

CSC 411: Lecture 17: Ensemble Methods I

Richard Zemel, Raquel Urtasun and Sanja Fidler

University of Toronto

- Ensemble Methods
- Bagging
- Boosting

Ensemble methods

- Typical application: classification

Ensemble methods

- Typical application: classification
- **Ensemble** of classifiers is a set of classifiers whose individual decisions are combined in some way to classify new examples

Ensemble methods

- Typical application: classification
- **Ensemble** of classifiers is a set of classifiers whose individual decisions are combined in some way to classify new examples
- Simplest approach:
 1. Generate multiple classifiers
 2. Each votes on test instance
 3. Take majority as classification

Ensemble methods

- Typical application: classification
- **Ensemble** of classifiers is a set of classifiers whose individual decisions are combined in some way to classify new examples
- Simplest approach:
 1. Generate multiple classifiers
 2. Each votes on test instance
 3. Take majority as classification
- Classifiers are different due to different sampling of training data, or randomized parameters within the classification algorithm

Ensemble methods

- Typical application: classification
- **Ensemble** of classifiers is a set of classifiers whose individual decisions are combined in some way to classify new examples
- Simplest approach:
 1. Generate multiple classifiers
 2. Each votes on test instance
 3. Take majority as classification
- Classifiers are different due to different sampling of training data, or randomized parameters within the classification algorithm
- Aim: take simple mediocre algorithm and transform it into a super classifier without requiring any fancy new algorithm

Ensemble methods: Summary

- Differ in training strategy, and combination method

Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets

Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets
 1. **Bagging** (bootstrap aggregation) – train separate models on overlapping training sets, average their predictions

Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets
 1. **Bagging** (bootstrap aggregation) – train separate models on overlapping training sets, average their predictions
 - ▶ Sequential training, iteratively re-weighting training examples so current classifier focuses on hard examples: **boosting**

Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets
 1. **Bagging** (bootstrap aggregation) – train separate models on overlapping training sets, average their predictions
 - ▶ Sequential training, iteratively re-weighting training examples so current classifier focuses on hard examples: **boosting**
 - ▶ Parallel training with objective encouraging division of labor: **mixture of experts**

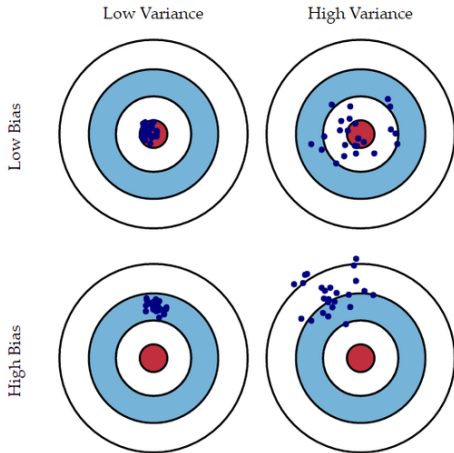
Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets
 1. **Bagging** (bootstrap aggregation) – train separate models on overlapping training sets, average their predictions
 - ▶ Sequential training, iteratively re-weighting training examples so current classifier focuses on hard examples: **boosting**
 - ▶ Parallel training with objective encouraging division of labor: **mixture of experts**
- Notes:
 - ▶ Also known as **meta-learning**

Ensemble methods: Summary

- Differ in training strategy, and combination method
 - ▶ Parallel training with different training sets
 1. **Bagging** (bootstrap aggregation) – train separate models on overlapping training sets, average their predictions
 - ▶ Sequential training, iteratively re-weighting training examples so current classifier focuses on hard examples: **boosting**
 - ▶ Parallel training with objective encouraging division of labor: **mixture of experts**
- Notes:
 - ▶ Also known as **meta-learning**
 - ▶ Typically applied to weak models, such as decision stumps (single-node decision trees), or linear classifiers

Bias-Variance Decomposition



Variance-bias Tradeoff

- Minimize two sets of errors:
 1. **Variance**: error from sensitivity to small fluctuations in the training set

Variance-bias Tradeoff

- Minimize two sets of errors:
 1. **Variance**: error from sensitivity to small fluctuations in the training set
 2. **Bias**: erroneous assumptions in the model

Variance-bias Tradeoff

- Minimize two sets of errors:
 1. **Variance**: error from sensitivity to small fluctuations in the training set
 2. **Bias**: erroneous assumptions in the model
- **Variance-bias decomposition** is a way of analyzing the generalization error as a sum of 3 terms: variance, bias and irreducible error (resulting from the problem itself)

Why do Ensemble Methods Work?

- Based on one of two basic observations:

Why do Ensemble Methods Work?

- Based on one of two basic observations:
 1. **Variance reduction**: if the training sets are completely independent, it will always help to average an ensemble because this will reduce variance without affecting bias (e.g., bagging)
 - ▶ reduce sensitivity to individual data points

Why do Ensemble Methods Work?

- Based on one of two basic observations:
 1. **Variance reduction**: if the training sets are completely independent, it will always help to average an ensemble because this will reduce variance without affecting bias (e.g., bagging)
 - ▶ reduce sensitivity to individual data points
 2. **Bias reduction**: for simple models, average of models has much greater capacity than single model (e.g., hyperplane classifiers, Gaussian densities).

Why do Ensemble Methods Work?

- Based on one of two basic observations:
 1. **Variance reduction**: if the training sets are completely independent, it will always help to average an ensemble because this will reduce variance without affecting bias (e.g., bagging)
 - ▶ reduce sensitivity to individual data points
 2. **Bias reduction**: for simple models, average of models has much greater capacity than single model (e.g., hyperplane classifiers, Gaussian densities).
 - ▶ Averaging models can reduce bias substantially by increasing capacity, and control variance by fitting one component at a time (e.g., boosting)

Ensemble Methods: Justification

- Ensemble methods more accurate than any individual members if:
 - ▶ Accurate (better than guessing)
 - ▶ Diverse (different errors on new examples)

Ensemble Methods: Justification

- Ensemble methods more accurate than any individual members if:
 - ▶ Accurate (better than guessing)
 - ▶ Diverse (different errors on new examples)
- Why?

