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

Lecture 18: Nash's Theorem and Von Neumann's Minimax Theorem

Lecturer: Jacob Abernethy *Scribes: Lili Su, Editors: Weiqing Yu and Andrew Melfi*

18.1 Review: On-line Learning with Experts (Actions)

Setting Given n experts (actions), the general on-line setting involves T rounds. For round $t = 1 \dots T$:

- The algorithm plays with the distribution $\mathbf{p}^t = \frac{\boldsymbol{\omega}^t}{\|\boldsymbol{\omega}^t\|_1} \in \Delta_n$.
- The i-th expert (action) suffers the loss $\ell_i^t \in [0, 1]$.
- The algorithm suffers the loss $\mathbf{p}^t \cdot \boldsymbol{\ell}^t$.

Theorem 18.1 (Regret Bound for EWA)**.**

$$\underbrace{\frac{1}{T}\sum_{t=1}^{T}\mathbf{p}^t \cdot \boldsymbol{\ell}^t}_{\mathcal{L}_{MA}^{t+1}} \leq \underbrace{\min_i \frac{1}{T}\sum_{t=1}^{T}\ell_i^t}_{\mathcal{L}_{IX}^{t+1}} + O\left(\sqrt{\frac{\log N}{T}}\right) = \min_{\mathbf{p}\in\Delta_n}\frac{1}{T}\sum_{t=1}^{T}\mathbf{p}\cdot\boldsymbol{\ell}^t + O\left(\sqrt{\frac{\log N}{T}}\right)$$

Note The distribution $\mathbf{p} = \mathbf{e}_i$ means the algorithm puts all mass on the i-th action.

18.2 Two Player Game

Definition 18.2 (Two Player Game)**.** *A **two player game** is defined by a pair of matrices $M, N \in [0, 1]^{n\times m}$.*

Definition 18.3 (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 18.4 (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 $\mathbf{p}^\top N\mathbf{q} = \sum_{i,j} p_i q_j N_{ij}$.*

Definition 18.5 (Zero-sum Game)**.** *A **zero-sum game** is a two player game, where the matrices M, N has the relation $M = -N$.*

18.3 Nash's Theorem

Definition 18.6 (Nash Equilibrium)**.** *In a two player game, a **Nash Equilibrium(Neq)**, in which P1 plays with the distribution $\widetilde{\mathbf{p}} \in \Delta_n$, and P2 plays with the distribution $\widetilde{\mathbf{q}} \in \Delta_m$, satisfies*

- *for all $\mathbf{p} \in \Delta_n$, $\widetilde{\mathbf{p}}^\top M\widetilde{\mathbf{q}} \geq \mathbf{p}^\top M\widetilde{\mathbf{q}}$*
- *for all $\mathbf{q} \in \Delta_m$, $\widetilde{\mathbf{p}}^\top N\widetilde{\mathbf{q}} \geq \widetilde{\mathbf{p}}^\top N\mathbf{q}$*

Theorem 18.7 (Nash's Theorem)**.** *Every two player game has a Nash Equilibrium(Neq). (Not all have pure strategy equilibria.)*

Lemma 18.8 (Brouwer's Fixed-point Theorem)**.** *Let $B \subseteq \mathcal{R}^d$ be a compact convex set, and a function $f : B \to B$ is continuous. Then there exists $x \in B$, such that $x = f(x)$.*

Proof Sketch of Nash's Theorem

1. Let $c_i(\mathbf{p},\mathbf{q}) = \max\left(0, \mathbf{e}_i^\top M\mathbf{q} - \mathbf{p}^\top M\mathbf{q}\right)$, and $d_i(\mathbf{p},\mathbf{q}) = \max\left(0, \mathbf{q}^\top M\mathbf{e}_j - \mathbf{p}^\top M\mathbf{q}\right)$.
2. Define a map $f : (\mathbf{p},\mathbf{q}) \to (\mathbf{p}',\mathbf{q}')$, $p_i' = \frac{p_i + c_i(\mathbf{p},\mathbf{q})}{1+\sum\limits_{i'\in[n]} c_{i'}(\mathbf{p},\mathbf{q})}$, and $q_i' = \frac{q_i + d_i(\mathbf{p},\mathbf{q})}{1+\sum\limits_{i'\in[m]} d_{i'}(\mathbf{p},\mathbf{q})}$.
3. By Brouwer's fixed-point theorem, there exists a fixed-point $(\widetilde{\mathbf{p}},\widetilde{\mathbf{q}})$, $f(\widetilde{\mathbf{p}},\widetilde{\mathbf{q}}) = (\widetilde{\mathbf{p}},\widetilde{\mathbf{q}})$.
4. Show the fixed-point $(\widetilde{\mathbf{p}},\widetilde{\mathbf{q}})$ is the Nash Equilibrium.

18.4 Von Neumann's Minimax Theorem

Theorem 18.9 (Von Neumann's Minimax Theorem)**.**

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

Proof by Nash's Theorem

- Exercise

Proof by the Exponential Weighted Average Algorithm

a) The "$\geq$" direction is straightforward. Let $\mathbf{p}_1 \in \Delta_n$, $\mathbf{q}_1 \in \Delta_m$ be the choices for $\min\limits_{\mathbf{p}\in\Delta_n}\max\limits_{\mathbf{q}\in\Delta_m} \mathbf{p}^\top M\mathbf{q} = \mathbf{p}_1^\top M\mathbf{q}_1$, and $\mathbf{p}_2 \in \Delta_n$, $\mathbf{q}_2 \in \Delta_m$ be the choices for $\max\limits_{\mathbf{q}\in\Delta_m}\min\limits_{\mathbf{p}\in\Delta_n} \mathbf{p}^\top M\mathbf{q} = \mathbf{p}_2^\top M\mathbf{q}_2$.

