we bet a fixed amount of money, with p_{it} being the proportion we spent on horse i at the t -th round. At the end of round t we receive a loss $\ell_{it} \in [0, 1]$ that we suffer on horse i . Note that the losses can be completely arbitrary (e.g., no i.i.d. assumption), although for simplicity we assume they are bounded. How should we place our bets (i.e., p_{it}) to hedge our risk? Needless to say, we must decide the proportions \mathbf{p}_t *before* we see the losses ℓ_t . On the other hand, the losses ℓ_t can be completely adversarial (i.e., depend on \mathbf{p}_t).

The Hedging algorithm gives a reasonable (in fact in some sense optimal) strategy: basically if a horse i is doing well (badly) on round t , then we spend a larger (smaller) proportion $p_{i,t+1}$ on it in the next round $t+1$. On a high level, this is similar to the idea behind perceptron (see Section 1).

Note that line 5 continues decreasing each entry in the weight vector (because both $\beta \in [0, 1]$ and $\ell \in [0, 1]$). To prevent the weights from underflowing, we can re-scale them by a positive number $c > 0$, as shown in the optional line 6. Clearly, **this scaling does not change \mathbf{p} hence the algorithm** (in any essential way).

Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139.

^aViewer discretion is advised; please do not try at home!

Alert 14.11: Multiplicative Updates

The Hedging algorithm belongs to the family of multiplicative updates, where we repeatedly multiplying, instead of adding (cf. Perceptron in Section 1), our weights by some correction terms. For multiplicative updates, it is important to start with **strictly positive** initializations, for a zero weight will remain as zero in multiplicative updates.

Theorem 14.12: Hedging Guarantee

For any nonempty $S \subseteq \{1, \dots, n\}$, we have

$$L \leq \frac{-\ln(\sum_{i \in S} p_{i1}) - (\ln \beta) \max_{i \in S} L_i}{1 - \beta},$$

where $L := \sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle$ is our total (expected) loss and $L_i := \sum_{t=1}^T \ell_{it}$ is the total loss on the i -th horse.

Proof: We first lower bound the sum of weights at the end:

$$\begin{aligned} \mathbf{1}^\top \mathbf{w}_{T+1} &= \sum_{i=1}^n w_{i,T+1} \geq \sum_{i \in S} w_{i,T+1} && // \text{weights remain nonnegative} \\ &= \sum_{i \in S} w_{i,1} \beta^{L_i} && // \text{line 5 of Algorithm 14.10} \end{aligned} \quad (14.2)$$

$$\geq \beta^{\max_{i \in S} L_i} \sum_{i \in S} w_{i,1} \quad // \beta \in [0, 1]. \quad (14.3)$$

Then, we upper bound the sum of weights:

$$\begin{aligned} \mathbf{1}^\top \mathbf{w}_{t+1} &= \sum_{i=1}^n w_{i,t+1} = \sum_{i=1}^n w_{i,t} \beta^{\ell_{i,t}} && // \text{line 5 of Algorithm 14.10} \\ &\leq \sum_{i=1}^n w_{i,t} (1 - (1 - \beta)\ell_{i,t}) && // x^\ell \leq 1^\ell + (x - 1)\ell: x^\ell \text{ concave when } \ell \in [0, 1] \text{ and } x \geq 0 \text{ (14.4)} \\ &= (\mathbf{1}^\top \mathbf{w}_t) [1 - (1 - \beta) \mathbf{p}_t^\top \boldsymbol{\ell}_t] && // \text{line 2 of Algorithm 14.10.} \end{aligned}$$

Thus, by telescoping:

$$\frac{\mathbf{1}^\top \mathbf{w}_{T+1}}{\mathbf{1}^\top \mathbf{w}_1} \leq \prod_{t=1}^T [1 - (1 - \beta) \mathbf{p}_t^\top \boldsymbol{\ell}_t] \leq \exp \left[-(1 - \beta) \sum_{t=1}^T \mathbf{p}_t^\top \boldsymbol{\ell}_t \right] = \exp[-(1 - \beta)L], \quad (14.5)$$

where we have used the elementary inequality $1 - x \leq \exp(-x)$ for any x .

Finally, combining the inequalities (14.3) and (14.5) completes the proof. ■

By inspecting the proof closely, we realize that the same conclusion still holds if we change the update to

$$\mathbf{w}_{t+1} = \mathbf{w}_t \odot U_\beta(\boldsymbol{\ell}_t)$$

as long as the function U_β satisfies

$$\beta^\ell \leq U_\beta(\ell) \leq 1 - (1 - \beta)\ell$$

(so that inequalities (14.2) and (14.4) still hold).