Ensemble Methods: Justification

- Ensemble methods more accurate than any individual members if:
 - ▶ Accurate (better than guessing)
 - ▶ Diverse (different errors on new examples)
- Why?
- Independent errors: prob k of N classifiers (independent error rate ϵ) wrong:

$$P(\text{num errors} = k) = \binom{N}{k} \epsilon^k (1 - \epsilon)^{N-k}$$

Ensemble Methods: Justification

- Ensemble methods more accurate than any individual members if:
 - ▶ Accurate (better than guessing)
 - ▶ Diverse (different errors on new examples)
- Why?
- Independent errors: prob k of N classifiers (independent error rate ϵ) wrong:

$$P(\text{num errors} = k) = \binom{N}{k} \epsilon^k (1 - \epsilon)^{N-k}$$

- Probability that majority vote wrong: error under distribution where more than $N/2$ wrong

Ensemble Methods: Justification

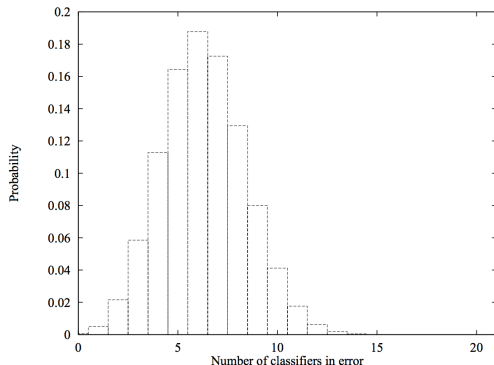


Figure : Example: The probability that exactly K (out of 21) classifiers will make an error assuming each classifier has an error rate of $\epsilon = 0.3$ and makes its errors independently of the other classifier. The area under the curve for 11 or more classifiers being simultaneously wrong is 0.026 (much less than ϵ).

[Credit: T. G Dietterich, Ensemble Methods in Machine Learning]

Ensemble Methods: Justification

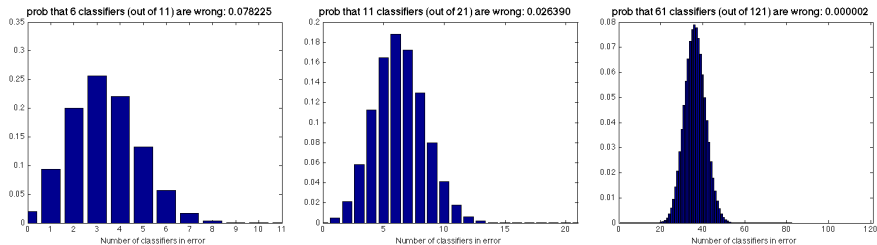


Figure : $\epsilon = 0.3$: (left) $N = 11$ classifiers, (middle) $N = 21$, (right) $N = 121$.

Ensemble Methods: Justification

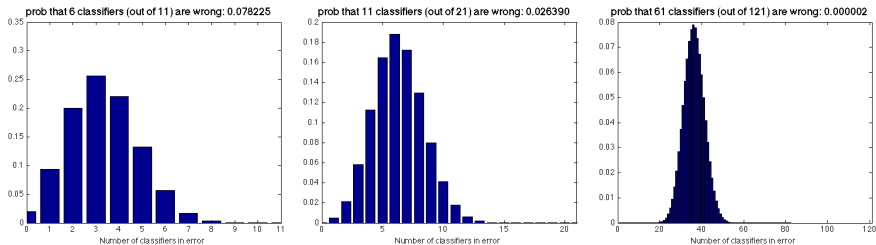


Figure : $\epsilon = 0.3$: (left) $N = 11$ classifiers, (middle) $N = 21$, (right) $N = 121$.

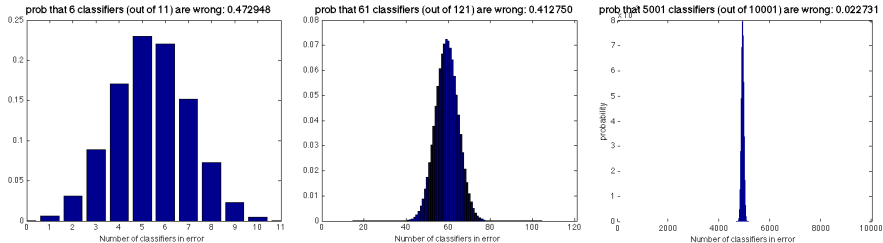


Figure : $\epsilon = 0.49$: (left) $N = 11$, (middle) $N = 121$, (right) $N = 10001$.

Ensemble Methods: Netflix

- Clear demonstration of the power of ensemble methods

Ensemble Methods: Netflix

- Clear demonstration of the power of ensemble methods
- Original progress prize winner (BellKor) was ensemble of 107 models!

