

# Lecture 17

Boosting

## Lecture 4: Combining predictors - Part I

#### Marina Meilă mmp@stat.washington.edu

Department of Statistics University of Washington

November, 2015

Averaginp Redictors

There is more than one way to average predictors

Bayesian averaging Bagging Stacking Mixtures of Experts Forward Fitting and Backfitting

#### Reducing bias: Boosting

AdaBoost Properties on the training set  $B = \{b\} \\ base class$ ``weak $<math display="block">f = \sum_{k=1}^{K} \beta_k b_k(x) \in F$   $b = b_k(x) \in F$   $b = b_k(x) \in F$ 

 $f(x) = \hat{y} - determine$ 

predicts y at x

**Reading HTF:** 15. Random Forests, 16. Ensembles of predictors, 8. Model inference and averaging, 10. Boosting, Murphy: 16.3, 16.3.1 Generalized Additive models (ignore the regularization, for "smoother" read "minimize loss function"), 16.4.1-5 [and optionally 8] Bosting, 3.2.4 Bayesian model averaging, 16.6.3, 16.6.1 Ensembles of predictors

## There is more than one way to average predictors

Classification will be the running exaple here, but most results hold for other prediction problems as well.

Denote  $\mathcal{B} = \{b\}$  a **base classifier** family Averaging:

$$f(x) = \sum_{k=1}^{M} \beta^{k} b^{k}(x)$$
 (1)

f is real-valued even if the  $b^k$ 's are  $\pm 1$  valued Why average predictors?

- to reduce bias
- to compensate for local optima (a form of bias)
- to reduce variance
- if b<sub>1</sub>, b<sub>2</sub>,... b<sub>M</sub> make independent errors, averaging reduces expected error (loss). We say that b<sub>1</sub> and b<sub>2</sub> make independent errors ⇔
   P(b<sub>1</sub> wrong | x) = P(b<sub>1</sub> wrong | x, b<sub>2</sub> wrong)
- because the b<sup>k</sup> functions are given: real world domain experts (weather prediction), a set of black-box classifiers (from a software package), a set of [expert designed] features (speech recognition, car/human recognition) each of them weakly informative of y
- $\blacktriangleright$  because  ${\cal B}$  is a set of (simple) "basis functions" and we need a more complex predictor in our task

## Averaging is not always the same thing

Depending how we choose  $\mathcal{B}, b^1, b^2, \dots b^M$  and  $\beta^1, \beta^2, \dots \beta^M$ , we can obtain very different effects.

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

We will examine

- Bayesian averaging (briefly)
- Mixtures of experts (briefly)
- Bagging (briefly)
- Backfitting (briefly)
- Stacking (see Murphy)
- Boosting

## Bayesian averaging

Assume any predictor  $b \in \mathcal{B}$  could be the "true" predictor with some **prior** probability  $P_0(b)$ . Learning means changing the probability distribution of b after seeing the data.

Before seeing data $P_0(b)$ prior probability of bAfter seeing  $\mathcal{D}$  $P(b|\mathcal{D})$ posterior probability of bBayes formula $P(b|\mathcal{D}) = \frac{P_0(b)P(\mathcal{D}|b)}{\sum_{b' \in \mathcal{B}} P_0(b')P(\mathcal{D}|b')}$ 

Classification of a new instance by Bayesian averaging:

$$f(x) = \int_{\mathcal{B}} b(x) dP(b|\mathcal{D})$$

or

$$P(y|x, D) = \int_{\mathcal{B}} \mathbf{1}_{b(x)=y} dP(b|D)$$

Hence classifiers (or more generally predictors) are weighted by their posterior probability.

Intuition: The likelihood becomes more concentrated when N increases

#### Bayesian averaging and model complexity



- model too simple: likelihood low, prior high
- model about right: likelihood high, prior not too low
- model too complex: likelihood high, prior very low

Bayesian averaging in practice.

$$\hat{P}(b^k|\mathcal{D}) = \frac{P_0(b^k)P(\mathcal{D}|b^k)}{\sum_{k'=1}^M P_0(b_{k'})P(\mathcal{D}|b_{k'})} \equiv \beta_k$$

 $b^{1:M}$  are either sampled from  $\mathcal{B}$ , or trained separately (e.g local minima of  $\hat{L}$ , models of different complexities) **Priors in practice** 

- non-informative
- complexity penalizing, sparsity inducing, etc

## Reducing variance: Bagging

What if we had several (independently sampled) training sets  $\mathcal{D}_1, \mathcal{D}_2, \dots \mathcal{D}_M$ ?

