# Introduction to Machine Learning (67577) Lecture 4 

## Shai Shalev-Shwartz

School of CS and Engineering, The Hebrew University of Jerusalem

Boosting

## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Weak Learnability

## Definition ( $\epsilon, \delta)$-Weak-Learnability)

A class $\mathcal{H}$ is $(\epsilon, \delta)$-weak-learnable if there exists a learning algorithm, $A$, and a training set size, $m \in \mathbb{N}$, such that for every distribution $\mathcal{D}$ over $\mathcal{X}$ and every $f \in \mathcal{H}$,

$$
\mathcal{D}^{m}\left(\left\{S: L_{\mathcal{D}, f}(A(S)) \leq \epsilon\right\}\right) \geq 1-\delta
$$

## Weak Learnability

## Definition (( $\epsilon, \delta)$-Weak-Learnability)

A class $\mathcal{H}$ is $(\epsilon, \delta)$-weak-learnable if there exists a learning algorithm, $A$, and a training set size, $m \in \mathbb{N}$, such that for every distribution $\mathcal{D}$ over $\mathcal{X}$ and every $f \in \mathcal{H}$,

$$
\mathcal{D}^{m}\left(\left\{S: L_{\mathcal{D}, f}(A(S)) \leq \epsilon\right\}\right) \geq 1-\delta .
$$

## Remarks:

- Almost identical to (strong) PAC learning, but we only need to succeed for specific $\epsilon, \delta$


## Weak Learnability

## Definition ( $\epsilon, \delta)$-Weak-Learnability)

A class $\mathcal{H}$ is $(\epsilon, \delta)$-weak-learnable if there exists a learning algorithm, $A$, and a training set size, $m \in \mathbb{N}$, such that for every distribution $\mathcal{D}$ over $\mathcal{X}$ and every $f \in \mathcal{H}$,

$$
\mathcal{D}^{m}\left(\left\{S: L_{\mathcal{D}, f}(A(S)) \leq \epsilon\right\}\right) \geq 1-\delta .
$$

## Remarks:

- Almost identical to (strong) PAC learning, but we only need to succeed for specific $\epsilon, \delta$
- Every class $\mathcal{H}$ is $(1 / 2,0)$-weak-learnable


## Weak Learnability

## Definition (( $\epsilon, \delta)$-Weak-Learnability)

A class $\mathcal{H}$ is $(\epsilon, \delta)$-weak-learnable if there exists a learning algorithm, $A$, and a training set size, $m \in \mathbb{N}$, such that for every distribution $\mathcal{D}$ over $\mathcal{X}$ and every $f \in \mathcal{H}$,

$$
\mathcal{D}^{m}\left(\left\{S: L_{\mathcal{D}, f}(A(S)) \leq \epsilon\right\}\right) \geq 1-\delta .
$$

## Remarks:

- Almost identical to (strong) PAC learning, but we only need to succeed for specific $\epsilon, \delta$
- Every class $\mathcal{H}$ is $(1 / 2,0)$-weak-learnable
- Intuitively, one can think of a weak learner as an algorithm that uses a simple 'rule of thumb' to output a hypothesis that performs just slightly better than a random guess


## Example of weak learner

- $\mathcal{X}=\mathbb{R}, \mathcal{H}$ is the class of 3 -piece classifiers, e.g.



## Example of weak learner

- $\mathcal{X}=\mathbb{R}, \mathcal{H}$ is the class of 3 -piece classifiers, e.g.

- Let $B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}$ be the class of Decision Stumps


## Example of weak learner

- $\mathcal{X}=\mathbb{R}, \mathcal{H}$ is the class of 3 -piece classifiers, e.g.

- Let $B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}$ be the class of Decision Stumps
- Claim: There is a constant $m$, such that $\mathrm{ERM}_{B}$ over $m$ examples is a (5/12, 1/2)-weak learner for $\mathcal{H}$


## Example of weak learner

- $\mathcal{X}=\mathbb{R}, \mathcal{H}$ is the class of 3 -piece classifiers, e.g.