- Clear demonstration of the power of ensemble methods
- Original progress prize winner (BellKor) was ensemble of 107 models!
 - ▶ *"Our experience is that most efforts should be concentrated in deriving substantially different approaches, rather than refining a simple technique."*

- Clear demonstration of the power of ensemble methods
- Original progress prize winner (BellKor) was ensemble of 107 models!
 - ▶ *"Our experience is that most efforts should be concentrated in deriving substantially different approaches, rather than refining a simple technique."*
 - ▶ *"We strongly believe that the success of an ensemble approach depends on the ability of its various predictors to expose different complementing aspects of the data. Experience shows that this is very different than optimizing the accuracy of each individual predictor."*

BELLKOR'S PRAGMATIC CHAOS WINS \$1 MILLION NETFLIX PRIZE BY MERE MINUTES



Bootstrap Estimation

- Repeatedly draw n samples from D

Bootstrap Estimation

- Repeatedly draw n samples from D
- For each set of samples, estimate a statistic

Bootstrap Estimation

- Repeatedly draw n samples from D
- For each set of samples, estimate a statistic
- The bootstrap estimate is the mean of the individual estimates

Bootstrap Estimation

- Repeatedly draw n samples from D
- For each set of samples, estimate a statistic
- The bootstrap estimate is the mean of the individual estimates
- Used to estimate a statistic (parameter) and its variance

Bootstrap Estimation

- Repeatedly draw n samples from D
- For each set of samples, estimate a statistic
- The bootstrap estimate is the mean of the individual estimates
- Used to estimate a statistic (parameter) and its variance
- Bagging: bootstrap aggregation (Breiman 1994)

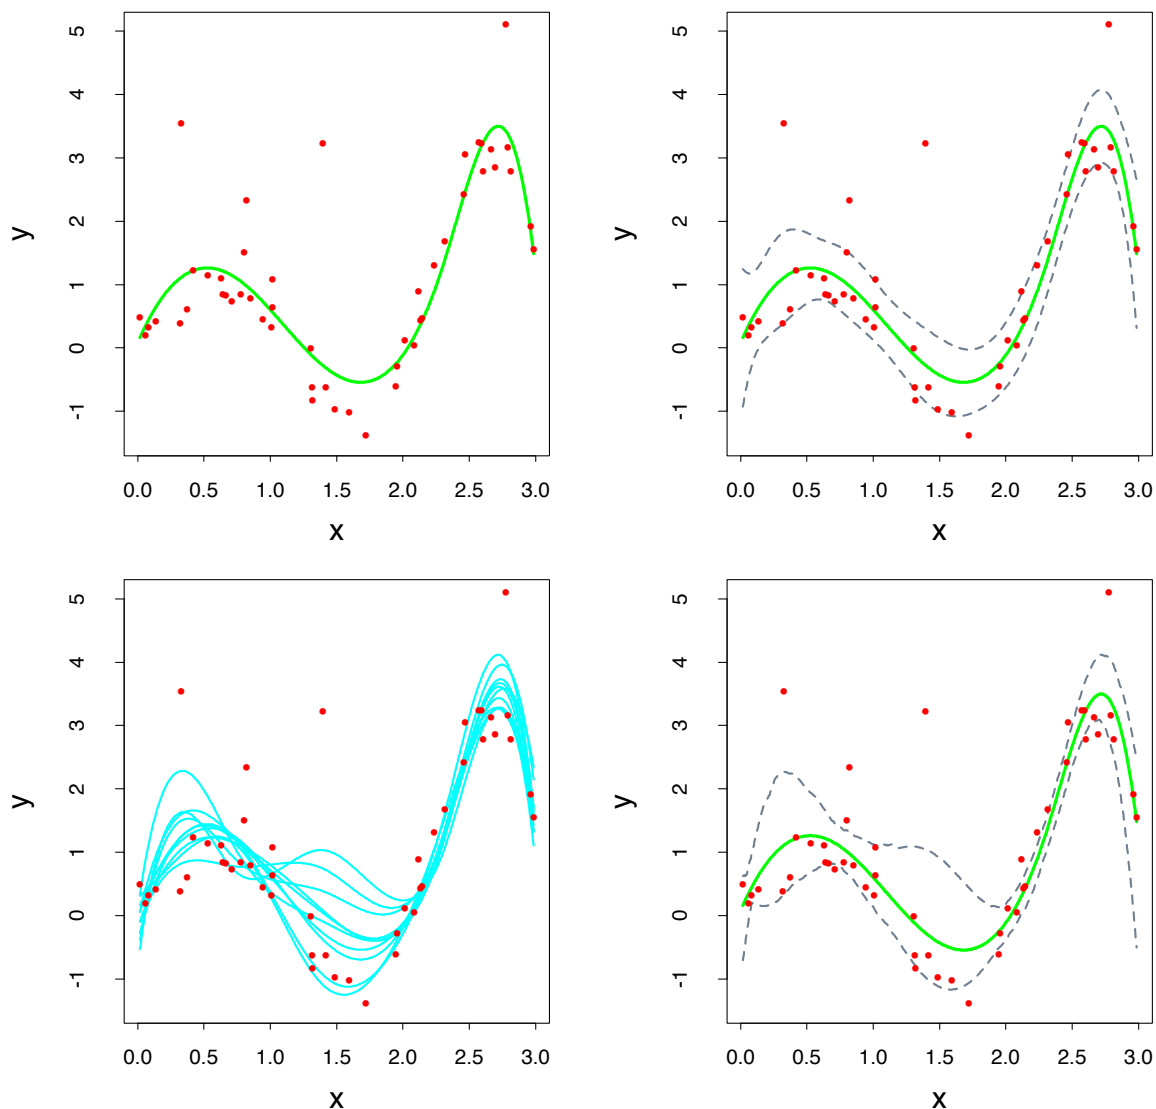


FIGURE 8.2. (Top left:) *B-spline smooth of data.* (Top right:) *B-spline smooth plus and minus $1.96 \times$ standard error bands.* (Bottom left:) *Ten bootstrap replicates of the B-spline smooth.* (Bottom right:) *B-spline smooth with 95% standard error bands computed from the bootstrap distribution.*

Bagging

- Simple idea: generate M bootstrap samples from your original training set. Train on each one to get y_m , and average them

$$y_{bag}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M y_m(\mathbf{x})$$

