COSC-6590/GSCS-6390

Games: Theory and Applications

Lecture 05 - Minimax Theorem

Luis Rodolfo Garcia Carrillo

School of Engineering and Computing Sciences
Texas A&M University - Corpus Christi, USA

Table of contents

Theorem Statement

Theorem Statement

Consider a game specified by an $m \times n$ matrix A.

- m actions for P_1, and n actions for P_2.

$$A = \underbrace{\begin{bmatrix} & \vdots & \\ \cdots & a_{ij} & \cdots \\ & \vdots & \end{bmatrix}}_{P_2 \text{ choices (maximizer)}} \left.\vphantom{\begin{bmatrix} \\ \\ \\ \end{bmatrix}}\right\} P_1 \text{ choices (minimizer)}$$

Theorem 5.1 (Minimax). For every matrix A, the **average security levels** of both players coincide, i.e.,

$$\underline{V}_m(A) := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z =: \bar{V}_m(A)$$

Module 05 is devoted to this result.

Theorem Statement

From Property **P4.3**

$$\underline{V}_m(A) \leq \bar{V}_m(A)$$

If the inequality were strict, there would be a constant such that

$$\underline{V}_m(A) < c < \bar{V}_m(A)$$

Proof of **Theorem 5.1** consists in showing this is not possible.

We will show that for any c, at least one of the players can guarantee a security level of c.

$$c \leq \underline{V}_m(A) \qquad \text{or} \qquad \bar{V}_m(A) \leq c$$

To achieve this we will use a key result in convex analysis.

Convex Hull

Convex Hull

Given k vectors v_1, v_2, ... , $v_k \in \mathbb{R}^n$, the **linear subspace generated by these vectors** is the set

$$\text{span}(v_1, v_2, \ldots, v_k) = \left\{ \sum_{i=1}^{k} \alpha_i v_i : \alpha_i \in \mathbb{R} \right\} \subset \mathbb{R}^n$$

which is represented graphically as

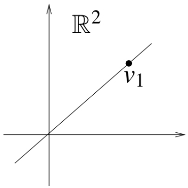

linear subspace of $\mathbb{R}^2$ generated by v_1

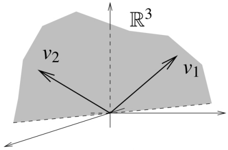

linear subspace of $\mathbb{R}^3$ generated by v_1,v_2

Convex Hull

The **(closed) convex hull generated by these vectors** is the set

$$\operatorname{co}(v_1, v_2, \ldots, v_k) = \left\{ \sum_{i=1}^{k} \alpha_i v_i : \alpha_i \geq 0, \sum_{i=1}^{k} \alpha_i = 1 \right\} \subset \mathbb{R}^n$$

which is represented graphically as

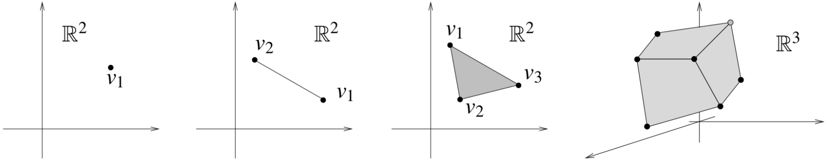

convex hull in $\mathbb{R}^2$ generated by v_1

convex hull in $\mathbb{R}^2$ generated by v_1, v_2

convex hull in $\mathbb{R}^2$ v_1, v_2, v_3

convex hull in $\mathbb{R}^3$

Convex Hull

The convex hull is always a **convex set** in the sense that

$$x_1, x_2 \in \text{co}(v_1, v_2, \ldots, v_k) \Rightarrow \frac{\lambda x_1 + (1-\lambda)x_2}{2} \in \text{co}(v_1, v_2, \ldots, v_k), \quad \forall \lambda \in [0, 1]$$

i.e., if x_1 and x_2 belong to the set, then any point in the line segment between x_1 and x_2 also belongs to the set.

Conclusion: all sets in previous figure are convex.

But this is not the case for the set below (a nonconvex set)

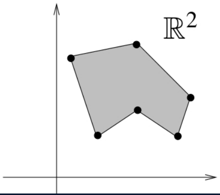

Separating Hyperplane Theorem

Separating Hyperplane Theorem

An **hyperplane** in $\mathbb{R}^n$ is a set of the form

$$\mathcal{P} := \left\{x \in \mathbb{R}^n : v'(x - x_0) = 0\right\}$$

- $x_0 \in \mathbb{R}^n$ is a point that belongs to the hyperplane
- $v \in \mathbb{R}^n$ a vector called the **normal to the hyperplane**.

