Scribe Notes on Minimax Theorems and Game Theory

Lecturers: Mitali Nanda, Vasiliki Mavrogeorgou
Scribe: Trevor Tidy

April 8th, 2026

1 Introduction

This lecture introduces the basics of game theory, which studies how agents make decisions in interactive settings, or *games*. One of the most fundamental settings in game theory is that of a *two-player zero-sum game*, which were extensively studied by John von Neumann [4], leading to the establishment of the field. In this seminal work, von Neumann introduced what would be the first of many *minimax theorems*, which prove when equations of the form

$$\max_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \max_{x \in X} f(x, y)$$

hold, given conditions on the sets X, Y and function f. In this lecture, we will see how such theorems are closely tied to convex analysis, give a strong duality (minimax) proof, and discuss algorithms for minimax problems.

2 Game Theory

As mentioned above, we begin the discussion with *two-player games*. In such games, you have two actors (or players) who each select actions (playing the game), and will receive a subsequent reward (payoff) as a result. In a *zero-sum game*, the total payoffs of the players will add to zero.

We can fit many real-life games into this framework, such as rock paper scissors. In rock paper scissors, each player has the choice of three actions: rock, paper, or scissors. The actions are played simultaneously, and the pair of actions picked decides the outcome. If both players select the same action, the game is tied, but otherwise, rock beats scissors, scissors beats paper, and paper beats rock. We can fully describe this interactive setting using a *payoff matrix*, detailing the action pairs and rewards:

P_1 \ P_2	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

Figure 1: The rock paper scissors payoff matrix. In this game, player 1 (P_1, column actions) aims to maximize the return value, and player 2 (P_2, row actions) aims to minimize the return value.

With P_1 looking to maximize return and P_2 looking to minimize return, we can see how this corresponds to strategic play in rock paper scissors via Figure 1. When P_1 "beats" P_2, i.e. when the actions selected by the players are (R, S), (P, R), or (S, P), the return is $+1$, which is good for P_1 and bad for P_2. The same is true when P_2 "beats" P_1, and when the result is a tie, both get 0. This is also a zero-sum game, as the amount gained for one player is exactly that lost for the other player (just think of the reward value for P_2 being the negative value of the return given by the matrix).

The actions of each player is governed by their *strategy.* An example of a strategy for a player could be picking rock every time. This is what we call a *pure strategy,* and we can label strategies by a vector with entries being the probability of picking each action. In this case, the vector with entries for actions (R, P, S) would be $(1, 0, 0)$. If a strategy does not deterministically follow the same choice every time, this is called a *mixed strategy.* For example, the strategy in which you randomly select any of the three actions with equal probability is a mixed strategy, and is given by the vector $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Rock paper scissors perhaps isn't the most strategic two-player zero-sum game out there, but we can see more interesting behaviours arise when we change up the reward structure of the game. Say we now have a game where both players want to maximize their return. In a more general setting, we can set the payoff values in the payoff matrix to be a tuple (a, b) where a is the payoff for player 1 and b is the payoff for player two. Take a game with the following payoff structure:

P_1 \ P_2	α_1	α_2
σ_1	$(4, -4)$	$(3, -3)$
σ_2	$(6, -6)$	$(-1, 1)$

Figure 2: A payoff matrix for a new game. σ_i are the action choices for P_1, while α_i are the action choices for P_2. This matrix setup is equally expressive to that in Figure 1 (and can be expressed this way using just the first of the tuple values in each entry), but this design is common for more general games.

The natural question here is: how do we expect these two players to play this game if they both wish to maximize their reward?

Putting ourselves into the shoes of player 2, we notice that no matter what action player 1 chooses, action α_2 yields a better reward. This means the strategy of picking α_2 *strictly dominates* that of picking α_1. Therefore, if both players are rational actors, P_2 will play α_2 every time. Player 1 will notice this as well, and now assumes that as player 2 is rational, this player can expect α_2 to be played. With this knowledge, σ_1 is a better choice for P_1.