Bagging

- Simple idea: generate M bootstrap samples from your original training set. Train on each one to get y_m , and average them

$$y_{bag}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M y_m(\mathbf{x})$$

- For regression: average predictions

Bagging

- Simple idea: generate M bootstrap samples from your original training set. Train on each one to get y_m , and average them

$$y_{bag}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M y_m(\mathbf{x})$$

- For regression: average predictions
- For classification: average class probabilities (or take the majority vote if only hard outputs available)

Bagging

- Simple idea: generate M bootstrap samples from your original training set. Train on each one to get y_m , and average them

$$y_{bag}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M y_m(\mathbf{x})$$

- For regression: average predictions
- For classification: average class probabilities (or take the majority vote if only hard outputs available)
- Bagging approximates the Bayesian posterior mean. The more bootstraps the better, so use as many as you have time for

Bagging

- Simple idea: generate M bootstrap samples from your original training set. Train on each one to get y_m , and average them

$$y_{bag}^M(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M y_m(\mathbf{x})$$

- For regression: average predictions
- For classification: average class probabilities (or take the majority vote if only hard outputs available)
- Bagging approximates the Bayesian posterior mean. The more bootstraps the better, so use as many as you have time for
- Each bootstrap sample is drawn with replacement, so each one contains some duplicates of certain training points and leaves out other training points completely

Boosting (AdaBoost): Summary

- Also works by manipulating training set, but **classifiers trained sequentially**

Boosting (AdaBoost): Summary

- Also works by manipulating training set, but **classifiers trained sequentially**
- Each classifier trained given knowledge of the performance of previously trained classifiers: **focus on hard examples**

Boosting (AdaBoost): Summary

- Also works by manipulating training set, but **classifiers trained sequentially**
- Each classifier trained given knowledge of the performance of previously trained classifiers: **focus on hard examples**
- Final classifier: weighted sum of component classifiers

Making Weak Learners Stronger

- Suppose you have a weak learning module (a **base classifier**) that can always get $(0.5 + \epsilon)$ correct when given a two-way classification task

Making Weak Learners Stronger

- Suppose you have a weak learning module (a **base classifier**) that can always get $(0.5 + \epsilon)$ correct when given a two-way classification task
 - ▶ That seems like a weak assumption but beware!

Making Weak Learners Stronger

- Suppose you have a weak learning module (a **base classifier**) that can always get $(0.5 + \epsilon)$ correct when given a two-way classification task
 - ▶ That seems like a weak assumption but beware!
- Can you apply this learning module many times to get a strong learner that can get close to zero error rate on the training data?

Making Weak Learners Stronger

- Suppose you have a weak learning module (a **base classifier**) that can always get $(0.5 + \epsilon)$ correct when given a two-way classification task
 - ▶ That seems like a weak assumption but beware!
- Can you apply this learning module many times to get a strong learner that can get close to zero error rate on the training data?
 - ▶ Theorists showed how to do this and it actually led to an effective new learning procedure (Freund & Shapire, 1996)

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.
- Then re-weight the training data to emphasize the hard cases and train a second model.

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.
- Then re-weight the training data to emphasize the hard cases and train a second model.
 - ▶ How do we re-weight the data?

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.
- Then re-weight the training data to emphasize the hard cases and train a second model.
 - ▶ How do we re-weight the data?
- Keep training new models on re-weighted data

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.
- Then re-weight the training data to emphasize the hard cases and train a second model.
 - ▶ How do we re-weight the data?
- Keep training new models on re-weighted data
- Finally, use a weighted committee of all the models for the test data.

Boosting (ADABOOST)

- First train the base classifier on all the training data with equal importance weights on each case.
- Then re-weight the training data to emphasize the hard cases and train a second model.
 - ▶ How do we re-weight the data?
- Keep training new models on re-weighted data
- Finally, use a weighted committee of all the models for the test data.
 - ▶ How do we weight the models in the committee?

How to Train Each Classifier

- Input: \mathbf{x} , Output: $y(\mathbf{x}) \in \{1, -1\}$

How to Train Each Classifier

- Input: \mathbf{x} , Output: $y(\mathbf{x}) \in \{1, -1\}$
- Target $t \in \{-1, 1\}$

How to Train Each Classifier

- Input: \mathbf{x} , Output: $y(\mathbf{x}) \in \{1, -1\}$
- Target $t \in \{-1, 1\}$
- Weight on example n for classifier m : \mathbf{w}_n^m

How to Train Each Classifier

- Input: \mathbf{x} , Output: $y(\mathbf{x}) \in \{1, -1\}$
- Target $t \in \{-1, 1\}$
- Weight on example n for classifier m : \mathbf{w}_n^m
- Cost function for classifier m

$$J_m = \sum_{n=1}^N w_n^m \underbrace{[y_m(\mathbf{x}^n) \neq t^{(n)}]}_{\substack{1 \text{ if error,} \\ 0 \text{ o.w.}}} = \sum \text{weighted errors}$$

How to weight each training case for classifier m

- Recall cost function is

$$J_m = \sum_{n=1}^N w_n^m \underbrace{[y_m(\mathbf{x}^n) \neq t^{(n)}]}_{\substack{1 \text{ if error,} \\ 0 \text{ o.w.}}} = \sum \text{weighted errors}$$

How to weight each training case for classifier m

- Recall cost function is

$$J_m = \sum_{n=1}^N w_n^m \underbrace{[y_m(\mathbf{x}^n) \neq t^{(n)}]}_{\substack{1 \text{ if error,} \\ 0 \text{ o.w.}}} = \sum \text{weighted errors}$$

- Weighted error rate of a classifier

$$\epsilon_m = \frac{J_m}{\sum w_n^m}$$

How to weight each training case for classifier m

- Recall cost function is

$$J_m = \sum_{n=1}^N w_n^m \underbrace{[y_m(\mathbf{x}^n) \neq t^{(n)}]}_{\substack{1 \text{ if error,} \\ 0 \text{ o.w.}}} = \sum \text{weighted errors}$$

