Lecture 9: Game Theory and Boosting

Lecturer: Jacob Abernethy *Scribes: Ezra Goss and Jad Salem*

9.1 Nash Equilibria and the Minimax Theorem

Definition 9.1 (Bimatrix Game) *A* two-player bimatrix game *with matrices $M \in \mathbb{R}^{n\times m}$ and $N \in \mathbb{R}^{n\times m}$, is a game where player 1 selects a distribution $p \in \Delta_n$ and player 2 selects a distribution $q \in \Delta_m$. The utility $U^k(p,q)$ received by player k is*

$$U^1(p,q) = \mathbb{E}_{\substack{i\sim p\\ j\sim q}}[M_{ij}] = p^\top M q \quad \text{and} \quad U^2(p,q) = \mathbb{E}_{\substack{i\sim p\\ j\sim q}}[N_{ij}] = p^\top N q \ .$$

In particular, if player 1 chooses $i \in [n]$ and player 2 chooses $j \in [m]$, then the utilities received by each player are

$$U^1(i,j) = M_{ij}, \quad U^2(i,j) = N_{ij} \ .$$

If $M + N = 0$, this is called a zero-sum game.

Note that if the strategy $p \in \Delta_n$ of Player 1 is known to Player 2, then Player 2 can choose an optimal strategy using linear programming. It is often of interest to study pairs (p,q) of strategies which are stable in the following sense:

1. given a strategy of p for Player 1, q is an optimal strategy for Player 2, and
2. given a strategy of q for Player 2, p is an optimal strategy for Player 1.

If such a pair of strategies is used, then neither player is incentivized to change strategies. This notion is formalized in the following definition.

Definition 9.2 (Nash Equilibrium) *Given $p \in \Delta_n$ and $q \in \Delta_m$, the pair (p,q) is a* Nash Equilibrium *if*

$$p^\top M q \geq \tilde{p}^\top M q \quad \forall \tilde{p} \in \Delta_n$$
$$p^\top N q \geq p^\top N \tilde{q} \quad \forall \tilde{q} \in \Delta_m$$

The natural next question is: does there always exist a Nash equilibrium? Nash answered this question in the affirmative:

Theorem 9.3 (Nash's Theorem) *There is a Nash Equilibrium*[1] *for every bimatrix game.*

Proof idea: Given any pair (p,q) of strategies, let $f(p,q)$ be the set of strategies $(\tilde{p},\tilde{q})$ for which $\tilde{p}$ is optimal with respect to q and $\tilde{q}$ is optimal with respect to p. Then a Nash Equilibrium is simply a point $(p,q) \in \Delta_n \times \Delta_m$ such that $(p,q) \in f(p,q)$. The existence of such a point is guaranteed by the Kakutani fixed-point theorem. ■

[1] It can be shown that the problem of finding (approximate) Nash Equilibria is complete for the class PPAD, as it reduces to finding (approximate) Brouwer fixed points.

Theorem 9.4 (Minimax Theorem, von Neumann) *Take a zero-sum game given by $M \in \mathbb{R}^{n\times m}$. Then*

$$\min_{p\in\Delta_n}\max_{q\in\Delta_m} p^\top Mq = \max_{q\in\Delta_m}\min_{p\in\Delta_n} p^\top Mq \ .$$

Note that Theorem 9.4 implies Nash's theorem. To see why, take $p_* \in \arg\min_{p\in\Delta_n}\max_{q\in\Delta_m} p^\top Mq$ and $q_* \in \arg\max_{q\in\Delta_m}\min_{p\in\Delta_n} p^\top Mq$. Then

$$\begin{aligned}\min_p \max_q p^\top Mq &= \max_q p_*^\top Mq \\ &\geq p_*^\top Mq_* && (9.1)\\ &\geq \min_p p^\top Mq_* && (9.2)\\ &= \max_q \min_p p^\top Mq\end{aligned}$$

By Theorem 9.4, the inequalities in (9.1) and (9.2) are equalities, and we get

$$\max_q p_*^\top Mq \overset{(a)}{=} p_*^\top Mq_* \overset{(b)}{=} \min_p p^\top Mq_*$$

The equality (a) implies that

$$p_*^\top Mq_* \geq p_* M\tilde{q} \ \ \forall \tilde{q} \in \Delta_m$$

Similarly, (b) implies that

$$p_*^\top(-M)q_* \geq \tilde{p}^\top(-M)q_* \ \ \forall \tilde{p} \in \Delta_n$$

Hence Theorem 9.4 implies Nash's Theorem. Now we prove Theorem 9.4.

Proof of Theorem 9.4: We have shown weak duality above, so it remains to show $(\leq)$ (strong duality). We do so by setting up an iterative game, and applying the Exponential Weights Algorithm. Consider the following game.

For $t = 1, \ldots, T$:

1. p_t is chosen by P1
2. q_t is chosen by P2
3. p_t observes a loss vector $\ell_t = Mq_t$
4. q_t observes a loss vector $h_t = -M^T p_t$
5. Players update their strategies to p_{t+1}, q_{t+1} via the Exponential Weights Algorithm, with respect to these loss vectors.