$$\min_{\mathbf{p}\in\Delta_n}\max_{\mathbf{q}\in\Delta_m} \mathbf{p}^\top M\mathbf{q} = \mathbf{p}_1^\top M\mathbf{q}_1 \geq \mathbf{p}_1^\top M\mathbf{q}_2 \geq \mathbf{p}_2^\top M\mathbf{q}_2 = \max_{\mathbf{q}\in\Delta_m}\min_{\mathbf{p}\in\Delta_n} \mathbf{p}^\top M\mathbf{q}.$$

An intuitive explanation for the first inequality is in $\min_{\mathbf{p}\in\Delta_n}\max_{\mathbf{q}\in\Delta_m} \mathbf{p}^\top M\mathbf{q}$, $\mathbf{q}$ is chosen to maximize $\mathbf{p}^\top M\mathbf{q}$ for any given $\mathbf{q}$, therefore, $\mathbf{p}_1^\top M\mathbf{q}_1 \geq \mathbf{p}_1^\top M\mathbf{q}$ for any $\mathbf{q} \neq \mathbf{q}_1$. Similar explanation goes for the second inequality.

b) Show the "$\leq$" direction holds up to $O\left(\frac{1}{\sqrt{t}}\right)$ approximation.

Setting Imagine playing a T-round game against a really hard adversary. For round $t = 1 \dots T$:

- Player 1 plays with the distribution $\mathbf{p}^t = \frac{\boldsymbol{\omega}^t}{\|\boldsymbol{\omega}^t\|_1} \in \Delta_n$.
- Player 2 plays with the distribution $\mathbf{q}^t = \arg\max\limits_{\mathbf{q}\in\Delta_m} \mathbf{p}^t M\mathbf{q}$.
- Let $\boldsymbol{\ell}^t = M\mathbf{q}^t$, and Player 1 suffers the loss $\mathbf{p}^t \cdot \boldsymbol{\ell}^t = \mathbf{p}^t \cdot M\mathbf{q}^t$.
- Let $\boldsymbol{\omega}^1 = (1 \dots 1)$, and update $\omega_i^{t+1} = \omega_i^t \exp(-\eta \ell_i^t)$.

Trick Analyze $\frac{1}{T}\sum\limits_{t=1}^{T} \mathbf{p}^t \cdot M\mathbf{q}^t$.

1. By Jensen's Inequality,

$$\frac{1}{T}\sum_{t=1}^{T} \mathbf{p}^t M\mathbf{q}^t = \frac{1}{T}\sum_{t=1}^{T} \max_{\mathbf{q}\in\Delta_m} \mathbf{p}^t M\mathbf{q} \geq \max_{\mathbf{q}\in\Delta_m} \left(\frac{1}{T}\sum_{t=1}^{T} \mathbf{p}^t\right) M\mathbf{q} \geq \min_{\mathbf{p}\in\Delta_n}\max_{\mathbf{q}\in\Delta_m} \mathbf{p}^\top M\mathbf{q}.$$

2. By the exponential weighted average algorithm,

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\mathbf{p}^t M\mathbf{q}^t &= \frac{1}{T}\sum_{t=1}^{T}\mathbf{p}^t\cdot\boldsymbol{\ell}^t \\
&\le \min_i \frac{1}{T}\sum_{t=1}^{T}\mathbf{p}^t\cdot\ell_i^t + \epsilon_T (:= \frac{\text{Regret}_T}{T}) \\
&= \min_{\mathbf{p}\in\Delta_n}\frac{1}{T}\sum_{t=1}^{T}\mathbf{p}\cdot\boldsymbol{\ell}^t + \epsilon_T \\
&= \min_{\mathbf{p}\in\Delta_n}\frac{1}{T}\sum_{t=1}^{T}\mathbf{p}\cdot M\mathbf{q}^t + \epsilon_T \\
&= \min_{\mathbf{p}\in\Delta_n}\mathbf{p}\cdot M(\frac{1}{T}\sum_{t=1}^{T}\mathbf{q}^t) + \epsilon_T \\
&\le \max_{\mathbf{q}\in\Delta_m}\min_{\mathbf{p}\in\Delta_n}\mathbf{p}^\top M\mathbf{q} + \epsilon_T.
\end{aligned}$$

Putting the results in 1. and 2. together, we have

$$\min_{\mathbf{p}\in\Delta_n}\max_{\mathbf{q}\in\Delta_m}\mathbf{p}^\top M\mathbf{q} \le \max_{\mathbf{q}\in\Delta_m}\min_{\mathbf{p}\in\Delta_n}\mathbf{p}^\top M\mathbf{q} + \epsilon_T.$$

T can be chosen as big as we wanted, and thus $\epsilon_T = O\left(\frac{1}{\sqrt{T}}\right)$ vanishes. It completes the prove of the "$\le$" direction

Theorem 18.10 (Generalization of Von Neumann's Minimax Theorem)**.** *Let $X \subseteq \mathcal{R}^n$, $Y \subseteq \mathcal{R}^m$ be compact convex sets. Let $f : X \times Y \to \mathcal{R}$ be some differentiable function with bounded gradients, where $f(\cdot, \mathbf{y})$ is convex in its first argument for all fixed $\mathbf{y}$, and $f(\mathbf{x}, \cdot)$ is concave for in its second argument for all fixed $\mathbf{x}$. Then*

$$\inf_{x\in X}\sup_{y\in Y} f(x,y) = \sup_{y\in Y}\inf_{x\in X} f(x,y).$$