Exercise 14.13: Optional re-scaling has no effect

Show that the same conclusion in Theorem 14.12 still holds even if we include the optional line 6 in Algorithm 14.10.

Corollary 14.14: Comparable to best “horse” in hindsight

If we choose $\mathbf{w}_1 = \mathbf{1}$ and $|S| = 1$, then

$$\sum_{t=1}^T \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle \leq \frac{\min_i L_i \ln \frac{1}{\beta} + \ln n}{1 - \beta}. \quad (14.6)$$

In addition, if we choose $\beta = \frac{1}{1 + \sqrt{(2m)/U}}$, where $U \geq L_{\min} := \min_i L_i$ and $m \geq \ln n$, then

$$\frac{1}{T} \sum_{t=1}^T \langle \mathbf{p}_t, \ell_t \rangle \leq \frac{L_{\min}}{T} + \frac{\sqrt{2mU}}{T} + \frac{\ln n}{T}. \quad (14.7)$$

If we also choose $m = \ln n$ and $U = T$ (in our choice of β), then

$$\frac{1}{T} \sum_{t=1}^T \langle \mathbf{p}_t, \ell_t \rangle \leq \frac{L_{\min}}{T} + \sqrt{\frac{2 \ln n}{T}} + \frac{\ln n}{T}. \quad (14.8)$$

Proof: Indeed, for $\beta \in [0, 1]$ we have

$$\begin{aligned} 2 \leq 2/\beta &\implies 2\beta \geq 2 \ln \beta + 2 && // 2\beta - 2 \ln \beta - 2 \text{ is decreasing and nonnegative at } \beta = 1 \\ &\implies \beta^2 \leq 1 + 2\beta \ln \beta && // \beta^2 - 2\beta \ln \beta - 1 \text{ is increasing and nonpositive at } \beta = 1 \\ &\iff \ln \frac{1}{\beta} \leq \frac{1 - \beta^2}{2\beta}. \end{aligned}$$

Plugging our choice of β into (14.6) and apply the above inequality we obtain (14.7). The bound (14.8) holds because we can choose $U = T$ (recall that we assume $\ell_{i,t} \in [0, 1]$ for all t and i). ■

Based on a result due to Vovk (1998), Freund and Schapire (1997) proved that the coefficients in (14.6) (i.e., $\frac{-\ln \beta}{1 - \beta}$ and $\frac{\ln n}{1 - \beta}$) cannot be simultaneously improved for any β , using any algorithm (not necessarily Hedging).

In our last setting, $\beta = \frac{1}{1 + \sqrt{(2 \ln n)/T}}$, indicating for a longer game (larger T) we should use a larger β (to discount less aggressively) while for more horses (larger n) we should do the opposite (although this effect is much smaller due to the log).

Vovk, V. (1998). “A Game of Prediction with Expert Advice”. *Journal of Computer and System Sciences*, vol. 56, no. 2, pp. 153–173.

Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139.

Remark 14.15: Appreciating the significance of Hedging

Corollary 14.14 implies that the *average* loss of Hedging on the left-hand side of (14.8) is no larger than the average loss of the “best horse”, plus a constant term proportional to $\sqrt{\frac{2 \ln n}{T}}$ (where we omit the higher order term $\frac{\ln n}{T}$). As the number of rounds T goes to infinity, Hedging can compete against the best horse in hindsight! However, we emphasize three important points:

- **It does not mean Hedging will achieve small loss**, simply because the best horse may itself achieve a *large* average loss, in which case being competitive against the best horse does not really mean much.
- **The loss vectors ℓ_t can be adversarial against the Hedging algorithm!** However, if the environment always tries to “screw up” the Hedging algorithm, then the guarantee in Corollary 14.14 implies that the environment inevitably also “screws up” the best horse.
- If the best horse can achieve very small loss, i.e. $L_{\min} \approx 0$, then by setting $U \approx 0$ we know from (14.7) that Hedging is off by at most a term proportional to $\frac{\ln n}{T}$, which is much smaller than the dominating term $\sqrt{\frac{2 \ln n}{T}}$ in (14.8).

Remark 14.16: History of Boosting

Schapire (1990) first formally proved the possibility to combine a few mediocre classifiers into a very accurate one, which was subsequently improved by Freund (1995) and eventually by Freund and Schapire (1997), which has since become the main source of the celebrated Adaboost algorithm. Freund and Schapire received the Gödel prize in 2003 for this seminal contribution.