This reasoning is enough to see that the optimal strategies are $z^* = (1, 0)$ for P_1 and $w_* = (0, 1)$ for P_2. This pair of strategies lands us in a *Nash equilibrium* (equivalent to a *saddle point* in two-player zero-sum games), as no player gains from changing their strategy (given the other player sticks with their strategy). In many games, equilibria and saddle points are not unique, and there may be many equally viable strategies.

3 Convex Analysis Approach

As we are just working with two-player zero-sum games, we can consider payoff matrices A like those seen in Figure 1, where P_1 aims to maximize the value while P_2 seeks to minimize it. Given sets of feasible strategies $\mathbb{Z}$ for P_1 and $\mathbb{W}$ for P_2 respectively, if P_1 picks strategy $z \in \mathbb{Z}$ and P_2 picks strategy $w \in \mathbb{W}$, then the return is

$$f(w,z) = \sum_{i=1}^{m}\sum_{j=1}^{n} z_i a_{ij} w_j = z^T A w$$

where a_{ij} denotes the entries of the $m \times n$ payoff matrix A.

This leads us to the primal value problem:

$$p_* = \inf_{w\in\mathbb{W}} \sup_{z\in\mathbb{Z}} f(w,z),$$

and the corresponding dual problem:

$$\partial^* = \sup_{z\in\mathbb{Z}} \inf_{w\in\mathbb{W}} f(w,z).$$

4 Saddle Points and Duality

Now that we have the background of Sections 1 and 2, we can tie together the ideas of minimax theorems: saddle points and duality. To formalize the notion of saddle points, a point $(w_*, z^*) \in \mathbb{W}\times\mathbb{Z}$ is called a *saddle point* of f if

$$f(w,z^*) \le f(w_*,z^*) \le f(w_*,z) \qquad \text{for all } w\in\mathbb{W},\ z\in\mathbb{Z}.$$

We introduced the notion of saddle points in Section 2, and in this section we will see how we can equate this notion with that of *strong duality* of the primal p_* and dual ∂^* defined above. Before concerning ourselves with strong duality (an equality statement), it is insightful to see that weak duality, $p_* \ge \partial^*$, always holds:

Proof. Let $w\in\mathbb{W}$, $z\in\mathbb{Z}$ be arbitrary. Then

$$\inf_{u\in W} f(u,z) \le f(w,z) \le \sup_{v\in Z} f(w,v)$$

holds by definition of inf and sup. Furthermore, as w, z are arbitrary, this will hold if we take the supremum of one and infimum of the other, giving us:

$$\partial^* = \sup_{y\in Z}\inf_{u\in W} f(u,y) \le \inf_{x\in W}\sup_{v\in Z} f(x,v) = p_*,$$

which is weak duality. □

We can now tie duality and saddle points together with the following proposition:

Proposition 1. *Given $\mathbb{W},\mathbb{Z}$ nonempty and f, p_*, ∂^* as defined above, then a point $(w_*,z^*)\in\mathbb{W}\times\mathbb{Z}$ is a saddle point if and only if $p_* = f(w_*,z^*) = \partial^*$, where w_* attains the primal infimum and z^* attains the dual supremum*[1]*.*

Proof. ($\Longrightarrow$) As (w_*,z^*) is a saddle point, we have

$$f(w_*,z) \le f(w_*,z^*) \le f(w,z^*)$$

for any $w\in\mathbb{W}, z\in\mathbb{Z}$. As this holds for all values, we can again take the left supremum and right infimum to get

$$\sup_{z\in\mathbb{Z}} f(w_*,z) \le f(w_*,z^*) \le \inf_{w\in\mathbb{W}} f(w,z^*).$$

[1] This is the same as ensuring $z^* \in \text{argmax}_z f(w_*,z)$ and $w_* \in \text{argmax}_w f(w,z^*)$, which were the conditions given in the lecture.