- Let $B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}$ be the class of Decision Stumps
- Claim: There is a constant $m$, such that $\mathrm{ERM}_{B}$ over $m$ examples is a (5/12, 1/2)-weak learner for $\mathcal{H}$
- Proof:
- Observe that there's always a decision stump with $L_{\mathcal{D}, f}(h) \leq 1 / 3$
- Apply VC bound for the class of decision stumps


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class H


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class $\mathcal{H}$
- Can we use $A$ to construct a strong learner ?


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class H
- Can we use $A$ to construct a strong learner ?
- If $A$ is computationally efficient, can we boost it efficiently ?


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class H
- Can we use $A$ to construct a strong learner ?
- If $A$ is computationally efficient, can we boost it efficiently ?
- Two questions:


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class H
- Can we use $A$ to construct a strong learner ?
- If $A$ is computationally efficient, can we boost it efficiently ?
- Two questions:
- Boosting the confidence


## The problem of boosting

- Suppose we have an $\left(\epsilon_{0}, \delta_{0}\right)$-weak-learner algorithm, $A$, for some class H
- Can we use $A$ to construct a strong learner ?
- If $A$ is computationally efficient, can we boost it efficiently ?
- Two questions:
- Boosting the confidence
- Boosting the accuracy


## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Boosting the confidence

- Suppose $A$ is an $\left(\epsilon_{0}, \delta_{0}\right)$-weak learner for $\mathcal{H}$ that requires $m_{0}$ examples


## Boosting the confidence

- Suppose $A$ is an $\left(\epsilon_{0}, \delta_{0}\right)$-weak learner for $\mathcal{H}$ that requires $m_{0}$ examples
- For any $\delta, \epsilon \in(0,1)$ we show how to learn $\mathcal{H}$ to accuracy $\epsilon_{0}+\epsilon$ with confidence $\delta$


## Boosting the confidence

- Suppose $A$ is an $\left(\epsilon_{0}, \delta_{0}\right)$-weak learner for $\mathcal{H}$ that requires $m_{0}$ examples
- For any $\delta, \epsilon \in(0,1)$ we show how to learn $\mathcal{H}$ to accuracy $\epsilon_{0}+\epsilon$ with confidence $\delta$
- Step 1: Apply $A$ on $k=\left\lceil\frac{\log (2 / \delta)}{\log \left(1 / \delta_{0}\right)}\right\rceil$ i.i.d. samples, each of which of $m_{0}$ examples, to obtain $h_{1}, \ldots, h_{k}$


## Boosting the confidence

- Suppose $A$ is an $\left(\epsilon_{0}, \delta_{0}\right)$-weak learner for $\mathcal{H}$ that requires $m_{0}$ examples
- For any $\delta, \epsilon \in(0,1)$ we show how to learn $\mathcal{H}$ to accuracy $\epsilon_{0}+\epsilon$ with confidence $\delta$
- Step 1: Apply $A$ on $k=\left\lceil\frac{\log (2 / \delta)}{\log \left(1 / \delta_{0}\right)}\right\rceil$ i.i.d. samples, each of which of $m_{0}$ examples, to obtain $h_{1}, \ldots, h_{k}$
- Step 2: Take additional validation sample of size $|V| \geq \frac{2 \log (4 k / \delta)}{\epsilon^{2}}$ and output $\hat{h} \in \operatorname{argmin}_{h_{i}} L_{V}\left(h_{i}\right)$


