ECE 273 Convex Optimization and Applications — Instructor: Jun-Kun Wang
Scribe: Yunzhou Yan — May 28, 2024
Editors/TAs: Marialena Sfyraki

Lecture 16: (Continue) Min-Max Optimization

1 Saddle points in min-max optimization

The goal of min-max optimzation

Consider the following optimization problem:

$$\inf_{x\in X}\sup_{y\in Y} g(x,y)$$

where $g : X\times Y \to \mathbb{R}$ is a given function, X and Y are sets over which the optimization is performed, inf denotes the infimum (or greatest lower bound), and sup denotes the supremum (or least upper bound).

Definition of saddle points

Definition 1 (Saddle Points/Nash Equilibrium)**.** *Let $x \in X$ and $y \in Y$ and $g(\cdot,\cdot) : X \times Y \to \mathbb{R}$. A pair of points $(x_*, y_*) \in X \times Y$ is a saddle point of $g(\cdot,\cdot)$ if*

$$g(x_*, y) \le g(x_*, y_*) \le g(x, y_*), \forall x \in X, y \in Y.$$

Remark. This condition implies that at the saddle point, $g(x_*, y_*)$ represents a Nash equilibrium in the sense that no player can unilaterally improve their payoff by changing their strategy from x_* or y_*.

Theorem 1. *Let $g : X \times Y \to \mathbb{R}$, where X and Y are non-empty sets. A point (x_*, y_*) is a saddle point of g if and only if the following conditions are satisfied:*

1. *The supremum in $\sup_{y\in Y}\inf_{x\in X} g(x,y)$ is attained at y_*.*
2. *The infimum in $\inf_{x\in X}\sup_{y\in Y} g(x,y)$ is attained at x_*.*
3. *Moreover, $\sup_{y\in Y}\inf_{x\in X} g(x,y) = \inf_{x\in X}\sup_{y\in Y} g(x,y)$.*

Remarks.

1. If inf sup and sup inf have different values, then there is no saddle point.

2. If a saddle point exists, then:
   - There might be multiple ones, all of them must have the same minimax value, i.e., $\sup_{y\in Y}\inf_{x\in X} g(x,y) = \inf_{x\in X}\sup_{y\in Y} g(x,y)$
   - The set of saddle points is the Cartesian product $X_* \times Y_*$ when nonempty.
   - The set x_* is the optimal solution to $\inf_{x\in X}\sup_{y\in Y} g(x,y)$.
   - The set y_* is the optimal solution to $\sup_{y\in Y}\inf_{x\in X} g(x,y)$.

Example of No Saddle Points

Consider the function $g(x,y) = (x-y)^2$ with $X = [-1,1]$ and $Y = [-1,1]$. Then, we evaluate the infimum and supremum as follows:

$$\inf_{x\in X}\sup_{y\in Y}(x-y)^2 = \inf_{x\in X}(1+|x|)^2 = 1,$$

where the infimum is taken over the maximum value the function can achieve for each x, realizing that the maximum occurs at the endpoints of Y. Similarly,

$$\sup_{y\in Y}\inf_{x\in X}(x-y)^2 = \sup_{y\in Y} 0 = 0,$$

where the infimum for each y is achieved when $x = y$, leading to a minimum value of 0 for all y.

This discrepancy between the infimum of the supremum and the supremum of the infimum indicates that there are no saddle points for $g(x,y) = (x-y)^2$ over the given domain.

2 Metric to measure the progress of min-max optimization

In the context of min-max optimization, it is paramount to quantify the progress of optimization from the perspectives of participating entities. For a given function $g : X \times Y \to \mathbb{R}$, where X and Y represent the strategy sets for two players within the optimization problem, we define two metrics, $\ell(x)$ and $h(y)$, to assess progress from the viewpoints of the x-player and y-player respectively.

