

Lecture 5: 2-Player Zero Sum Games

Lecturer: Yishay Mansour

Scribe: Ilan Cohen, Natan Rubin, Ophir Bleiberg*

1

5.1 2-Player Zero Sum Games

In this lecture we will discuss 2-player zero sum games. Such games are completely competitive, where whatever one player wins, the other loses. Examples of such games include chess, checkers, backgammon, etc. We will show that in such games:

- An equilibrium always exists (although not necessarily a pure one);
- All equilibrium points yield the same payoff for the players;
- The set of equilibrium points is actually the cartesian product of independent sets of equilibrium strategies per player.

We will also show applications of the theory of 2-Players zero sum games.

Definition Let G be the game defined by $\langle N, (A_i), (u_i) \rangle$ where N is the number of players, A_i is the set of possible pure strategies for player i , and u_i is the payoff function for player i . Let A be the Cartesian product $A = \prod_{i=1}^n A_i$. Then G is a *zero sum game* if and only if:

$$\forall \vec{a} \in A, \sum_{i=1}^n u_i(\vec{a}) = 0 \quad (5.1)$$

In other words, a *zero sum game* is a game in which, for any joint action, the sum of payoffs to all players is zero.

We naturally extend the definition of u_i to any probability distribution \vec{p} over A by defining $u_i(\vec{p}) = E_{\vec{a} \sim \vec{p}}(u_i(\vec{a}))$. The following is immediate due to the linearity of the expectation and the zero sum constraint:

Corollary 5.1 *Let G be a zero sum game, and $\Delta(A)$ the set of probability distributions over A . Then*

$$\forall \vec{p} \in \Delta(A), \sum_{i=1}^n u_i(\vec{p}) = 0 \quad (5.2)$$

¹The scribe notes are based in part on the scribe notes of Yair Ha-Levi and Daniel Deutsch, 2004

Specifically, this will also hold for any probability distribution that is the product of N independent distributions, one per player, which applies to our mixed strategies.

A *2-player zero sum game* is a zero sum game with $N = 2$. In this case, 5.1 may be written as

$$\forall a_1 \in A_1, a_2 \in A_2, \quad u_1(a_1, a_2) = -u_2(a_1, a_2) \quad (5.3)$$

Such a game is completely competitive. There is no motivation for cooperation between the players.

A two person zero sum game may also be described by a single function $\pi : A_1 \times A_2 \rightarrow R$ describing the payoff value for player I, or the loss value for player II. The goal of player I is to maximize π , while the goal of player II is to minimize π . We say that $\pi(i, j)$ is the *value of the game for strategies i and j* or simply the *payoff for i and j* .

Given a certain ordering of the pure strategies of both players, we can also represent a finite 2-player zero sum game using a real matrix $A_{m \times n}$ (the *payoff matrix*), where m is the number of pure strategies for player I and n is the number of pure strategies for player II. The element a_{ij} in the i th row and j th column of A is the payoff (for player I) assuming player I chooses his i th strategy and player II chooses his j th strategy.

When a 2-player zero sum game is represented as a matrix A , a deterministic Nash equilibrium for the game is a saddle point of A , or a pair of strategies i, j so that

$$a_{ij} = \max_k a_{kj}$$

$$a_{ij} = \min_l a_{il}$$

However, such an equilibrium, does not necessarily exist (e.g. consider "Paper, scissors, rock").

5.2 Nash Equilibria

The Nash equilibria of a 2-player zero sum game have several interesting properties. First, they all exhibit the same value. Second, they are *interchangeable*, meaning that given 2 Nash equilibrium points, it is possible to replace a strategy for one of the players in the first point by the strategy of the same player in the second point and obtain another Nash equilibrium. Formally:

Theorem 5.2 *Let G be a 2-player zero sum game defined by $\langle (A_1, A_2), \pi \rangle$. Let (τ_1, τ_2) and (σ_1, σ_2) be two Nash equilibria for G . Then*

1. Both (σ_1, τ_2) and (τ_1, σ_2) are Nash equilibria of G .
2. $\pi(\tau_1, \tau_2) = \pi(\tau_1, \sigma_2) = \pi(\sigma_1, \tau_2) = \pi(\sigma_1, \sigma_2)$

Proof: (σ_1, σ_2) is a Nash equilibrium. Therefore, for the first player (who plays to maximize π), we have

$$\pi(\sigma_1, \sigma_2) \geq \pi(\tau_1, \sigma_2)$$

However, (τ_1, τ_2) is a Nash equilibrium as well. Therefore, for the second player (who plays to minimize π) we have

$$\pi(\tau_1, \sigma_2) \geq \pi(\tau_1, \tau_2)$$

Combining these two inequalities we get

$$\pi(\sigma_1, \sigma_2) \geq \pi(\tau_1, \sigma_2) \geq \pi(\tau_1, \tau_2)$$

Similarly,

$$\pi(\sigma_1, \sigma_2) \leq \pi(\sigma_1, \tau_2) \leq \pi(\tau_1, \tau_2)$$

From the last two inequalities we obtain

$$\pi(\sigma_1, \sigma_2) = \pi(\tau_1, \tau_2) = \pi(\sigma_1, \tau_2) = \pi(\tau_1, \sigma_2)$$

