Lecture 7: Le Cam's two-point method

Lecturer: Yanjun Han

April 19, 2021

The idea of testing

$$\inf_T \sup_{\theta\in\Theta} \mathbb{E}_\theta[L(\theta, T)] \geq \inf_T \sup_{\theta\in\{\theta_1,\cdots,\theta_m\}} \mathbb{E}_\theta[L(\theta, T)]$$

$m = 2$: two-point method

- point vs. point P_{θ_1} vs. P_{θ_2}
- point vs. mixture P_{θ_1} vs. $\mathbb{E}[P_{\theta_2}]$
- mixture vs. mixture $\mathbb{E}[P_{\theta_1}]$ vs. $\mathbb{E}[P_{\theta_2}]$

$m > 2$: testing multiple hypotheses

- Fano and Assouad
- Global Fano with covering and packing

Today's plan

Two-point method:

- derivation of tools from the first principle
- example in (robust) parameter estimation
- example in functional estimation
- example in bandits
- example in stochastic optimization

Start from the first principle

$$\max_{\theta \in \{\theta_0, \theta_1\}} \mathbb{E}_\theta [L(\theta, T)] = \max \left\{ \sum_{x \in \mathcal{X}} L(\theta_0, T(x)) \cdot P_{\theta_0}(x), \sum_{x \in \mathcal{X}} L(\theta_1, T(x)) P_{\theta_1}(x) \right\}$$

$$\geq \frac{1}{2} \sum_{x \in \mathcal{X}} L(\theta_0, T(x)) P_{\theta_0}(x) + \frac{1}{2} \sum_{x \in \mathcal{X}} L(\theta_1, T(x)) \cdot P_{\theta_1}(x)$$

$$= \frac{1}{2} \sum_{x \in \mathcal{X}} \left(\underbrace{L(\theta_0, T(x)) \cdot P_{\theta_0}(x) + L(\theta_1, T(x)) \cdot P_{\theta_1}(x)}_{} \right)$$

$L(\theta_0, a) + L(\theta_1, a) \geq \Delta$ separation condition

$$\geq \frac{1}{2} \sum_{x \in \mathcal{X}} \Delta \cdot \min \{ P_{\theta_0}(x), P_{\theta_1}(x) \}$$

$$= \frac{\Delta}{2} \left(1 - \| P_{\theta_0} - P_{\theta_1} \|_{TV} \right)$$

Two-point method: statement

Two-point method

Fix any $\theta_0, \theta_1 \in \Theta$. Suppose that the following separation condition holds:

$$\min_{a \in \mathcal{A}} L(\theta_0, a) + L(\theta_1, a) \geq \Delta > 0.$$

Then

indistinguishability

$$\inf_T \max_{\theta \in \{\theta_0, \theta_1\}} \mathbb{E}_\theta[L(\theta, T(X))] \geq \frac{\Delta}{2} \left(1 - \|P_{\theta_0} - P_{\theta_1}\|_{\mathsf{TV}}\right).$$

- in many scenarios, further upper bound TV by Hellinger, KL, χ^2, etc

Two-point method: discussions

- the simplest tool in testing-based lower bound
- works better if the target to be estimated is low-dimensional
- key difficulty when applying: finding the appropriate two points

Example I: normal mean estimation

- $X \sim \mathcal{N}(\theta, \sigma^2)$ with known σ
- target: estimate θ under quadratic loss $L(\theta, a) = (\theta - a)^2$
- choice of two points: $\theta_0 = 0, \theta_1 = \delta$
- two-point method:

$L(\theta_0, a) + L(\theta_1, a) = a^2 + (\delta - a)^2 \geq \frac{\delta^2}{2} \quad (:= \Delta)$

$\| P_{\theta_0} - P_{\theta_1} \|_{TV} = \quad 2\Phi\left(\frac{|\delta|}{2\sigma}\right) - 1$

$$R_\sigma^\star \geq \underline{\frac{\delta^2}{2}\left(1 - \Phi\left(\frac{|\delta|}{2\sigma}\right)\right)}$$

- maximizing over δ gives $R_\sigma^\star \geq \underline{0.332\sigma^2}$

$(R_\sigma^* = \sigma^2)$

Example II: robust normal mean estimation

- Huber's ε-contaminated model:

$$X_1, \cdots, X_n \sim P_{(\theta,Q)} = (1-\varepsilon)\cdot\mathcal{N}(\theta,1) + \varepsilon\cdot Q$$