For the x-Player

Define $\ell(x)$ as the supremum of $g(x, y)$ over all $y \in Y$:

$$\ell(x) := \sup_{y \in Y} g(x, y).$$

From the x-player's perspective, the progress is measured as:

$$\ell(x) - \inf_{x \in X} \ell(x).$$

For the y-Player

Define $h(y)$ as the infimum of $g(x, y)$ over all $x \in X$:

$$h(y) := \inf_{x \in X} g(x, y).$$

For the y-player's perspective, the progress is captured by:

$$\sup_{y \in Y} h(y) - h(y).$$

Let $g : X \times Y \to \mathbb{R}$ be a given function, and $\hat{x} \in X$, $\hat{y} \in Y$ represent specific selections within their respective domains. By the definition of sup and inf, the following relation holds:

$$\sup_{y \in Y} g(\hat{x}, y) \geq g(\hat{x}, \hat{y}) \geq \inf_{x \in X} g(x, \hat{y}).$$

Combining the optimality gap of each player, we have that

$$\begin{aligned} \text{Gap}(\hat{x}, \hat{y}) &:= \ell(\hat{x}) - \inf_{x \in X} \ell(x) + \sup_{y \in Y} h(y) - h(\hat{y}) \\ &= \sup_{y \in Y} g(x, y) - \inf_{x \in X} \sup_{y \in Y} g(x, y) + \sup_{y \in Y} \inf_{x \in X} g(x, y) - \inf_{x \in X} g(x, y) \\ &= \sup_{y \in Y} g(\hat{x}, y) - \inf_{x \in X} g(x, \hat{y}), \end{aligned}$$

where the second-to-the-last line is by assuming the existence of a saddle point.

Definition 2 (Duality Gap)**.** *The duality gap $Gap(\hat{x}, \hat{y})$ is defined as:*

$$Gap(\hat{x}, \hat{y}) := \sup_{y \in Y} g(\hat{x}, y) - \inf_{x \in X} g(x, \hat{y}),$$

Remark. Duality gap is always non-negative even if the saddle point does not exist. By the definition of sup and inf, we have

$$\sup_{y \in Y} g(\hat{x}, y) \geq g(\hat{x}, \hat{y}) \geq \inf_{x} g(x, \hat{y})$$

Therefore,

$$\begin{aligned}
\mathrm{Gap}(\hat{x}, \hat{y}) &:= \sup_{y \in Y} g(\hat{x}, y) - \inf_{x \in X} g(x, \hat{y}) \\
&= \sup_{y \in Y} g(\hat{x}, y) - g(\hat{x}, \hat{y}) + g(\hat{x}, \hat{y}) - \inf_{x} g(x, \hat{y}) \\
&\geq 0.
\end{aligned}$$

ϵ-equilibrium / ϵ-saddle point

Assume a saddle point of $g(\cdot, \cdot)$ exists. Let us define the value v_* as follows:

$$v_* = \inf_{x \in X} \sup_{y \in Y} g(x, y) = \sup_{y \in Y} \inf_{x \in X} g(x, y).$$

Definition 3 (ϵ-equilibrium / ϵ-saddle point)**.** *A pair $(\hat{x}, \hat{y}) \in X \times Y$ is an ε-equilibrium or ε-saddle point if*

$$v_* - \varepsilon \leq \inf_{x \in X} g(x, \hat{y}) \leq v_* \leq \sup_{y \in Y} g(\hat{x}, y) \leq v_* + \varepsilon.$$

Remark. This definition extends the concept of a saddle point by introducing a margin of ε, allowing for a near-optimal equilibrium within an ε range of the optimal value v_*. Using the following inequality,

$$\sup_{y \in Y} g(\hat{x}, y) \geq g(\hat{x}, \hat{y}) \geq \inf_{x \in X} g(x, \hat{y}),$$

we can derive the following two inequalities

$$v_* - \varepsilon \leq \inf_{x \in X} g(x, \hat{y}) \leq v_* \leq g(\hat{x}, \hat{y})$$

$$g(\hat{x}, \hat{y}) \leq \sup_{y \in Y} g(\hat{x}, y) \leq v_* + \varepsilon$$

Thus, the above definition implies that

$$v_* - \epsilon \leq g(\hat{x}, \hat{y}) \leq v_* + \epsilon.$$

Lemma 1. *Given that the duality gap $Gap(\hat{x}, \hat{y}) \leq \varepsilon$ and assuming the existence of a saddle point, it follows that the pair $(\hat{x}, \hat{y}) \in X \times Y$ constitutes an ε-equilibrium or ε-saddle point.*

Proof. By definition of the duality gap

$$\begin{aligned} Gap(\hat{x}, \hat{y}) := & \sup_{y \in Y} g(\hat{x}, y) - \inf_{x \in X} g(x, \hat{y}) \leq \varepsilon \\ & \Leftrightarrow \sup_{y \in Y} g(\hat{x}, y) \leq \inf_{x \in X} g(x, \hat{y}) + \varepsilon. \end{aligned}$$

