Algorithmic Excursions: Topics in Computer Science II — Lecture - Spring 2016

Week 10: MW Applications

Lecturer: Kasturi Varadarajan — *Scribe: Leixin Zhou*

1 MW when gains considered

There are situations where it makes more sense for the vector $m^{(t)}$ to specify gains for each expert rather than losses. Now our goal is to get as much total expected payoff as possible in comparison to the total payoff of the best expert. We can get an algorithm for this case simply by running the Multiplicative Weights algorithm using the cost vector $-m^{(t)}$. The resulting algorithm is identical, and the following theorem follows directly from Claim 9.2 and Corollary 9.3 by simply negating the quantities. Under the scenario of gains, we have following Theorem and Corollary. Assume there are n experts, MW updates T rounds and $\eta \leq \frac{1}{2}$.

Theorem 10.1. For any expert i,

$$\sum_{t=1}^{T} m^{(t)} \cdot \mathbf{p}^{(t)} \geq \eta \sum_{t=1}^{T} |m_i^{(t)}| - \frac{\ln n}{\eta}$$

Corollary 10.1. For any distribution $\mathbf{p}$ over experts,

$$\sum_{t=1}^{T} m^{(t)} \cdot \mathbf{p}^{(t)} \geq \sum_{t=1}^{T} (m^{(t)} - \eta |m^{(t)}|) \cdot \mathbf{p} - \frac{\ln n}{\eta}$$

2 Learning a linear classifier: Winnow

Input: m labelled examples: $(a_1, l_1), (a_2, l_2), \cdots, (a_m, l_m)$, where $a_j \in R^n$ is a feature vector and each $l_j \in \{-1, +1\}$ is the corresponding label.

Goal: Find vector $X \in R^n$ with $X_i \geq 0$ for each i and $sign(a_j \cdot X) = l_j$ or $l_j(a_j \cdot X) \geq 0$ for each j.

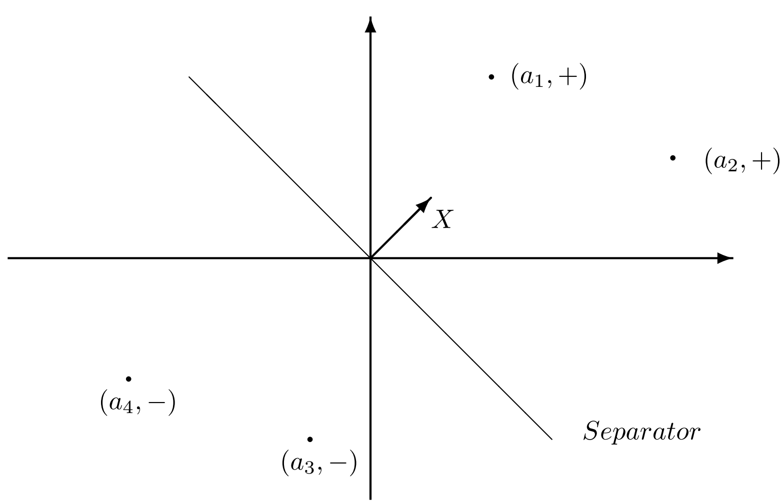

The above figure shows a simple example when $n = 2$. X is a goal vector candidate and the corresponding separator is the line orthogonal to X.

For notational convenience, if we redefine a_j to be $l_j a_j$, then the problem reduces to finding a solution to the following LP: find $X \in R^n$ s.t $\forall j = 1, 2, \cdots, m$, $a_j \cdot X \geq 0$, $\mathbf{1} \cdot \mathbf{X} = \mathbf{1}$, $\forall i : X_i \geq 0$.

Assumption: there exists a large-margin solution X^*, i.e, there is $\varepsilon > 0$ s.t $a_j \cdot X^* \geq \varepsilon$ for each j.

Algorithm: Run MW, each feature is an expert. Pick $\eta = \frac{\varepsilon}{2\rho}$. At each time, MW algorithm has a probability distribution $\mathbf{p}^{(t)}$ over experts / features.

— If $a_j \cdot \mathbf{p}^{(t)} \geq 0$ for each j, we stop

— If not, pick any j s.t $a_j \cdot \mathbf{p}^{(t)} < 0$. Let gain vector $m^{(t)} = \frac{1}{\rho} \times a_j$, where $\rho = \max_{i,j} |a_{i,j}|$.