- target: estimate θ under quadratic loss $L(\theta, a) = (\theta - a)^2$
- claim:

$$R^{\star}_{\sigma,\varepsilon} = \Omega\left(\frac{1}{n} + \varepsilon^2\right)$$

- idea of new lower bound: find (θ_0, θ_1) and (Q_0, Q_1) such that

$$P_{(\theta_0,Q_0)} = P_{(\theta_1,Q_1)}$$

Construction of two points

Lemma

Let P_0, P_1 be two probability distributions with $\|P_0 - P_1\|_{\mathsf{TV}} \le \varepsilon/(1-\varepsilon)$. Then there exist distributions Q_0, Q_1 such that

$$(1-\varepsilon)P_0 + \varepsilon Q_0 = (1-\varepsilon)P_1 + \varepsilon Q_1.$$

$\forall x: \quad Q_0(x) - Q_1(x) = \frac{1-\varepsilon}{\varepsilon}(P_1(x) - P_0(x))$

$$\begin{cases} Q_0(x) = \frac{1-\varepsilon}{\varepsilon}\max\{P_1(x) - P_0(x), 0\} \\ Q_1(x) = \frac{1-\varepsilon}{\varepsilon}\max\{P_0(x) - P_1(x), 0\} \end{cases}$$

$\frac{1-\varepsilon}{\varepsilon}\cdot\|P_0 - P_1\|_{TV} = \sum_x Q_0(x) = \sum_x Q_1(x) \le 1$

Original problem: find the maximum $\delta > 0$ with

$$\|\mathcal{N}(0,1) - \mathcal{N}(\delta,1)\|_{\mathsf{TV}} \le \varepsilon/(1-\varepsilon) \Longrightarrow \delta \asymp \varepsilon$$

Example III: nonparametric estimation at a point

$s \le 1$: $|f(x) - f(y)| \le L \cdot |x-y|^s$, $\forall x, y \in [0,1]$

$s > 1$: $|f^{(m)}(x) - f^{(m)}(y)| \le L|x-y|^{s-m}$. $m = \lfloor s \rfloor$

- $X_1, \cdots, X_n \sim f$ supported on $[0,1]$ with Hölder smoothness s
- target: estimate $f(1/2)$ under loss $L(f,T) = |T - f(1/2)|$ $\quad \|f^{(s)}\|_\infty \le L$.
- two-point construction: $f_0 \equiv 1$, $f_1(x) = 1 + c_0 \cdot h^s g((x-1/2)/h)$

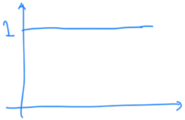

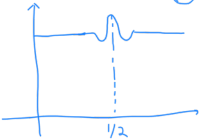

g: g is smooth on $\mathbb{R}$
supported on $[-1,1]$
$\int g = 0$
$g(0) \neq 0$

- separation $\Theta(h^s)$
- KL divergence:

$$D_{\mathrm{KL}}(f_1^{\otimes n} \| f_0^{\otimes n}) \le n \int_0^1 (f_1(x) - 1)^2 dx \asymp nh^{2s+1} = O(1)$$

- choosing $h \asymp n^{-1/(2s+1)}$ gives $R_{n,s}^\star = \Omega(n^{-s/(2s+1)})$

Example IV: entropy estimation

- $X_1, \cdots, X_n \sim P = (p_1, \cdots, p_k)$
- target: estimate the entropy $H(P) = \sum_{i=1}^k -p_i \log p_i$ under quadratic loss
- local asymptotic minimax theorem implies (HW2)

$$R^\star_{n,k} = \Omega\left(\frac{(\log k)^2}{n} \cdot (1 + o_n(1))\right)$$

- claim: non-asymptotically in (n, k), it holds that

$$R^\star_{n,k} = \Omega\left(\frac{(\log k)^2}{n}\right)$$

Tight minimax rate

$$R^\star_{n,k} = \Theta\left(\left(\frac{k}{n \log n}\right)^2 + \frac{(\log k)^2}{n}\right), \quad n \gtrsim \frac{k}{\log k}.$$

Guideline for two-point construction

- let's start from scratch and choose two points P_0, P_1
- separation condition: maximize $|H(P_0) - H(P_1)|$
- indistinguishability condition:

$$n \cdot D_{\mathrm{KL}}(P_0 \| P_1) = D_{\mathrm{KL}}(P_0^{\otimes n} \| P_1^{\otimes n}) = O(1)$$