Which proves part 2 of the theorem. To prove part 1 we observe that because (σ_1, σ_2) is a Nash equilibrium for player I,

$$\forall \alpha'_1 \in A_1, \quad \pi(\alpha'_1, \sigma_2) \leq \pi(\sigma_1, \sigma_2) = \pi(\tau_1, \sigma_2)$$

Where the right-hand equation is due to part 2 of the theorem which has already been proven. Similarly, because (τ_1, τ_2) is a Nash equilibrium for player II,

$$\forall \alpha'_2 \in A_2, \quad \pi(\tau_1, \alpha'_2) \geq \pi(\tau_1, \tau_2) = \pi(\tau_1, \sigma_2)$$

Which means that (τ_1, σ_2) is a Nash equilibrium as well. The proof is similar for (σ_1, τ_2) . \square

Theorem 5.2 holds with the same proof for both the deterministic and the nondeterministic case.

We define the *equilibrium strategies* of a player as the set of all strategies played by the player in any equilibrium point. For player I, this is given by

$$\{\sigma_1 \in A_1 \mid \exists \sigma_2 \in A_2, (\sigma_1, \sigma_2) \text{ is an eq. pt.}\}$$

Corollary 5.3 *The set of Nash equilibrium points of a 2-player zero sum game is the Cartesian product of the equilibrium strategies of each player.*

5.3 Payoff Bounds

For the case of pure strategies, player I can guarantee a payoff lower bound by choosing a pure strategy for which the minimal payoff is maximized. This assumes player II is able to know player I's choice and will play the worst possible strategy for player I (note that in a 2-player zero sum game this is also player II's best response to player I's chosen strategy).

We denote this "gain-floor" by V'_I :

$$V'_I = \max_i \min_j a_{ij}$$

Similarly, player II can guarantee a loss upper bound by choosing the pure strategy for which the maximal payoff is minimal. We denote this "loss-ceiling" by V'_{II} :

$$V'_{II} = \min_j \max_i a_{ij}$$

Lemma 5.4 *For any function $F : X \times Y \rightarrow R$, for which all the relevant minima and maxima exist:*

1. $\max_{x \in X} \min_{y \in Y} F(x, y) \leq \min_{y \in Y} \max_{x \in X} F(x, y)$
2. Equality holds iff:

$$\exists x_0 \in X, y_0 \in Y, F(x_0, y_0) = \min_{y \in Y} F(x_0, y) = \max_{x \in X} F(x, y_0)$$

Proof: The proof of this lemma is trivial and is not shown here. □

Applying Lemma 5.4 to our case proves the intuitive fact that player I's gain-floor cannot be greater than player II's loss-ceiling,

$$V'_I \leq V'_{II}$$

and that equality holds iff we have a saddle point and thus an equilibrium.

5.4 Mixed Strategies

For a finite 2-player zero sum game denoted as a matrix $A_{m \times n}$, we denote a mixed strategy for a player I (II) by a stochastic vector of length m (n), where the i th element in the vector is the probability for the i th pure strategy of this player (using the same order used to generate the payoff matrix).

Vectors in this text are always row vectors. We will typically use x for player I mixed strategies, and y for player II mixed strategies. We shall denote by Δ_d the set of stochastic vectors in R^d .

For a 2-player zero sum game given by matrix $A_{m \times n}$, and given mixed strategies x for player I and y for player II, the expected payoff is given by

$$A(x, y) = \sum_{i=1}^m \sum_{j=1}^n x_i a_{ij} y_j = xAy^T \quad (5.4)$$

Once again, if player I chose strategy x , the minimum gain (which is also player II's best response loss) is

$$v_{II}(x) = \min_{y \in \Delta_n} xAy^T, \quad (5.5)$$

which is player II best response to what player I strategy. It is easily shown that this minimum must be reachable in at least one pure strategy of player II.

Lemma 5.5

$$\forall x \in \Delta_m, \quad v_{II}(x) = \min_{y \in \Delta_n} xAy^T = \min_{1 \leq j \leq n} xAe_j^T$$

Proof: The proof is trivial given the fact that xAy^T is a convex combination of xAe_j^T , so xAy^T can never be less than all of xAe_j^T , and on the other hand, e_j is also in Δ_n , so $v_{II}(x) \leq xAe_j^T$. □

Therefore we can write 5.5 as

$$v_{II}(x) = \min_{1 \leq j \leq n} xAe_j^T = \min_{1 \leq j \leq n} \sum_{i=1}^m x_i a_{ij} \quad (5.6)$$

Which means that player I can guarantee the following lower bound on his payoff (gain-floor)

$$V_I = \max_{x \in \Delta_m} \min_{y \in \Delta_n} xAy^T = \max_{x \in \Delta_m} \min_{1 \leq j \leq n} xAe_j^T = \max_{x \in \Delta_m} \min_{1 \leq j \leq n} \sum_{i=1}^m x_i a_{ij} \quad (5.7)$$

Such a mixed strategy x that maximizes $v_{II}(x)$ is a *maximin* strategy for player I. Once again, this maximum exists due to compactness and continuity.