An **(open) half-space** in $\mathbb{R}^n$ n is a set of the form

$$\mathcal{H} := \left\{x \in \mathbb{R}^n : v'(x - x_0) > 0\right\}$$

- $x_0 \in \mathbb{R}^n$ is a point in the boundary of $\mathcal{H}$
- $v \in \mathbb{R}^n$ is the **inwards-pointing normal** to the half-space.

Each hyperplane partitions the whole space $\mathbb{R}^n$ into two half-spaces with symmetric normals.

Separating Hyperplane Theorem

Theorem 5.2 (Separating Hyperplane). For every convex set $\mathcal{C}$ and a point x_0 not in $\mathcal{C}$, there exists an hyperplane $\mathcal{P}$ that contains x_0 but does not intersect $\mathcal{C}$. Consequently, the set $\mathcal{C}$ is fully contained in one of the half spaces defined by $\mathcal{P}$.

Theorem (a key result in convex analysis) is illustrated below

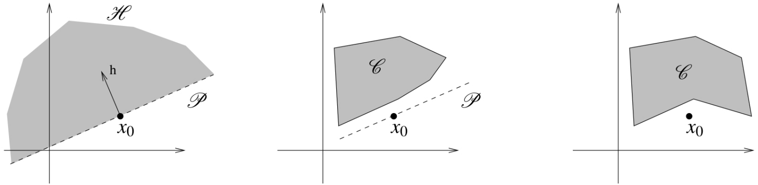

hyperplane and half-space | theorem correctly applied to a convex set | theorem fails for a non-convex set

On the Way to Prove the Minimax Theorem

On the Way to Prove the Minimax Theorem

Prove that for any number c, we either have

$$\underline{V}_m(A) := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z \geq c \qquad \text{or} \qquad \bar{V}_m(A) := \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z \leq c$$

To prove this, show that either there exists a $z^* \in \mathcal{Z}$ such that

$$y' A z^* \geq c, \quad \forall y \in \mathcal{Y} \Rightarrow \min_{y \in \mathcal{Y}} y' A z^* \geq c \Rightarrow \underline{V}_m(A) := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z \geq c$$

or there exists a $y^* \in \mathcal{Y}$ such that

$$y^{*\prime} A z \leq c, \quad \forall z \in \mathcal{Z} \Rightarrow \max_{z \in \mathcal{Z}} y^{*\prime} A z \leq c \Rightarrow \bar{V}_m(A) := \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z \leq c$$

The **Theorem of the Alternative for Matrices**, proves exactly this for the special case $c = 0$.

On the Way to Prove the Minimax Theorem

Theorem 5.3 (Theorem of the Alternative for Matrices).
For every $m \times n$ matrix M, one of the following statements must necessarily hold:

1. there exists some $y^* \in \mathcal{Y}$ such that $y^{*\prime} M z \leq 0, \forall z \in \mathcal{Z}$
2. there exists some $z^* \in \mathcal{Z}$ such that $y' M z^* \geq 0, \forall y \in \mathcal{Y}$

Note. We can regard

- y^* as a policy for P_1 that guarantees an outcome no larger than zero (since it guarantees $\bar{V}_m(A) \leq 0$)
- z^* as a policy for P_2 that guarantees an outcome no smaller than zero (since it guarantees $\underline{V}_m(A) \geq 0$.)

On the Way to Prove the Minimax Theorem

Proof. Consider separately the cases of whether or not the vector 0 belongs to the convex hull $\mathcal{C}$ of the columns of $[-M_{m\times n} \;\; I_m]$ where I_m denotes the identity matrix in $\mathbb{R}^m$

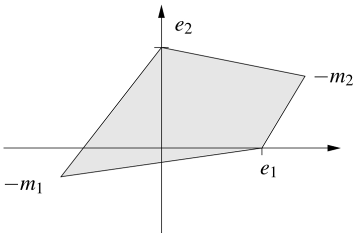

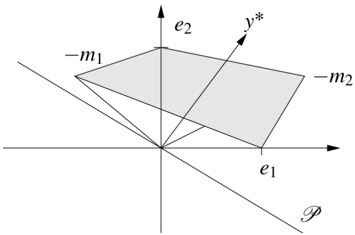

0 in the convex hull $\mathcal{C}$

0 not in the convex hull $\mathcal{C}$, with separating hyperplane $\mathcal{P}$ and inner-pointing normal y^*

On the Way to Prove the Minimax Theorem

Suppose that 0 belongs to the convex hull $\mathcal{C}$ of the columns of $[-M_{m\times n} \;\; I_m]$, therefore there exist scalars $\bar{z}_j$, $\bar{\eta}_j$ such that

$$[-M_{m\times n} \;\; I_m]\begin{bmatrix} \bar{z}_1 \\ \vdots \\ \bar{z}_n \\ \bar{\eta}_1 \\ \vdots \\ \bar{\eta}_m \end{bmatrix} = 0 \qquad \bar{z}_j \geq 0, \;\; \bar{\eta}_j \geq 0, \;\; \sum_j \bar{z}_j + \sum_j \bar{\eta}_j = 1.$$

Note that $\sum_j \bar{z}_j \neq 0$ since otherwise all the $\bar{z}_j$ would have to be exactly equal to zero and then so would all the $\bar{\eta}_j$, because of the left-hand side equality.