- Weighted error rate of a classifier

$$\epsilon_m = \frac{J_m}{\sum w_n^m}$$

- The quality of the classifier is

$$\alpha_m = \ln \left(\frac{1 - \epsilon_m}{\epsilon_m} \right)$$

It is zero if the classifier has weighted error rate of 0.5 and infinity if the classifier is perfect

How to weight each training case for classifier m

- Recall cost function is

$$J_m = \sum_{n=1}^N w_n^m \underbrace{[y_m(\mathbf{x}^n) \neq t^{(n)}]}_{\substack{1 \text{ if error,} \\ 0 \text{ o.w.}}} = \sum \text{weighted errors}$$

- Weighted error rate of a classifier

$$\epsilon_m = \frac{J_m}{\sum w_n^m}$$

- The quality of the classifier is

$$\alpha_m = \ln \left(\frac{1 - \epsilon_m}{\epsilon_m} \right)$$

It is zero if the classifier has weighted error rate of 0.5 and infinity if the classifier is perfect

- The weights for the next round are then

$$w_n^{m+1} = \exp \left(-\frac{1}{2} t^{(n)} \sum_{i=1}^m \alpha_i y_i(\mathbf{x}^{(n)}) \right) = w_n^m \exp \left(-\frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)}) \right)$$

How to make predictions using a committee of classifiers

- Weight the binary prediction of each classifier by the quality of that classifier:

$$y_M(\mathbf{x}) = \text{sign} \left(\sum_{m=1}^M \frac{1}{2} \alpha_m y_m(\mathbf{x}) \right)$$

- This is how to do inference, i.e., how to compute the prediction for each new example.

AdaBoost Algorithm

- Input: $\{\mathbf{x}^{(n)}, t^{(n)}\}_{n=1}^N$, and **WeakLearn**: learning procedure, produces classifier $y(\mathbf{x})$
- Initialize example weights: $w_n^m(\mathbf{x}) = 1/N$
- For $m=1:M$

- ▶ $y_m(\mathbf{x}) = \text{WeakLearn}(\{\mathbf{x}\}, \mathbf{t}, \mathbf{w})$, fit classifier by minimizing

$$J_m = \sum_{n=1}^N w_n^m [y_m(\mathbf{x}^n) \neq t^{(n)}]$$

- ▶ Compute unnormalized error rate

$$\epsilon_m = \frac{J_m}{\sum w_n^m}$$

- ▶ Compute classifier coefficient $\alpha_m = \log \frac{1-\epsilon_m}{\epsilon_m}$
- ▶ Update data weights

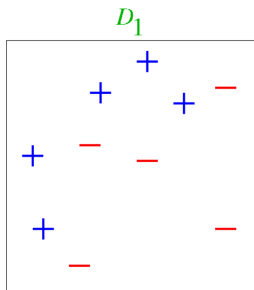
$$w_n^{m+1} = w_n^m \exp\left(-\frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right)$$

- Final model

$$Y(\mathbf{x}) = \text{sign}(y_M(\mathbf{x})) = \text{sign}\left(\sum_{m=1}^M \alpha_m y_m(\mathbf{x})\right)$$

AdaBoost Example

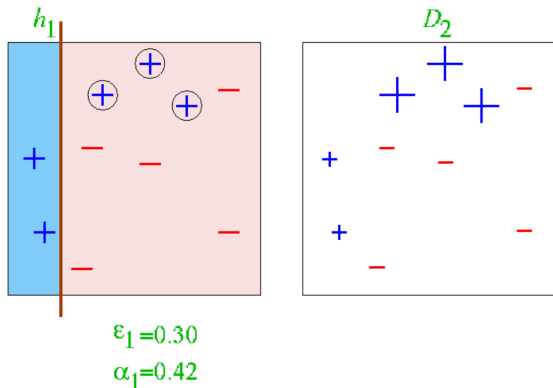
- Training data



[Slide credit: Verma & Thrun]

AdaBoost Example

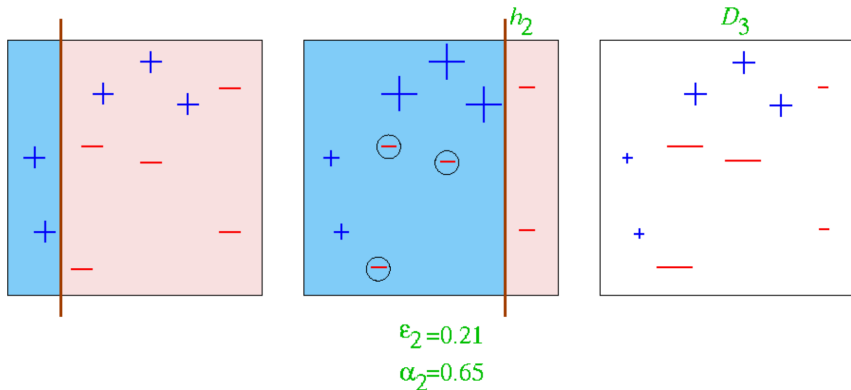
- Round 1



[Slide credit: Verma & Thrun]

