From Convex Analysis to Learning, Prediction, and Elicitation*
Lecture 9: Blackwell Approachability and Online Calibration

Lunjia Hu

1 Online Multi-objective Optimization

We consider the following online learning problem specified by three sets: learner's action set X.

In each round $t = 1, \ldots, T$:

1. Learner chooses $x_t \in X$;
2. Adversary reveals $y_t \in Y$.

The goal of the learner is to minimize the following quantity, where $Z \subseteq \mathbb{R}^d$ is a set of distinguishers, and $u : X \times Y \to \mathbb{R}^d$ is some fixed function known to the learner:

$$L(x_1, \ldots, x_T; y_1, \ldots, y_T) := \sup_{z \in Z} \left\langle \frac{1}{T} \sum_{t=1}^{T} u(x_t, y_t), z \right\rangle$$

Assumption. for every $z \in Z$, there exists $x \in X$ such that $\sup_{y \in Y} \langle u(x, y), z \rangle \le w$, where $w \in \mathbb{R}$ is some fixed and known threshold.

Algorithm 1. Online Multi-objective Optimization.

1. Use a low-regret algorithm (e.g. FTRL) to choose $z_t \in Z$.
2. Play $x_t \in X$ such that $\sup_{y \in Y} \langle u(x_t, y), z_t \rangle \le t$.
3. Observe $y_t \in Y$ from the adversary.

Suppose z_t's are chosen so that the following low-regret guarantee is satisfied:

$$\sup_{z \in Z} \left\langle \frac{1}{T} \sum_{t=1}^{T} u(x_t, y_t), z \right\rangle - \left\langle \frac{1}{T} \sum_{t=1}^{T} u(x_t, y_t), z_t \right\rangle \le \varepsilon.$$

Now we have

$$\begin{aligned} L(x_1, \ldots, x_T; y_1, \ldots, y_T) &= \sup_{z \in Z} \left\langle \frac{1}{T} \sum_{t=1}^{T} u(x_t, y_t), z \right\rangle \\ &\le \left\langle \frac{1}{T} \sum_{t=1}^{T} u(x_t, y_t), z_t \right\rangle + \varepsilon \\ &\le w + \varepsilon. \end{aligned}$$

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

Remark 1. *In many cases, the functions $f_z(x,y) := \langle u(x,y), z\rangle$ have the minimax property:*

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

2 Online Calibration

In each round $t = 1, \ldots, T$:

1. Predictor chooses distribution a τ_t of predictions $p \in [0,1]$;
2. Adaptive adversary reveals outcome $y_t \in \{0,1\}$;
3. Predictor's prediction p_t is sampled from τ_t.

Consider making discretized predictions among $1/m, 2/m, \ldots, 1$. Our goal is to achieve

$$\mathbb{E}[\mathsf{ECE}(p_1,\ldots,p_T; y_1,\ldots,y_T)] = O\left(\sqrt{\frac{m}{T}} + \frac{1}{m}\right).$$

Choosing $m \approx T^{1/3}$ gives $\mathsf{ECE} = O(T^{-1/3})$. ECE can be calculated as follows:

$$\begin{aligned}
\mathsf{ECE}(p_1,\ldots,p_T; y_1,\ldots,y_T) &= \frac{1}{T}\sum_{i=1}^m \left|\sum_{t=1}^T (y_t - \frac{i}{m})\mathbb{I}[p_t = \frac{i}{m}]\right| \\
&= \frac{1}{T}\sum_{i=1}^m \sum_{t=1}^T \sup_{z_i\in[-1,1]} (y_t - \frac{i}{m})\mathbb{I}[p_t = \frac{i}{m}]z_i \\
&= \sup_{z\in[-1,1]^m} \frac{1}{T}\sum_{t=1}^T \sum_{i=1}^m (y_t - \frac{i}{m})\mathbb{I}[p_t = \frac{i}{m}]z_i \\
&= \sup_{z\in[-1,1]^m} \left\langle \frac{1}{T}\sum_{t=1}^T u_t, z\right\rangle,
\end{aligned}$$

where

$$u_t = \left(\left(y_t - \frac{1}{m}\right)\mathbb{I}\left[p_t = \frac{1}{m}\right], \left(y_t - \frac{2}{m}\right)\mathbb{I}\left[p_t = \frac{2}{m}\right], \ldots, (y_t - 1)\,\mathbb{I}\left[p_t = 1\right]\right).$$

For $x \in \Delta_m$ and $y \in \{0,1\}$, define

$$u(x,y) := \left(\left(y - \frac{1}{m}\right)x_1, \left(y - \frac{2}{m}\right)x_2, \ldots, (y-1)x_m\right).$$

We have $\|u(x,y)\|_1 \le 1$. We set $Z = [-1,1]^m$.

