Lecture 17: Min-max vs. max-min techniques

Lecturer: Yanjun Han

May 24, 2021

Today's plan

most previous lower bound approaches:

$$\inf_{T} \sup_{\theta\in\Theta} \mathbb{E}_\theta[L(\theta, T)] \ge \inf_{T} \mathbb{E}_{\theta\sim\pi}\mathbb{E}_\theta[L(\theta, T)]$$

- choose the prior π first (usually the uniform distribution over two or multiple points), then argue that any T incurs a large average risk
- however, in principle this process could be reversed: first fix any T, then argue that there is some $\theta \in \Theta$ with a large pointwise risk
- latter approach could be useful if the observations depend on the mechanism design, e.g.
  - sequential experiments (bandits)
  - distributed estimation and testing
- three examples today, with no golden rule

Example I: dynamic pricing

- setting (posted-price auction): one seller with a single good, and n customers with iid valuations $v_1, \cdots, v_n \in [0,1]$
- at each time $t \in [n]$:
  - seller posts a price $p_t \in [0,1]$
  - customer t comes, and buys the good iff $v_t \geq p_t$
  - seller only observes the binary outcome $1(v_t \geq p_t)$
- seller's target: minimize the regret

$$R_n(\pi) = n \cdot \max_{p^\star \in [0,1]} p^\star D(p^\star) - \mathbb{E}\left[\sum_{t=1}^{n} p_t D(p_t)\right]$$

where $D(p) = \mathbb{P}(v_1 \geq p)$ denotes the unknown demand curve

- an additional assumption: there is a constant $\beta > 0$ such that

$$p^\star D(p^\star) - pD(p) \geq \beta (p^\star - p)^2$$

The main difficulty

Theorem (Kleinberg and Leighton, 2003)

$$R_n^\star = \Theta(\sqrt{n})$$

- try two-point analysis and find two demand curves D_1, D_2
- separation condition: if the maximizers $p_1^\star$ and $p_2^\star$ differ by ε, then separation condition holds with parameter $\Omega(n\varepsilon^2)$
- indistinguishability condition: KL divergence at time t

$$D_{\mathrm{KL}}(\mathrm{Bern}(D_1(p_t)) \| \mathrm{Bern}(D_2(p_t))) \asymp (D_1(p_t) - D_2(p_t))^2$$

- problem: the above quantity is $O(\varepsilon^4)$ when $p_t \in [p_1^\star, p_2^\star]$, but could be very large when p_t is far from optimal

Previous analysis

- solution: when p_t is far from optimal, the regret is large too
- argument of Kleinberg and Leighton: define "conditional regret" and "knowledge", and show that the former is lower bounded by the latter
- however, this does not fit into two-point analysis again

$$\mathbb{E}_{\theta_1}[L(\theta_1, T)] + \mathbb{E}_{\theta_2}[L(\theta_2, T)]$$
$$\not\gtrsim (1 - \|P_{\theta_1} - P_{\theta_2}\|_{\mathrm{TV}}) \cdot \mathbb{E}_{\theta}[L(\theta_1, T) + L(\theta_2, T)]$$

$L(\theta_1, a) P_{\theta_1}(x) + L(\theta_2, a) P_{\theta_2}(x) \geq \min\{P_{\theta_1}(x), P_{\theta_2}(x)\}(L(\theta_1, a) + L(\theta_2, a))$

$\min_a L(\theta_1, a) + L(\theta_2, a)$

- very involved arguments to make the insights work

Our new analysis

- new analysis: same intuition, but satisfied with the smallest separation parameter $\Omega(n\varepsilon^2)$
- key idea: argue that the KL divergence must be small to ensure a small regret under D_1
- in other words, will prove something like

$$R_n^\star \geq f(R_n^\star)$$

- construction of demand curves:

$$pD_{\pm}(p) = \begin{cases} \left[1 - 4(\frac{1}{4} \pm \varepsilon)^2\right] p & \text{if } p \leq \frac{1}{4} \\ \frac{1}{4} - (\frac{1}{2} \pm \varepsilon - p)^2 & \text{if } \frac{1}{4} < p \leq \frac{3}{4} \\ \frac{1}{16} + \left[\frac{3}{4} - 4(\frac{1}{4} \mp \varepsilon)^2\right](1 - p) & \text{if } p > \frac{3}{4} \end{cases}$$