AdaBoost Example

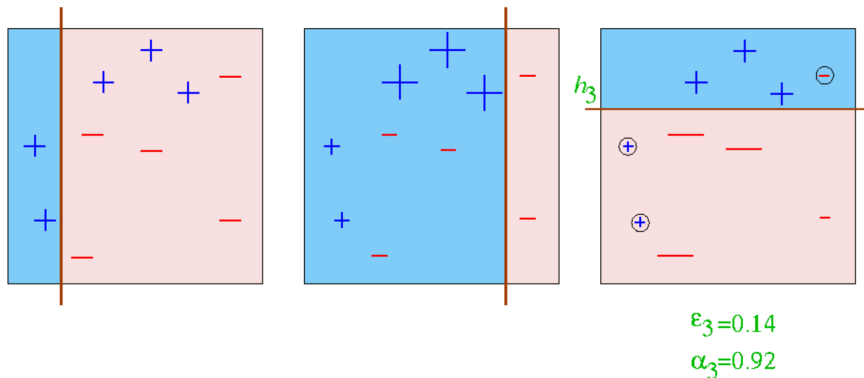
- Round 2



[Slide credit: Verma & Thrun]

AdaBoost Example

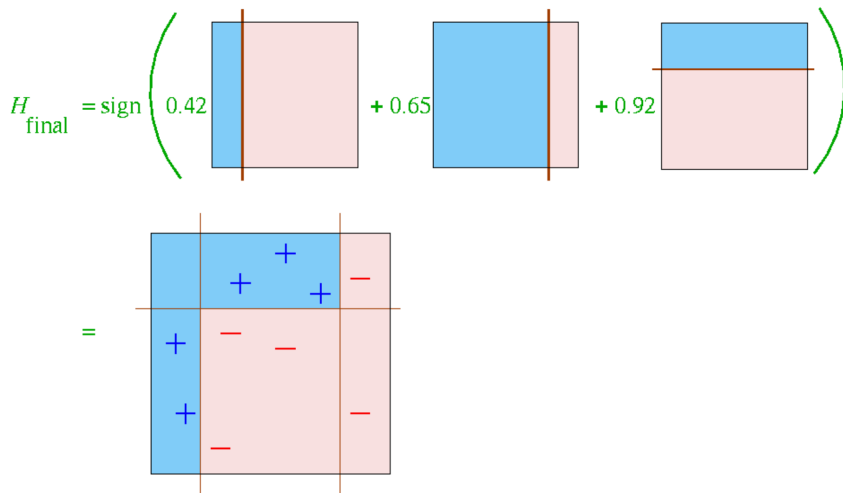
- Round 3



[Slide credit: Verma & Thrun]

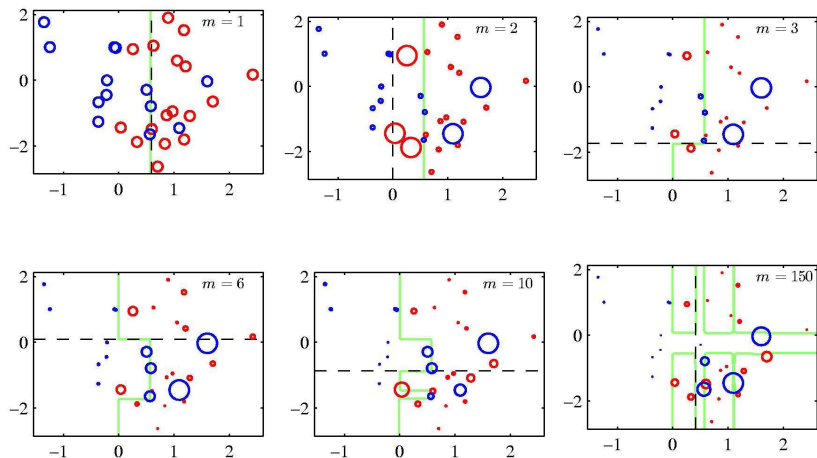
AdaBoost Example

- Final classifier



[Slide credit: Verma & Thrun]

AdaBoost example



- Each figure shows the number m of base learners trained so far, the decision of the most recent learner (dashed black), and the boundary of the ensemble (green)

AdaBoost Applet: <http://cseweb.ucsd.edu/~yfreund/adaboost/>

An alternative derivation of ADABOOST

- Just write down the right cost function and optimize each parameter to minimize it
 - ▶ [stagewise additive modeling](#) (Friedman et. al. 2000)

An alternative derivation of ADABOOST

- Just write down the right cost function and optimize each parameter to minimize it
 - ▶ **stagewise additive modeling** (Friedman et. al. 2000)
- At each step employ the exponential loss function for classifier m

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

An alternative derivation of ADAboost

- Just write down the right cost function and optimize each parameter to minimize it
 - ▶ **stagewise additive modeling** (Friedman et. al. 2000)
- At each step employ the exponential loss function for classifier m

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

- Real-valued prediction by committee of models up to m

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x})$$

An alternative derivation of ADaboost

- Just write down the right cost function and optimize each parameter to minimize it
 - ▶ **stagewise additive modeling** (Friedman et. al. 2000)

- At each step employ the exponential loss function for classifier m

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

- Real-valued prediction by committee of models up to m

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x})$$

- We want to minimize E w.r.t. α_m and the parameters of the classifier $y_m(\mathbf{x})$

An alternative derivation of ADaboost

- Just write down the right cost function and optimize each parameter to minimize it
 - ▶ **stagewise additive modeling** (Friedman et. al. 2000)

- At each step employ the exponential loss function for classifier m

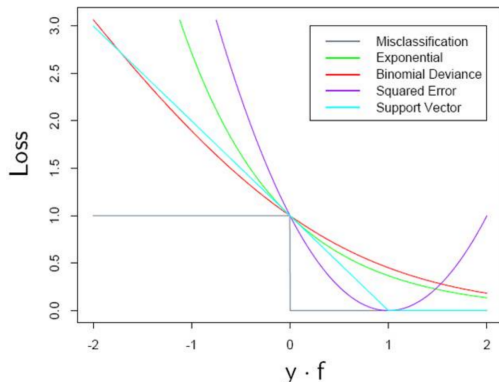
$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