However, as $w_* \in \mathbb{W}$ and $z^* \in \mathbb{Z}$, we also have $f(w_*, z^*) \leq \sup_{z\in\mathbb{Z}} f(w_*, z)$ and $\inf_{w\in\mathbb{W}} f(w, z^*) \leq f(w_*, z^*)$ respectively. Therefore:

$$\begin{aligned}\sup_{z\in\mathbb{Z}} f(w_*, z) =& f(w_*, z^*) = \inf_{w\in\mathbb{W}} f(w, z^*) \\ \implies \inf_{w\in\mathbb{W}} \sup_{z\in\mathbb{Z}} f(w, z) \leq& f(w_*, z^*) \leq \sup_{z\in Z} \inf_{w\in\mathbb{W}} f(w, z) \\ \implies p_* \leq& f(w_*, z^*) \leq \partial^*.\end{aligned}$$

By weak duality, we always have $\partial^* \leq p_*$, and thus

$$\begin{aligned}\partial^* \leq p_* \leq& f(w_*, z^*) \leq \partial^* \\ \implies p_* =& f(w_*, z^*) = \partial^*.\end{aligned}$$

($\Longleftarrow$) We now suppose that $p_* = f(w_*, z^*) = \partial^*$. Thus for every $w \in \mathbb{W}$ and $z \in \mathbb{Z}$ we have

$$f(w, z^*) \geq \inf_{u\in\mathbb{W}} f(u, z^*) = \partial^* = f(w_*, z^*)$$

and

$$f(w_*, z) \leq \sup_{v\in\mathbb{Z}} f(w_*, v) = p_* = f(w_*, z^*)$$

respectively. Therefore

$$f(w_*, z) \leq f(w_*, z^*) \leq f(w, z^*) \qquad \text{for all } w \in \mathbb{W},\ z \in \mathbb{Z},$$

proving (w_*, z^*) is a saddle point, as desired. □

Now that we have conditions from which we can find a saddle point, notice what the equivalence $p_* = \partial^*$ explicitly gives us:

$$p_* = \inf_{w\in\mathbb{W}} \sup_{z\in\mathbb{Z}} f(w, z) = \sup_{z\in\mathbb{Z}} \inf_{w\in\mathbb{W}} f(w, z) = \partial^*.$$

This tells us we can swap the infimum and supremum, and gives us a *minimax*-style result. This is of a slightly more general form than that of von Neumann's original theorem [4]:

Theorem 1 (von Neumann). *Let $\mathbb{W}$ and $\mathbb{Z}$ be the sets of possible mixed strategies for players 1 and 2*[2]*, and $A \in \mathbb{R}^{m\times n}$ be the payoff matrix of a (finite) zero-sum game. Taking $f(w, z) = w^T Az$, we have:*

$$\min_{w\in\mathbb{W}} \max_{z\in\mathbb{Z}} f(w, z) = \max_{z\in\mathbb{Z}} \min_{w\in\mathbb{W}} f(w, z).$$

What this theorem tells us is that optimal strategies (w_*, z^*) exist such that no one player can improve by changing their strategy (saddle point), and we reach this common equilibrium value (known as the *value* of the game) that realizes the minimax equality. Compared to our duality statement, we notice that supremum and infimum are generalizations of the usual maximum and minimum we see in minimax theorems. Though such statements often collapse to max and min (through realizing the supremum and infimum, which actually does happen above), there exist more general minimax theorems. One notable example is Sion's minimax theorem [2]:

Theorem 2 (Sion). *Let X be a convex subset of a linear topological space, and let Y be a compact convex subset of a linear topological space. Let*

$$f : X \times Y \to \mathbb{R}$$

be a real-valued function such that:

- *for every fixed $y \in Y$, the function*

  $$x \mapsto f(x, y)$$

  is upper semicontinuous and quasi-concave on X;

[2] We take these to be the simplices Δ^n and Δ^m.