Given the optimal value

$$v_* = \inf_{x \in X} \sup_{y \in Y} g(x, y),$$

it follows from the definition that

$$v_* \leq \sup_{y \in Y} g(\hat{x}, y).$$

Therefore, we can establish the chain of inequalities

$$v_* = \inf_{x \in X} \sup_{y \in Y} g(x, y) \leq \sup_{y \in Y} g(\hat{x}, y) \leq \inf_{x \in X} g(x, \hat{y}) + \varepsilon.$$

This sequence demonstrates the relationship between the optimal value v_*, the supremum over y for a fixed $\hat{x}$, and the adjusted infimum over x for a fixed $\hat{y}$ by an ε margin, reflecting the bounds within which v_* is situated.
The duality gap for a pair $(\hat{x}, \hat{y})$ is defined as:

$$Gap(\hat{x}, \hat{y}) := \sup_{y \in Y} g(\hat{x}, y) - \inf_{x \in X} g(x, \hat{y}) \leq \varepsilon$$

This can be equivalently expressed as:

$$\sup_{y \in Y} g(\hat{x}, y) \leq \inf_{x \in X} g(x, \hat{y}) + \varepsilon \leq \sup_{y \in Y} \inf_{x \in X} g(x, y) + \varepsilon = v_* + \varepsilon$$

Using similar arguments, we can prove the left side of the chain of inequalities.
Therefore, we have proven that

$$v_* - \varepsilon \leq \inf_{x \in X} g(x, \hat{y}) \leq v_* \leq \sup_{y \in Y} g(\hat{x}, y) \leq v_* + \varepsilon.$$

□

Definition 4. *Given a pair $(\hat{x}, \hat{y}) \in X \times Y$, it is considered to be an ε-equilibrium or ε-saddle point if the following condition holds:*

$$v_* - \varepsilon \leq \inf_{x \in X} g(x, \hat{y}) \leq v_* \leq \sup_{y \in Y} g(\hat{x}, y) \leq v_* + \varepsilon.$$

3 The algorithmic aspect of min-max optimization

Review of online convex optimization

Algorithm 1 Online convex optimization

1: **for** $t = 1, 2, \ldots$ **do**
2: Commit a point z_t with its convex decision space $Z \subset \mathbb{R}^d$.
3: Receive a loss function $\ell_t(\cdot) : Z \to \mathbb{R}$ and incurs a loss $\ell_t(z_t)$.
4: **end for**

The goal of online convex optimization is to learn to be competitive with the best-fixed predictor from the convex set S, which is captured by minimizing the regret. Formally, the regret of the algorithm relative to any fixed benchmark z^* in Z when running on a sequence of T examples is defined as

$$\text{Regret}_T(z_*) = \sum_{t=1}^{T} l_t(z_t) - \sum_{t=1}^{T} l_t(z_*).$$

The regret of the algorithm relative to a convex set Z is defined as

$$\text{Regret}_T(Z) = \operatorname*{argmax}_{z_* \in Z} Regret_T(z_*)$$

The x-Player Perspective

Consider the x-player who, at each time step t, plays a strategy $x_t \in X$. Upon choosing this strategy, the x-player receives a loss function defined as:

$$\ell_t(x) := g(x, y_t),$$

where $g : X \times Y \to \mathbb{R}$ is a given function that determines the loss based on the player's choice x_t and the strategy y_t chosen by the opponent at time t.

The y-Player Perspective

From the perspective of the y-player, the game proceeds as follows: at each time step t, the y-player selects a strategy $y_t \in Y$. Upon making this selection, the y-player receives a loss function, which is defined as:

$$h_t(y) := -g(x_t, y),$$

where $g : X \times Y \to \mathbb{R}$ is the function determining the outcome based on the strategy x_t chosen by the opponent and the y-player's own choice y at time t.

