From Convex Analysis to Learning, Prediction, and Elicitation*
Lecture 5: Saddle Points, KKT Conditions, and Constructive Minimax Theorem

Lunjia Hu

Suppose we have a function $f : X \times Y \to \mathbb{R}$ that satisfies the minimax theorem:

$$\inf_{x \in X} \sup_{y \in Y} f(x, y) = \sup_{y \in Y} \inf_{x \in X} f(x, y). \tag{1}$$

Let OPT denote the quantity equal to both sides of (1).

Assume the infimum over $x \in X$ on the left side of (1) and supremum over $y \in Y$ on the right side of (1) can be both be attained. That is, there exists an optimal primal solution $x_0 \in X$ such that

$$\sup_{y \in Y} f(x_0, y) = \mathsf{OPT} = \inf_{x \in X} \sup_{y \in Y} f(x, y), \tag{2}$$

and an optimal dual solution $y_0 \in Y$ such that

$$\inf_{x \in X} f(x, y_0) = \mathsf{OPT} = \sup_{y \in Y} \inf_{x \in X} f(x, y). \tag{3}$$

Now suppose someone else gives us $(x_0, y_0) \in X \times Y$ and claims that they satisfy (2) and (3). How do we verify that? We could directly verify (2) and (3), but that may not be easy. Condition (2) states that x_0 minimizes $\sup_{y \in Y} f(x, y)$, so it is an optimality condition on the supremum function $\sup_{y \in Y} f(x, y)$. Similarly, (3) is an optimality condition on the infimum function $\inf_{x \in X} f(x, y)$. Can we verify the optimality of (x_0, y_0) simply using optimality conditions on f itself?

We will show that the answer is "yes". The optimality conditions (2) and (3) together are equivalent to the condition that (x_0, y_0) is a *saddle point*, which is defined using optimality conditions on f itself (see Definition 1 below).

The notion of saddle point not only allows us to *verify* conditions (2) and (3), but also helps us *construct* points (x_0, y_0) satisfying, or approximately satisfying (2) and (3). In particular, in Section 3 below we use a no-regret online learning algorithm to construct such (x_0, y_0), giving a constructive proof of the minimax theorem.

*https://lunjiahu.com/convex-analysis/

1 Saddle Points

Definition 1. *Let X, Y be non-empty sets, and let $f : X \times Y \to \mathbb{R}$ be an arbitrary function. For $\varepsilon_1, \varepsilon_2 \geq 0$, we say $(x_0, y_0) \in X \times Y$ is an $(\varepsilon_1 + \varepsilon_2)$-saddle point if*

$$f(x_0, y_0) - \inf_{x \in X} f(x, y_0) \leq \varepsilon_1,$$
$$\sup_{y \in Y} f(x_0, y) - f(x_0, y_0) \leq \varepsilon_2.$$

The following two inequality chains are important for understanding the notion of saddle point:

$$\inf_x f(x, y_0) \quad \leq \quad f(x_0, y_0) \quad \leq \quad \sup_y f(x_0, y), \tag{4}$$
$$\inf_x f(x, y_0) \leq \sup_y \inf_x f(x, y) \leq \inf_x \sup_y f(x, y) \leq \sup_y f(x_0, y). \tag{5}$$

These inequality chains imply the following two theorems:

Theorem 1. *Let X, Y be non-empty sets, and let $f : X \times Y \to \mathbb{R}$ be an arbitrary function. Suppose $(x_0, y_0) \in X_0 \times Y_0$ is an ε-saddle point for some $\varepsilon \geq 0$. Then*

$$\sup_{y \in Y} f(x_0, y) - \inf_{x \in X} \sup_{y \in Y} f(x, y) \leq \varepsilon, \qquad \text{(Primal optimality of } x_0\text{)}$$
$$\sup_{y \in Y} \inf_{x \in X} f(x, y) - \inf_{x \in X} f(x, y_0) \leq \varepsilon, \qquad \text{(Dual optimality of } y_0\text{)}$$
$$\inf_{x \in X} \sup_{y \in Y} f(x, y) - \sup_{y \in Y} \inf_{x \in X} f(x, y) \leq \varepsilon. \qquad \text{(Approximate minimax condition)}$$

Theorem 2. *Let X, Y be non-empty sets, and let $f : X \times Y \to \mathbb{R}$ be an arbitrary function. For an error bound $\varepsilon \geq 0$, suppose f satisfies the following approximate minimax condition:*

$$\inf_{x \in X} \sup_{y \in Y} f(x, y) \leq \sup_{y \in Y} \inf_{x \in X} f(x, y) + \varepsilon.$$

Suppose $x_0 \in X$ and $y_0 \in Y$ are approximately optimal primal and dual solutions within error $\varepsilon_1, \varepsilon_2 \geq 0$:

$$\sup_{y \in Y} f(x_0, y) - \inf_{x \in X} \sup_{y \in Y} f(x, y) \leq \varepsilon_1,$$
$$\sup_{y \in Y} \inf_{x \in X} f(x, y) - \inf_{x \in X} f(x, y_0) \leq \varepsilon_2.$$

Then (x_0, y_0) is an $(\varepsilon + \varepsilon_1 + \varepsilon_2)$-saddle point.