We define $v_I(y)$ in a similar fashion as player I's most harmful response (to player II) to strategy y of player II (this is also player I's best response to y). Then, player II can guarantee the following upper bound on his loss (loss-ceiling)

$$V_{II} = \min_{y \in \Delta_n} \max_{x \in \Delta_m} xAy^T = \min_{y \in \Delta_n} \max_{1 \leq i \leq m} e_i Ay^T = \min_{y \in \Delta_n} \max_{1 \leq i \leq m} \sum_{j=1}^n y_j a_{ij} \quad (5.8)$$

Such a mixed strategy y that maximizes $v_I(y)$ is a *minimax* strategy for player II.

V_I and V_{II} are called the *values* of the game for players I and II, respectively.

5.5 The Minimax Theorem

Applying Lemma 5.4 to the maximin and minimax values of the game we obtain

$$V_I \leq V_{II} \quad (5.9)$$

We will show the following fundamental property of 2-player zero sum games

Theorem 5.6 (*The Minimax Theorem*)

$$V_I = V_{II}$$

The proof outline:

- We shall first prove the following lemma:

Lemma 5.7 (*Supporting Hyperplane Theorem*) *Let $B \subseteq \mathbb{R}^d$ be a closed convex set and $\vec{x} \notin B$ then $\vec{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_d)$ and α_{d+1} exist such that*

$$\vec{\alpha} \cdot \vec{x} = \sum_{i=1}^d \alpha_i x_i = \alpha_{d+1} \quad (5.10)$$

$$\forall y \in B, \quad \vec{\alpha} \cdot \vec{y} = \sum_{i=1}^d \alpha_i y_i > \alpha_{d+1} \quad (5.11)$$

In other words, given a convex closed set B and a point outside the set \vec{x} , the lemma claims that we can pass a hyperplane through \vec{x} such that B lies entirely on one side of the hyperplane. This lemma and its proof are schematically shown in figure 5.1.

- Using Lemma 5.7, we shall prove the following claim:

Claim 5.8 *For every matrix M (representing a 2 player zero-sum game) one of the following must be true:*

- $\exists y, My^t \leq 0$ (Where y is a stochastic vector).
- $\exists x, xM > 0$ (and x is a stochastic vector).

- Finally, using this claim we can easily prove the minimax theorem, as follows:

Proof of the minimax theorem:

According to Claim 5.8 one of two conditions must be true:

- If the first item is true then $\exists y, My^t \leq 0$. Since y is stochastic, it can be viewed as a strategy for player 2, which guarantees: $V_I(y) \leq 0$ and hence $V_{II} \leq 0$. Recall that $V_I \leq V_{II}$ meaning that $V_I \leq V_{II} \leq 0$.
- If the second item is true then $\exists x, xM > 0$ from which we conclude that $0 < V_I \leq V_{II}$.

From the above we conclude that it is impossible to have $V_I < 0 < V_{II}$.

Now assume that V_I, V_{II} are the appropriate values for matrix M . Consider the matrix obtained by subtracting $(V_I + V_{II})/2$ from each coordinate in the matrix M , denoted by \bar{M} . Thus we have

$$\forall 1 \leq i \leq m, \forall 1 \leq j \leq n, \bar{m}_{ij} = m_{ij} - (V_I + V_{II})/2.$$

Denote the relevant values for \bar{M} by \bar{V}_I and \bar{V}_{II} . It is easy to see that the strategies for players I and II have not changed, since the utility function of the game was only changed by a constant. Therefore:

- $\bar{V}_I = (V_I - V_{II})/2$
- $\bar{V}_{II} = (V_{II} - V_I)/2$

Since we have shown that $V_I \leq V_{II}$ we immediately conclude that if $V_I \neq V_{II}$ then $\bar{V}_I < 0 < \bar{V}_{II}$. Since we have already established that this is impossible, we conclude that $V_I = V_{II}$ □

Proof of the supporting hyperplane theorem: Let $\vec{z} \in B$ be the point in B nearest to \vec{x} . Such a point exists because B is closed, and the distance function is both continuous and bounded from below by 0. We define

$$\begin{aligned} \vec{\alpha} &= \vec{z} - \vec{x} \\ \alpha_{d+1} &= \vec{\alpha} \cdot \vec{x} \end{aligned}$$

Identity 5.10 holds immediately. We shall prove 5.11. Note that $\vec{\alpha} \neq 0$ because $\vec{z} \in B$ and $\vec{x} \notin B$. Thus,

$$\vec{\alpha} \cdot \vec{z} - \alpha_{d+1} = \vec{\alpha} \cdot \vec{z} - \vec{\alpha} \cdot \vec{x} = \vec{\alpha} \cdot (\vec{z} - \vec{x}) = \vec{\alpha} \cdot \vec{\alpha} = |\vec{\alpha}|^2 > 0$$