- ▶ we could train classifiers  $\{b^1, b^2, \dots b^M\}$  on the respective  $\mathcal{D}_1, \mathcal{D}_2, \dots \mathcal{D}_M$
- we could estimate  $E_{P(b)}[b] \sim f = \frac{1}{M} \sum_{k=1}^{M} b^k$
- f has always lower variance than b<sup>k</sup>

Idea of **bagging**: sample  $\mathcal{D}_1, \mathcal{D}_2, \dots \mathcal{D}_M$  from the given  $\mathcal{D}$  and estimate  $b^k$  on  $\mathcal{D}_k$ 

$$f(x) = \operatorname{sgn} \frac{1}{M} \sum_{k=1}^{M} b^{k}(x)$$

Thus, bagging is a form of boostrap.

## Bagging reduces variances

It was shown theoretically and empirically that bagging reduces variance. Bagging is good for

- base classifiers with high variance (complex)
- unstable classifiers (decision trees, decision lists, neural networks)
- noisy data

## Example

**Random Forests** A large ensemble of decision trees is fitted to the same data set, introducing randomness in various ways, such as (1) resampling the data set, (2) taking random splits, with probabilities that favor "good" splits, etc. The output predictor is the (unweighted) average of all the trees.

can be extended to more than classification (regression, feature selection)

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

- training can be done in parallel
- computing f(x) on new example is fast enough
- VERY POPULAR in industry

## Stacking

- very general method (any kind of predictors or costs)
- for complex base classifiers
- Fit predictors b<sup>1</sup>,...b<sup>M</sup> to the data D. The predictors can be from different model classes (i.e neural networks, CART, nearest neighbors, logistic regressions) or use different sets of features. For k = 1 : M, for data point i = 1 : N
   train b<sup>k</sup><sub>k</sub>; from the model class of b<sup>k</sup> on D \ {(x<sup>i</sup>, y<sup>i</sup>)}
- 2. Fit coefficients  $\beta_{1:M}$  by minimizing the leave-one-out (loo) empirical cost  $\hat{L}^{loo}$

$$\hat{L}^{loo}(\beta_{1:M} = \frac{1}{N} \sum_{i=1}^{N} L(y^{i}, \sum_{k=1}^{M} \beta_{k} b_{-i}^{k}(x^{i}))$$
(2)

(or by cross-validation).

Note that if *L* is a non-linear function, the minimization in (2) is a non-linear minimization, and in particular for convex losses this is a convex optimization problem in  $\beta_{1:M}$ .

## Mixtures of Experts

Here  $\beta_k = \beta_k(x)$ , and  $\sum_k \beta^k(x) = 1$  for all x in the domain of the inputs. Idea each  $b^k$  is an "expert" in one region of the input space. Wanted  $f \approx b^k(x)$  in the region of expertise of expert  $b^k$ The vector function  $\beta(x) = [\beta^1(x) \dots \beta^M(x)]$  is sometimes called the gating function.

#### Example

Suppose the true function to learn is  $f^*(x) = |x|, x \in \mathbb{R}$ . This can be well approximated by two linear experts  $f^1(x) = x$ ,  $f^2(x) = -x$  with weights  $\beta^1 = \phi(x), \beta^1 = \phi(-x)$ , where  $\phi$  is the sigmoid function (hence  $\beta^1 + \beta^2 = 1$  everywhere.

The example highlights that, by using a mixture of experts we can construct a more complex classifier by from simple classifiers (linear). "Simple" can mean easy to fit, or low complexity (aka simple decision region), or both. The effective sample size for each  $b^k$  is smaller than N, and corresponds to the data for which  $\beta^k(x^i)$  is away from 0.

For more than two experts, a natural gating function is the softmax function

$$\beta^{k}(x) = \frac{e^{v_{k}^{T}x}}{\sum_{l=1}^{M} e^{v_{l}^{T}x}} \text{ with } v_{l} \in \mathbb{R}^{n} \text{ a vector of parameters}$$
(3)

Training By descent methods. Often the experts  $f^{1:M}$  and gating functions  $\beta^{1:M}$  are trained simultaneously (to maximize log-likelihood) by steepest descent, or EM algorithm (which we'll study later).