2 KKT Conditions

Consider the following general optimization problem specified by a domain $E \subseteq \mathbb{R}^m$, functions $f_i : E \to \mathbb{R}$ for $i = 0, \ldots, n$, a matrix $A \subseteq \mathbb{R}^{n' \times m}$, and a set $S \subseteq \mathbb{R}^{n'}$:

$$\underset{x \in E}{\text{minimize}} \quad f_0(x) \tag{P1}$$
$$\text{s.t.} \quad f_i(x) \leq 0 \quad \text{for every } i = 1, \ldots, n, \tag{6}$$
$$Ax \in S. \tag{7}$$

The Lagrangian L of the optimization problem is a function of $(x, s, h, h') \in E \times S \times \mathbb{R}^n_{\geq 0} \times \mathbb{R}^{n'}$:

$$L(x, s; h, h') := f_0(x) + \sum_{i=1}^{n} h_i f_i(x) + \langle Ax - s, h' \rangle. \tag{8}$$

Suppose there exists $x_0 \in E$ that attains the optimal objective value OPT of (P1) while satisfying the two constraints (6) and (7). This is equivalent to the following condition, where we define $s_0 := Ax_0$:

$$\sup_{(h,h') \in \mathbb{R}^n_{\geq 0} \times \mathbb{R}^{n'}} L(x_0, s_0; h, h') \leq \mathsf{OPT}. \tag{9}$$

Lagrange duality is the following condition: there exists $(h_0, h') \in \mathbb{R}^n_{\geq 0} \times \mathbb{R}^{n'}$ such that

$$\inf_{(x,s) \in E \times S} L(x, s; h_0, h'_0) \geq \mathsf{OPT}. \tag{10}$$

If (9) and (10) both hold, then $(x_0, s_0; h_0, h'_0)$ is a 0-saddle point and both (9) and (10) become equalities. This is because all inequalities below must be equalities:

$$0 \leq \sup_{h,h'} L(x_0, s_0; h, h') - \inf_{x,s} L(x, s; h_0, h') \leq \mathsf{OPT} - \mathsf{OPT} = 0.$$

In particular, by (5), the minimax theorem holds for L:

$$\sup_{h,h'} \inf_{x,s} L(x, s; h, h') = \inf_{x,s} \sup_{h,h'} L(x, s; h, h') = \mathsf{OPT}. \tag{11}$$

The following theorem is a direct corollary of the inequality chain (4):

Theorem 3 (KKT condition)**.** *Consider the Lagrangian L (defined in Equation* (8)*) of the optimization problem* (P1)*. For every $(x_0, s_0, h_0, h'_0) \in E \times S \times \mathbb{R}^n_{\geq 0} \times \mathbb{R}^{n'}$, the following two statements are equivalent:*

1. *There exists $\mathsf{OPT} \in \mathbb{R}$ such that* (9) *and* (10) *both hold (in which case the value of OPT is given by* (11)*, and the two inequalities* (9) *and* (10) *both become equalities).*

2. *(x_0, s_0, h_0, h'_0) satisfies the following saddle-point condition (i.e., the* KKT conditions*):*

$$L(x_0, s_0; h_0, h'_0) = \min_{x,s} L(x, s; h_0, h'_0), \tag{12}$$

$$L(x_0, s_0; h_0, h'_0) = \max_{h,h'} L(x_0, s_0; h, h'). \tag{13}$$

The second KKT condition (13) can be simplified to the following equivalent form: for every $i = 1, \ldots, n$,

$$(h_0)_i > 0 \Longrightarrow f_i(x_0) = 0, \qquad \text{(Complementary slackness)}$$

$$f_i(x_0) \leq 0 \text{ and } Ax_0 = s_0. \qquad \text{(Primal feasibility)}$$

3 Constructive Minimax Theorem from No-Regret Learning

Theorem 1 allows us to prove that a function f satisfies the approximate minimax condition by showing the existence of a saddle point. We prove Lemma 4 below which allows us to construct a saddle point from two sequences satisfying a low-regret condition (14). Combining these two results, we can prove that a function f satisfies the approximate minimax condition by constructing low-regret sequences. This leads to *constructive* proofs of the minimax theorem (see the proof of Theorem 5 below) in contract to the *non-constructive* proofs we have seen in Lecture 2.