Therefore,

$$\vec{\alpha} \cdot \vec{z} > \alpha_{d+1}$$




Figure 5.1: Supporting Hyperplane

Now, assume that there exists $\vec{y} \in B$ such that

$$\vec{a} \cdot \vec{y} \leq \alpha_{d+1}$$

As B is convex, for any $0 \leq \lambda \leq 1$,

$$\vec{w}_\lambda = \lambda \vec{y} + (1 - \lambda) \vec{z} \in B$$

The square of the distance between \vec{x} and \vec{w}_λ is given by

$$D^2(\vec{x}, \vec{w}_\lambda) = |\vec{x} - \lambda \vec{y} - (1 - \lambda) \vec{z}|^2 = \sum_{i=1}^d (x_i - \lambda y_i - (1 - \lambda) z_i)^2$$

Deriving by λ we obtain

$$\frac{\partial D^2}{\partial \lambda} = 2(\vec{x} - \lambda \vec{y} - (1 - \lambda) \vec{z}) \cdot (\vec{z} - \vec{y})$$

$$\begin{aligned}
&= 2(\vec{z} - \vec{x}) \cdot \vec{y} - 2(\vec{z} - \vec{x}) \cdot \vec{z} + 2\lambda(\vec{z} - \vec{y})^2 \\
&= 2\vec{\alpha} \cdot \vec{y} - 2\vec{\alpha} \cdot \vec{z} + 2\lambda(\vec{z} - \vec{y})^2
\end{aligned}$$

Evaluating at $\lambda = 0$ we get

$$\frac{\partial D^2}{\partial \lambda} = 2\vec{\alpha} \cdot \vec{y} - 2\vec{\alpha} \cdot \vec{z}$$

But according to our assumption the first term $\vec{\alpha} \cdot \vec{y} \leq \alpha_{d+1}$ and we showed that the second term $\vec{\alpha} \cdot \vec{z} > \alpha_{d+1}$, and therefore

$$\left. \frac{\partial D^2}{\partial \lambda} \right|_{\lambda=0} < 0$$

Hence, for λ close enough to 0 we must have

$$D^2(\vec{x}, \vec{w}_\lambda) < D^2(\vec{x}, \vec{z})$$

But \vec{z} was chosen to minimize the distance to \vec{x} , so we have a contradiction. Therefore for all $\vec{y} \in B$, and inequality 5.11 must hold. □

Theorem of the Alternative for Matrices: Recall that we wish to prove Claim 5.8. That is: For a matrix M , one of the following must hold:

- $\exists y, My^t \leq 0$ (Where y is a stochastic vector).
- $\exists x, xM > 0$ (and x is a stochastic vector).

Let $A = (a_{ij})$ be an $m \times n$ real matrix. Let $\{\vec{a}_i\}_{i=1}^m = (a_{i1}, a_{i2}, \dots, a_{in})$ be the rows of the matrix.

1. Suppose that the point $\vec{0}$ in R^n is in the convex hull of the $m+n$ points $\{\vec{a}_i\}_{i=1}^m \cup \{\vec{e}_i\}_{i=1}^n$ where \vec{e}_i is the i th elementary vector in R^n . Then there exists a stochastic vector v of size $m+n$, representing a linear combination of the above points, that produces $\vec{0}$. Denote the vectors y and y' to be the vectors composed of the first m elements in v and the last n elements, respectively. Formally, we have:

$$Ay^t + Iy'^t = \vec{0}$$

and therefore:

$$Ay^t = -Iy'^t$$

Notice that y cannot be 0, because then y' is also zero and v is known to be nonzero. Additionally, $y' \geq 0$ and therefore $-Iy'^t \leq 0$, and

$$Ay \geq 0$$

Since y is not zero, and has only non-negative elements (v is stochastic), y can be normalized to a stochastic vector s.t.

$$My^t > 0$$

2. Alternatively, suppose 1 does not hold. Denote the convex hull mentioned in 1 by C . By Lemma 5.7, there exist $\vec{\alpha} \in R^n$ and α_{n+1} such that

$$\vec{\alpha} \cdot \vec{0} = \alpha_{n+1}$$

(which means that $\alpha_{n+1} = 0$, of course) and

$$\forall \vec{y} \in C, \quad \vec{\alpha} \cdot \vec{y} > 0$$

In particular, this will hold if \vec{y} is any of the vectors \vec{a}_i or \vec{e}_i . Thus

$$\begin{aligned} \vec{a}_i \cdot \vec{\alpha} &> 0 && \text{for all } 1 \leq i \leq m, \\ \alpha_j &> 0 && \text{for all } 1 \leq j \leq n. \end{aligned}$$

Since $\forall 1 \leq j \leq n, \alpha_j > 0$ we have $\sum_{j=1}^n \alpha_j > 0$, so we can scale by the sum and define

$$x_j = \alpha_j \left/ \sum_{j=1}^n \alpha_j \right.$$

Therefore

$$\begin{aligned} \sum_{j=1}^n x_j &= 1 \\ \forall 1 \leq j \leq n, \quad x_j &> 0 \\ \forall 1 \leq i \leq m, \quad \vec{a}_i \cdot \vec{x} = \sum_{j=1}^n a_{ij} x_j &> 0 \end{aligned}$$

In other words, there exists a stochastic vector $\vec{x} = (x_1, \dots, x_n) \in R^n$ s.t. $xM > 0$.