- Real-valued prediction by committee of models up to m

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x})$$

- We want to minimize E w.r.t. α_m and the parameters of the classifier $y_m(\mathbf{x})$
- We do this in a sequential manner, one classifier at a time

Loss Functions



- Misclassification: 0/1 loss
- Exponential loss: $\exp(-t \cdot f(x))$ (AdaBoost)
- Squared error: $(t - f(x))^2$
- Soft-margin support vector (hinge loss): $\max(0, 1 - t \cdot y)$

Learning classifier m using exponential loss

- At iteration m , the energy is computed as

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

Learning classifier m using exponential loss

- At iteration m , the energy is computed as

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

with

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x}) = \frac{1}{2} \alpha_m y_m(\mathbf{x}) + \frac{1}{2} \sum_{i=1}^{m-1} \alpha_i y_i(\mathbf{x})$$

Learning classifier m using exponential loss

- At iteration m , the energy is computed as

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

with

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x}) = \frac{1}{2} \alpha_m y_m(\mathbf{x}) + \frac{1}{2} \sum_{i=1}^{m-1} \alpha_i y_i(\mathbf{x})$$

- We can compute the part that is relevant for the m -th classifier

$$E_{\text{relevant}} = \sum_{n=1}^N \exp\left(-t^{(n)} f_{m-1}(\mathbf{x}^{(n)}) - \frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right)$$

Learning classifier m using exponential loss

- At iteration m , the energy is computed as

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

with

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x}) = \frac{1}{2} \alpha_m y_m(\mathbf{x}) + \frac{1}{2} \sum_{i=1}^{m-1} \alpha_i y_i(\mathbf{x})$$

- We can compute the part that is relevant for the m -th classifier

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N \exp\left(-t^{(n)} f_{m-1}(\mathbf{x}^{(n)}) - \frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right) \\ &= \sum_{n=1}^N w_n^m \exp\left(-\frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right) \end{aligned}$$

Learning classifier m using exponential loss

- At iteration m , the energy is computed as

$$E = \sum_{n=1}^N \exp\{-t^{(n)} f_m(\mathbf{x}^{(n)})\}$$

with

$$f_m(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^m \alpha_i y_i(\mathbf{x}) = \frac{1}{2} \alpha_m y_m(\mathbf{x}) + \frac{1}{2} \sum_{i=1}^{m-1} \alpha_i y_i(\mathbf{x})$$

- We can compute the part that is relevant for the m -th classifier

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N \exp\left(-t^{(n)} f_{m-1}(\mathbf{x}^{(n)}) - \frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right) \\ &= \sum_{n=1}^N w_n^m \exp\left(-\frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right) \end{aligned}$$

with $w_n^m = \exp(-t^{(n)} f_{m-1}(\mathbf{x}^{(n)}))$

Continuing the derivation

$$E_{\text{relevant}} = \sum_{n=1}^N w_n^m \exp\left(-t^{(n)} \frac{\alpha_m}{2} y_m(\mathbf{x}^{(n)})\right)$$

Continuing the derivation

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N w_n^m \exp\left(-t^{(n)} \frac{\alpha_m}{2} y_m(\mathbf{x}^{(n)})\right) \\ &= e^{-\frac{\alpha_m}{2}} \sum_{\text{right}} w_n^m + e^{\frac{\alpha_m}{2}} \sum_{\text{wrong}} w_n^m \end{aligned}$$

Continuing the derivation

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N w_n^m \exp\left(-t^{(n)} \frac{\alpha_m}{2} y_m(\mathbf{x}^{(n)})\right) \\ &= e^{-\frac{\alpha_m}{2}} \sum_{\text{right}} w_n^m + e^{\frac{\alpha_m}{2}} \sum_{\text{wrong}} w_n^m \\ &= \underbrace{\left(e^{\frac{\alpha_m}{2}} - e^{-\frac{\alpha_m}{2}}\right)}_{\text{multiplicative constant}} \underbrace{\sum_n w_n^m [t^{(n)} \neq y_m(\mathbf{x}^{(n)})]}_{\text{wrong cases}} + e^{-\frac{\alpha_m}{2}} \underbrace{\sum_n w_n^m}_{\text{unmodifiable}} \end{aligned}$$

Continuing the derivation

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N w_n^m \exp\left(-t^{(n)} \frac{\alpha_m}{2} y_m(\mathbf{x}^{(n)})\right) \\ &= e^{-\frac{\alpha_m}{2}} \underbrace{\sum_{\text{right}} w_n^m}_{\text{right}} + e^{\frac{\alpha_m}{2}} \underbrace{\sum_{\text{wrong}} w_n^m}_{\text{wrong}} \\ &= \underbrace{\left(e^{\frac{\alpha_m}{2}} - e^{-\frac{\alpha_m}{2}}\right)}_{\text{multiplicative constant}} \underbrace{\sum_n w_n^m [t^{(n)} \neq y_m(\mathbf{x}^{(n)})]}_{\text{wrong cases}} + e^{-\frac{\alpha_m}{2}} \underbrace{\sum_n w_n^m}_{\text{unmodifiable}} \end{aligned}$$

- The second term is constant w.r.t. $y_m(\mathbf{x})$

Continuing the derivation

$$\begin{aligned} E_{\text{relevant}} &= \sum_{n=1}^N w_n^m \exp\left(-t^{(n)} \frac{\alpha_m}{2} y_m(\mathbf{x}^{(n)})\right) \\ &= e^{-\frac{\alpha_m}{2}} \sum_{\text{right}} w_n^m + e^{\frac{\alpha_m}{2}} \sum_{\text{wrong}} w_n^m \\ &= \underbrace{\left(e^{\frac{\alpha_m}{2}} - e^{-\frac{\alpha_m}{2}}\right)}_{\text{multiplicative constant}} \underbrace{\sum_n w_n^m [t^{(n)} \neq y_m(\mathbf{x}^{(n)})]}_{\text{wrong cases}} + e^{-\frac{\alpha_m}{2}} \underbrace{\sum_n w_n^m}_{\text{unmodifiable}} \end{aligned}$$