Schapire, R. E. (1990). “The strength of weak learnability”. *Machine Learning*, vol. 5, no. 2, pp. 197–227.

Freund, Y. (1995). “Boosting a Weak Learning Algorithm by Majority”. *Information and Computation*, vol. 121, no. 2, pp. 256–285.

Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139.

Algorithm 14.17: Adaboost (Freund and Schapire 1997)

Algorithm: Adaptive Boosting.

Input: initial weight vector $\mathbf{w}_1 \in \mathbb{R}_{++}^n$, training set $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \subseteq \mathbb{R}^d \times \{0, 1\}$
Output: meta-classifier $\bar{h} : \mathbb{R}^d \rightarrow \{0, 1\}$, $\mathbf{x} \mapsto \lfloor \sum_{t=1}^T (\ln \frac{1}{\beta_t})(h_t(\mathbf{x}) - \frac{1}{2}) \geq 0 \rfloor$

```

1 for  $t = 1, 2, \dots, T$  do
2    $\mathbf{p}_t = \mathbf{w}_t / \mathbf{1}^\top \mathbf{w}_t$  // normalization
3    $h_t \leftarrow \text{WeakLearn}(\mathcal{D}_n, \mathbf{p}_t)$  //  $t$ -th weak classifier  $h_t : \mathbb{R}^d \rightarrow [0, 1]$ 
4    $\forall i, \ell_{it} = 1 - |h_t(\mathbf{x}_i) - y_i|$  // loss is higher if prediction is more accurate!
5    $\epsilon_t = 1 - \langle \mathbf{p}_t, \boldsymbol{\ell}_t \rangle = \sum_{i=1}^n p_{it} |h_t(\mathbf{x}_i) - y_i|$  // weighted (expected) error  $\epsilon_t \in [0, 1]$  of  $h_t$ 
6    $\beta_t = \epsilon_t / (1 - \epsilon_t)$  // adaptive discounting parameter  $\beta_t \leq 1 \iff \epsilon_t \leq \frac{1}{2}$ 
7    $\mathbf{w}_{t+1} = \mathbf{w}_t \odot \beta_t^{\boldsymbol{\ell}_t}$  // element-wise product  $\odot$  and power
8   optional scaling:  $\mathbf{w}_{t+1} \leftarrow c_{t+1} \mathbf{w}_{t+1}$  //  $c_{t+1} > 0$  can be arbitrary
```

Provided that $\epsilon_t \leq \frac{1}{2}$ hence $\beta_t \in [0, 1]$ (so the classifier h_t is better than random guessing), if h_t predicts correctly on a training example \mathbf{x}_i , then we suffer a larger loss ℓ_{it} so that in the next iteration we assign less weight to \mathbf{x}_i . In other words, each classifier is focused on hard examples that are misclassified by the previous classifier. On the other hand, if $\epsilon_t > \frac{1}{2}$ then $\beta_t > 1$ and we do the opposite.

For the meta-classifier \bar{h} , we first perform a weighted aggregation of the confidences of individual (weak) classifiers and then threshold. Alternatively, we could threshold each weak classifier first and then perform (weighted) majority voting. The former approach, which we adopt here, is found to work better in practice. Note that, a classifier h_t with lower (training) error ϵ_t will be assigned a higher weight $\ln \frac{1}{\beta_t} = \ln(\frac{1}{\epsilon_t} - 1)$ in the final aggregation, making intuitive sense.

Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139.

Exercise 14.18: Implicitly fixing “bad” classifiers

If $\epsilon_t \geq \frac{1}{2}$ (hence $\beta_t \geq 1$), the t -th (weak) classifier is in fact worse than random guessing! In our binary setting here, a natural idea is to discard h_t but use instead $\tilde{h}_t := 1 - h_t$. Prove the following (all tilde quantities are w.r.t. the flipped classifier \tilde{h}_t):

- $\tilde{\ell}_t = 1 - \ell_t$
- $\tilde{\epsilon}_t = 1 - \epsilon_t$
- $\tilde{\beta}_t = 1/\beta_t$
- The Adaboost algorithm is not changed in any essential way even if we flip all “bad” classifiers. In particular, we arrive at the same meta-classifier.

In other words, Adaboost implicitly fixes all “bad” classifiers!