## Boosting the confidence

- Suppose $A$ is an $\left(\epsilon_{0}, \delta_{0}\right)$-weak learner for $\mathcal{H}$ that requires $m_{0}$ examples
- For any $\delta, \epsilon \in(0,1)$ we show how to learn $\mathcal{H}$ to accuracy $\epsilon_{0}+\epsilon$ with confidence $\delta$
- Step 1: Apply $A$ on $k=\left\lceil\frac{\log (2 / \delta)}{\log \left(1 / \delta_{0}\right)}\right\rceil$ i.i.d. samples, each of which of $m_{0}$ examples, to obtain $h_{1}, \ldots, h_{k}$
- Step 2: Take additional validation sample of size $|V| \geq \frac{2 \log (4 k / \delta)}{\epsilon^{2}}$ and output $\hat{h} \in \operatorname{argmin}_{h_{i}} L_{V}\left(h_{i}\right)$
- Claim: W.p. at least $1-\delta$, we have $L_{\mathcal{D}}(\hat{h}) \leq \epsilon_{0}+\epsilon$


## Proof

- First, by the validation procedure guarantees

$$
\mathbb{P}\left[L_{\mathcal{D}}(\hat{h})>\min _{i} L_{\mathcal{D}}\left(h_{i}\right)+\epsilon\right] \leq \delta / 2 .
$$

## Proof

- First, by the validation procedure guarantees

$$
\mathbb{P}\left[L_{\mathcal{D}}(\hat{h})>\min _{i} L_{\mathcal{D}}\left(h_{i}\right)+\epsilon\right] \leq \delta / 2 .
$$

- Second,

$$
\begin{aligned}
\mathbb{P}\left[\min _{i} L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] & =\mathbb{P}\left[\forall_{i} L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] \\
& =\prod_{i=1}^{k} \mathbb{P}\left[L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] \\
& \leq \delta_{0}^{k} \leq \delta / 2
\end{aligned}
$$

## Proof

- First, by the validation procedure guarantees

$$
\mathbb{P}\left[L_{\mathcal{D}}(\hat{h})>\min _{i} L_{\mathcal{D}}\left(h_{i}\right)+\epsilon\right] \leq \delta / 2 .
$$

- Second,

$$
\begin{aligned}
\mathbb{P}\left[\min _{i} L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] & =\mathbb{P}\left[\forall_{i} L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] \\
& =\prod_{i=1}^{k} \mathbb{P}\left[L_{\mathcal{D}}\left(h_{i}\right)>\epsilon_{0}\right] \\
& \leq \delta_{0}^{k} \leq \delta / 2 .
\end{aligned}
$$

- Apply the union bound to conclude the proof.


## Boosting a learner that succeeds on expectation

- Suppose that $A$ is a learner that guarantees:

$$
\underset{S \sim \mathcal{D}^{m}}{\mathbb{E}}\left[L_{\mathcal{D}}(A(S))\right] \leq \min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)+\epsilon
$$

## Boosting a learner that succeeds on expectation

- Suppose that $A$ is a learner that guarantees:

$$
\underset{S \sim \mathcal{D}^{m}}{\mathbb{E}}\left[L_{\mathcal{D}}(A(S))\right] \leq \min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)+\epsilon
$$

- Denote $\theta=L_{\mathcal{D}}(A(S))-\min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)$, so we obtain

$$
\underset{S \sim D^{m}}{\mathbb{E}}[\theta] \leq \epsilon
$$

## Boosting a learner that succeeds on expectation

- Suppose that $A$ is a learner that guarantees:

$$
\underset{S \sim D^{m}}{\mathbb{E}}\left[L_{\mathcal{D}}(A(S))\right] \leq \min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)+\epsilon
$$

- Denote $\theta=L_{\mathcal{D}}(A(S))-\min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)$, so we obtain

$$
\underset{S \sim \mathcal{D}^{m}}{\mathbb{E}}[\theta] \leq \epsilon
$$

- Since $\theta$ is a non-negative random variable, we can apply Markov's inequality to obtain

$$
\mathbb{P}[\theta \geq 2 \epsilon] \leq \frac{\mathbb{E}[\theta]}{2 \epsilon} \leq \frac{1}{2}
$$

## Boosting a learner that succeeds on expectation

- Suppose that $A$ is a learner that guarantees:

$$
\underset{S \sim D^{m}}{\mathbb{E}}\left[L_{\mathcal{D}}(A(S))\right] \leq \min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)+\epsilon
$$