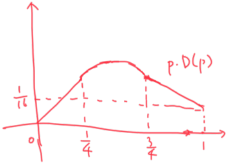

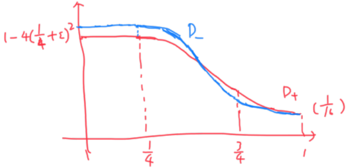

- separation parameter: $\Omega(n\varepsilon^2)$
- both D_+, D_- bounded away from zero and one

Detailed analysis (cont'd)

- indistinguishability condition:

$$D_{\mathrm{KL}}(D_+^{\otimes n} \| D_-^{\otimes n}) \asymp \sum_{t=1}^{n} \mathbb{E}_{D_+}(D_+(p_t) - D_-(p_t))^2$$

$\lesssim \varepsilon^4 + \varepsilon^2 (p_t - \frac{1}{2})^2$

$$\lesssim n\varepsilon^4 + \varepsilon^2 \cdot \sum_{t=1}^{n} \mathbb{E}_{D_+}(p_t - \frac{1}{2})^2$$

- regret lower bound under D_+:

$$R_n^\star \geq \sum_{t=1}^{n} (p^\star D_+(p^\star) - \mathbb{E}_{D_+}[p_t D_+(p_t)])$$

$$\geq \sum_{t=1}^{n} \mathbb{E}_{D_+}(p_t - \frac{1}{2} - \varepsilon)^2 \geq \frac{1}{2} \sum_{t=1}^{n} \mathbb{E}_{D_+}(p_t - \frac{1}{2})^2 - n\varepsilon^2$$

- two-point method:

$$R_n^\star \gtrsim n\varepsilon^2 \exp\left(-c(n\varepsilon^4 + \varepsilon^2 R_n^\star)\right)$$

$\varepsilon \sim n^{-1/4}$

Example II: batched bandit problem

- usual K-armed bandit problem, but with at most M batches
- batch constraint: design a grid $0 = t_0 < t_1 < \cdots < t_M = T$, need to determine the actions $(\pi_t)_{t_{j-1}+1 \le t \le t_j}$ in each batch simultaneously
- two types of grids:
  - static grid: $\{t_1, \cdots, t_M\}$ is fixed ahead of time
  - adaptive grid: t_j could be determined by all outcomes up to time t_{j-1}

Claim (Gao, Han, Ren, and Zhou, 2019)

$$R^\star_{M,K,T} \gtrsim \begin{cases} \sqrt{K} \cdot T^{\frac{1}{2-2^{1-M}}} & \text{for static grid,} \\ \frac{\sqrt{K}}{M^2} \cdot T^{\frac{1}{2-2^{1-M}}} & \text{for adaptive grid.} \end{cases}$$

Lower bound under static grid

- same construction as classical bandit:

$$\begin{aligned}
\mu_1 &= (\Delta, 0, 0, \cdots, 0) \\
\mu_2 &= (\Delta, 2\Delta, 0, \cdots, 0) \\
&\vdots \\
\mu_K &= (\Delta, 0, 0, \cdots, 2\Delta)
\end{aligned}$$

- tree-based inequality gives (see Lecture 10)

$t_{j-1} < t \le t_j$

$\Rightarrow$ # obs $\le t_{j-1}$

$$R^\star_{M,K,T} \gtrsim \Delta \sum_{j=1}^{M} (t_j - t_{j-1}) \exp\left(-\frac{t_{j-1}\Delta^2}{K}\right)$$

- for a static grid, choosing $\Delta = \sqrt{K/t_j}$ with $j = 1, \cdots, M-1$ gives

$$R^\star_{M,K,T} \gtrsim \sqrt{K} \cdot \max\left\{t_1, \frac{t_2}{\sqrt{t_1}}, \cdots, \frac{T}{\sqrt{t_{M-1}}}\right\} \gtrsim \sqrt{K} \cdot T^{\frac{1}{2-2^{1-M}}}$$

How about adaptive grid?