On the Way to Prove the Minimax Theorem

Defining

$$z^* := \frac{1}{\sum_j \bar{z}_j} \begin{bmatrix} \bar{z}_1 \\ \vdots \\ \bar{z}_n \end{bmatrix} \qquad \eta^* := \frac{1}{\sum_j \bar{\eta}_j} \begin{bmatrix} \bar{\eta}_1 \\ \vdots \\ \bar{\eta}_m \end{bmatrix}$$

we conclude $z^* \in \mathcal{Z}$ and $Mz^* = \eta^*$.

Then, for every $y \in \mathcal{Y}$

$$y'Mz^* = y'\eta^* \geq 0$$

which shows that **Statement 2** holds

- recall that all entries of y and η^* are non negative.

On the Way to Prove the Minimax Theorem

Suppose the 0 vector does not belong to the convex hull $\mathcal{C}$ of the columns of $[-M_{m\times n} \;\; I_m]$.

Use the **Separating Hyperplane Theorem** to conclude that there must exist an half space $\mathcal{H}$ with 0 in its boundary that fully contains $\mathcal{C}$.

Denoting by y^* the inwards-pointing normal to $\mathcal{H}$, we have

$$\mathcal{H} = \left\{x \in \mathbb{R}^m : y^{*\prime}x > 0\right\} \supset \mathcal{C}$$

Therefore, for every x in the convex hull $\mathcal{C}$ of the columns of $[-M_{m\times n} \;\; I_m]$, we have

$$y^* x > 0$$

On the Way to Prove the Minimax Theorem

We conclude that for every $\bar{z}_j \geq 0, \quad \bar{\eta}_j \geq 0, \quad \sum_j \bar{z}_j + \sum_j \bar{\eta}_j = 1$

$$y^* [-M_{m\times n} \quad I_m] \begin{bmatrix} \bar{z}_1 \\ \vdots \\ \bar{z}_n \\ \bar{\eta}_1 \\ \vdots \\ \bar{\eta}_m \end{bmatrix} > 0$$

In particular, for convex combinations with all the $\eta_j = 0$, we obtain

$$y^* M \bar{z} < 0, \qquad \forall \bar{z} \in \mathcal{Z}$$

On the Way to Prove the Minimax Theorem

On the other hand, from the convex combinations with $\eta_j = 1$ and all other coefficients equal to zero, we conclude that

$$y_j^* > 0, \qquad \forall j$$

In case $\sum_j y_j^* = 1$, then $y^* \in \mathcal{Y}$ (which is the hyperplane normal) provides the desired vector y^* for **Statement 1**.

Otherwise, we simply need to rescale the normal by a positive constant to get $\sum_j y_j^* = 1$.

Note: rescaling by a positive constant does not change the validity of $y^* M \bar{z} < 0$, $\forall \bar{z} \in \mathcal{Z}$.

Proof of the Minimax Theorem

Proof of the Minimax Theorem

Pick c, and show that either there exists a $z^* \in \mathcal{Z}$ such that

$$y'Az^* \geq c, \quad \forall y \in \mathcal{Y} \Rightarrow \min_{y \in \mathcal{Y}} y'Az^* \geq c \Rightarrow \underline{V}_m(A) := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y'Az \geq c$$

or there exists a $y^* \in \mathcal{Y}$ such that

$$y^{*\prime}Az \leq c, \quad \forall z \in \mathcal{Z} \Rightarrow \max_{z \in \mathcal{Z}} y^{*\prime}Az \leq c \Rightarrow \bar{V}_m(A) := \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y'Az \leq c$$

This is achieved by applying **Theorem 5.3** to the matrix

$$M = A - c\mathbf{1}$$

- where $\mathbf{1}$ denotes a $m \times n$ matrix with all entries equal to 1.

Property: $y'\mathbf{1}z = 1$ for every $y \in \mathcal{Y}$ and $z \in \mathcal{Z}$.

Proof of the Minimax Theorem

When **Statement 1** in **Theorem 5.3** holds, there exists some $y^* \in \mathcal{Y}$ such that

$$y^{*\prime}(A - c\mathbf{1})z = y^{*\prime}Az - c \leq 0, \quad \forall z \in \mathcal{Z}$$

and the previous $\bar{V}_m(A)$ equation holds.