## Forward Fitting and Backfitting

 $(b^k, \beta^k)$  are fitted iteratively (sequentially), one k at a time. In some cases, the weights  $\beta^k$  can be absorbed into  $b^k$ .  $f^t$  is f at iteration t The residual  $r \in \mathbb{R}^N$  is defined as  $r^t(x^i) = y^i - f^t(x^i)$ .

FORWARDFITTING ALGORITHM Input M, labeled training set  $\mathcal{D}$ Initialize f = 0repeat for k = 1, 2, ..., Mfit k-th predictor  $\beta^k, b^k = \operatorname{argmin} \hat{L}(f + \beta b)$ update  $f = f + b^k \beta^k$ until change in  $\hat{L}$  small enough (or, change in  $b^k$  small enough) Output  $f(x) = \sum_{k=1}^M \beta^k b^k(x)$ Note that M does not have to be set in advance (just like in boosting).

◆□▶ ◆□▶ ◆三▶ ◆三▶ ◆□ ◆ ◇◇◇

## Backfitting

Set M at the beginning, and cycle through the M predictors, updating predictor k while keeping the others fixed. Denote by

$$f^{-k}(x) = \sum_{l \neq k} \beta^k b^k(x) \tag{4}$$

the combined predictor f "minus" the k-th base predictor  $b^k$ .

BACKFITTING ALGORITHM Input *M*, labeled training set  $\mathcal{D}$ Initialize  $b^{1:M} = 0$ ,  $[\beta^{1:M} = 0$  if there are coefficients  $\beta$ ] repeat for k = 1, 2, ..., Mcalculate  $r^k(x^i) = y^i - f^{-k}(x^i)$ , i = 1 : Noptimize  $\hat{L}$  w.r.t k-th base predictor  $\beta^k$ ,  $b^k = \operatorname{argmin} \hat{L}(r^k + \beta b)$ until change in  $\hat{L}$  small enough (or, change in  $b^{1:M}$  small enough) Output  $f(x) = \sum_{k=1}^{M} \beta^k b^k(x)$ Often these methods are used when  $b^k$  are assumed to be simple, weak, and the predictor f is built from the cooperation of several b's. Thus, they are generally bias reduction method.

## Example

Least squares regression In this problem, it is useful to denote

$$r^{k}(x^{i}) = y^{i} - f^{-k}(x^{i})$$
 (5)

the residual of  $f^{-k}$  at data point  $x^i$ . Then,  $L_{LS}(y^i, f(x^i)) = y^i - f^{-k}(x^i) - \beta^k b^k(x^i)$ . Hence, optimizing  $\hat{L}_{LS}$  w.r.t.  $\beta^k, b^k$  can be expressed as

$$\beta^{k}, b^{k} = \operatorname{argmin}_{\beta, b} \sum_{i=1}^{N} (r^{i} - \beta b(x^{i}))^{2}.$$
(6)

In other words, each step of backfitting is a least squares regression problem, where the output variable values  $y^i$  are replaced with the current residuals  $r^i$ .

See also Additive Models



Base classifier family  $\underline{\mathcal{B}}$  has large bias (e.g. linear classifier, decision stumps) but learning always produces *b* that is better (on the training set) than random guessing. Preconditions for boosting

1. Learning algorithm accepts weighted data sets. Training minimizes

$$\hat{L}_{01}^{w}(b) = \sum_{i=1}^{W} w_i L_{01}(y^i, b(x^i)) \text{ with } \sum_{i=1}^{W} w_i = 1.$$

2.  $\mathcal{B}$  is a weak classifier family. For any  $\mathcal{D}$  and any weights  $w_{1:N}$  there can be found  $b \in \mathcal{B}$  such that the training error of b on  $\mathcal{D}$  is bounded below one half.

$$0 < \hat{L}_{01}^w(b) \leq \frac{\delta}{a} < \frac{1}{2}$$

Idea of boosting: train a classifier  $b^1$  on D, then train a  $b^2$  to correct the errors of  $b^1$ , then  $b^3$  to correct the errors of  $b^2$ , etc.

$$\hat{L}(b_1 \otimes) = \frac{1}{W} \sum_{i=1}^{N} L(y^i, b(x^i)) \underset{\text{Loss}}{\text{min}}$$

(日) (同) (三) (三) (三) (○) (○)



## Example

#### Boosting with stumps

Stumps are decision trees with a single split.