- significantly more complicated. why?
- previous analysis breaks down even for randomized grid:

$$R^\star_{M,K,T} \gtrsim \Delta \cdot \mathbb{E}\left[\sum_{j=1}^{M}(t_j - t_{j-1})\exp\left(-\frac{t_{j-1}\Delta^2}{K}\right)\right]$$

not convex in $(t_1, \cdots, t_M)$

we don't know how to choose Δ

- even worse for general adaptive grid:

$\mathbb{E}_\Delta$

$$R^\star_{M,K,T} \gtrsim \Delta \sum_{t=1}^{T}\exp\left(-\frac{\mathbb{E}[\text{last grid point before } t]\cdot\Delta^2}{K}\right)$$

  - problem I: expectation depends on Δ
  - problem II: even for $M = 2$, if $t_1 \sim \mathrm{Unif}([T])$, resulting lower bound only of the order $\Omega(\sqrt{KT})$

High-level idea

- if t_1 is too large: large regret in first batch
- if t_1 is too small and t_2 is too large: large regret in second batch
- $\cdots$
- if t_{M-1} is too small: large regret in last batch
- the above events form a partition of the entire space

Carrying out the idea

- choosing $T_0 = 0, T_1, \cdots, T_M = T$ appropriately, define M disjoint events:
$$A_j = \{t_{j-1} \leq T_{j-1} \text{ and } t_j \geq T_j\}, \quad j \in [M].$$
- if any A_j holds, choose hypotheses $(P_{j,1}, \cdots, P_{j,M})$ depending on j
  - note: the lower bound will depend on $P_{j,1}(A_j)$
- argue that at least one of A_j will hold with large probability
  - note: not direct pigeon-hole principle as we need to deal with $P_{j,1}(A_j)$
- choice of hypothesis:
$$\mu_{j,1} = (\Delta_j + \Delta_M, 0, \cdots, 0, \Delta_M)$$
$$\mu_{j,k} = (\Delta_j + \Delta_M, 0, \cdots, 0, 3\Delta_j, 0, \cdots, \Delta_M)$$
- perform hypothesis testing after fixing the policy

First step: if one of A_j holds

Lemma

If $P_{j,1}(A_j) \geq 1/(2M)$, then the target lower bound holds for π.

- again the tree-based lower bound
- a different lower bound on the TV distance: for $t \leq T_j$

$$\begin{aligned}
\sum_x \min\{P_{j,1}^t(x), P_{j,k}^t(x)\} &\geq \sum_x \min\{P_{j,1}^{T_j}(x), P_{j,k}^{T_j}(x)\} \\
&\geq \sum_{x \in A_j} \min\{P_{j,1}^{T_j}(x), P_{j,k}^{T_j}(x)\} \\
&= \sum_{x \in A_j} \min\{P_{j,1}^{T_{j-1}}(x), P_{j,k}^{T_{j-1}}(x)\} \\
&\geq P_{j,1}(A_j) - \frac{3}{2}\|P_{j,1}^{T_{j-1}} - P_{j,k}^{T_{j-1}}\|_{\mathrm{TV}}
\end{aligned}$$

data-processing

$x \in A_j = \{t_{j-1} \leq T_{j-1},\ t_j \geq T_j\}$

$\min\{x, y\} = \frac{x+y-|x-y|}{2}$

Second step: A_j's roughly form a partition

- a slight change in the construction: for $\ell \in [K-1]$, define

$$\mu_{j,1}^{\ell} = (0, \cdots, 0, \Delta_j + \Delta_M, 0, \cdots, 0, \Delta_M)$$

- previous lemma remains unchanged for each ℓ

lemma

$$\frac{1}{K-1} \sum_{\ell=1}^{K-1} \sum_{j=1}^{M} P_{j,1}^{\ell}(A_j) \geq \frac{1}{2}.$$

- key argument:

$\mu_M = (0, 0, \cdots, 0, \Delta_M)$

$$|P_{j,1}^{\ell}(A_j) - P_M(A_j)| \leq \|P_{j,1}^{\ell, T_{j-1}} - P_M^{T_{j-1}}\|_{\mathrm{TV}} \leq \sqrt{\frac{(\Delta_j + \Delta_M)^2}{4} \cdot \mathbb{E}_{P_M}[\tau_{j,\ell}]}$$

$t_{j-1} \leq T_{j-1}$, $t_j \geq T_j$

- average over ℓ and use $\sum_{\ell} \tau_{j,\ell} \leq T_{j-1}$

Example III: distributed uniformity testing

- same setting as Lecture 14
- each of n nodes holds $X_i \sim p = (p_1, \cdots, p_d)$, has k bits to communicate to the central server, and would like to test whether $p = \mathrm{Unif}_d$ or $\|p - \mathrm{Unif}_d\|_{\mathrm{TV}} \geq \delta$
- table of sample complexities:

SMP w/ public coin	$\frac{\sqrt{d}}{\delta^2} \cdot \sqrt{\frac{d}{\min\{d, 2^k\}}}$
SMP w/ private coin	$\frac{\sqrt{d}}{\delta^2} \cdot \frac{d}{\min\{d, 2^k\}}$
SMP w/ s-bit shared randomness	$\frac{\sqrt{d}}{\delta^2} \cdot \sqrt{\frac{d}{\min\{d, 2^k\}}} \cdot \sqrt{\frac{d}{\min\{d, 2^{k+s}\}}}$

High-level idea

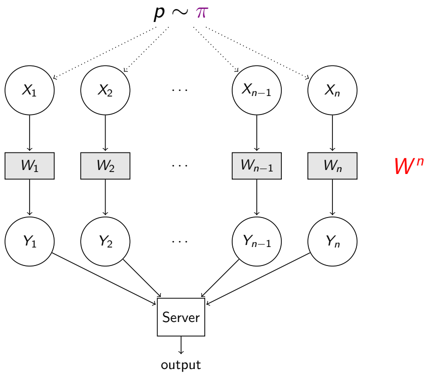

Learner: choose communication channel $W^n = (W_1, \cdots, W_n)$ to **perform** constrained inference.

Adversary: choose prior π on the underlying distribution p to **confuse** the learner.

Role of shared randomness:

- without shared randomness: W^n is a product channel;
- with shared randomness: W^n is a mixture of product channels.

Semi-max-min information

- Let $I(W^n \to \pi)$ be a suitable notion of "information" provided by a given channel W^n to a given prior π.
- Semi-max-min information:

$$\underline{\overline{I}} = \max_{\mathcal{W}:|\mathcal{W}|=2^s} \min_{\pi} \mathbb{E}_{W^n \sim \mathrm{Unif}(\mathcal{W})}[I(W^n \to \pi)].$$

- $s = 0$ gives the max-min information for private randomness:

$$\underline{\overline{I}} \geqslant \underline{I} = \max_{W^n} \min_{\pi} I(W^n \to \pi).$$

- $s = \infty$ gives the min-max information for public randomness:

$$\underline{\overline{I}} \leqslant \overline{I} = \min_{\pi} \max_{W^n} I(W^n \to \pi).$$

More details

- choice of the prior (simplified): for $v \sim \mathrm{Unif}(\{\pm 1\}^{d/2})$, set

$$p_v = \left(\frac{1+\delta u_1 v_1}{d}, \frac{1-\delta u_1 v_1}{d}, \cdots, \frac{1+\delta u_{d/2} v_{d/2}}{d}, \frac{1-\delta u_{d/2} v_{d/2}}{d} \right)$$

- the choice of the information $I(W^n \to \pi)$:

$$\begin{aligned} I(W^n \to \pi) &= \log(1 + \chi^2(\mathbb{E}_v[W^n \circ p_v], W^n \circ p_0)) \\ &\lesssim \frac{n^2\delta^4}{d^2} \cdot (u^\top H(W^n) u)^2 \end{aligned}$$

- semi-max-min information:

$$\begin{aligned} \underline{I} &\lesssim \frac{n^2\delta^4}{d^2} \max_{\mathcal{W}:|\mathcal{W}|=2^s} \min_u \mathbb{E}_{W^n \sim \mathrm{Unif}(\mathcal{W})} (u^\top H(W^n) u)^2 \\ &\le \frac{n^2\delta^4}{d^2} \max_{W^n} \min_u 2^s \cdot (u^\top H(W^n) u)^2 = \frac{n^2\delta^4 2^s}{d^2} \max_{W^n} \lambda_{\min}(H(W^n))^2 \end{aligned}$$

References

- Robert Kleinberg, and Tom Leighton. "The value of knowing a demand curve: Bounds on regret for online posted-price auctions." In *44th Annual IEEE Symposium on Foundations of Computer Science*, 2003.
- Zijun Gao, Yanjun Han, Zhimei Ren, and Zhengqing Zhou. "Batched multi-armed bandits problem." *NeurIPS*, 2019.
- David Simchi-Levi, and Yunzong Xu. "Phase Transitions and Cyclic Phenomena in Bandits with Switching Constraints." *arXiv preprint* arXiv:1905.10825v4 (2021).
- Jayadev Acharya, Clément L. Canonne, Yanjun Han, Ziteng Sun, and Himanshu Tyagi. "Domain compression and its application to randomness-optimal distributed goodness-of-fit." In *Conference on Learning Theory*, pp. 3-40. PMLR, 2020.

Next (final) lecture: adaptation lower bounds