Lemma 4 (Saddle point from no-regret sequences)**.** *Let $f : X \times Y \to \mathbb{R}$ be a convex-concave function. Let $x_1, \dots, x_T \in X$ and $y_1, \dots, y_T \in Y$ be two sequences satisfying the following assumption:*

$$\sup_{y \in Y} \frac{1}{T} \sum_{t=1}^{T} f(x_t, y) - \varepsilon_2 \leq \frac{1}{T} \sum_{t=1}^{T} f(x_t, y_t) \leq \inf_{x \in X} \frac{1}{T} \sum_{t=1}^{T} f(x, y_t) + \varepsilon_1. \tag{14}$$

Then $(\bar{x}, \bar{y})$ is an $(\varepsilon_1 + \varepsilon_2)$-saddle point, where $\bar{x} := \frac{1}{T} \sum_{t=1}^{T} x_t$ and $\bar{y} := \frac{1}{T} \sum_{t=1}^{T} y_t$.

Proof. By the assumption that f is convex-concave, we can apply Jensen's inequality and get

$$\inf_{x \in X} \frac{1}{T} \sum_{t=1}^{T} f(x, y_t) \leq \inf_{x \in X} f(x, \bar{y}),$$

$$\sup_{y \in Y} \frac{1}{T} \sum_{t=1}^{T} f(x_t, y) \geq \sup_{y \in Y} f(\bar{x}, y).$$

Plugging these inequality into (14), we get

$$\sup_{y \in Y} f(\bar{x}, y) - \inf_{x \in X} f(x, \bar{y}) \leq \varepsilon_1 + \varepsilon_2.$$ □

We now demonstrate the power of Lemma 4 by giving a *constructive* proof of the following minimax theorem:

Theorem 5. *Let $X, Y \subseteq \mathbb{R}^d$ be compact convex sets. Then*

$$\inf_{x \in X} \sup_{y \in Y} \langle x, y \rangle = \sup_{y \in Y} \inf_{x \in X} \langle x, y \rangle. \tag{15}$$

Theorem 5 is a special case of the minimax theorems we have learned in Lecture 2. However, the proofs we have seen are *non-constructive*, in that they do not directly give us solutions to the primal and dual problems. We now give a constructive proof of Theorem 5 using Theorem 1 and Lemma 4.

Proof of Theorem 5. Consider a sequential game between a learner and an adversary. In each round $t = 1, \dots, T$, the learner chooses $x_t \in X$, and then the adversary chooses $y_t \in Y$. Based on what we learned in the previous lecture, the learner has a strategy that guarantees low regret regardless of the adversary's strategy. Specifically, the learner can ensure that the second inequality in (14) holds with $\varepsilon_1 = O(\sqrt{1/T})$. Now we let the adversary choose y_t as the *best response* to x_t. That is, $y_t := \arg\max_{y \in Y} \langle x_t, y \rangle$. This ensures that (14) holds with $\varepsilon_2 = 0$. By Lemma 4,

$$(\bar{x}, \bar{y}) \text{ is an } \varepsilon_1\text{-saddle point.}$$

By Theorem 1,

$$\inf_{x\in X}\sup_{y\in Y}\langle x,y\rangle - \sup_{y\in Y}\inf_{x\in X}\langle x,y\rangle \le \varepsilon_1 = O\left(\sqrt{1/T}\right).$$

Sending $T \to +\infty$ proves (15). □

Remark 1. *In the proof above, y_t is the best response to x_t. This is stronger than the low-regret condition needed in Lemma 4 and allows us to construct a "simpler" saddle point than $(\bar{x},\bar{y})$. In particular, let $t^* := \arg\min_{t=1,\dots,T} f(x_t,y_t)$. We have*

$$\sup_{y\in Y} f(x_{t^*},y) = f(x_{t^*},y_t) \le \frac{1}{T}\sum_{t=1}^{T} f(x_t,y_t).$$

Our earlier argument shows, for $\varepsilon_1 = O(\sqrt{1/T})$,

$$\frac{1}{T}\sum_{t=1}^{T} f(x_t,y_t) \le \inf_{x\in X}\frac{1}{T}\sum_{t=1}^{T} f(x,y_t) + \varepsilon_1 \le \inf_{x\in X} f(x,\bar{y}) + \varepsilon_1.$$

Combining the two equations above, we know that $(x_{t^*},\bar{y})$ is an ε_1-saddle point:

$$\sup_{y\in Y} f(x_{t^*},y) \le \inf_{x\in X} f(x,\bar{y}) + \varepsilon_1.$$

Remark 2. *All our analysis extends to the case where y_t is an* approximate *best response. That is, for some $\varepsilon_2 > 0$, we have $\langle x_t,y_t\rangle \ge \sup_{y\in Y}\langle x_t,y\rangle - \varepsilon_2$. In this case, $(\bar{x},\bar{y})$ and $(x_{t^*},\bar{y})$ are both $(\varepsilon_1+\varepsilon_2)$-saddle points.*