- *for every fixed $x \in X$, the function*

$$y \mapsto f(x, y)$$

is lower semicontinuous and quasi-convex on Y.

Then

$$\sup_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \sup_{x \in X} f(x, y).$$

Minimax theorems have uses beyond game theory, in fields such as optimization and machine learning. Being able to prove the existence of saddle points is useful in many settings (such as the games seen prior), and can be applied in practice for tasks such as robust optimization. For a direct tie to our coursework on convex analysis, we will analyze how Sion's minimax theorem relates to Langrangian duality in the following subsection.

4.1 Connection to Lagrangian Duality

Consider the convex optimization problem

$$p^* = \inf_{x \in X} f_0(x)$$

that is subject to constraints

$$f_i(x) \leq 0, \qquad i = 1, ..., n.$$

We define the *Lagrangian* L to be

$$L(x, \lambda) = f_0(x) + \sum_{i=1}^{n} \lambda_i f_i(x),$$

where $\lambda_i \geq 0$ are the *multipliers*, which serve to penalize violations of the constraints. Given this function L, notice that if we fix x, we see that

$$\sup_{\lambda \geq 0} L(x, \lambda) = \begin{cases} f_0(x), & \text{if } f_i(x) \leq 0 \quad \forall i \\ +\infty, & \text{else} \end{cases}$$

and we can rewrite our constrained optimization problem as

$$p^* = \inf_{x \in X} \sup_{\lambda \geq 0} L(x, \lambda).$$

This problem has now taken a familiar form. Similarly to what we introduced earlier, the Lagrangian dual problem reverses the order:

$$d^* = \sup_{\lambda \geq 0} \inf_{x \in X} L(x, \lambda),$$

and *Lagrangian strong duality* will be obtained when $p^* = d^*$, or:

$$\inf_{x \in X} \sup_{\lambda \geq 0} L(x, \lambda) = \sup_{\lambda \geq 0} \inf_{x \in X} L(x, \lambda).$$

We have arrived at a minimax problem once again, and using the minimax theorems given in the previous subsection now gives us a path to Lagrangian strong duality. By Sion's theorem, if $L(x, \lambda)$ is convex and lower semicontinuous in x, concave and upper semicontinuous in λ, and the feasible sets satisfy the conditions of Sion's theorem, then we have Lagrangian strong duality automatically. Furthermore, this connection tells us that when the primal infimum and dual supremum are attained, strong duality is equivalent to the existence of a saddle point, (x^*, λ^*) for the Lagrangian[3].

[3]For a longer, but also clear treatise on Lagrangian duality and the minimax formulation, see [1].

5 Algorithms for Minimax Problems

The guarantees provided by minimax theorems are great in theory, but in practice it is useful to have methods for finding optimal solutions to minimax problems (i.e. saddle points). In this section, we will go over three of the most common algorithms: alternating minimax, Uzawa's algorithm, and gradient descent-ascent (GDA).

5.1 The Alternating Minimax Algorithm

The alternating minimax algorithm is given in Algorithm 1:

Algorithm 1 Alternating Minimax

Input: $(w_0, z_0) \in (\mathbb{W} \times \mathbb{Z}) \cap \operatorname{dom} f$

for $t = 0, 1, 2, \ldots$ **do**

$$w_{t+1} \leftarrow \arg\min_{w \in \mathbb{W}} f(w, z_t)$$

$$z_{t+1} \leftarrow \arg\max_{z \in \mathbb{Z}} f(w_{t+1}, z)$$

end for

Intuitively speaking, this algorithm works as if two perfect players are playing this game against each other, playing the optimal response at each turn. It should be noted that if f is convex in w and concave in z, and if $\mathbb{W}$ and $\mathbb{Z}$ are convex sets, then each step of Algorithm 1 is a convex optimization problem:

$$w_{t+1} \in \arg\min_{w \in \mathbb{W}} f(w, z_t),$$

and

$$z_{t+1} \in \arg\max_{z \in \mathbb{Z}} f(w_{t+1}, z) = \arg\min_{z \in \mathbb{Z}} \big(-f(w_{t+1}, z)\big).$$