Theorem 14.19: Adaboost Guarantee

The following bound on the **training error** holds for the Adaboost Algorithm 14.17:

$$\sum_{i=1}^n p_{i1} [\bar{h}(\mathbf{x}_i) \neq y_i] \leq \prod_{t=1}^T \sqrt{4\epsilon_t(1-\epsilon_t)}.$$

Proof: The proof closely follows what we have seen in the proof of Theorem 14.12. As before, we upper bound the sum of weights in the same manner:

$$\frac{\|\mathbf{w}_{T+1}\|_1}{\|\mathbf{w}_1\|_1} \leq \prod_{t=1}^T [1 - (1 - \beta_t) \mathbf{p}_t \ell_t] = \prod_{t=1}^T [1 - (1 - \beta_t)(1 - \epsilon_t)].$$

To lower bound the sum of weights, let $S = \{i : \bar{h}(\mathbf{x}_i) \neq y_i\}$, where recall that \bar{h} is the meta-classifier constructed in Algorithm 14.17. Proceed as before:

$$\frac{\|\mathbf{w}_{T+1}\|_1}{\|\mathbf{w}_1\|_1} \geq \sum_{i \in S} \frac{w_{i,T+1}}{\|\mathbf{w}_1\|_1} = \sum_{i \in S} p_{i1} \prod_{t=1}^T \beta_t^{\ell_{i,t}} \geq \sum_{i \in S} p_{i1} \left[\prod_{t=1}^T \beta_t \right]^{1/2},$$

where the last inequality follows from the definition of S and the meta-classifier \bar{h} :

$$i \in S \implies 0 \geq \text{sign}(2y_i - 1) \sum_t (\ln \frac{1}{\beta_t}) (h_t(\mathbf{x}_i) - \frac{1}{2}) = \sum_t (\ln \frac{1}{\beta_t}) (\frac{1}{2} - |h_t(\mathbf{x}_i) - y_i|) = \sum_t (\ln \frac{1}{\beta_t}) (\ell_{i,t} - \frac{1}{2}).$$

Combine the upper bound and the lower bound:

$$\sum_{i \in S} p_{i1} \prod_{t=1}^T \sqrt{\beta_t} \leq \prod_{t=1}^T [1 - (1 - \beta_t)(1 - \epsilon_t)], \quad \text{i.e.,} \quad \sum_{i \in S} p_{i1} \leq \prod_{t=1}^T \frac{1 - (1 - \beta_t)(1 - \epsilon_t)}{\sqrt{\beta_t}}.$$

Optimizing w.r.t. β_t we obtain $\beta_t = \epsilon_t / (1 - \epsilon_t)$. Plugging it back in we complete the proof. ■

Importantly, we observe that the training error of the meta-classifier \bar{h} is upper bounded by the errors of all classifiers h_t : **improving any individual classifier leads to a better bound**. Moreover, the symmetry on the right-hand side confirms again that in the binary setting a very “inaccurate” classifier (i.e., large ϵ_t) is as good as a very accurate classifier (i.e., small ϵ_t).

Corollary 14.20: Exponential decay of training error

Assume $|\epsilon_t - \frac{1}{2}| > \gamma_t$, then

$$\sum_{i=1}^n p_{i1} [\bar{h}(\mathbf{x}_i) \neq y_i] \leq \prod_{t=1}^T \sqrt{1 - 4\gamma_t^2} \leq \exp\left(-2 \sum_{t=1}^T \gamma_t^2\right).$$

In particular, if $\gamma_t \geq \gamma$ for all t , then

$$\sum_{i=1}^n p_{i1} [\bar{h}(\mathbf{x}_i) \neq y_i] \leq \exp(-2T\gamma^2).$$

Thus, to achieve ϵ (weighted) **training error**, we need to combine at most

$$T = \lceil \frac{1}{2\gamma^2} \ln \frac{1}{\epsilon} \rceil$$

weak classifiers, each of which is slightly better than random guessing (by a margin of γ). ■

Remark 14.21: Generalization Error of Adaboost

It is possible to bound the generalization error of Adaboost as well. For instance, Freund and Schapire (1997) bounded the VC dimension of the (family of) meta-classifiers constructed by Adaboost. A standard application of the VC theory then relates the generalization error with the training error. More refined analysis can be found in Schapire et al. (1998), Koltchinskii and Panchenko (2002), Koltchinskii et al. (2003), Koltchinskii and Panchenko (2005), Freund et al. (2004), and Rudin et al. (2007) (just to give a few pointers).