Meta-algorithm for solving min-max problems

Algorithm 2 Meta-algorithm for Solving Min-Max Problems

Initialize OAlg^x (OCO Algorithm for x) and OAlg^y (OCO Algorithm for y).
Define weight sequence $\alpha_1, \alpha_2, \dots, \alpha_T$.
for $t = 1, 2, \dots, T$ **do**
    x plays $x_t \leftarrow \text{OAlg}^x(\alpha_1 \ell_1, \alpha_2 \ell_2, \dots, \alpha_{t-1} \ell_{t-1})$
    y plays $y_t \leftarrow \text{OAlg}^y(\alpha_1 h_1, \alpha_2 h_2, \dots, \alpha_{t-1} h_{t-1})$
    x receives $\alpha_t \ell_t(x) := \alpha_t g(x, y_t)$
    y receives $\alpha_t h_t(y) := -\alpha_t g(x_t, y)$
end for
Output the average strategies x_T and y_T, where:

$$x_T := \frac{\sum_{t=1}^T \alpha_t x_t}{A_T}, \quad y_T := \frac{\sum_{t=1}^T \alpha_t y_t}{A_T},$$
$$\text{with} \quad A_T := \sum_{t=1}^T \alpha_t.$$

From the x-player perspective:

- Play $x_t \in X$.
- Receives the loss function at t, $\alpha_t \ell_t(x) := \alpha_t g(x, y_t)$.

(Weighted) Regret of the x-player:

$$\alpha\text{-}Regret^x := \sum_{t=1}^T \alpha_t \ell_t(x_t) - \inf_{x \in X} \sum_{t=1}^T \alpha_t \ell_t(x).$$

(Weighted) Average regret of the x-player:

$$\overline{\alpha\text{−Regret}^x} := \frac{\alpha\text{-}Regret^x}{A_T},$$

where $A_T := \sum_{t=1}^T \alpha_t$.

From the y-player perspective:

- Play $y_t \in Y$.
- Receives the loss function at t, $h_t(y) := -\alpha_t g(x_t, y)$.

(Weighted) Regret of the y-player:

$$\alpha\text{-}Regret^y := \sum_{t=1}^{T} \alpha_t h_t(y_t) - \inf_{y \in Y} \sum_{t=1}^{T} \alpha_t h_t(y).$$

(Weighted) Average regret of the y-player:

$$\overline{\alpha\text{−Regret}^y} := \frac{\alpha\text{-}Regret^y}{A_T},$$

where $A_T := \sum_{t=1}^{T} \alpha_t$.

Guarantees of the meta-algorithm

Theorem 2. *Let $g(x, y)$ be convex w.r.t x and concave w.r.t. y. The output $(\overline{x}_T, \overline{y}_T)$ of the meta-algorithm is an ϵ-equilibrium of $g(\cdot, \cdot)$, where*

$$\epsilon := \overline{\alpha\text{−Regret}^x} + \overline{\alpha\text{−Regret}^y}.$$

Also, the duality gap is bounded as

$$\text{Gap}(\overline{x}_T, \overline{y}_T) := \sup_{y \in Y} g(\overline{x}_T, y) - \inf_{x \in X} g(x, \overline{y}_T) \leq \overline{\alpha\text{−Regret}^x} + \overline{\alpha\text{−Regret}^y}.$$

x-perspective $\ell_t(x) = g(x, y_t)$

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) = \frac{1}{A_T} \sum_{t=1}^{T} \alpha_t \ell_t(x_t)$$

This expression can further be decomposed into the infimum over x in X of the weighted outcomes, adjusted by the weighted regret for the x-player, and be simplified by using the definition of $\alpha\text{-}Regret^x$ and $\overline{\alpha\text{-}Regret^x}$:

$$= \inf_{x \in X} \left(\sum_{t=1}^{T} \frac{1}{A_T} \alpha_t g(x, y_t) \right) + \frac{\alpha\text{-}Regret^x}{A_T}$$

$$= \inf_{x \in X} \left(\sum_{t=1}^{T} \frac{1}{A_T} \alpha_t g(x, y_t) \right) + \overline{\alpha\text{-}Regret^x} \qquad (1)$$

Using the Jensen's inequality, we have

$$\leq \inf_{x \in X} g\left(x, \sum_{t=1}^{T} \frac{\alpha_t}{A_t} y_t\right) + \overline{\alpha\text{-}Regret^x} \tag{2}$$

$$\leq \sup_{y \in Y} \inf_{x \in X} g(x, y) + \overline{\alpha\text{-}Regret^x} \tag{3}$$

y-perspective $h_t(y) = -g(x_t, y)$

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) = \frac{1}{A_T} - \sum_{t=1}^{T} \alpha_t h_t(y_t)$$