- The second term is constant w.r.t. $y_m(\mathbf{x})$
- Thus we minimize the weighted number of wrong examples

AdaBoost Algorithm

- Input: $\{\mathbf{x}^{(n)}, t^{(n)}\}_{n=1}^N$, and **WeakLearn**: learning procedure, produces classifier $y(\mathbf{x})$
- Initialize example weights: $w_n^m(\mathbf{x}) = 1/N$
- For $m=1:M$

- ▶ $y_m(\mathbf{x}) = \text{WeakLearn}(\{\mathbf{x}\}, \mathbf{t}, \mathbf{w})$, fit classifier by minimizing

$$J_m = \sum_{n=1}^N w_n^m [y_m(\mathbf{x}^n) \neq t^{(n)}]$$

- ▶ Compute unnormalized error rate

$$\epsilon_m = \frac{J_m}{\sum w_n^m}$$

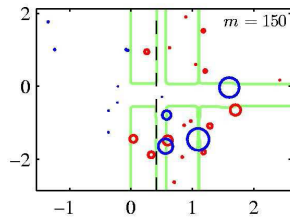
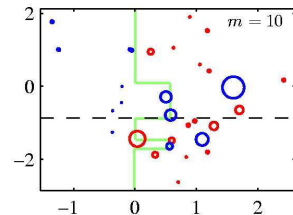
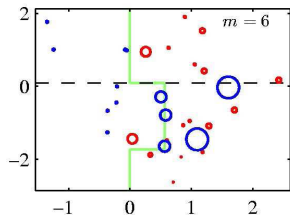
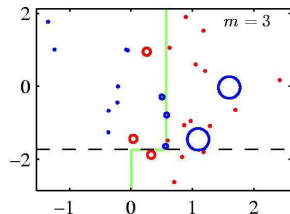
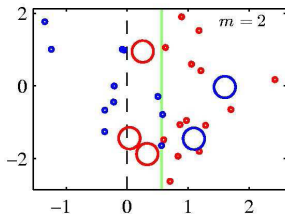
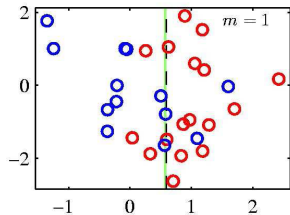
- ▶ Compute classifier coefficient $\alpha_m = \log \frac{1-\epsilon_m}{\epsilon_m}$
- ▶ Update data weights

$$w_n^{m+1} = w_n^m \exp\left(-\frac{1}{2} t^{(n)} \alpha_m y_m(\mathbf{x}^{(n)})\right)$$

- Final model

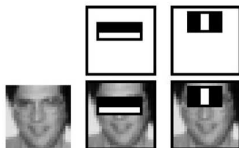
$$Y(\mathbf{x}) = \text{sign}(y_M(\mathbf{x})) = \text{sign}\left(\sum_{m=1}^M \alpha_m y_m(\mathbf{x})\right)$$

AdaBoost Example



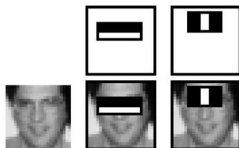
An impressive example of boosting

- Viola and Jones created a very fast face detector that can be scanned across a large image to find the faces.



An impressive example of boosting

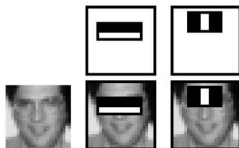
- Viola and Jones created a very fast face detector that can be scanned across a large image to find the faces.



- The base classifier/weak learner just compares the total intensity in two rectangular pieces of the image.

An impressive example of boosting

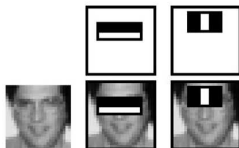
- Viola and Jones created a very fast face detector that can be scanned across a large image to find the faces.



- The base classifier/weak learner just compares the total intensity in two rectangular pieces of the image.
 - ▶ There is a neat trick for computing the total intensity in a rectangle in a few operations.

An impressive example of boosting

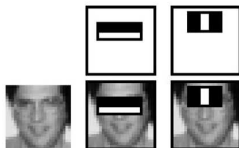
- Viola and Jones created a very fast face detector that can be scanned across a large image to find the faces.



- The base classifier/weak learner just compares the total intensity in two rectangular pieces of the image.
 - ▶ There is a neat trick for computing the total intensity in a rectangle in a few operations.
 - ▶ So its easy to evaluate a huge number of base classifiers and they are very fast at runtime.

An impressive example of boosting

- Viola and Jones created a very fast face detector that can be scanned across a large image to find the faces.



- The base classifier/weak learner just compares the total intensity in two rectangular pieces of the image.
 - ▶ There is a neat trick for computing the total intensity in a rectangle in a few operations.
 - ▶ So its easy to evaluate a huge number of base classifiers and they are very fast at runtime.
 - ▶ The algorithm adds classifiers greedily based on their quality on the weighted training cases.

AdaBoost in Face Detection


- Famous application of boosting: detecting faces in images
- Two twists on standard algorithm

AdaBoost in Face Detection

- Famous application of boosting: detecting faces in images
- Two twists on standard algorithm
 - ▶ Pre-define weak classifiers, so optimization=selection

AdaBoost in Face Detection

- Famous application of boosting: detecting faces in images
- Two twists on standard algorithm
 - ▶ Pre-define weak classifiers, so optimization=selection
 - ▶ Change loss function for weak learners: false positives less costly than misses



AdaBoost Face Detection Results

