



Slides adapted from Mohri

Classification

Jordan Boyd-Graber
University of Maryland
WEIGHTED MAJORITY

Beyond Binary Classification

- Before we've talked about combining weak predictor (boosting)

Beyond Binary Classification

- Before we've talked about combining weak predictor (boosting)
 - What if you have strong predictors?

Beyond Binary Classification

- Before we've talked about combining weak predictor (boosting)
 - What if you have strong predictors?
- How do you make inherently binary algorithms multiclass?
- How do you answer questions like ranking?

General Online Setting

- For $t = 1$ to T :
 - Get instance $x_t \in X$
 - Predict $\hat{y}_t \in Y$
 - Get true label $y_t \in Y$
 - Incur loss $L(\hat{y}_t, y_t)$
- Classification: $Y = \{0, 1\}$, $L(y, y') = |y' - y|$
- Regression: $Y \subset \mathbb{R}$, $L(y, y') = (y' - y)^2$

General Online Setting

- For $t = 1$ to T :
 - Get instance $x_t \in X$
 - Predict $\hat{y}_t \in Y$
 - Get true label $y_t \in Y$
 - Incur loss $L(\hat{y}_t, y_t)$
- Classification: $Y = \{0, 1\}$, $L(y, y') = |y' - y|$
- Regression: $Y \subset \mathbb{R}$, $L(y, y') = (y' - y)^2$
- **Objective:** Minimize total loss $\sum_t L(\hat{y}_t, y_t)$

Prediction with Expert Advice

- For $t = 1$ to T :
 - Get instance $x_t \in X$ and advice $a_t, i \in Y, i \in [1, N]$
 - Predict $\hat{y}_t \in Y$
 - Get true label $y_t \in Y$
 - Incur loss $L(\hat{y}_t, y_t)$

Prediction with Expert Advice

- For $t = 1$ to T :
 - Get instance $x_t \in X$ and **advice** $a_t, i \in Y, i \in [1, N]$
 - Predict $\hat{y}_t \in Y$
 - Get true label $y_t \in Y$
 - Incur loss $L(\hat{y}_t, y_t)$
- **Objective:** Minimize regret, i.e., difference of total loss vs. best expert

$$\text{Regret}(T) = \sum_t L(\hat{y}_t, y_t) - \min_i \sum_t L(a_{t,i}, y_t) \quad (1)$$

Mistake Bound Model

- Define the maximum number of mistakes a learning algorithm L makes to learn a concept c over any set of examples (until it's perfect).

$$M_L(c) = \max_{x_1, \dots, x_T} |\text{mistakes}(L, c)| \quad (2)$$

- For any concept class C , this is the max over concepts c .

$$M_L(C) = \max_{c \in C} M_L(c) \quad (3)$$

Mistake Bound Model

- Define the maximum number of mistakes a learning algorithm L makes to learn a concept c over any set of examples (until it's perfect).

$$M_L(c) = \max_{x_1, \dots, x_T} |\text{mistakes}(L, c)| \quad (2)$$

- For any concept class C , this is the max over concepts c .

$$M_L(C) = \max_{c \in C} M_L(c) \quad (3)$$

- In the expert advice case, assumes some expert matches the concept (realizable)

Halving Algorithm

```

 $H_1 \leftarrow H;$ 
for  $t \leftarrow 1 \dots T$  do
  Receive  $x_t$ ;
   $\hat{y}_t \leftarrow \text{Majority}(H_t, \vec{a}_t, x_t);$ 
  Receive  $y_t$ ;
  if  $\hat{y}_t \neq y_t$  then
     $H_{t+1} \leftarrow \{a \in H_t : a(x_t) = y_t\};$ 
return  $H_{T+1}$ 

```

Algorithm 1: The Halving Algorithm (Mitchell, 1997)

Halving Algorithm Bound (Littlestone, 1998)

- For a finite hypothesis set

$$M_{\text{Halving}(H)} \leq \lg |H| \quad (4)$$

- After each mistake, the hypothesis set is reduced by at least by half

Halving Algorithm Bound (Littlestone, 1998)