This expression can further be decomposed into the infimum over y in Y of the weighted outcomes, adjusted by the weighted regret for the y-player, and be simplified by using the definition of $\alpha\text{-}Regret^y$ and $\overline{\alpha\text{-}Regret^y}$:

$$= -\inf_{y \in Y} \left(\sum_{t=1}^{T} \frac{1}{A_T} \alpha_t g(x, y_t)\right) - \frac{\alpha\text{-}Regret^y}{A_T}$$

$$= \sup_{y \in Y} \left(\sum_{t=1}^{T} \frac{1}{A_T} \alpha_t g(x_t, y)\right) - \overline{\alpha\text{-}Regret^y}$$

Using the Jensen's inequality, we have

$$\geq \sup_{y \in Y} g\left(\sum_{t=1}^{T} \frac{\alpha_t}{A_t} x_t, y\right) - \overline{\alpha\text{-}Regret^y} \tag{4}$$

$$\geq \inf_{x \in X} \inf_{y \in Y} g(x, y) - \overline{\alpha\text{-}Regret^y} \tag{5}$$

Thus, from (2) and (4), we have

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) \leq \inf_{x \in X} g\left(x, \sum_{t=1}^{T} \frac{\alpha_t}{A_t} y_t\right) + \overline{\alpha\text{-}Regret^x},$$

and

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) \geq \sup_{y \in Y} g\left(\sum_{t=1}^{T} \frac{\alpha_t}{A_t} x_t, y\right) - \overline{\alpha\text{-}Regret^y},$$

which implies that

$$\mathrm{Gap}(\bar{x}_T, \bar{y}_T) := \sup_{y \in Y} g(\bar{x}_T, y) - \inf_{x \in X} g(x, \bar{y}_T) \leq \overline{\alpha - \mathrm{Regret}^x} + \overline{\alpha - \mathrm{Regret}^y}.$$

First implication

Recall the Theorem:

Theorem 3. *Let $g(x, y)$ be convex w.r.t x and concave w.r.t. y. The output $(\overline{x}_T, \overline{y}_T)$ of the meta-algorithm is an ϵ-equilibrium of $g(\cdot, \cdot)$, where*

$$\epsilon := \overline{\alpha\text{−Regret}^x} + \overline{\alpha\text{−Regret}^y}.$$

Also, the duality gap is bounded as

$$\text{Gap}(\overline{x}_T, \overline{y}_T) := \sup_{y \in Y} g(\overline{x}_T, y) - \inf_{x \in X} g(x, \overline{y}_T) \leq \overline{\alpha\text{−Regret}^x} + \overline{\alpha\text{−Regret}^y}.$$

We have the following implication:
Let $g(x, y)$ be convex w.r.t x and concave w.r.t. y. If the descision space X and Y are convex and compact and $g(\cdot, \cdot)$ is Lipschitz continuous, then we know there are sublinear regret algorithms. This implies our second implication.

Second implication

Theorem 4. *Let X, Y be compact convex subsets of $\mathbb{R}^n$ and $\mathbb{R}^m$ respectively. Let $g(x, y) : X \times Y \to \mathbb{R}$ be convex in its first argument and concave in its second, and Lipschitz with respect to both. Then,*

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

Proof. From (3) and (5), we have

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) \leq \sup_{y \in Y} \inf_{x \in X} g(x, y) + \overline{\alpha\text{-}Regret^x}$$

and

$$\frac{1}{A_T} \sum_{t=1}^{T} \alpha_t g(x_t, y_t) \geq \inf_{x \in X} \inf_{y \in Y} g(x, y) - \overline{\alpha\text{-}Regret^y}$$

we can derive that

$$\sup_y \inf_x g(x, y) + \overline{\alpha\text{-}Regret^x} \geq \inf_x \sup_y g(x, y) - \overline{\alpha\text{-}Regret^y}$$

$$\Leftrightarrow \sup_y \inf_x g(x, y) + \overline{\alpha\text{-}Regret^x} + \overline{\alpha\text{-}Regret^y} \geq \inf_x \sup_y g(x, y)$$

Recall the following lemma in the last lecture:

Lemma 2. *Let $g(\cdot,\cdot) : X \times Y \to \mathbb{R}$, where X and Y are not empty. Then,*

$$\inf_{x\in X}\sup_{y\in Y} g(x,y) \geq \sup_{y\in Y}\inf_{x\in X} g(x,y)$$

Therefore, we get

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

□

The above result together with the following theorem that we saw in the last lecture imply that a saddle point exsits for when $g(x, y)$ is convex w.r.t x and concave w.r.t. y, $g(\cdot,\cdot)$ is Lipschitz continuous, and the descision space X and Y are convex and compact.

