Algorithmic Game Theory and Applications

Lecture 4: 2-player zero-sum games, and the Minimax Theorem

Kousha Etessami

School of Informatics

2-person zero-sum games

A finite 2-person zero-sum (2p-zs) strategic game Γ, is a strategic game where:

- For players $i \in \{1,2\}$, the *payoff functions* $u_i : S \mapsto \mathbb{R}$ are such that for all $s = (s_1, s_2) \in S$,

$$u_1(s) + u_2(s) = 0$$

I.e., $u_1(s) = -u_2(s)$.

$u_i(s_1, s_2)$ can conveniently be viewed as a $m_1 \times m_2$ payoff matrix A_i, where:

$$A_1 = \begin{bmatrix} u_1(1,1) & \ldots\ldots & u_1(1,m_2) \\ \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots \\ u_1(m_1,1) & \ldots\ldots & u_1(m_1,m_2) \end{bmatrix}$$

Note, $A_2 = -A_1$. Thus we may assume only one function $u(s_1, s_2)$ is given, as one matrix, $A = A_1$. Player 1 wants to maximize $u(s)$, while Player 2 wants to minimize it (i.e., to maximize its negative).

matrices and vectors

As just noted, a 2p-zs game can be described by an $m_1 \times m_2$ matrix:

$$A = \begin{bmatrix} a_{1,1} & \ldots\ldots & a_{1,m_2} \\ \vdots & \vdots & \vdots \\ \vdots & a_{i,j} & \vdots \\ \vdots & \vdots & \vdots \\ a_{m_1,1} & \ldots\ldots & a_{m_1,m_2} \end{bmatrix}$$

where $a_{i,j} = u(i,j)$.

For any $(n_1 \times n_2)$-matrix A we'll either use $a_{i,j}$ or $(A)_{i,j}$ to denote the entry in the i'th row and j'th column of A.

For $(n_1 \times n_2)$ matrices A and B, let

$$A \geq B$$

denote that for all i, j, $a_{i,j} \geq b_{i,j}$.
Let

$$A > B$$

denote that for all i, j, $a_{i,j} > b_{i,j}$.

For a matrix A, let $A \geq 0$ denote that every entry is ≥ 0. Likewise, let $A > 0$ mean every entry is > 0.

more review of matrices and vectors

Recall matrix multiplication: given $(n_1 \times n_2)$-matrix A and $(n_2 \times n_3)$-matrix B, the product AB is an $(n_1 \times n_3)$-matrix C, where

$$c_{i,j} = \sum_{k=1}^{n_2} a_{i,k} * b_{k,j}$$

Fact: matrix multiplication is "associative": i.e.,

$$(AB)C = A(BC)$$

(Note: for the multiplications to be defined, the dimensions of the matrices A, B, and C need to be "consistent": $(n_1 \times n_2)$, $(n_2 \times n_3)$, and $(n_3 \times n_4)$, respectively.)

Fact: For matrices A, B, C, of appropriate dimensions, if $A \geq B$, and $C \geq 0$, then

$AC \geq BC$, and likewise, $CA \geq CB$.

(C's dimensions might be different in each case.)

more on matrices and vectors

For a $(n_1 \times n_2)$ matrix B, let B^T denote the $(n_2 \times n_1)$ **transpose** matrix, where $(B^T)_{i,j} := (B)_{j,i}$.

We can view a column vector, $y = \begin{bmatrix} y(1) \\ \vdots \\ \vdots \\ y(m) \end{bmatrix}$, as a $(m \times 1)$-matrix. Then, y^T would be a $(1 \times m)$-matrix, i.e., a row vector.

Typically, we think of "vectors" as column vectors and explicitly transpose them if we need to. We'll call a length m vector an m-vector.

Multiplying a $(n_1 \times n_2)$-matrix A by a n_2-vector y is just a special case of matrix multiplication:
Ay is a n_1-vector.

Likewise, pre-multiplying A, by a n_1-row vector y^T, is also just a special case of matrix multiplication: $y^T A$ is a n_2-row vector.