- For a finite hypothesis set

$$M_{\text{Halving}(H)} \leq \lg |H| \quad (4)$$

- After each mistake, the hypothesis set is reduced by at least by half
- Consider the optimal mistake bound $\text{opt}(H)$. Then

$$\text{VC}(H) \leq \text{opt}(H) \leq M_{\text{Halving}(H)} \leq \lg |H| \quad (5)$$

- For a fully shattered set, form a binary tree of mistakes with height $\text{VC}(H)$

Halving Algorithm Bound (Littlestone, 1998)

- For a finite hypothesis set

$$M_{\text{Halving}(H)} \leq \lg |H| \quad (4)$$

- After each mistake, the hypothesis set is reduced by at least by half
- Consider the optimal mistake bound $\text{opt}(H)$. Then

$$\text{VC}(H) \leq \text{opt}(H) \leq M_{\text{Halving}(H)} \leq \lg |H| \quad (5)$$

- For a fully shattered set, form a binary tree of mistakes with height $\text{VC}(H)$
- What about non-realizable case?

Weighted Majority (Littlestone and Warmuth, 1998)

```

for  $i \leftarrow 1 \dots N$  do
  |  $w_{1,i} \leftarrow 1$ ;
for  $t \leftarrow 1 \dots T$  do
  | Receive  $x_t$ ;
  |  $\hat{y}_t \leftarrow \mathbb{1} \left[ \sum_{a_{t,i}=1} w_t \geq \sum_{a_{t,i}=0} w_t \right]$ ;
  | Receive  $y_t$ ;
  | if  $\hat{y}_t \neq y_t$  then
  |   | for  $i \leftarrow 1 \dots N$  do
  |     | if  $a_{t,i} \neq y_t$  then
  |       |  $w_{t+1,i} \leftarrow \beta w_{t,i}$ ;
  |     | else
  |       |  $w_{t+1,i} \leftarrow w_{t,i}$ 
return  $w_{T+1}$ 
  
```

- Weights for every expert
- Classifications in favor of side with higher total weight ($y \in \{0, 1\}$)
- Experts that are wrong get their weights decreased ($\beta \in [0, 1]$)
- If you're right, you stay unchanged

Weighted Majority (Littlestone and Warmuth, 1998)

```

for  $i \leftarrow 1 \dots N$  do
  |  $w_{1,i} \leftarrow 1$ ;
for  $t \leftarrow 1 \dots T$  do
  | Receive  $x_t$ ;
  |  $\hat{y}_t \leftarrow \mathbb{1} \left[ \sum_{a_{t,i}=1} w_t \geq \sum_{a_{t,i}=0} w_t \right]$ ;
  | Receive  $y_t$ ;
  | if  $\hat{y}_t \neq y_t$  then
  |   | for  $i \leftarrow 1 \dots N$  do
  |     | if  $a_{t,i} \neq y_t$  then
  |       |  $w_{t+1,i} \leftarrow \beta w_{t,i}$ ;
  |     | else
  |       |  $w_{t+1,i} \leftarrow w_{t,i}$ 
  |   |
  |   return  $w_{T+1}$ 

```

- Weights for every expert
- Classifications in favor of side with higher total weight ($y \in \{0, 1\}$)
- Experts that are wrong get their weights decreased ($\beta \in [0, 1]$)
- If you're right, you stay unchanged

Weighted Majority (Littlestone and Warmuth, 1998)

```

for  $i \leftarrow 1 \dots N$  do
  |  $w_{1,i} \leftarrow 1$ ;
for  $t \leftarrow 1 \dots T$  do
  | Receive  $x_t$ ;
  |  $\hat{y}_t \leftarrow \mathbb{1} \left[ \sum_{a_{t,i}=1} w_t \geq \sum_{a_{t,i}=0} w_t \right]$ ;
  | Receive  $y_t$ ;
  | if  $\hat{y}_t \neq y_t$  then
  |   | for  $i \leftarrow 1 \dots N$  do
  |     | if  $a_{t,i} \neq y_t$  then
  |       |  $w_{t+1,i} \leftarrow \beta w_{t,i}$ ;
  |     | else
  |       |  $w_{t+1,i} \leftarrow w_{t,i}$ 
return  $w_{T+1}$ 
  
```

- Weights for every expert
- Classifications in favor of side with higher total weight ($y \in \{0, 1\}$)
- Experts that are wrong get their weights decreased ($\beta \in [0, 1]$)
- If you're right, you stay unchanged