Theorem 5. *Let $g(x, y) : X \times Y \to \mathbb{R}$, where X and Y are not empty. A point (x_*, y_*) is a saddle point if and only if*

- *The supremum in $\sup_{y\in Y}\inf_{x\in X} g(x,y)$ is attained at y_* & the infimum in $\inf_{x\in X}\sup_{y\in Y} g(x,y)$ is attained at x_*.*
- *Also, $\sup_{y\in Y}\inf_{x\in X} g(x,y) = \inf_{x\in X}\sup_{y\in Y} g(x,y)$.*

4 Applications of the min-max theorem

Boosting as a bilinear game

Denote the training set $\{z_j \in \mathbb{R}^d, l_j = \{+1,-1\}\}_{j=1}^m$. Let $H := \{h_i(\cdot)\}_{i=1}^n$ be a set of prediction functions, i.e.,

$$h_i(\cdot) : \mathbb{R}^d \to \{+1,-1\}.$$

We can construct the misclassification matrix as

$$A_{i,j} = \begin{cases} 1 & \text{if } h_i(z_j) \neq l_j, \\ 0 & \text{otherwise.} \end{cases}$$

We have that

$$\min_{x\in\Delta_n}\max_{y\in\Delta_m} x^\top Ay := \min_{x\in\Delta_n}\max_{y\in\Delta_m} \sum_{i=1}^n\sum_{j=1}^m x[i]y[j]\mathbb{I}\{h_i(z_j) \neq l_j\}$$

Assume the existence of a weak learning oracle, i.e.,

$$\sum_{j=1}^{m} y[j]\mathbb{I}\{h_{i_*}(z_j) \neq l_j\} \leq \frac{1}{2} - \gamma,$$

where $\gamma > 0$. Here, i_* is the index of the predictor that gives a y-weighted error better than chance. Furthermore, for any $y \in \Delta_m$,

$$\min_{x\in\Delta_n} x^\top Ay \leq e_{i_*}^\top Ay \leq \frac{1}{2} - \gamma.$$

Recall $v_* = \max_{y\in\Delta_m} \min_{x\in\Delta_n} x^\top Ay$. These imply that

$$v_* \leq \frac{1}{2} - \gamma < \frac{1}{2}.$$

Thus,

$$\max_{y\in\Delta_m} \min_{x\in\Delta_n} x^\top Ay = v_* \leq \frac{1}{2} - \gamma.$$

As we know the Nash equilibrium/Saddle points (x_*, y_*) exist,

$$x^{*\top} Ay_* = v_* \leq \frac{1}{2} - \gamma.$$

The above implies that there exists $x_* \in \Delta_n$ such that

$$\forall j \in [m] : \sum_{i=1}^{n} x_*[i]\mathbb{I}\{h_i(z_j) \neq l_j\} = x^{*\top} Ae_j \leq v_* \leq \frac{1}{2} - \gamma < \frac{1}{2}.$$

Less than half of the base predictors misclassify when weighted by $x_*[i]$ for each sample $j \in [m]$. The above implies that

$$\sum_{i=1}^{n} x_*[i]\mathbb{I}\{h_i(z_j) \neq l_j\} = x^{*\top} Ae_j \leq v_* \leq \frac{1}{2} - \gamma < \frac{1}{2}.$$

We can correctly classify all the samples using a weighted majority vote.

5 Meta-algorithm for solving min-max problems (Simultaneously Play)

Instance of the meta-algorithm

$$OAlg^x = FTRL, OMD, OptimisticMD, ...$$

Algorithm 3 Meta-algorithm for solving min-max problems (Simultaneously Play)

