LECTURE #21

APRIL 19, 2016

LECTURER: ELAD HAZAN SCRIBE: XINYI CHEN

1 LP Duality

1.1 LP

Linear Program:

$$\begin{aligned} \min_x \quad & c^T x \\ \text{s.t.} \quad & Ax \geq b \\ & x \geq 0 \end{aligned}$$

Simplex algorithm:
The set of linear constraints define a polyhedron, and we can use the simplex algorithm to solve a linear program. The simplex algorithm relies on the fact that under the optimal solution, some constraints must be tight. Therefore, we start at some vertex of the polyhedron, and then move to an adjacent vertex. At each vertex, the objective value might be improved.

1.2 Dual LP and Strong Duality

Suppose the minimum is attained at λ, then $\exists$ an automatic way to show that the optimal value is $\geq \lambda$: by maximizing the lower bound of the objective function.The dual LP is given by

$$\begin{aligned} \max_y \quad & y^T b \\ \text{s.t.} \quad & A^T y \leq c \\ & y \geq 0 \end{aligned}$$

Theorem 1.1. *Suppose the optimal value of the primal LP is λ and the optimal value of the dual LP is μ, where λ and μ are bounded, then $\mu = \lambda$.*

Von Neumann presented a topological argument to prove the strong duality theorem using the fixed point theorem, and usually the proof involves Farkas Lemma, but here we consider an equivalence to Zero-Sum Games.

Definition 1.2. Zero-Sum Games (ZSG): we have a set of players, and each player has a set of strategies. We define a game to be a mapping from all strategy profiles to real numbers.

Example 1.3. Rock-Paper-Scissors
Player 1: strategies = $\{$R, P, S$\}$; Player 2: strategies = $\{$R, P, S$\}$
Payoff:

	R	P	S
R	(0, 0)	(-1, 1)	(1, -1)
P	(1, -1)	(0, 0)	(-1, 1)
S	(-1, 1)	(1, -1)	(0, 0)

which we denote M, where $M_{i,j}[1]$ = payoff of player 1 if player 1 plays i and player 2 plays j, and $M_{i,j}[2]$ = payoff of player 2 if player 2 plays j and player 1 plays i. In a 2-player ZSG, the payoff to player 1 = -payoff to player 2, so we only need a single matrix to represent the payoff to both players.

Payoff to player 1:

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

Example 1.4. Non-ZSG: Prisoner's Dilemma
Payoff:

	C	N
C	(5, 5)	(0, 10)
N	(10, 0)	(3, 3)

where C denotes "Collaborate with Police" and N denotes "Not Collaborate with Police".

Definition 1.5. Equilibrium of ZSG: a pair of mixed strategies (i, j) such that there is no motivation for either player to change strategy. Mixed strategy: a distribution over strategies. For Rock-Paper-Scissors, the equilibrium is $[(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})]$, and there is no equilibrium for Prisoner's Dilemma.

Consider a game with payoff matrix $M \in R^{m \times n}$. The mixed strategy of player 1 is given by $p \in \triangle_m$, and the mixed strategy of player 2 is given by $q \in \triangle_n$. If (p, q) is an equilibrium, then for every $i \in m$, $p^T Mq \geq e_i^T Mq$, and for every $j \in n$, $p^T Mq \leq p^T Me_j$. The "best" non-assuming strategy for player 1 is given by:

$$p^* = argmax_{p \in \triangle_m} \, min_{q \in \triangle_n} \, p^T Mq \tag{1}$$
$$\lambda_R = min_{q \in \triangle_n} \, p^{*T} Mq \tag{2}$$

and for player 2 is given by:

$$q^* = argmin_{q \in \triangle_n} \, max_{p \in \triangle_m} \, p^T Mq \tag{3}$$
$$\lambda_C = max_{p \in \triangle_m} \, p^T Mq^* \tag{4}$$

Theorem 1.6 (Von Neumann Minimax Thm)**.** *For every 2-player ZSG, we have $\lambda_C = \lambda_R$, and (p^*, q^*) is an equilibrium. The set of all equilibria for a given game is continuous and convex.*

How do we compute p^*? We can solve the following LP:

$$\begin{array}{ll} \max & \lambda_R \\ \text{s.t.} & \forall i \in [n], \; p^T Me_i \geq \lambda_R \\ & \sum p_i = 1 \\ & p_i \geq 0 \end{array}$$

which has dual:

$$\begin{aligned}
\min \quad & \lambda_C \\
\text{s.t.} \quad & \forall j \in [m],\ e_j^T M q \le \lambda_C \\
& \sum q_i = 1 \\
& q_i \ge 0
\end{aligned}$$

Now we show an equivalence between the Von Neumann Minimax Theorem and LP Duality.

Proof.
$\Rightarrow$): We've already shown ZSG $\Rightarrow$ LP.
$\Leftarrow$): LP $\Rightarrow$ ZSG:
Consider repeated game:
$p_1 = \overrightarrow{1} * \frac{1}{m}$, which is the uniform distribution in $\triangle_m$
$q_{t+1} = q_t e^{-\varepsilon p_t^T M}$ + normalization.
$l_t = M q_t$
$p_{t+1}(i) = \frac{p_t(i) e^{\varepsilon l_t(i)}}{\sum p_{t+1}(i)}$ (or any regret minimization algorithm, for example $p_{t+1} = \Pi(p_t - \eta l_t)$)
An easy observation: by weak duality,

$$\begin{aligned}
\lambda_C = \min_q \max_p p^T M q &= \max_p p^T M q^* \\
&\ge \max_p \min_q p^T M q = \lambda_R
\end{aligned}$$

So we have $\lambda_C \ge \lambda_R$.

Let $\bar{q} = \frac{1}{T} \sum q_t$, $\bar{p} = \frac{1}{T} \sum p_t$.

$$\lambda_C = \min_q \max_p p^T M q \le \max_p p^T M \bar{q} = \max_p \frac{1}{T} \sum_{t=1}^{T} p M q_t \tag{5}$$

$$\sum_t p_t^T M q_t \ge \max_{p^*} \sum_t p^{*T} M q_t + O(\sqrt{T log n}) \tag{6}$$

$$\sum_t p_t^T M q_t \le \min_{q^*} \sum_t p_t^T M q^* + O(\sqrt{T log n}) \tag{7}$$

$$\begin{aligned}
\min_q \max_p p^T M q \le \max_p p^T M \bar{q} = \max_p \frac{1}{T} \sum_t p^T M q_t \le \frac{1}{T} \sum_t p_t^T M q_t + O(\frac{1}{\sqrt{T}}) & \qquad \text{by 6} \\
\le \min_{q^*} \frac{1}{T} \sum_t p_t^T M q^* + O(\frac{1}{\sqrt{T}}) & \qquad \text{by 7} \\
= \min_{q^*} \bar{p}^T M q^* + O(\frac{1}{\sqrt{T}}) & \\
\le \max_p \min_q p^T M q + O(\frac{1}{\sqrt{T}}) &
\end{aligned}$$

So we have $\lambda_C \le \lambda_R$.
As $T \to \infty$, LHS $= \max_p \min_q p^T M q$ $\square$