Claim 1. *For every $z \in [-1,1]^m$, there exists $x \in \Delta_m$ such that*

$$\sup_{y\in\{0,1\}} \langle u(x,y), z\rangle \le \frac{1}{m}.$$

Proof. Let $z \in [-1,1]^m$ be arbitrary. By the minimax theorem, it suffices to prove that for every distribution π on $\{0,1\}$, there exists $x \in \Delta_m$ such that

$$\mathbb{E}_{y\sim\pi}\langle u(x,y), z\rangle \le \frac{1}{m}.$$

Let i/m be the value closest to $\mathbb{E}_\pi[y]$ among $\{1/m, 2/m, \dots, 1\}$. We simply choose $x = \mathbf{e}_i := (0,\dots,0,1,0,\dots,0)$ where the value 1 is at the i-th coordinate. Now we have

$$\mathbb{E}_{y\sim\pi}\langle u(x,y), z\rangle = \mathbb{E}_{y\sim\pi}\left[\left(y - \frac{i}{m}\right) z_i\right] = \left(\mathbb{E}_\pi[y] - \frac{i}{m}\right) z_i \le \left|\mathbb{E}_\pi[y] - \frac{i}{m}\right| \cdot |z_i| \le \frac{1}{m}\cdot 1 = \frac{1}{m}.$$

□

Claim 2. *There is an (efficient) low-regret online algorithm for choosing $z_1,\dots,z_T \in [-1,1]^m$ that guarantees the following regret bound, regardless of how $x_1,\dots,x_T \in \Delta_m$ and $y_1,\dots,y_T \in \{0,1\}$ are chosen:*

$$\sup_{z\in Z}\left\langle \frac{1}{T}\sum_{t=1}^T u(x_t,y_t), z\right\rangle - \left\langle \frac{1}{T}\sum_{t=1}^T u(x_t,y_t), z_t\right\rangle \le O(\sqrt{m/T}).$$

Proof. Claim 2 is a standard regret bound for online linear optimization (OLO), where the learner's action z comes from $[-1,1]^m$, and the adversary's action $u(x,y)$ comes from $\bar{B}_{\ell_1}(\mathbf{0},1)$. Specifically, consider running Follow the Regularized Leader (FTRL) with regularizer $\varphi(z) = \frac{1}{2}\|z\|_2^2$. Clearly, φ is bounded between 0 and $m/2$ on $[-1,1]^m$. Moreover, it is 1-strongly convex w.r.t. the ℓ_2 norm, and thus 1-strongly convex also w.r.t. the ℓ_∞ norm (which is the dual of the ℓ_1 norm in which $u(x,y)$ is bounded). Therefore, the total regret of T rounds of FTRL with learning rate η is at most

$$\frac{m}{2\eta} + \frac{\eta T}{2}.$$

Choosing $\eta = \sqrt{m/T}$ gives a regret bound of $\sqrt{mT}$. Thus the average regret over T rounds is at most $\sqrt{mT}/T = \sqrt{m/T}$. □

Combining Claim 1 and Claim 2, we know that Algorithm 1 guarantees

$$L(x_1,\dots,x_T; y_1,\dots,y_T) := \sup_{z\in[-1,1]^m}\left\langle \frac{1}{T}\sum_{t=1}^T u(x_t,y_t), z\right\rangle = O\left(\sqrt{\frac{m}{T}} + \frac{1}{m}\right).$$

Now in each round t, we choose τ_t to be the distribution of $p \in [0,1]$ corresponding to x_t. Concretely, we set $\Pr_{p\sim\tau_t}[p = i/m]$ to be x_i for every $i = 1,\dots,m$. Since each p_t is drawn from τ_t, by standard martingale concentration inequalities, we can show that

$$\mathbb{E}|\mathsf{ECE}(p_1,\dots,p_T; y_1,\dots,y_T) - L(x_1,\dots,x_T; y_1,\dots,y_T)| = O\left(\sqrt{\frac{m}{T}}\right).$$

Combining the two equations above, we get

$$\mathbb{E}[\mathsf{ECE}(p_1,\dots,p_T; y_1,\dots,y_T)] = O\left(\sqrt{\frac{m}{T}} + \frac{1}{m}\right).$$