1: OAlg^x (OCO Alg. of x) and OAlg^y (OCO Alg. of y).
2: Weight sequence $\alpha_1, \alpha_2, \dots, \alpha_T$.
3: **for** $t = 1, 2, \dots, T$ **do**
4: $\begin{cases} \text{x plays } x_t \leftarrow \text{OAlg}^x(\alpha_1 \ell_1, \alpha_2 \ell_2, \dots, \alpha_{t-1}\ell_{t-1}) \\ \text{y plays } y_t \leftarrow \text{OAlg}^y(\alpha_1 h_1, \alpha_2 h_2, \dots, \alpha_{t-1} h_{t-1}) \end{cases}$
5: $\begin{cases} \text{x receives } \alpha_t \ell_t(x) := \alpha_t g(x, y_t) \\ \text{y receives } \alpha_t h_t(y) := -\alpha_t g(x_t, y) \end{cases}$
6: **end for**
7: Output: $\left(\overline{x}_T := \frac{\sum_{t=1}^T \alpha_t x_t}{A_T}, \overline{y}_T := \frac{\sum_{t=1}^T \alpha_t y_t}{A_T} \right)$, where $A_T := \sum_{t=1}^T \alpha_t$.

$$OAlg^y = FTRL, OMD, OptimisticMD, ...$$

Assume that $\alpha_t = 1$ and $\overline{x_T}$ and $\overline{y_T}$ are $\epsilon-$equilibrium points

$$\epsilon = \frac{\textbf{Regret}_T(\text{OMD})}{T} + \frac{\textbf{Regret}_T(\text{OMD})}{T} = \frac{\mathcal{O}(\sqrt{T})}{T} \to 0, \text{ as } T \to \infty$$

Question: Can we get a better rate than $\mathcal{O}(\frac{1}{\sqrt{T}})$? Yes!

Recall Online Mirror Descent

The function $\ell_t(z)$ is convex but not necessarily differentiable. $g_t \in \partial \ell_t(z_t)$ is the subgradient of $\ell_t(\cdot)$ at z_t.

Algorithm 4 Online Mirror Descent

1: **for** $t = 1, 2, \dots$ **do**
2: $z_{t+1} = \arg\min_{z \in C} \langle g_t, z - z_t \rangle + \frac{1}{\eta} D^\phi_{z_t}(z)$.
3: **end for**

Mirror Descent has

$$\sum_{t=1}^T \ell_t(z_t) - \ell_t(z^*) \le \frac{1}{\eta} D^\phi_{z_1}(z^*) + \sum_{t=1}^T \frac{\eta}{2} \|g_t\|_*^2,$$

for any benchmark $z^* \in Z$.
If the loss $\ell_t(\cdot)$ is scaled by α_t,

$$\alpha\text{-Regret}_z(z^*) \le \frac{1}{\eta} D^\phi_{z_1}(z^*) + \sum_{t=1}^T \frac{\eta}{2} \|\alpha_t g_t\|_*^2,$$

for any benchmark $z^* \in Z$.
Assume there is a good guess m_t of g_t.

Algorithm 5 Optimistic Mirror Descent

1: **for** $t = 1, 2, \ldots$ **do**
2: $\quad z_{t-\frac{1}{2}} = \arg\min_{z \in C} \alpha_{t-1}\langle g_{t-1}, z\rangle + \frac{1}{\eta} D^{\phi}_{z_{t-\frac{3}{2}}}(z)$.
3: $\quad z_t = \arg\min_{z \in C} \alpha_t \langle m_t, z\rangle + \frac{1}{\eta} D^{\phi}_{z_{t-\frac{1}{2}}}(z)$.
4: **end for**

We have that

$$\alpha\text{-Regret}^z(z^*) \leq \frac{1}{\eta} D^{\phi}_{z_1}(z_*) + \sum_{t=1}^{T} \frac{\eta}{2} \|\alpha_t(g_t - m_t)\|_*^2,$$

for any benchmark $z^* \in Z$.
By putting two Optimistic Mirror Descent against each other, we can get $\mathcal{O}(\frac{1}{T})$ in a min-max problem, see e.g., [3] for details.

6 Bibliographic notes

More materials about min-max optimization can be found in [1],[2],[3],[4].

References

[1] Francesco Orabona, *A Modern Introduction to Online Learning*, Chapter 11.

[2] Jun-Kun Wang, Jacob Abernethy, and Kfir Y. Levy, *No-Regret Dynamics in the Fenchel Game: A Unified Framework for Algorithmic Convex Optimization*, Mathematical Programming, 2023.

[3] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire, *Fast Convergence of Regularized Learning in Games*, NeurIPS 2015.

[4] Robert E. Schapire and Yoav Freund, *Boosting: Foundations and Algorithms*, MIT Press, 2012.