While this algorithm does have a good logical basis, it is not guaranteed to converge. For example, consider the problem

$$\min_{w \in [-1,1]} \max_{z \in [-1,1]} wz,$$

i.e. we have $f(w, z) = wz$, and $\mathbb{W}, \mathbb{Z} = [-1, 1]$. It is easy to see that $(0, 0)$ is a saddle point here, but what happens when we run the alternating algorithm? Following the algorithm, we will have

$$w_{t+1} = -\operatorname{sign}(z_t), \qquad z_{t+1} = \operatorname{sign}(w_{t+1}),$$

which just oscillates when $z_0 \neq 0$. What we have encountered here is a saddle point $(0, 0)$ that is *unstable*, and the alternating algorithm does not converge.

This is not always the case, and often the alternating algorithm will quickly converge to the saddle point. For example, take the problem

$$\min_{w \in [-1,1]} \max_{z \in [-1,1]} ze^w.$$

Since $e^w > 0$, the maximizing player chooses $z = 1$. The minimizing player then chooses $w = -1$, thus the saddle point is

$$(w_*, z^*) = (-1, 1).$$

and we see Algorithm 1 converges in just two steps[4]. In this example, $(-1, 1)$ is a *stable* saddle point, hence the convergence of the algorithm.

[4] It is important to note that we used some of our own mathematical knowledge of the given functions to reach this conclusion. In practice, finding the exact arg max and arg min is usually a complicated task to solve algorithmically.

5.2 Uzawa's Algorithm

Uzawa's algorithm treats the minimax problem as minimizing the value $\sup_{z\in\mathbb{Z}} f(w,z)$ iteratively [3].

Algorithm 2 Uzawa's Algorithm

1: **Input:** $(w_0, z_0) \in (\mathbb{W} \times \mathbb{Z}) \cap \operatorname{dom} f$
2: **for** $t = 0, 1, 2, \dots$ **do**
3:
$$z_t \leftarrow \arg\max_{z\in\mathbb{Z}} f(w_t, z)$$
4: Compute a subgradient
$$g_t \in \partial_w f(w_t, z_t)$$
5: Choose step size $\eta_t > 0$
6: **Optional:**
$$g_t \leftarrow \frac{g_t}{\|g_t\|}$$
7:
$$w_{t+1} \leftarrow P_{\mathbb{W}}(w_t - \eta_t g_t)$$
8: **end for**

There are frequent scenarios in which $\sup_{z\in\mathbb{Z}} f(w,z)$ is non-smooth, which prompts the use of subgradients in the algorithm. The technique at play boils down to projective subgradient descent, realizing worst-case responses z_t and adapting w_t to reduce this value. Uzawa's algorithm is not guaranteed to converge, but will always do so when the following conditions are met[5]:

- The step sizes η_t diminish.
- $\mathbb{W}, \mathbb{Z}$ are closed, convex sets.
- f is smooth, as well as convex in w and concave in z.
- Subgradients are bounded, and the supremum is realized.

Algorithm 2 is more applicable than Algorithm 1 when the outer minimization problem is difficult to solve in an exact way. For example, say you have the problem

$$\min_{w\in\mathbb{W}} \max_{z\in[n]} \ell(w,z)$$

where the loss function is

$$\ell(w,z) = a_z^T w + b_z,$$

i.e. you are attempting to minimize your worst case loss for all scenarios z[6]. Here we can see the inner minimization problem is easy, as for a fixed w we can just check each of the z scenarios to see which loss is maximal. The outer minimization problem, however, is likely to be much more difficult to solve in practice. The minimization is now over the maximum of multiple linear functions, which is now potentially piecewise and non-smooth[7]. Using subgradient steps for the outer minimization is the perfect approach here, making Uzawa's algorithm the most applicable.