- the final program:

$$\begin{aligned} \text{maximize} \quad & |H(P_0) - H(P_1)| \\ \text{subject to} \quad & D_{\mathrm{KL}}(P_0 \| P_1) \le c/n \end{aligned}$$

Some attempts

- Attempt 1:

$$P_0 = (\frac{1}{k}, \cdots, \frac{1}{k}), \quad P_1 = (\frac{1+\varepsilon}{k}, \frac{1-\varepsilon}{k}, \cdots, \frac{1+\varepsilon}{k}, \frac{1-\varepsilon}{k})$$

  Result: $|H(P_0) - H(P_1)| \asymp \varepsilon^2, D_{\mathrm{KL}}(P_0 \| P_1) \asymp \varepsilon^2$

- Attempt 2:

$$P_0 = (\frac{1}{k}, \cdots, \frac{1}{k}), \quad P_1 = (\frac{1-\varepsilon}{k}, \frac{1-\varepsilon}{k}, \cdots, \frac{1-\varepsilon}{k}, \frac{1+(k-1)\varepsilon}{k})$$

  Result: $|H(P_0) - H(P_1)| \asymp \min\{\varepsilon \log k, k\varepsilon^2\}, \ D_{\mathrm{KL}}(P_0 \| P_1) \asymp k\varepsilon^2$

- Attempt 3: let $k' = k - 1$ and

$$P_0 = (\frac{1}{2k'}, \cdots, \frac{1}{2k'}, \frac{1}{2}), \quad P_1 = (\frac{1-\varepsilon}{2k'}, \cdots, \frac{1-\varepsilon}{2k'}, \frac{1+\varepsilon}{2})$$

  Result: $|H(P_0) - H(P_1)| \asymp \varepsilon \log k, \ D_{\mathrm{KL}}(P_0 \| P_1) \asymp \varepsilon^2$

  $\varepsilon \asymp \frac{1}{\sqrt{n}}$

Example V: two-armed bandit

- two arms associated with unknown mean reward $(\mu_1, \mu_2) \in [0,1]^2$
- learner pulls arm $\pi_t \in \{1,2\}$ at time t
- learner observes a Gaussian random reward $r_t \sim \mathcal{N}(\mu_{\pi_t}, 1)$
- learner's cumulative regret (loss function):

$$R_T(\pi) = T \max\{\mu_1, \mu_2\} - \sum_{t=1}^{T} \mu_{\pi_t}$$

Claim

- Without additional assumptions, $R_T^\star = \Omega(\sqrt{T})$;
- If it is additionally known that $|\mu_1 - \mu_2| \ge \Delta$, then

$$R_T^\star = \Omega\left(\frac{\max\{1, \log(T\Delta^2)\}}{\Delta} \wedge T\Delta\right).$$

First construction of two points

- a natural choice: $\theta_1 = (\Delta, 0), \theta_2 = (0, \Delta)$
- instantaneous regret at each time:

$$R_T(\pi) = \sum_{t=1}^{T} (\max\{\mu_1, \mu_2\} - \mu_{\pi_t})$$

- separation condition: a gap of Δ at each time
- indistinguishability condition:

$$D_{\mathrm{KL}}(P_1^t \| P_2^t) = \frac{t\Delta^2}{2}$$

- choosing $\Delta \asymp T^{-1/2}$ gives $R_T^\star = \Omega(\sqrt{T})$

Second construction of two points

- failure of first construction:

$$R_T^\star = \Omega\left(\sum_{t=1}^{T} \Delta \exp(-t\Delta^2)\right) = \Omega\left(\frac{1-\exp(-T\Delta^2)}{\Delta}\right)$$

- a less natural choice: $\theta_0 = (\Delta, 0)$, $\theta_1 = (\Delta, 2\Delta)$
- separation condition: same as the first construction
- indistinguishability condition: could be shown that (HW2)

\# of pulling arm 2

$$D_{\mathrm{KL}}(P_1^T \| P_2^T) = 2\Delta^2 \cdot \mathbb{E}_1[T_2]$$

- two-point method:

$$\mathbb{E}[R_T(\pi)] = \Omega(T\Delta \cdot \exp(-2\Delta^2 \cdot \mathbb{E}_1[T_2]))$$

- on the other hand, $\mathbb{E}[R_T(\pi)] = \Omega(\Delta \cdot \mathbb{E}_1[T_2])$
- find $\mathbb{E}_1[T_2]$ to achieve the best tradeoff...