It remains to show that $\min_{p\in\Delta_n}\max_{q\in\Delta_m} p^\top Mq \leq \max_{q\in\Delta_m}\min_{p\in\Delta_n} p^\top Mq + \varepsilon$ for all $\varepsilon > 0$. To that end, consider

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^T p_t^\top Mq_t = \frac{1}{T}\sum p_t^\top \ell_t = \min_{p\in\Delta_n}\frac{1}{T}\sum_{t=1}^T p^\top \ell_t + \underbrace{\frac{\text{Regret}_T^p}{T}}_{=\varepsilon_T^p} \quad &\text{(by definition)}\\ = \min_p \frac{1}{T}\sum_{t=1}^T p^\top Mq_t + \varepsilon_T^p &\\ = \min_p p^\top M\overline{q_T} + \varepsilon_T^p &\\ \leq \max_q \min_p p^\top Mq + \varepsilon_T^p &\end{aligned}$$

where $\overline{q_T}$ is an average of the q_ts. Note that $\varepsilon_T^p \to 0$ since Regret is sublinear. Similarly

$$\begin{aligned}
-\frac{1}{T}\sum p_t^\top M q_t &= \frac{1}{T}\sum_{t=1}^{T} q_t^\top h_t \\
&= \min_{q\in\Delta_m} \frac{1}{T}\sum q^\top h_t + \underbrace{\frac{\text{Regret}_T^q}{T}}_{=:\varepsilon_T^q} \\
&= \min_q -\Big(\frac{1}{T}\sum p_t^\top\Big) M q + \varepsilon_T^q \\
&= -\max_q \overline{p_T}^\top M q + \varepsilon_T^q \\
&\le -\min_p \max_q p^\top M q + \varepsilon_T^q
\end{aligned}$$

Combining, we get

$$\begin{aligned}
\min_p \max_q p^\top M q - \varepsilon_T^q \le \max_q \overline{p_T}^\top M q - \varepsilon_T^q \\
&= \frac{1}{T}\sum p_t^\top M q_t \\
&= \min_p p^\top M \overline{q_T} + \varepsilon_T^p \\
&\le \max_q \min_p p^\top M q + \varepsilon_T^p
\end{aligned}$$

Sending $T \to \infty$, we get that $\varepsilon_T^p + \varepsilon_T^q \to 0$. ■

Corollary 9.5 *$(\overline{p_T}, \overline{q_T})$ is an ε-approximate Nash equilibrium, where $\varepsilon = \varepsilon_T^p + \varepsilon_T^q$.*

Remark. This proof also outlines an algorithm for finding approximate Nash equilibria.

9.2 Boosting

We are given data $\{(x_1, y_1), \ldots, (x_n, y_n)\} \subset \mathcal{X} \times \{-1, 1\}$, and access to a class of *weak learners* $\mathcal{H}$. Each $h \in \mathcal{H}$ is a map $\mathcal{X} \to \{-1, 1\}$. By "weak," we mean that we shouldn't expect to find a "good" $h \in \mathcal{H}$ on any particular problem, but we can find an $h \in \mathcal{H}$ which does slightly better than a random coin toss. One goal is to combine many h's from $\mathcal{H}$ to get an ensemble:

$$F(x) = \text{sign} \sum_{i=1}^{m} q_i h_i(x)$$

which is highly correlated with the true classifier.

We now formally define the Weak Learning Hypothesis (WLH) and the Strong Learning Hypothesis (SLH).

Definition 9.6 (Weak Learning Hypothesis $(\gamma > 0)$.) *For any distribution $p \in \Delta_n$, there is some $h \in \mathcal{H}$ such that*

$$\mathbb{P}_{(x_i,y_i)\sim p}[h(x_i) = y_i] = \sum p_i \mathbb{1}[h(x_i) = y_i] \ge \frac{1}{2} + \gamma$$

Definition 9.7 (Strong Learning Hypothesis) *There is a distribution $q \in \Delta(\mathcal{H})$ such that for every (x, y) in the sample,*

$$\sum_{h \in \mathcal{H}} q(h)h(x_i)y_i > 0 \ .$$

Equivalently, $F(x) = y$ for all (x, y) in the sample, where $F(x) = \mathrm{sign}(\sum q(h)h(x))$.

Theorem 9.8 *The weak learning hypothesis is equivalent to the strong learning hypothesis.*

Proof: Assume $\mathcal{H} = \{h_1, \ldots, h_m\}$. Then the WLH($\gamma$) implies that for all $p \in \Delta_n$, there is a $j \in [m]$ such that $p^\top M e_j \geq 2\gamma$, where $M_{ij} = h_j(x_i)y_i$. This implies that

$$\min_{p \in \Delta_n} \max_{j \in [m]} p^\top M e_j \geq 2\gamma$$

By Theorem 9.4, we have that

$$\max_q \min_p p^\top M q \geq 2\gamma > 0$$

So

$$\max_q \min_{i \in [n]} e_i^\top M q > 0$$

So, there is a $q \in \Delta_m$ such that for all (x_i, y_i), $\sum q_j h_j(x_i) y_i > 0$. ■