Analysis: note that if $\mathbf{p}^{(t)}$ fails to satisfy constraints, $m^{(t)} \cdot \mathbf{p}^{(t)} = \frac{1}{\rho} a_j \cdot \mathbf{p}^{(t)} < 0$, whereas $m^{(t)} \cdot X^* = \frac{1}{\rho} a_j \cdot X^* \geq \frac{\varepsilon}{\rho}$. Suppose we have T fail iterations. Then using Corollary 1:

$$0 > \sum_{t=1}^{T} m^{(t)} \cdot p^{(t)} \geq \sum_{t=1}^{T} (m^{(t)} - \eta |m^{(t)}|) \cdot \mathbf{p} - \frac{\ln n}{\eta}$$

$$\geq \frac{\varepsilon}{\rho} T - \eta T - \frac{\ln n}{\eta}$$

$$= \frac{\varepsilon}{\rho} T - \frac{\varepsilon}{2\rho} T - \frac{2\rho \ln n}{\varepsilon} = \frac{\varepsilon}{2\rho} T - \frac{2\rho \ln n}{\varepsilon}$$

$$\Rightarrow 0 \geq \frac{\varepsilon}{2\rho} T - \frac{2\rho \ln n}{\varepsilon}$$

$$\Rightarrow T \leq \frac{4\rho^2 \ln n}{\varepsilon^2}$$

3 Solving zero-sum games approximately

Here is a very simple zero-sum game example demonstrated by table 1. There are two player i and j. Row player (i.e i) has two strategies, each of which correspond to one row of the table. Column player (i.e j) has three strategies, each of which corresponds to one column. $A(i, j)$ denotes payoff from row player to column

$i \backslash j$	1	2	3
1	9	3	4
2	1	7	6

Table 1: Example 1 with 2 players i and j.

player. For example, if row player uses strategy one and column player picks strategy one, the payoff is $A(1, 1) = 9$. Player i wants to minimize payoff and player j wants to maximize payoff.

If row player goes first, his optimal play is to picks a row i that minimizes $\max_{j'} A(i, j')$. For the given example, the optimal is 7. If column player goes first, his best play is to pick a column j, that maximize $\min_{i'} A(i', j)$. In this case, the optimal is 4.

Claim 10.1.

$$\min_i \max_j A(i, j) \geq \max_j \min_i A(i, j)$$

Proof. Exercise ☐

An example is matching pennies, which is shown in table 2. In this case, the left side is 1 and the right side is 0, i.e $1 > 0$.

	H	T
H	0	1
T	1	0

Table 2: Matching pennies game.

Now suppose that a strategy of row players is a distribution $\mathbf{p}$ over rows. Strategy of column play is distribution $\mathbf{q}$ over columns. Given $\mathbf{p}$, expected payoff for column j is defined as:

$$A(\mathbf{p}, j) := E_{i \sim \mathbf{p}}[A(i,j)]$$

For example, if we consider the example 1 and $\mathbf{p} = (\frac{1}{4}, \frac{3}{4})$, then $A(\mathbf{p}, 1) = \frac{1}{4} \times 9 + \frac{3}{4} \times 1$.

Under the probability scenario, if row player goes first, optimal for column player is : $\max_j A(\mathbf{p}, j)$; optimal for row player is : $\min_{\mathbf{p}} \max_j A(\mathbf{p}, j)$. If column player goes first, optimal for column player is : $\max_{\mathbf{q}} \min_i A(i, \mathbf{q})$. Right now, **Claim 10.1** becomes:

$$\min_{\mathbf{p}} \max_j A(\mathbf{p}, j) \geq \max_{\mathbf{q}} \min_i A(i, \mathbf{q})$$

Von Neumann's min-max theorem: Equality holds. e.g in the case of example 2, $\mathbf{p}_{optimal} = \mathbf{q}_{optimal} = (\frac{1}{2}, \frac{1}{2})$ and the corresponding payoff are both $\frac{1}{2}$.

Given A, each entry is in $[0, 1]$ range and let $\lambda^* = \min_{\mathbf{p}} \max_j A(\mathbf{p}, j)$. We will find a distribution $\tilde{\mathbf{p}}$ s.t :

$$\max_j A(\tilde{\mathbf{p}}, j) \leq \lambda^* + \varepsilon$$

Apply MW. Each row is an expert. So at each time $t = 1, 2, \cdots, T$, algorithm has a probability distribution $\mathbf{p}^{(t)}$ over rows. Let j^t be a column that maximizes $A(\mathbf{p}^{(t)}), j)$. Column j^t is loss vector at time t. By MW corollary, for any distribution $\mathbf{p}$,

$$\begin{aligned}
\sum_{t=1}^{T} A(\mathbf{p}^{(t)}, j^{(t)}) &\leq (1+\eta) \sum_{t=1}^{T} A(\mathbf{p}, j^{(t)}) + \frac{\ln n}{\eta}) \\
&\leq \sum_{t=1}^{T} A(\mathbf{p}, j^{(t)}) + \eta \sum_{t=1}^{T} \mathbf{p} \cdot |m^{(t)}| + \frac{\ln n}{\eta} \\
&\leq \sum_{t=1}^{T} A(\mathbf{p}, j^{(t)}) + \eta \sum_{t=1}^{T} 1 + \frac{\ln n}{\eta} \\
&\leq \sum_{t=1}^{T} A(\mathbf{p}, j^{(t)}) + \eta T + \frac{\ln n}{\eta}]
\end{aligned}$$

Dividing by T,

$$\frac{1}{T} \sum_{t=1}^{T} A(\mathbf{p}^{(t)}, j^{(t)}) \leq \frac{1}{T} \sum_{t=1}^{T} A(\mathbf{p}, j^{(t)}) + \eta + \frac{\ln n}{\eta T}$$