Example VI: multi-armed bandit

- same setting as two-armed bandit, but now with K arms
- learner's cumulative regret:

$$R_T(\pi) = T \max_{i \in [K]} \mu_i - \sum_{t=1}^{T} \mu_{\pi_t}$$

- target: use two-point method to show that

$$R^{\star}_{K,T} = \Omega(\sqrt{KT})$$

Wait... two points suffice?

- construction of the first point: $\theta_1 = (\Delta, 0, 0, \cdots, 0)$
- try second point: $\theta_2 = (\Delta, 0, \cdots, 0, 2\Delta, 0, \cdots, 0)$
- if learner knows θ_2 in advance, then problem reduces to two arms
- observation: separation Δ, indistinguishability

$$D_{\mathrm{KL}}(P_1^T \| P_2^T) = 2\Delta^2 \cdot \mathbb{E}_1[T_i]$$

- idea: choose an appropriate i after knowing learner's policy
- pigeonhole principle: there must exist i such that $\mathbb{E}_1[T_i] \le T/(K-1)$
- two-point method:

$\sum_{i=2}^{K} \mathbb{E}_1[T_i] \le T.$

$$R_T^\star = \Omega\left(T\Delta \cdot \exp\left(-\frac{2\Delta^2 T}{K-1}\right)\right)$$

$\Delta \approx \sqrt{\frac{K}{T}}$

message: choose the points depending on the policy!

Example VII: stochastic optimization

- $\mathcal{F} = \{f : [-1,1] \to [-1,1], f$ convex and 1-Lipschitz$\}$
- at time $t = 1, 2, \cdots, T$:
  - learner queries $x_t \in [-1,1]$
  - learner receives $y_t, z_t \in [-1,1]$ such that

$$\mathbb{E}[y_t] = f(x_t), \quad \mathbb{E}[z_t] = f'(x_t)$$

  - learner adds (x_t, y_t, z_t) to her history
- learner outputs some $\widehat{x} \in [-1,1]$
- loss function:

$$L(f, \widehat{x}) = f(\widehat{x}) - \min_{x \in [-1,1]} f(x)$$

- target: show that

$$\min_{\widehat{x}} \max_{f \in \mathcal{F}} \mathbb{E}_f[L(f, \widehat{x})] \geq \frac{c}{\sqrt{T}}$$

Choice of two points

- choice of two points:

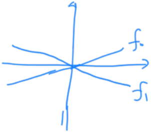

$$f_0(x) = \frac{x}{\sqrt{T}}$$

$$f_1(x) = -\frac{x}{\sqrt{T}}$$

$\mathbb{E}[y_t] = f(x_t)$

$\mathbb{E}[z_t] = f'(x_t)$

$$y_t = \begin{cases} 1 & \text{w.p. } \frac{1+f(x_t)}{2} \\ -1 & \text{w.p. } \frac{1-f(x_t)}{2} \end{cases}$$

$$z_t = \begin{cases} 1 & \text{w.p. } \frac{1+f'(x_t)}{2} \\ -1 & \text{w.p. } \frac{1-f'(x_t)}{2} \end{cases}$$

Verification of conditions

- separation condition:

$$L(f_0, \hat{x}) + L(f_1, \hat{x}) = f_0(\hat{x}) - \min_x f_0(x) + f_1(\hat{x}) - \min_x f_1(x)$$

$$\geq \underbrace{\min_{\hat{x}} (f_0 + f_1)(\hat{x}) - \min_x f_0(x) - \min_x f_1(x)} \gtrsim \frac{1}{\sqrt{T}}$$

- indistinguishability condition:

$$\|f_0 - f_1\|_\infty \leq \frac{2}{\sqrt{T}}$$

$$\|f_0' - f_1'\|_\infty \leq \frac{2}{\sqrt{T}}$$

References

- David L. Donoho and Richard C. Liu. "Geometrizing rates of convergence, II." *The Annals of Statistics*, 19: 668–701, 1991.
- Yury Polyanskiy and Yihong Wu. "Dualizing Le Cams method, with applications to estimating the unseens." *arXiv preprint* arXiv:1902.05616 (2019).
- Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet. "Bounded regret in stochastic multi-armed bandits." *Conference on Learning Theory*, 2013.
- Xi Chen, Yanjun Han, and Yining Wang. "Adversarial combinatorial bandits with general non-linear reward functions." *arXiv preprint* arXiv:2101.01301 (2021).

Next lecture: point vs. mixture, Ingster–Suslina method