When **Statement 2** in **Theorem 5.3** holds, there exists some $z^* \in \mathcal{Z}$ such that

$$y'(A - c\mathbf{1})z^* = y'Az^* - c \geq 0, \quad \forall y \in \mathcal{Y}$$

and the previous $\underline{V}_m(A)$ equation holds.

Proof of the Minimax Theorem

If the Minimax Theorem did not hold, we could pick c such that

$$\max_{z\in\mathcal{Z}}\min_{y\in\mathcal{Y}} y'Az < c < \min_{y\in\mathcal{Y}}\max_{z\in\mathcal{Z}} y'Az$$

which contradicts the equations for $\underline{V}_m(A)$ and $\bar{V}_m(A)$.

Therefore there must **not be a gap** between the $\max_z \min_y$ and the $\min_y \max_z$.

Consequences of the Minimax Theorem

Consequences of the Minimax Theorem

Combining **Theorem 4.1** with the **Minimax Theorem 5.1** we conclude:

Corollary 5.1. Consider a game defined by a matrix A:

P5.1 A mixed saddle-point equilibrium always exist and

$$\underline{V}_m(A) := \max_{z\in\mathcal{Z}} \min_{y\in\mathcal{Y}} y' A z = \min_{y\in\mathcal{Y}} \max_{z\in\mathcal{Z}} y' A z =: \bar{V}_m(A)$$

P5.2 If y^* and z^* are mixed security policies for P_1 and P_2, then (y^*, z^*) is a mixed saddle-point equilibrium and its value $y^{*\prime} A z^*$ is equal to the equation in **P5.1**.

Consequences of the Minimax Theorem

P5.3 If (y^*, z^*) is a mixed saddle-point equilibrium then y^* and z^* are mixed security policies for P_1 and P_2, and the equation in **P5.1** is equal to the mixed saddle-point value $y^{*\prime}Az^*$.

P5.4 If (y_1^*, z_1^*) and (y_2^*, z_2^*) are mixed saddle-point equilibria then (y_1^*, z_2^*) and (y_2^*, z_1^*) are also mixed saddle-point equilibria and

$$y_1^{*\prime}Az_1^* = y_2^{*\prime}Az_2^* = y_1^{*\prime}Az_2^* = y_2^{*\prime}Az_1^*$$

Practice Exercise

Practice Exercise

5.1 (Symmetric games).

A game defined by a matrix A is called symmetric if A is skew symmetric, i.e., if $A' = -A$.
For such games, show that the following statements hold:

- $V_m(A) = 0$
- If y^* is a mixed security policy for P_1 , then y^* is also a security policy for P_2 and vice-versa.
- If (y^*, z^*) is a mixed saddle-point equilibrium then (z^*, y^*) is also a mixed saddle-point equilibrium.

Hint: Make use of the two facts below:

$$\max_x f(x) = -\min_x(-f(x)), \qquad \min_w \max_x f(x) = -\max_w \min_x(-f(x))$$

Note that the Rock-Paper-Scissors game is symmetric.

Practice Exercise

Solution to Exercise 5.1.

1. Denoting by y^* a mixed security policy for P_1, we have that

$$V_m(A) := \min_y \max_z y'Az = \max_z y^{*\prime}Az$$

Since $y'Az$ is a scalar and $A' = -A$, we conclude that

$$y'Az = (y'Az)' = z'A'y = -z'Ay, \quad y^{*\prime}Az = \cdots = -z'Ay^*$$

Using this in $V_m(A)$, we conclude that

$$V_m(A) = \min_y \max_z(-z'Ay) = \max_z(-z'Ay^*)$$

Practice Exercise

We now use the **hint** to obtain

$$V_m(A) = -\max_y \min_z z'Ay = -\min_z z'Ay^*$$

However, in mixed games $\max_y \min_z z'Ay$ is also equal to $V_m(A)$ and therefore we have

$$V_m(A) = -V_m(A) \Rightarrow V_m(A) = 0$$

Practice Exercise

2. On the other hand, since we just saw in using the hint in $V_m(A)$ that

$$\max_y \min_z z' A y = \min_z z' A y^*$$

we have that y^* is indeed a mixed security policy for P_2. To prove the converse, we follow a similar reasoning starting from the previous equation

$$V_m(A) := \min_y \max_z y' A z = \max_z y^{*\prime} A z$$

but with $\max_y \min_z$ instead of $\min_z \max_y$. This results in the proof that if y^* is a mixed security policy for P_2 then it is also a mixed security policy for P_1.

Practice Exercise

3. If (y^*, z^*) is a mixed saddle-point equilibrium then both y^* and z^* are mixed security policies for P_1 and P_2, respectively.

However, from the previous results we conclude that these are also security policies for P_2 and P_1, respectively, which means that (z^*, y^*) is indeed a mixed saddle-point equilibrium.

End of Lecture

05 - Minimax Theorem

Questions?