□

5.6 Results

We have shown that in a 2-player zero sum game the gain-ceiling for player I is equal to the loss-floor for player II. We denote this value simply by V and call it the *value of the game*.

Part 2 of Lemma 5.4 tells us that $V_I = V_{II}$ means that we have a Nash equilibrium point. It is easy to see that the payoff in this equilibrium is exactly the value of the game. Theorem

5.2 tells us that all Nash equilibria will have this value, and that the set of all Nash equilibria is actually a cartesian product of the equilibrium strategies of each player.

A strategy x for player I satisfying

$$\forall 1 \leq j \leq n, \quad \sum_{i=1}^m x_i a_{ij} \geq V \quad (5.12)$$

is optimal for player I in the sense that this strategy guarantees a payoff of V against every strategy of player II, and there is no strategy that guarantees a higher payoff against every strategy of player II. Similarly, a strategy y for player II satisfying

$$\forall 1 \leq i \leq m, \quad \sum_{j=1}^n y_j a_{ij} \leq V \quad (5.13)$$

is optimal for player II. It is clear that

$$xAy^T = V$$

otherwise one of 5.12 or 5.13 will not hold. It is easy to see that (x, y) is a Nash equilibrium. Also, any Nash equilibrium must satisfy 5.12 and 5.13.

To summarize

Theorem 5.9 *Let G be a 2-player zero sum game. Then*


1. *The gain-floor for player I and loss-ceiling for player II are equal (the value of the game, V).*
2. *There is at least one mixed Nash equilibrium.*
3. *The set of equilibria for the game is the Cartesian product of the sets of equilibrium strategies for each player.*
4. *The value of the game in all equilibria is V .*

5.7 Infinite 2-Players Zero Sum Games

We have just proven that $V_I = V_{II}$ for finite 2-players zero sum games. The following example shows that in the infinite case this result does not hold.

Consider the following game (described schematically by the illustration above):

- The action of player 1 is, $A_1 = \{x | x \in [0, 1]\}$.
- The action of player 2 is, $A_2 = \{y | y \in [0, 1]\}$



- The utility - $K(x, y) = u_1(x, y) = -u_2(x, y)$, where:
 - $K(x, y) = 1$, if $y < x$ or $y > x + 1/2$.
 - $K(x, y) = 0$, if $y = x$ or $y = x + 1/2$.
 - $K(x, y) = -1$, if $(x + 1/2) > y > x$.

We will show:

1. For any mixed strategy for player 1, x , player 2 has a pure strategy y that guarantees a maximum loss of $1/3$. Hence player 1 cannot guarantee more than $1/3$. Hence $V_I \leq 1/3$.
2. For any mixed strategy for player 2, y , player 1 has a pure strategy x that guarantees a minimum payoff of $3/7$. Hence $V_{II} \geq 3/7$.

Therefore, $V_{II} \geq 3/7 > 1/3 \geq V_I$, which is in contrast to Theorem 5.6.

Proof:

To prove (1), we split the analysis into two cases:

- $Pr[x \in [0, 1/2]] > 1/3$, then y can always choose $1/2 - \epsilon$ for arbitrary small $\epsilon > 0$. Then $E[K(x, y)] \leq 2/3 - 1/3$ (at least $1/3$ of the cases will result the score -1 , at most $2/3$ of the cases will result in the score 1)
- $Pr[x \in [0, 1/2]] \leq 1/3$ then y can always choose 1 , hence $E[K(x, y)] \leq 1/3$ (at most $1/3$ of the cases the score is 1)

To prove (2), we analyze three cases:

- $Pr[y = 1] \leq 4/7$ then x can always choose 1 , hence $E[K(x, y)] \geq 3/7$ (at least 3/7 of the cases the score is 1, otherwise the score is 0)
- $Pr[y = 1] > 4/7$ and $Pr[y \in (0, 1/2)] \leq 1/7$ then x can always choose 0 , hence $E[K(x, y)] \geq 4/7 - 1/7$. (at least 4/7 of the cases the score is 1, at most 1/7 of the cases the score is -1)
- $Pr[y = 1] > 4/7$ and $Pr[y \in (0, 1/2)] > 1/7$ then x can always choose $1/2 - \epsilon$, hence $E[K(x, y)] \geq 3/7$ ($K(1/2 - \epsilon, 1) = 1$, also if $y \in (0, 1/2)$ then $K(1/2 - \epsilon, y) = 1$ at least 5/7 of the cases the score is 1, at most 2/7 of the cases the score is -1)

□

5.8 Application of Zero Sum Games in Computer Science

5.8.1 Deterministic vs. Random Algorithms

In this example $\Omega = \{A_i\}$ is a finite set of deterministic algorithms that can take as input any element of the finite input set $\Lambda = \{x_j\}$. We will denote by Δ_S the set of probability distributions over the set S , for any set S .