- Denote $\theta=L_{\mathcal{D}}(A(S))-\min _{h \in \mathcal{H}} L_{\mathcal{D}}(h)$, so we obtain

$$
\underset{S \sim \mathcal{D}^{m}}{\mathbb{E}}[\theta] \leq \epsilon
$$

- Since $\theta$ is a non-negative random variable, we can apply Markov's inequality to obtain

$$
\mathbb{P}[\theta \geq 2 \epsilon] \leq \frac{\mathbb{E}[\theta]}{2 \epsilon} \leq \frac{1}{2}
$$

- Corollary: $A$ is $(2 \epsilon, 1 / 2)$-weak learner.


## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Boosting the accuracy

## Problem raised in 1988 by Kearns and Valiant



Solved in 1990 by Robert Schapire, then a graduate student at MIT


> In 1995, Schapire \& Freund proposed the AdaBoost algorithm

## AdaBoost ('Adaptive Boosting')

- Input: $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, where for each $i, y_{i}=f\left(\mathbf{x}_{i}\right)$


## AdaBoost ('Adaptive Boosting')

- Input: $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, where for each $i, y_{i}=f\left(\mathbf{x}_{i}\right)$
- Output: hypothesis $h$ with small error on $S$


## AdaBoost ('Adaptive Boosting')

- Input: $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, where for each $i, y_{i}=f\left(\mathbf{x}_{i}\right)$
- Output: hypothesis $h$ with small error on $S$
- We'll later analyze $L_{(\mathcal{D}, f)}(h)$ as well


## AdaBoost ('Adaptive Boosting')

- Input: $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, where for each $i, y_{i}=f\left(\mathbf{x}_{i}\right)$
- Output: hypothesis $h$ with small error on $S$
- We'll later analyze $L_{(\mathcal{D}, f)}(h)$ as well
- AdaBoost calls the weak learner on distributions over $S$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :
- invoke weak learner $h_{t}=\mathrm{WL}\left(\mathbf{D}^{(t)}, S\right)$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :
- invoke weak learner $h_{t}=\mathrm{WL}\left(\mathbf{D}^{(t)}, S\right)$
- compute $\epsilon_{t}=L_{\mathbf{D}^{(t)}}\left(h_{t}\right)=\sum_{i=1}^{m} D_{i}^{(t)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :
- invoke weak learner $h_{t}=\mathrm{WL}\left(\mathbf{D}^{(t)}, S\right)$
- compute $\epsilon_{t}=L_{\mathbf{D}^{(t)}}\left(h_{t}\right)=\sum_{i=1}^{m} D_{i}^{(t)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}$
- let $w_{t}=\frac{1}{2} \log \left(\frac{1}{\epsilon_{t}}-1\right)$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :
- invoke weak learner $h_{t}=\mathrm{WL}\left(\mathbf{D}^{(t)}, S\right)$
- compute $\epsilon_{t}=L_{\mathbf{D}^{(t)}}\left(h_{t}\right)=\sum_{i=1}^{m} D_{i}^{(t)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}$
- let $w_{t}=\frac{1}{2} \log \left(\frac{1}{\epsilon_{t}}-1\right)$
- update $D_{i}^{(t+1)}=\frac{D_{i}^{(t)} \exp \left(-w_{t} y_{i} h_{t}\left(\mathbf{x}_{i}\right)\right)}{\sum_{j=1}^{m} D_{j}^{(t)} \exp \left(-w_{t} y_{j} h_{t}\left(\mathbf{x}_{j}\right)\right)}$ for all $i=1, \ldots, m$