For a column (row) vector y, we use $(y)_j$ to denote the entry $(y)_{j,1}$ (respectively, $(y)_{1,j}$).

A matrix view of zero-sum games

Suppose we have a 2p-zs game given by a $(m_1 \times m_2)$-matrix, A.

Suppose Player 1 chooses a mixed strategy x_1, and Player 2 chooses mixed strategy x_2 (assume x_1 and x_2 are given by column vectors). Consider the product

$$x_1^T A x_2$$

If you do the calculation,

$$x_1^T A x_2 = \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} (x_1(i) * x_2(j)) * a_{i,j}$$

But note that $(x_1(i) * x_2(j))$ is precisely the probability of the pure combination $s = (i,j)$. Thus, for the mixed profile $x = (x_1, x_2)$

$$x_1^T A x_2 = U_1(x) = -U_2(x)$$

where $U_1(x)$ is the expected payoff (which Player 1 is trying to maximize, and Player 2 is trying to minimize).

"minmaximizing" strategies

Suppose Player 1 chooses a mixed strategy $x_1^* \in X_1$, by trying to maximize the "worst that can happen". The worst that can happen would be for Player 2 to choose x_2 which minimizes $(x_1^*)^T A x_2$.

Definition: $x_1^* \in X_1$ is a **minmaximizer** for Player 1 if

$$\min_{x_2 \in X_2} (x_1^*)^T A x_2 = \max_{x_1 \in X_1} \min_{x_2 \in X_2} (x_1)^T A x_2$$

Similarly, $x_2^* \in X_2$ is a **maxminimizer** for Player 2 if

$$\max_{x_1 \in X_1} (x_1)^T A x_2^* = \min_{x_2 \in X_2} \max_{x_1 \in X_1} x_1^T A x_2$$

Note that

$$\min_{x_2 \in X_2} (x_1^*)^T A x_2 \leq (x_1^*)^T A x_2^* \leq \max_{x_1 \in X_1} x_1^T A x_2^*$$

Amazingly, von Neumann (1928) showed equality holds!

The Minimax Theorem

Theorem(von Neumann) Let a 2p-zs game Γ be given by an $(m_1 \times m_2)$-matrix A of real numbers. There exists a unique value $v^* \in \mathbb{R}$, such that there exists $x^* = (x_1^*, x_2^*) \in X$ such that

1. $((x_1^*)^T A)_j \geq v^*$, for $j = 1, \ldots, m_2$.

2. $(Ax_2^*)_j \leq v^*$, for $j = 1, \ldots, m_1$.

3. And (thus) $v^* = (x_1^*)^T A x_2^*$ and

$$\max_{x_1 \in X_1} \min_{x_2 \in X_2} (x_1)^T A x_2 = v^* = \min_{x_2 \in X_2} \max_{x_1 \in X_1} x_1^T A x_2$$

4. In fact, the above conditions all hold precisely when $x^* = (x_1^*, x_2^*)$ is any Nash Equilibrium.

   Equivalently, they hold precisely when x_1^* is any minmaximizer and x_2^* is any maxminimizer.

some remarks

Note:

$(1.)$ says x_1^* guarantees Player 1 at least expected profit v^*, and
$(2.)$ says x_2^* guarantees Player 2 at most expected "loss" v^*.

We call any such $x^* = (x_1^*, x_2^*)$ a **minimax profile**.

We call the unique v^* the **minimax value** of game Γ.

It is obvious that the maximum profit that Player 1 can guarantee for itself should be $\leq$ the minimum loss that Player 2 can guarantee for itself, i.e., that

$$\max_{x_1 \in X_1} \min_{x_2 \in X_2} (x_1)^T A x_2 \leq \min_{x_2 \in X_2} \max_{x_1 \in X_1} x_1^T A x_2$$

What is not obvious at all is why these two values should be the same!

Proof of the Minimax Theorem

The Minimax Theorem follows directly from Nash's Theorem (but historically, it predates Nash).