5.3 Gradient Descent Ascent

Algorithm 3 rounds out the minimax algorithms covered, and is called *Gradient Descent-Ascent*. Looking over the two algorithms seen so far, the techniques are:

[5] These are not the most general conditions for convergence, but are more so a standard and trivial selection.

[6] In this example, $\mathbb{Z} = \{1, ..., n\}$. We take the values $a_z \in \mathbb{R}^d$ and $b_z \in \mathbb{R}$.

[7] In our case, piecewise linear with potential corners with each change in w. This problem can also take on a more challenging shape when ℓ isn't just the simple linear function given.

- Algorithm 1 solves both the maximization and minimization problems exactly.
- Algorithm 2 solves the inner maximization problem exactly, but uses a subgradient on the outer problem with minimizing in w.

GDA removes the second exactness condition, giving us an algorithm that takes gradient steps for both optimizations:

Algorithm 3 Gradient Descent–Ascent (GDA)

Input: $(w_0, z_0) \in \operatorname{dom} f \cap (\mathbb{W} \times \mathbb{Z})$
for $t = 0, 1, 2, \ldots$ **do**
  Choose step size $\eta_t > 0$

$$w_{t+1} \leftarrow P_{\mathbb{W}}(w_t - \eta_t \nabla_w f(w_t, z_t))$$

$$z_{t+1} \leftarrow P_{\mathbb{Z}}(z_t + \eta_t \nabla_z f(w_{t+1}, z_t))$$

end for

This is the most computationally feasible of the algorithms when problems become increasingly complex in nature. As a result, this approach is the most used in applied settings such as machine learning. Many variations can be implemented using the general algorithm, such as:

- Using different step sizes η_t for w and z.
- Use w_{t+1} in the z update, or vice versa[8].
- Use of stochastic gradients in either GD or GA.
- After every update in w, perform k updates in z, or vice versa.

Convergence of GDA varies based on the problem and finer implementation details such as those given above. Going back to an early example, we can see GDA actually fails on the bilinear game seen in Subsection 5.1, as with

$$f(w, z) = wz,$$

the updates become

$$w_{t+1} = w_t - \eta z_t, \qquad z_{t+1} = z_t + \eta w_t.$$

Therefore, if we compute the squared distance to the saddle point $(0, 0)$, we find

$$w_{t+1}^2 + z_{t+1}^2 = (1 + \eta^2)(w_t^2 + z_t^2),$$

which grows whenever $\eta > 0$ (as $\eta + 1 > 1$). Therefore, the iterates "spiral away" from the saddle point, and do not converge.

Bilinear games (and those having unstable saddle points) are one example where GDA may fail, but other practical considerations such as step size and the convexity of the landscape are also common major factors in convergence. Even in problems where the standard GDA iterates may diverge, the averaged iterates

$$\bar{w}_T = \frac{1}{T}\sum_{t=1}^{T} w_t, \qquad \bar{z}_T = \frac{1}{T}\sum_{t=1}^{T} z_t$$

often converge to a saddle point in convex-concave problems. However, this is not a universal fix. Though these three algorithms are useful for a wide variety of problems, sometimes more sophisticated methods are needed to find solutions for minimax problems.

[8]This idea is analogous to that of gradient descent with momentum, or accelerated gradient descent

References

[1] Knowles, D. Lagrangian duality for dummies. *URL: https://wwwcs. stanford. edu/˜ david-knowles/lagrangian_duality. pdf (date of access: 24.10. 2024)* (2010).

[2] Sion, M. On general minimax theorems.

[3] Uzawa, H. Iterative methods for concave programming. *Studies in linear and nonlinear programming 6* (1958), 154–165.

[4] Von Neumann, J. On the theory of games of strategy. *Contributions to the Theory of Games 4*, 13-42 (1959), 41.