## The AdaBoost Algorithm

- input: training set $S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)$, weak learner WL, number of rounds $T$
- initialize $\mathbf{D}^{(1)}=\left(\frac{1}{m}, \ldots, \frac{1}{m}\right)$
- for $t=1, \ldots, T$ :
- invoke weak learner $h_{t}=\mathrm{WL}\left(\mathbf{D}^{(t)}, S\right)$
- compute $\epsilon_{t}=L_{\mathbf{D}^{(t)}}\left(h_{t}\right)=\sum_{i=1}^{m} D_{i}^{(t)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}$
- let $w_{t}=\frac{1}{2} \log \left(\frac{1}{\epsilon_{t}}-1\right)$
- update $D_{i}^{(t+1)}=\frac{D_{i}^{(t)} \exp \left(-w_{t} y_{i} h_{t}\left(\mathbf{x}_{i}\right)\right)}{\sum_{j=1}^{m} D_{j}^{(t)} \exp \left(-w_{t} y_{j} h_{t}\left(\mathbf{x}_{j}\right)\right)}$ for all $i=1, \ldots, m$
- output the hypothesis $h_{s}(\mathbf{x})=\operatorname{sign}\left(\sum_{t=1}^{T} w_{t} h_{t}(\mathbf{x})\right)$.


## Intuition: AdaBoost forces WL to focus on problematic examples

- Claim: The error of $h_{t}$ w.r.t. $\mathbf{D}^{(t+1)}$ is exactly $1 / 2$
- Proof:

$$
\sum_{i=1}^{m} D_{i}^{(t+1)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}=\frac{\sum_{i=1}^{m} D_{i}^{(t)} e^{-w_{t} y_{i} h_{t}\left(\mathbf{x}_{i}\right)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}}{\sum_{j=1}^{m} D_{j}^{(t)} e^{-w_{t} y_{j} h_{t}\left(\mathbf{x}_{j}\right)}}
$$

## Intuition: AdaBoost forces WL to focus on problematic examples

- Claim: The error of $h_{t}$ w.r.t. $\mathbf{D}^{(t+1)}$ is exactly $1 / 2$
- Proof:

$$
\begin{aligned}
& \sum_{i=1}^{m} D_{i}^{(t+1)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}=\frac{\sum_{i=1}^{m} D_{i}^{(t)} e^{-w_{t} y_{i} h_{t}\left(\mathbf{x}_{i}\right)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}}{\sum_{j=1}^{m} D_{j}^{(t)} e^{-w_{t} y_{j} h_{t}\left(\mathbf{x}_{j}\right)}} \\
& =\frac{e^{w_{t}} \epsilon_{t}}{e^{w_{t}} \epsilon_{t}+e^{-w_{t}}\left(1-\epsilon_{t}\right)}=\frac{\epsilon_{t}}{\epsilon_{t}+e^{-2 w_{t}}\left(1-\epsilon_{t}\right)}
\end{aligned}
$$

## Intuition: AdaBoost forces WL to focus on problematic examples

- Claim: The error of $h_{t}$ w.r.t. $\mathbf{D}^{(t+1)}$ is exactly $1 / 2$
- Proof:

$$
\begin{aligned}
& \sum_{i=1}^{m} D_{i}^{(t+1)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}=\frac{\sum_{i=1}^{m} D_{i}^{(t)} e^{-w_{t} y_{i} h_{t}\left(\mathbf{x}_{i}\right)} \mathbb{1}_{\left[y_{i} \neq h_{t}\left(\mathbf{x}_{i}\right)\right]}}{\sum_{j=1}^{m} D_{j}^{(t)} e^{-w_{t} y_{j} h_{t}\left(\mathbf{x}_{j}\right)}} \\
& =\frac{e^{w_{t}} \epsilon_{t}}{e^{w_{t}} \epsilon_{t}+e^{-w_{t}}\left(1-\epsilon_{t}\right)}=\frac{\epsilon_{t}}{\epsilon_{t}+e^{-2 w_{t}}\left(1-\epsilon_{t}\right)} \\
& =\frac{\epsilon_{t}}{\epsilon_{t}+\frac{\epsilon_{t}}{1-\epsilon_{t}}\left(1-\epsilon_{t}\right)}=\frac{1}{2} .
\end{aligned}
$$