Proof: Let $x^* = (x_1^*, x_2^*) \in X$ be a NE of the 2-player zero-sum game Γ, with matrix A.

Let $v^* := (x_1^*)^T A x_2^* = U_1(x^*) = -U_2(x^*)$.

Since x_1^* and x_2^* are "best responses" to each other, we know that for $i \in \{1, 2\}$

$$U_i(x^*_{-i}; \pi_{i,j}) \leq U_i(x^*)$$

But

1. $U_1(x^*_{-1}; \pi_{1,j}) = (Ax_2^*)_j$. Thus,

$$(Ax_2^*)_j \leq v^* = U_1(x^*)$$

   for all $j = 1, \ldots, m_1$.

2. $U_2(x^*_{-2}; \pi_{2,j}) = -((x_1^*)^T A)_j$. Thus,

$$((x_1^*)^T A)_j \geq v^* = -U_2(x^*)$$

   for all $j = 1, \ldots, m_2$.

3. $\max_{x_1 \in X_1} (x_1)^T A x_2^* \leq v^*$ because $(x_1)^T A x_2^*$ is a "weighted average" of $(Ax_2^*)_j$'s.

   Similarly, $v^* \leq \min_{x_2 \in X_2} (x_1^*)^T A x_2$ because $(x_1^*)^T A x_2$ is a "weighted average" of $((x_1^*)^T A)_j$'s. Thus

$$\max_{x_1 \in X_1} (x_1)^T A x_2^* \leq v^* \leq \min_{x_2 \in X_2} (x_1^*)^T A x_2$$

   We earlier noted the opposite inequalities, so,

$$\min_{x_2 \in X_2} \max_{x_1 \in X_1} x_1^T A x_2 = v^* = \max_{x_1 \in X_1} \min_{x_2 \in X_2} (x_1)^T A x_2$$

4. We didn't assume anything about the particular Nash Equilibrium we chose. So, for every NE, x^*, letting $v' = (x_1^*)^T A x_2^*$,

$$\max_{x_1 \in X_1} \min_{x_2 \in X_2} (x_1)^T A x_2 = v' = v^* = \min_{x_2 \in X_2} \max_{x_1 \in X_1} x_1^T A x_2$$

   Moreover, if $x^* = (x_1^*, x_2^*)$ satisfies conditions (1) and (2) for some v^*, then x^* must be a Nash Equilibrium. (This will be a homework exercise.)

Q.E.D. (Minimax Theorem)

remarks and food for thought

- Thus, for 2-player zero-sum games, Nash Equilibria and Minimax profiles are the same thing.

- Let us note here
  Useful Corollary for Minimax: In a minimax profile $x^* = (x_1^*, x_2^*)$,

  1. if $x_2^*(j) > 0$ then $((x_1^*)^T A)_j = (x_1^*)^T A x_2^* = v^*$.
  2. if $x_1^*(j) > 0$ then $(A x_2^*)_j = (x_1^*)^T A x_2^* = v^*$.

  This is an immediate consequence of the Useful Corollary for Nash Equilibria.

- If you were playing a 2-player zero-sum game (say, as player 1) would you always play a minmaximizer strategy?

- What if you were convinced your opponent is an idiot?

- Notice, we have no clue yet how to compute the minimax value and a minimax profile.

  That is about to change.

minimax as an optimization problem

Consider the following "optimization problem":

Maximize v

Subject to constraints:

$(x_1^T A)_j \geq v$ for $j = 1, \ldots, m_2$,

$x_1(1) + \ldots + x_1(m_1) = 1$,

$x_1(j) \geq 0$ for $j = 1, \ldots, m_1$

It follows from the minimax theorem that an optimal solution (x_1^*, v^*) would give precisely the minimax value v^*, and a minmaximizer x_1^* for Player 1.

We are optimizing a "linear objective",
under "linear constraints" (or "linear inequalities").

That's what Linear Programming is.

Fortunately, we have good algorithms for it.

Next time, we start Linear Programming.