Definition $\mathbf{Time(A, x)}$ is the time complexity (measured, as usual in complexity, in the means of number of commands) of running the deterministic algorithm A with the input x . Also denoted $T(A, x)$.

Definition A **Random Algorithm** is a probability distribution over the deterministic algorithms, $\vec{p} \in \Delta_\Omega$. We denote the probability for algorithm A_i by p_i .

Definition $\mathbf{RTime}(\vec{p}, x)$ is time complexity of the random algorithm defined by distribution \vec{p} for fixed input x . It is defined as the expected deterministic time complexity for the fixed input x :

$$RTime(\vec{p}, x) = \sum_i p_i \cdot T(A_i, x)$$

Definition $\mathbf{AvgTime(A, \vec{q})}$ is the time complexity of deterministic algorithm A given distribution \vec{q} over inputs. This is in essence an average-case complexity analysis for A .

It is defined as the expected time complexity for the deterministic algorithm A with input distributed according to \vec{q} :

$$\text{AvgTime}(A, \vec{q}) = \sum_j q_j \cdot T(A, x_j)$$

Complexity Analysis

Corollary 5.10 *Deterministic worst-case time complexity is $\min_i \max_j T(A_i, x_j)$.*

Proof: The complexity of the problem is the minimum complexity over all relevant algorithms (Ω). We must choose the deterministic algorithm before knowing the input. Thus, the complexity of deterministic algorithm A_i is analyzed for the worst input, which yields complexity $\max_j T(A, x_j)$, and then the complexity of the problem is the complexity of the best algorithm, which results in complexity $\min_i \max_j T(A_i, x_j)$. \square

Corollary 5.11 *Non-deterministic worst-case time complexity is $\max_j \min_i T(A_i, x_j)$*

Proof: For non-deterministic algorithms we can guess the best deterministic algorithm given the input. Thus, for input x_j , the complexity is $\min_i T(A_i, x_j)$. We now analyze for the worst case input, which yields complexity $\max_j \min_i T(A_i, x_j)$. \square

Corollary 5.12 *Random worst-case time complexity is*

$$RTime(\Omega) = \min_{\vec{p} \in \Delta_\Omega} \max_j RTime(\vec{p}, x_j)$$

Theorem 5.13 (Yao's Lemma) *For any distribution \vec{p} on the algorithms and \vec{q} on inputs*

$$\max_j E_{i \sim \vec{p}} [T(A_i, x_j)] \geq \min_i E_{j \sim \vec{q}} [T(A_i, x_j)]$$

Proof: We can view the complexity analysis as a 2-player zero sum game in the following way. The max player pure strategies are the possible inputs, Λ . The min player pure strategies are the deterministic algorithms Ω . The payoff is the time complexity $T(A_i, x_j)$.

Given such a game, the Minimax Theorem states,

$$\min_{\vec{p} \in \Delta_\Omega} \max_{\vec{q} \in \Delta_\Lambda} E_{i \sim \vec{p}, j \sim \vec{q}} [T(A_i, x_j)] = \max_{\vec{q} \in \Delta_\Lambda} \min_{\vec{p} \in \Delta_\Omega} E_{i \sim \vec{p}, j \sim \vec{q}} [T(A_i, x_j)] \quad (5.14)$$

As in the previous game analysis, it is easily shown that the internal maximum and minimum are obtained in deterministic points:

$$\max_{\vec{q} \in \Delta_\Lambda} E_{i \sim \vec{p}, j \sim \vec{q}} [T(A_i, x_j)] = \max_j E_{i \sim \vec{p}} [T(A_i, x_j)] \quad (5.15)$$

$$\min_{\vec{p} \in \Delta_\Omega} E_{i \sim \vec{p}, j \sim \vec{q}} [T(A_i, x_j)] = \min_i E_{j \sim \vec{q}} [T(A_i, x_j)] \quad (5.16)$$

Using 5.14, and substituting using 5.15 and 5.16 we obtain

$$\min_{\vec{p} \in \Delta_\Omega} \max_j E_{i \sim \vec{p}} [T(A_i, x_j)] = \max_{\vec{q} \in \Delta_\Lambda} \min_i E_{j \sim \vec{q}} [T(A_i, x_j)] \quad (5.17)$$

Hence for any $\vec{p} \in \Delta_\Omega$

$$\max_j E_{i \sim \vec{p}} [T(A_i, x_j)] = \max_{\vec{q} \in \Delta_\Lambda} \min_i E_{j \sim \vec{q}} [T(A_i, x_j)]$$

Thus for any $\vec{p} \in \Delta_\Omega$ and $\vec{q} \in \Delta_\Lambda$

$$\max_j E_{i \sim \vec{p}} [T(A_i, x_j)] = \min_i E_{j \sim \vec{q}} [T(A_i, x_j)]$$

□

Note that Yao's Lemma is actually a result of the weaker inequality established in Lemma 5.4.