## Theorem

If $W \mathrm{~L}$ is $(1 / 2-\gamma, \delta)$ weak learner then, with probability at least $1-\delta T$,

$$
L_{S}\left(h_{s}\right) \leq \exp \left(-2 \gamma^{2} T\right)
$$

## Theorem

If $W \mathrm{~L}$ is $(1 / 2-\gamma, \delta)$ weak learner then, with probability at least $1-\delta T$,

$$
L_{S}\left(h_{s}\right) \leq \exp \left(-2 \gamma^{2} T\right)
$$

Remarks:

- For any $\epsilon>0$ and $\gamma \in(0,1 / 2)$, if $T \geq \frac{\log (1 / \epsilon)}{2 \gamma^{2}}$, then AdaBoost will output a hypothesis $h_{s}$ with $L_{S}\left(h_{s}\right) \leq \epsilon$.


## Theorem

If $W \mathrm{~L}$ is $(1 / 2-\gamma, \delta)$ weak learner then, with probability at least $1-\delta T$,

$$
L_{S}\left(h_{s}\right) \leq \exp \left(-2 \gamma^{2} T\right)
$$

## Remarks:

- For any $\epsilon>0$ and $\gamma \in(0,1 / 2)$, if $T \geq \frac{\log (1 / \epsilon)}{2 \gamma^{2}}$, then AdaBoost will output a hypothesis $h_{s}$ with $L_{S}\left(h_{s}\right) \leq \epsilon$.
- Setting $\epsilon=1 /(2 m)$ the hypothesis $h_{s}$ must have a zero training error


## Theorem

If $W \mathrm{~L}$ is $(1 / 2-\gamma, \delta)$ weak learner then, with probability at least $1-\delta T$,

$$
L_{S}\left(h_{s}\right) \leq \exp \left(-2 \gamma^{2} T\right)
$$

## Remarks:

- For any $\epsilon>0$ and $\gamma \in(0,1 / 2)$, if $T \geq \frac{\log (1 / \epsilon)}{2 \gamma^{2}}$, then AdaBoost will output a hypothesis $h_{s}$ with $L_{S}\left(h_{s}\right) \leq \epsilon$.
- Setting $\epsilon=1 /(2 m)$ the hypothesis $h_{s}$ must have a zero training error
- Since the weak learner is invoked on a distribution over $S$, in many cases $\delta$ can be 0 . In any case, by "boosting the confidence", we can assume w.l.o.g. that $\delta$ is very small.


## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost

4 AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## AdaBoost as a Learner for Halfspaces++

- Let $B$ be the set of all hypotheses the WL may return


## AdaBoost as a Learner for Halfspaces++

- Let $B$ be the set of all hypotheses the WL may return
- Observe that AdaBoost outputs a hypothesis from the class

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}\left(\sum_{t=1}^{T} w_{t} h_{t}(x)\right): \mathbf{w} \in \mathbb{R}^{T}, \forall t, \quad h_{t} \in B\right\}
$$

## AdaBoost as a Learner for Halfspaces++

- Let $B$ be the set of all hypotheses the WL may return
- Observe that AdaBoost outputs a hypothesis from the class

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}\left(\sum_{t=1}^{T} w_{t} h_{t}(x)\right): \mathbf{w} \in \mathbb{R}^{T}, \forall t, \quad h_{t} \in B\right\}
$$

- Since WL is invoked only on distributions over $S$ we can assume w.l.o.g. that $B=\left\{g_{1}, \ldots, g_{d}\right\}$ for some $d \leq 2^{m}$.