(below,  $c_1 \ldots c_4$  denote the coefficients  $\beta_{1:4}$ ).



▲□▶ ▲□▶ ▲三▶ ▲三▶ 三三 のへで

AdaBoost Algorithm  

$$k \neq i$$
  
 $k \neq i$   
 $k \neq i$   

#### Remarks

- If b(x) ∈ {±1} then y<sup>i</sup>b(x<sup>i</sup>) ∈ {±1}, and <sup>1-y<sup>i</sup>b(x<sup>i</sup>)</sup>/<sub>2</sub> = 1 if an error occurs and 0 otherwise. Thus, ε<sup>k</sup> in step 2 adds up the weights of the errors. If b(x) ∈ [-1,1] then the errors contribute different amounts to the loss depending on their margin.
- 2. In both cases,  $\varepsilon^k \in [0, \delta], \delta < 0.5$  by the weak learner property of  $\mathcal{B}$
- 3.  $\beta^k > 0$  whenever  $\varepsilon^k < 1/2$ .
- If b ∈ {±1}, then step 4 can be written equivalently (up to a multiplicative constant)

$$w_{i}^{k+1} = \begin{cases} \frac{1}{Z^{k}} w_{i}^{k} & \text{if } b^{k}(x_{i}) = y_{i} \\ \frac{1}{Z^{k}} w_{i}^{k} e^{2\beta^{k}} & \text{if } b^{k}(x_{i}) \neq y_{i} \end{cases}$$
(7)

This form corresponds to the DISCRETEADABOOST algorithm, the first AdaBoost algorithm published, which assumed  $b(x) \in \{\pm 1\}$ . As we shall see later, modern boosting algorithms dispense with the assumption  $b \in [-1, 1]$  too.



• weight of example *i* at step *k* is proportional to  $e^{-y_i f^{k-1}(x_i)}$  the exponential of its negative margin

イロト イポト イヨト イヨト

-

- Examples that have been hard to classify get exponentially high weights.
- Examples that are classified with high margins get vanishingly small weights.

## 2. The normalization constant is an average loss

If we sum both sides of (8) over i we obtain

or

where

$$\phi(z) = e^{-z}. \tag{11}$$

Lou

(9)

and

$$L_{\phi}(y, f(x)) = \phi(yf(x))$$
(12)

 $\sum_{i} W_{i}^{kH} = 1 = \sum_{i} e^{-y_{i}f^{k}(x_{i})}$   $\prod_{i} Z^{k} = \frac{\sum_{i} e^{-y_{i}f^{k}(x_{i})}}{N} \equiv \hat{L}_{\phi}(f^{k}) = \frac{1}{N} \sum_{i} \varphi(y_{i}^{k})$ 

Hence, the r.h.s of (10) is the average over the data set of the **exponential loss**  $L_{\phi}$ . The function  $\phi$  decreases with the margin, thus decreasing  $\hat{L}_{\phi}$  will produce a better classifier (on the training set). In this sense,  $L_{\phi}$  is an alternative loss function for classification.

## $\mathbf{\mathcal{F}}_{L_{\phi}}$ decreases exponentially with M

where  $\gamma < \mathbf{1}$  depends on  $\delta$  the maximum error. It follows that

# The training set error $\hat{L}_{01}$ decreases exponentially with M

Note that  $\phi(z) \ge 1_{z<0}$  for all z (see also figure ??). Therefore

$$\hat{L}(f^k) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{[y_i f^k(x_i) < 0]}$$
(18)

$$\leq \frac{1}{N} \sum_{i=1}^{N} e^{-y_i f^k(x_i)} = \hat{L}_{\phi}(f^k) \leq \gamma^k$$
 (19)

In other words, the training error  $\hat{L}(f^k)$  is bounded by a decaying exponential. Moreover, since  $\hat{L}(f^k) \in \{0, 1/N, 2/N, \dots 1\}$ , it follows that after a finite number of steps, when  $\gamma^{k^0} < 1/N$ , the training error will become 0 and the training data will be perfectly classified!

## The test set error and overfitting

- ▶ Do NOT take  $M = k^0$ . The number of steps M for good generalization error is often much larger than  $k^0$  (and sometimes smaller).
- Below is a typical plot of  $\hat{L}$  and L (which can be estimated from an independent sample) vs the number of boosting iterations.



▲□▶ ▲圖▶ ★ 国▶ ★ 国▶ - 国 - のへで