Corollary 5.14 *In order to prove a lower bound for the worst-case expected time complexity of any random algorithm for a given problem, it is sufficient to prove a lower bound for any deterministic algorithm on some distribution over the input.*

Proof: Using

$$\begin{aligned} E_{i \sim \vec{p}} [T(A_i, x_j)] &= RTime(\vec{p}, x_j) \\ E_{j \sim \vec{q}} [T(A_i, x_j)] &= AvgTime(A_i, \vec{q}) \end{aligned}$$

we can write Yao's Lemma as

$$\max_j RTime(\vec{p}, x_j) \geq \min_i AvgTime(A_i, \vec{q})$$

So given a lower bound B on the complexity of any deterministic algorithm on some input distribution \vec{q} , we obtain

$$B \leq \min_i AvgTime(A_i, \vec{q}) \leq \max_j RTime(\vec{p}, x_j)$$

So B is a lower bound on the worst-case complexity of any random algorithm. □

Example - Sorting a List of Numbers

We wish to lower-bound the complexity of a random algorithm for sorting n numbers (comparison based sort). We can describe any deterministic comparison based sort algorithm as a decision tree, where each internal node corresponds to a comparison the algorithm performs, with 2 possible outcomes (we assume all elements are different). For a specific input, the execution of the algorithm corresponds to a path from the root to a leaf. It is impossible for 2 different permutations to result in the same path. The running time for the algorithm over an input is the length of the path.

Therefore, the decision tree must have at least $n!$ leaves. Thus the depth of the tree is at least $\log(n!) = O(n \log n)$ nodes. The number of leaves whose depth is not greater than l is $\leq 2^{l+1}$.

Thus, for any deterministic algorithm A , at least one half of the permutations are in depth greater than l , where $l + 1 = \log(n!/2)$ (since then the number of leaves whose depth is less than l is $\leq 2^{\log(n!/2)} = n!/2$). $l + 1 = \log(n!/2) \implies l = \log(n!) - 2 = O(n \log n)$.

We shall choose a uniform distribution \vec{q} over the possible inputs (all permutations of n numbers), and fix a deterministic algorithm A . The running time of A over this distribution is simply the average of the depths of the leaves for all possible inputs. But at least $n!/2$ inputs are of depth at least $\log(n!) - 1$, so the average running time will be at least

$$\frac{\frac{n!}{2} \cdot (\log(n!) - 1)}{n!} = \Omega(n \log n)$$

And using Yao's lemma, the complexity of any random algorithm is also $\Omega(n \log n)$.

5.8.2 Weak vs. Strong Learning

Given a weak learner for a binary classification problem we will show that strong learning is possible.

The model: f is the target function, H a function family.

$$f : X \longrightarrow \{0, 1\}$$

$$\forall h \in H, \quad h : X \longrightarrow \{0, 1\}$$

X is finite, and as a consequence H is finite ($|H| \leq 2^{|X|}$)

The ϵ -WL (weak learning) assumption: For every distribution D on X there exists $h \in H$ and $\epsilon > 0$ such that

$$[Pr_D[h(x) = f(x)] \geq 1/2 + \epsilon] \tag{5.18}$$

Question: How will f be approximated by functions in H ?

We represent the problem as a 2-player zero sum game as follows. The max player pure strategies are the inputs X . The min player pure strategies are the functions H . The payoff is an error indicator:

$$M(h, x) = \begin{cases} 0 & \text{if } f(x) = h(x) \text{ (no error)} \\ 1 & \text{if } f(x) \neq h(x) \text{ (error)} \end{cases}$$

Note that $M(h, x) = |f(x) - h(x)|$. The max player is trying to select a distribution over X that will maximize the expected error, while the min player is trying to select a distribution over H that will minimize it.

The WL proposition implies that for every D there exists a hypothesis and $h \in H$ such that

$$1/2 - \epsilon \geq Pr_D[h(x) \neq f(x)] = Pr_D[M(h, x) = 1] = E_{x \sim D}[M(h, x)]$$

Thus

$$\min_h E_{x \sim D}[M(h, x)] \leq 1/2 - \epsilon.$$

This means that

$$V_{X \text{ Player}} \leq 1/2 - \epsilon.$$

Since in a zero sum game, the values of both players are equal, we conclude that

$$V_{H \text{ Player}} < 1/2 - \epsilon.$$

Hence

$$\min_q \max_x E_{h \sim q}[M(h, x)] \leq 1/2 - \epsilon.$$

Therefore there exists a distribution q (the one which attains the minimum) such that

$$\max_x E_{h \sim q}[M(h, x)] \leq 1/2 - \epsilon.$$

Thus

$$\forall x \in X, \quad E_{h \sim q}[M(h, x)] \leq 1/2 - \epsilon.$$

In other words, for this q , for every $x \in X$,

$$1/2 - \epsilon \geq \sum_{h \in H} q(h) M(h, x) = \sum_{h \in H} q(h) |f(x) - h(x)|$$

We define an approximation $G(x) = \sum_{h \in H} q(h) \cdot h(x)$. Now, for all $x \in X$,

$$|f(x) - G(x)| = \left| \sum_{h \in H} q(h) [f(x) - h(x)] \right| \leq \sum_{h \in H} q(h) |f(x) - h(x)| < 1/2.$$

So by rounding $G(x)$ to either 0 or 1 we obtain $f(x)$, for all $x \in X$.

