EECS 598-005: Theoretical Foundations of Machine Learning **Fall 2015**

Lecture 17: Online learning with EWA and introduction to game theory

Lecturer: Jacob Abernethy *Scribes: Henry Oskar Singer, Editors: Weiqing Yu and Andrew Melfi*

17.1 Exponential Weights Algorithm

Given a loss function $\ell(\hat{y}, y) \in [0,1]$ that is convex in $\hat{y}$, with $\eta > 0$. Let $\mathbf{w}^1 = \langle 1, \ldots, 1 \rangle$,

---

for $t = 1, 2, \ldots, T$ **do**
  Algorithm receives prediction $f_i^t \in \{0,1\}$ from expert i
  Algorithm predicts $\hat{y}^t = \frac{\sum_i w_i^t f_i^t}{\sum_j w_j^t}$
  Nature reveals $y^t \in \{0,1\}$
  Algorithm loss increases: $L_{MA}^{t+1} = L_{MA}^t + \ell(\hat{y}^t, y^t)$
  $w_i^{t+1} = w_i^t \exp(-\eta \ell(f_i^t, y^t))$
end for

NOTE: f_i^t and y^t can be real-valued, but we are assuming for simplicity that they are binary.

Theorem 17.1. *For any sequence of* $\{y^t\}_t$, $\{f_i^t\}_{i,t}$ *we have*

$$L_{MA} \leq \frac{\eta L_i^{t+1} + \log N}{1 - \exp(-\eta)}$$

for all i *where* $L_i^{t+1} = \sum_{s=1}^t \ell(f_i^s, y^s)$.

Corollary 17.2. *With* η *tuned appropriately*

$$L_{MA} \leq L_{i^*}^{T+1} + \log N + \sqrt{2 L_{i^*}^T \log N}$$

where i^* *is the index of the "best expert". Notice that*

$$\frac{L_{MA}}{T} \leq \frac{L_{i^*}^{T+1}}{T} + \epsilon_T$$

where ϵ_T *is approaching* 0 *at a rage of about* $O\left(\frac{1}{\sqrt{T}}\right)$, *since* $L_{i^*}^T$ *is at most* T.

17.2 Hedge Setting

Theorems in this part are proposed by Freund and Schapire, 95[1].

Assuming that we have N actions (or bets), we do the following algorithm.

---

for $t = 1, \ldots, T$ **do**
  Alg chooses distribution $\mathbf{p}^t \in \triangle_N$
  Alg samples $i_t \in \mathbf{p}^t$
  Nature/adversary reveals $\ell^t \in [0,1]^N$
  Alg suffers $\ell^t_{i_t}$, but in expectation, $L_{MA} = \sum_i \ell_i^t p_i^t$
end for

Theorem 17.3. *The hedge setting gives the same bound as the exponential weights algorithm when you choose*

$$\mathbf{p}^t = \frac{\mathbf{w}^t}{\sum_j w_j^t}.$$

Proof: For this proof, we will need to call on the following inequality that holds for all $s \in \mathbb{R}$:

$$\log \mathbb{E} \exp(sX) \leq (e^s - 1)\, \mathbb{E} X.$$

Assume X is a random variable taking values in $[0,1]$ on round t. Let $X^t = \ell(f_i^t, y^t)$ w.p. $\frac{w_i^t}{\sum_{j=1}^N w_j^t}$. Let

$$\Phi_t = -\log \sum_{i=1}^N w_i^t = -\log \sum_{i=1}^N \exp\left(-\eta L_i^t\right).$$

Then

$$\begin{aligned}
\Phi_{t+1} - \Phi_t &= -\log\left(\frac{\sum_i w_i^{t+1}}{\sum_j w_j^t}\right) \\
&= -\log\left(\frac{\sum_i w_i^t \exp(-\eta \ell(f_i^t, y^t))}{\sum_j w_j^t}\right) \\
&= -\log \mathbb{E} \exp(-\eta x^t) \\
&\geq -(e^{-\eta} - 1)\mathbb{E} X^t \\
&= (1 - e^{-\eta}) \frac{\sum_i w_i^t \ell(f_i^t, y^t)}{\sum_j w_j^t} \\
&\geq (1 - e^{-\eta}) \ell\left(\frac{\sum_i w_i^t f_i^t}{\sum_j w_j^t}, y^t\right) \\
&= (1 - e^{-\eta}) \ell(\hat{y}^t, y^t)
\end{aligned}$$

[1] A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting.

Recall that the loss of the algorithm on t is $\ell(\frac{\sum_i w_i^t f_i^t}{\sum_j w_j^t}, y^t)$. This is required for the last step of the sequence of inequalities and equations above.

Whence,

$$\begin{aligned}(1-e^{-\eta})L_{MA}^{T+1} &= \sum_{t=1}^{T}(\Phi_{t+1}-\Phi_t) \\ &= -\log\sum_i \exp\left(-\eta L_i^{T+1}\right)+\log N \\ &\le -\log\left(\exp\left(-\eta L_i^{T+1}\right)\right)+\log N \\ &= \eta L_i^{T+1}+\log N,\end{aligned}$$

which implies that

$$L_{MA} \le \frac{\eta L_i^{T+1}+\log N}{1-e^{-\eta}}$$

■

17.3 Zero-sum games

We are given n strategies/actions for $P1$ and m for $P2$, and the payoff matrix $M \in [-1,+1]^{n\times m}$. Simultaneously,

$$\begin{aligned}&P1 \text{ chooses } i \in [n] \\ &P2 \text{ chooses } j \in [m].\end{aligned}$$

As a result, $P1$ earns M_{ij}, and $P2$ earns $-M_{ij}$.

<u>**Example:**</u> Rock-Paper-Scissors

$$M = \begin{bmatrix} 0 & -1 & +1 \\ +1 & 0 & -1 \\ -1 & +1 & 0 \end{bmatrix}$$

Definition 17.4 (Pure Strategy)**.** *With a* ***pure strategy*** *in a two player game, P1 chooses an action* $i \in [n]$*, and P2 chooses an action* $j \in [m]$*. P1 thus earns* M_{ij}*, and P2 earns* N_{ij}*.*

Definition 17.5 (Mixed Strategy)**.** *With a* ***mixed strategy*** *in a two player game, P1 plays with a distribution* $\mathbf{p} \in \Delta_n$*, and P2 plays with a distribution* $\mathbf{q} \in \Delta_m$*. P1 thus earns* $\mathbf{p}^\top M\mathbf{q} = \sum_{i,j} p_i q_j M_{ij}$*, and P2 earns* $\mathbf{p}^\top N\mathbf{q} = \sum_{i,j} p_i q_j N_{ij}$*.*

17.4 Quick View on Von Neumann's Minimax Theorem

$$\min_{\mathbf{q}}\max_{\mathbf{p}} \mathbf{p}^\top M\mathbf{q} = \max_{\mathbf{p}}\min_{\mathbf{q}} \mathbf{p}^\top M\mathbf{q}$$

The minimizer gets to see the maximizer's strategy before picking his/her own, so the right side will clearly be less than or equal to the left. The other way is more difficult.