- Freund, Y. and R. E. Schapire (1997). “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139.
- Schapire, R. E., Y. Freund, P. Bartlett, and W. S. Lee (1998). “Boosting the margin: a new explanation for the effectiveness of voting methods”. *The Annals of Statistics*, vol. 26, no. 5, pp. 1651–1686.
- Koltchinskii, V. and D. Panchenko (2002). “Empirical Margin Distributions and Bounding the Generalization Error of Combined Classifiers”. *The Annals of Statistics*, vol. 30, no. 1, pp. 1–50.
- Koltchinskii, V., D. Panchenko, and F. Lozano (2003). “Bounding the generalization error of convex combinations of classifiers: balancing the dimensionality and the margins”. *The Annals of Applied Probability*, vol. 13, no. 1, pp. 213–252.
- Koltchinskii, V. and D. Panchenko (2005). “Complexities of convex combinations and bounding the generalization error in classification”. *The Annals of Statistics*, vol. 33, no. 4, pp. 1455–1496.
- Freund, Y., Y. Mansour, and R. E. Schapire (2004). “Generalization bounds for averaged classifiers”. *The Annals of Statistics*, vol. 32, no. 4, pp. 1698–1722.
- Rudin, C., R. E. Schapire, and I. Daubechies (2007). “Analysis of boosting algorithms using the smooth margin function”. *The Annals of Statistics*, vol. 35, no. 6, pp. 2723–2768.

Alert 14.22: Does Adaboost Overfit? (Breiman 1999; Grove and Schuurmans 1998)

It has long been observed that Adaboost, even after decreasing the training error to zero, continues to improve test error. In other words, Adaboost does not seem to overfit even when we combine many many (weak) classifiers.

A popular explanation, due to Schapire et al. (1998), attributes Adaboost’s resistance against overfitting to margin maximization: after decreasing the training error to 0, Adaboost continues to improve the margin (i.e. $yh(\mathbf{x})$), which leads to better generalization. However, Breiman (1999) and Grove and Schuurmans (1998) later designed the LPboost that explicitly maximizes the margin but observed inferior generalization.

- Breiman, L. (1999). “Prediction Games and Arcing Algorithms”. *Neural Computation*, vol. 11, no. 7, pp. 1493–1517.
- Grove, A. J. and D. Schuurmans (1998). “Boosting in the Limit: Maximizing the Margin of Learned Ensembles”. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pp. 692–699.
- Schapire, R. E., Y. Freund, P. Bartlett, and W. S. Lee (1998). “Boosting the margin: a new explanation for the effectiveness of voting methods”. *The Annals of Statistics*, vol. 26, no. 5, pp. 1651–1686.


Algorithm 14.23: Face Detection (Viola and Jones 2004)

Viola and Jones (2004) applied the Adaboost algorithm to real-time face detection and are among the first few people who demonstrated the power of Adaboost in real challenging applications.

- Viola, P. and M. J. Jones (2004). “Robust Real-Time Face Detection”. *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154.

Remark 14.24: Comparison

The following figure illustrates the similarity and difference between bagging and boosting:



To summarize:

- Both bagging and boosting train an **ensemble** of (weak) classifiers, which are combined in the end to produce a “meta-classifier”;
- Bagging is amenable to parallelization while boosting is strictly sequential;
- Bagging **resamples** training examples while boosting **reweights** training examples;
- As suggested in Breiman (2004), we can think of bagging as averaging *independent* classifiers (in order to reduce variance), while boosting averages dependent classifiers which may be analyzed through dynamic system and ergodic theory;
- We can of course combine bagging with boosting. For instance, each classifier in boosting can be obtained through bagging (called bag-boosting by Bühlmann and Yu, see the discussion of Friedman et al. (2000)).

Breiman, L. (2004). “Population theory for boosting ensembles”. *The Annals of Statistics*, vol. 32, no. 1, pp. 1–11.

Friedman, J., T. Hastie, and R. Tibshirani (2000). “Additive logistic regression: a statistical view of boosting (With discussion and a rejoinder by the authors)”. *The Annals of Statistics*, vol. 28, no. 2, pp. 337–407.

Exercise 14.25: Diversity

Recall that $\epsilon_t(h) := \sum_{i=1}^n p_{it} |h(\mathbf{x}_i) - y_i|$ is the weighted error of classifier h at iteration t . Assume the (weak) classifiers are always binary-valued (and $\beta_t \in (0, 1)$ for all t). Prove:

$$\epsilon_{t+1}(h_t) \equiv \frac{1}{2}.$$

In other words, **Adaboost would never choose the same weak classifier twice in a row.**