5.8.3 Correlated Equilibrium Existence

Consider a general strategic game of the form $\langle [N], (A_i)_{i=1}^N, (u_i)_{i=1}^N \rangle$. Here N denotes the number of players, A_i denotes the set of strategies allowed for player $i \in [N]$ and $u_i : \times_{i=1}^N A_i \rightarrow R$ denotes the utility function of player $i \in [N]$.

Recall, that *correlated equilibrium* (CE) in strategic game is a random distribution Q over $\times_{i=1}^N A_i$ such that:

$$\forall \alpha \in A_i, \alpha \in BR(Q|_{a_i=\alpha})$$

Equivalently,

$$\forall \alpha \in A_i, E[u_i(a) - u_i(a^{-i}, \beta) \mid a_i = \alpha] \geq 0$$

The expectation is taken w.r.t. Q .

Theorem 5.15 *Every finite strategic game has at least one correlated equilibrium.*

Proof: Let G be a finite strategic game $\langle [N], (A_i)_{i=1}^N, (u_i)_{i=1}^N \rangle$. Define the following zero-sum game, played by two players: the row-player and the column player. Assume that the row player's objective is to maximize the utility value and the column player's objective is to minimize.

Strategies of the row player are from $A = \times_{i=1}^N A_i$, namely the set of all joint strategies in the original game. Strategies of the column player are all possible triplets (i, b_1, b_2) , such that $b_1, b_2 \in A_i$.

The utility function is given as a matrix M , where the rows correspond to the strategies of the row-player and the columns correspond to the strategies of column player. M is defined in the following way:

$$M(a, (i, b_1, b_2)) = u_i(a) - u_i(a^{-i}, a')$$

where

$$a'_i = \begin{cases} b_1 & \text{if } a_i = b_2 \\ a_i & \text{otherwise} \end{cases}$$

In terms of the original strategic game $\langle [N], (A_i)_{i=1}^N, (u_i)_{i=1}^N \rangle$, $M(a, (i, b_1, b_2))$ describes the gain of player i from choosing strategy $a_i = b_2$, when the joint action is a , instead of switching to b_1 , assuming the strategies of the rest of the players remain unchanged. Clearly, the existence of a CE for $\langle [N], (A_i)_{i=1}^N, (u_i)_{i=1}^N \rangle$ is equivalent to a non-negative game value.

Observe that the game value is at most 0, since the column player can choose a strategy of the form (i, b, b) . We are going to prove that the game value is exactly 0.

According to the Minimax theorem, it is sufficient to prove that the column player cannot guarantee his loss to be less than 0. Indeed, let D be a mixed strategy for the column player. For all $1 \leq i \leq N$, let D_i denote the stochastic matrix, such that:

$$D_i(b_1, b_2) = Pr(a'_i = b_1 \mid a_i = b_2, (i, *, *) \text{ chosen by column player}).$$

If the probability of choosing $(i, *, *)$ under D is zero, w.l.o.g. assume that $D_i(b_1, b_2) = \frac{1}{|A_i|}$ uniformly.

Denote Π_i the stationary distribution vector of D_i , such that $D_i \Pi_i^T = \Pi_i^T$. Now, define a mixed strategy for the row player, that is distribution Π over A , as the cartesian product of all Π_i , such that each player's strategy a_i is chosen independently from A_i according to Π_i .

Note, that if $1 \leq i \leq N$ fixed, row player chooses $a \sim \Pi$ and column player chooses $(j, b_1, b_2) \sim D_{|(j, b_1, b_2) = (i, *, *)}$, then a_i and a'_i have the same distribution by the definition of Π_i as the stationary distribution of D_i .

The resulting game value is:

$$\sum_{i=1}^N E_{a \sim \Pi, (j, b_1, b_2) \sim D} [U_i(a) - U_i(a^{-i}, a'_i) \mid j = i] \cdot Pr_D((j, b_1, b_2) = (i, *, *))$$

which is zero, because $\forall 1 \leq i \leq N$,

$$\begin{aligned} & E_{a \sim \Pi, (j, b_1, b_2) \sim D} [U_i(a) - U_i(a^{-i}, a'_i) \mid j = i] = \\ & = E_{a \sim \Pi, (j, b_1, b_2) \sim D} [U_i(a) \mid j = i] - E_{a \sim \Pi, (j, b_1, b_2) \sim D} [U_i(a^{-i}, a'_i) \mid j = i] = 0 \end{aligned}$$

since a_i and a'_i have the same distribution. □

Bibliography

- [1] Guillermo Owen, Game Theory, Academic Press, 1992.