Weighted Majority

- Let m_t be the number of mistakes made by WM until time t
- Let m_t^* be the best expert's mistakes until time t
- N is the number of experts

$$m_t \leq \frac{\log N + m_t^* \log \frac{1}{\beta}}{\log \frac{2}{1+\beta}} \quad (6)$$

- Thus, mistake bound is $O(\log N)$ plus the best expert
- Halving algorithm $\beta = 0$

Proof: Potential Function

- Potential function is the sum of all weights

$$\Phi_t \equiv \sum_i w_{t,i} \quad (7)$$

- We'll create sandwich of upper and lower bounds

Proof: Potential Function

- Potential function is the sum of all weights

$$\Phi_t \equiv \sum_i w_{t,i} \quad (7)$$

- We'll create sandwich of upper and lower bounds
- For any expert i , we have lower bound

$$\Phi_t \geq w_{t,i} = \beta^{m_{t,i}} \quad (8)$$

Proof: Potential Function

- Potential function is the sum of all weights

$$\Phi_t \equiv \sum_i w_{t,i} \quad (7)$$

- We'll create sandwich of upper and lower bounds
- For any expert i , we have lower bound

$$\Phi_t \geq w_{t,i} = \beta^{m_{t,i}} \quad (8)$$

Weights are nonnegative, so $\sum_i w_{t,i} \geq w_{t,i}$

Proof: Potential Function

- Potential function is the sum of all weights

$$\Phi_t \equiv \sum_i w_{t,i} \quad (7)$$

- We'll create sandwich of upper and lower bounds
- For any expert i , we have lower bound

$$\Phi_t \geq w_{t,i} = \beta^{m_{t,i}} \quad (8)$$

Each error multiplicatively reduces weight by β

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta\Phi_t}{2} \tag{9}$$

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta \Phi_t}{2} \tag{9}$$

Half (at most) of the experts by weight were right

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta \Phi_t}{2} \tag{9}$$

Half (at least) of the experts by weight were wrong

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta\Phi_t}{2} = \left[\frac{1+\beta}{2} \right] \Phi_t \quad (9)$$

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta\Phi_t}{2} = \left[\frac{1+\beta}{2} \right] \Phi_t \quad (9)$$

- Initially potential function sums all weights, which start at 1

$$\Phi_1 = N \quad (10)$$

Proof: Potential Function (Upper Bound)

- If an algorithm makes an error at round t

$$\Phi_{t+1} \leq \frac{\Phi_t}{2} + \frac{\beta\Phi_t}{2} = \left[\frac{1+\beta}{2} \right] \Phi_t \quad (9)$$

- Initially potential function sums all weights, which start at 1

$$\Phi_1 = N \quad (10)$$

- After m_T mistakes after T rounds

$$\Phi_T \leq \left[\frac{1+\beta}{2} \right]^{m_T} N \quad (11)$$

Weighted Majority Proof

- Put the two inequalities together, using the best expert

$$\beta^{m_T^*} \leq \Phi_T \leq \left[\frac{1 + \beta}{2} \right]^{m_T} N \quad (12)$$

Weighted Majority Proof

- Put the two inequalities together, using the best expert

$$\beta^{m_T^*} \leq \Phi_T \leq \left[\frac{1+\beta}{2} \right]^{m_T} N \quad (12)$$

- Take the log of both sides

$$m_T^* \log \beta \leq \log N + m_T \log \left[\frac{1+\beta}{2} \right] \quad (13)$$

Weighted Majority Proof

- Put the two inequalities together, using the best expert

$$\beta^{m_T^*} \leq \Phi_T \leq \left[\frac{1+\beta}{2} \right]^{m_T} N \quad (12)$$

- Take the log of both sides

$$m_T^* \log \beta \leq \log N + m_T \log \left[\frac{1+\beta}{2} \right] \quad (13)$$

- Solve for m_T

$$m_T \leq \frac{\log N + m_T^* \log \frac{1}{\beta}}{\log \left[\frac{2}{1+\beta} \right]} \quad (14)$$

Weighted Majority Recap

- Simple algorithm
- No harsh assumptions (non-realizable)
- Depends on best learner
- Downside: Takes a long time to do well in worst case (but okay in practice)
- Solution: Randomization