## AdaBoost as a Learner for Halfspaces++

- Let $B$ be the set of all hypotheses the WL may return
- Observe that AdaBoost outputs a hypothesis from the class

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}\left(\sum_{t=1}^{T} w_{t} h_{t}(x)\right): \mathbf{w} \in \mathbb{R}^{T}, \forall t, \quad h_{t} \in B\right\}
$$

- Since WL is invoked only on distributions over $S$ we can assume w.l.o.g. that $B=\left\{g_{1}, \ldots, g_{d}\right\}$ for some $d \leq 2^{m}$.
- Denote $\psi(x)=\left(g_{1}(x), \ldots, g_{d}(x)\right)$. Therefore:

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}(\langle w, \psi(x)\rangle): \mathbf{w} \in \mathbb{R}^{d},\|\mathbf{w}\|_{0} \leq T\right\}
$$

where $\|\mathbf{w}\|_{0}=\left|\left\{i: w_{i} \neq 0\right\}\right|$.

## AdaBoost as a Learner for Halfspaces++

- Let $B$ be the set of all hypotheses the WL may return
- Observe that AdaBoost outputs a hypothesis from the class

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}\left(\sum_{t=1}^{T} w_{t} h_{t}(x)\right): \mathbf{w} \in \mathbb{R}^{T}, \forall t, \quad h_{t} \in B\right\}
$$

- Since WL is invoked only on distributions over $S$ we can assume w.l.o.g. that $B=\left\{g_{1}, \ldots, g_{d}\right\}$ for some $d \leq 2^{m}$.
- Denote $\psi(x)=\left(g_{1}(x), \ldots, g_{d}(x)\right)$. Therefore:

$$
L(B, T)=\left\{x \mapsto \operatorname{sign}(\langle w, \psi(x)\rangle): \mathbf{w} \in \mathbb{R}^{d},\|\mathbf{w}\|_{0} \leq T\right\}
$$

where $\|\mathbf{w}\|_{0}=\left|\left\{i: w_{i} \neq 0\right\}\right|$.

- That is, AdaBoost learns a composition of the class of halfspaces with sparse coefficients over the mapping $x \mapsto \psi(x)$


## Expressiveness of $L(B, T)$

- Suppose $\mathcal{X}=\mathbb{R}$ and $B$ is Decision Stumps,

$$
B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}
$$

## Expressiveness of $L(B, T)$

- Suppose $\mathcal{X}=\mathbb{R}$ and $B$ is Decision Stumps,

$$
B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}
$$

- Let $\mathcal{G}_{T}$ be the class of piece-wise constant functions with $T$ pieces


## Expressiveness of $L(B, T)$

- Suppose $\mathcal{X}=\mathbb{R}$ and $B$ is Decision Stumps,

$$
B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}
$$

- Let $\mathcal{G}_{T}$ be the class of piece-wise constant functions with $T$ pieces
- Claim: $\mathcal{G}_{T} \subseteq L(B, T)$


## Expressiveness of $L(B, T)$

- Suppose $\mathcal{X}=\mathbb{R}$ and $B$ is Decision Stumps,

$$
B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}
$$

- Let $\mathcal{G}_{T}$ be the class of piece-wise constant functions with $T$ pieces
- Claim: $\mathcal{G}_{T} \subseteq L(B, T)$


## Expressiveness of $L(B, T)$

- Suppose $\mathcal{X}=\mathbb{R}$ and $B$ is Decision Stumps,

$$
B=\{x \mapsto \operatorname{sign}(x-\theta) \cdot b: \quad \theta \in \mathbb{R}, b \in\{ \pm 1\}\}
$$

- Let $\mathcal{G}_{T}$ be the class of piece-wise constant functions with $T$ pieces
- Claim: $\mathcal{G}_{T} \subseteq L(B, T)$


Composing halfspaces on top of simple classes can be very expressive !

## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Bias-complexity

Recall:


- We have argued that the expressiveness of $L(B, T)$ grows with $T$


## Bias-complexity

Recall:


- We have argued that the expressiveness of $L(B, T)$ grows with $T$
- In other words, the approximation error decreases with $T$


## Bias-complexity

## Recall:



- We have argued that the expressiveness of $L(B, T)$ grows with $T$
- In other words, the approximation error decreases with $T$
- We'll show that the estimation error increases with $T$


## Bias-complexity

## Recall:



- We have argued that the expressiveness of $L(B, T)$ grows with $T$
- In other words, the approximation error decreases with $T$
- We'll show that the estimation error increases with $T$
- Therefore, the parameter $T$ of AdaBoost enables us to control the bias-complexity tradeoff


## The Estimation Error of $L(B, T)$

- Claim:

$$
\operatorname{VCdim}(L(B, T)) \leq \tilde{O}(T \cdot \operatorname{VCdim}(B))
$$

## The Estimation Error of $L(B, T)$

- Claim:

$$
\operatorname{VCdim}(L(B, T)) \leq \tilde{O}(T \cdot \operatorname{VCdim}(B))
$$

- Corollary: if $m \geq \tilde{\Omega}\left(\frac{\log (1 / \delta)}{\gamma^{2} \epsilon}\right)$ and $T=\log (m) /\left(2 \gamma^{2}\right)$, then w.p. of at least $1-\delta$,

$$
L_{(\mathcal{D}, f)}\left(h_{s}\right) \leq \epsilon .
$$

## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Weak Learnability and Separability with Margin

- We have essentially shown: if $\mathcal{H}$ is weak learnable, then $\mathcal{H} \subseteq L(B, \infty)$


## Weak Learnability and Separability with Margin

- We have essentially shown: if $\mathcal{H}$ is weak learnable, then $\mathcal{H} \subseteq L(B, \infty)$
- What about the other direction ?


## Weak Learnability and Separability with Margin

- We have essentially shown: if $\mathcal{H}$ is weak learnable, then $\mathcal{H} \subseteq L(B, \infty)$
- What about the other direction ?
- Using von Neumanns minimax theorem, it can be shown that if $L(B, \infty)$ separates a training set with $\ell_{1}$ margin $\gamma$ then $\mathrm{ERM}_{B}$ is a $\gamma$ weak learner for $\mathcal{H}$


## Weak Learnability and Separability with Margin

- We have essentially shown: if $\mathcal{H}$ is weak learnable, then $\mathcal{H} \subseteq L(B, \infty)$
- What about the other direction ?
- Using von Neumanns minimax theorem, it can be shown that if $L(B, \infty)$ separates a training set with $\ell_{1}$ margin $\gamma$ then $\mathrm{ERM}_{B}$ is a $\gamma$ weak learner for $\mathcal{H}$
- This is beyond the scope of the course


## Outline

(1) Weak learnability
(2) Boosting the confidence
(3) Boosting the accuracy using AdaBoost
4) AdaBoost as a learner for Halfspaces++
(5) AdaBoost and the Bias-Complexity Tradeoff
(6) Weak Learnability and Separability with Margin
(7) AdaBoost for Face Detection

## Face Detection

- Classify rectangles in an image as face or non-face


## Weak Learner for Face Detection

Rules of thumb:

- "eye region is often darker than the cheeks"
- "bridge of the noise is brighter than the eyes"


## Weak Learner for Face Detection

Rules of thumb:

- "eye region is often darker than the cheeks"
- "bridge of the noise is brighter than the eyes"

Goal:

- We want to combine few rules of thumb to obtain a face detector
- "Sparsity" reflects both small estimation error but also speed !


## Weak Learner for Face Detection

Each hypothesis in the base class is of the form $h(x)=f(g(x))$, where $f$ is a decision stump and $g: \mathbb{R}^{24,24} \rightarrow \mathbb{R}$ is parameterized by:

- An axis-aligned rectangle $R$. Since each image is of size $24 \times 24$, there are at most $24^{4}$ axis-aligned rectangles.
- A type, $t \in\{A, B, C, D\}$. Each type corresponds to a mask:



D

## AdaBoost for Face Detection

The first and second features selected by AdaBoost, as implemented by Viola and Jones.


## Summary

- Boosting the confidence using validation
- Boosting the accuracy using AdaBoost
- The power of composing halfspaces over simple classes
- The bias-complexity tradeoff
- AdaBoost works in many practical problems !

