A Reduction from Linear Contextual Bandits Lower Bounds to Estimations Lower Bounds

July 15, 2022

Linear Contextual Bandits

- Online sequential decision making with personalized information.
  - K alternative actions (arms).
  - Context $(\boldsymbol{x}_{t1}, \boldsymbol{x}_{t2}, \ldots, \boldsymbol{x}_{tK})$.
  - Reward $y_{ta} = \boldsymbol{x}_{ta}^\top \boldsymbol{\theta}^* + \xi_{ta}$, where $\boldsymbol{\theta}^*$ is unknown.
- Lots of variants: sparsity, privacy constraints, batch constraints, etc..

Literature

- Basic model: Auer (2002), Chu et al. (2011), L. Li et al. (2010), Y. Li, Y. Wang, and Y. Zhou (2019), Goldenshluger and Zeevi (2013), Bastani, Bayati, and Khosravi (2021), Kannan et al. (2018).
- Sparsity: Bastani and Bayati (2020), Kim and Paik (2019), Oh, Iyengar, and Zeevi (2021), Ren and Z. Zhou (2020), D. Wang and Xu (2019), W. Li, Barik, and Honorio (2021).
- Privacy constraints: Shariff and Sheffet (2018), Zheng et al. (2020), Yuxuan Han et al. (2021).
- Batch constraints: Yanjun Han et al. (2020), Ren and Z. Zhou (2020).

Regret

- Problem incident $(\boldsymbol{\theta}^*, F)$:
  - $\boldsymbol{\theta}^*$: $y_{ta} = \boldsymbol{x}_{ta}^\top \boldsymbol{\theta}^* + \xi_{ta}$
  - F: $(\boldsymbol{x}_{t1}, \boldsymbol{x}_{t2}, \ldots, \boldsymbol{x}_{tK}) \sim F$.
- Single period regret $R_t^\pi(\boldsymbol{\theta}^*, F) = E\left[\max_a \boldsymbol{x}_{ta}^\top \boldsymbol{\theta}^* - \boldsymbol{x}_{ta_t^\pi} \boldsymbol{\theta}^*\right]$,

  Cumulative regret $R^\pi(\boldsymbol{\theta}^*, F) = \sum_{t=1}^{T} R_t^\pi(\boldsymbol{\theta}^*, F)$,
- For a class of problem $(\boldsymbol{\Theta}^*, \mathcal{F}_K)$,
  - the worst-case regret $R^\pi(\boldsymbol{\Theta}^*, \mathcal{F}) = \sup_{(\boldsymbol{\theta}^*, F) \in (\boldsymbol{\Theta}^*, \mathcal{F})} R^\pi(\boldsymbol{\theta}^*, F)$
  - the minimum attainable worst case regret $R(\boldsymbol{\Theta}^*, \mathcal{F}) = \inf_{\pi \in \Pi} R^\pi(\boldsymbol{\Theta}^*, \mathcal{F})$.

Overview

- To develop algorithm and find upper bounds:

$$\begin{array}{rcl} \text{Estimator } \boldsymbol{\theta} & \rightsquigarrow & \text{Algorithm } \pi^{\boldsymbol{\theta}} \\ Loss(\boldsymbol{\theta}) & \succeq & Regret(\pi^{\boldsymbol{\theta}}) \end{array}$$

- Our work:

$$\begin{array}{rcl} \text{Estimator } \boldsymbol{\theta}^{\pi} & \leftsquigarrow & \text{Algorithm } \pi \\ Loss(\boldsymbol{\theta}^{\pi}) & \preceq & Regret(\pi) \end{array}$$

The 2-Armed Problems

- In period t, an algorithm is a binary classifier.
- The ground truth is linear: $(\boldsymbol{x}_{t1}, \boldsymbol{x}_{t2}) \mapsto 1 \iff \boldsymbol{z}_t^\top \boldsymbol{\theta}^* \geq 0$, where $\boldsymbol{z}_t = \boldsymbol{x}_{t1} - \boldsymbol{x}_{t2}$.
- $\boldsymbol{\theta}_t^\pi$: the best linear approximation, maximizer of

$$P\left(a_t^\pi = 1, \boldsymbol{z}_t^\top \boldsymbol{\theta} > 0\right) + P\left(a_t^\pi = 2, \boldsymbol{z}_t^\top \boldsymbol{\theta} \leq 0\right).$$

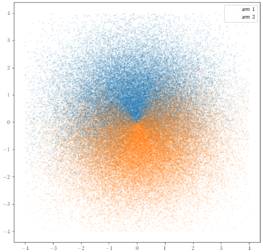

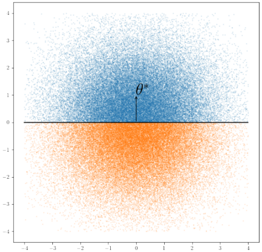

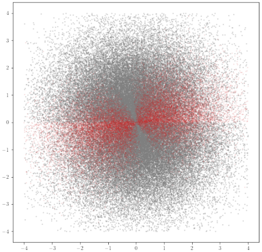

Label by UCB

True Label

Wrong vs. Correct

The 2-Armed Problems

Proposition

For any $h \geq 0$,

$$R_t^{\pi}(\boldsymbol{\theta}^*, F) \geq \frac{h}{2} P\left(\left(\boldsymbol{z}_t^\top \boldsymbol{\theta}^*\right)\left(\boldsymbol{z}_t^\top \boldsymbol{\theta}_t^{\pi}\right) < 0\right) - E\left[\left(h - \left|\boldsymbol{z}_t^\top \boldsymbol{\theta}^*\right|\right)^+ 1_{\{(\boldsymbol{z}_t^\top \boldsymbol{\theta}^*)(\boldsymbol{z}_t^\top \boldsymbol{\theta}_t^{\pi}) \leq 0\}}\right].$$

Proposition

For any $F \in \mathcal{F}_2$ such that $\boldsymbol{z}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$,

$$R_t^{\pi}(\boldsymbol{\theta}^*, F) = \Omega\left(\|\boldsymbol{\theta}^*\|_2 \cdot \left\|\frac{\boldsymbol{\theta}^*}{\|\boldsymbol{\theta}^*\|_2} - \frac{\boldsymbol{\theta}_t^{\pi}}{\|\boldsymbol{\theta}_t^{\pi}\|_2}\right\|_2^2\right).$$

The K-Armed Problems

Problem #1: $\begin{pmatrix} \boxed{\boldsymbol{x}_{t1}(1)^\top \boldsymbol{\theta}^*(1)} \\ \boldsymbol{x}_{t2}(1)^\top \boldsymbol{\theta}^*(1) \end{pmatrix} = \begin{pmatrix} \boxed{\boldsymbol{x}_{t1}(1)^\top} \\ \boldsymbol{x}_{t2}(1)^\top \end{pmatrix} \begin{pmatrix} \boldsymbol{\theta}^*(1) \end{pmatrix}$

Problem #2: $\begin{pmatrix} \boldsymbol{x}_{t1}(2)^\top \boldsymbol{\theta}^*(2) \\ \boxed{\boldsymbol{x}_{t2}(2)^\top \boldsymbol{\theta}^*(2)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{x}_{t1}(2)^\top \\ \boxed{\boldsymbol{x}_{t2}(2)^\top} \end{pmatrix} \begin{pmatrix} \boldsymbol{\theta}^*(2) \end{pmatrix}$

Problem #3: $\begin{pmatrix} \boxed{\boldsymbol{x}_{t1}(3)^\top \boldsymbol{\theta}^*(3)} \\ \boldsymbol{x}_{t2}(3)^\top \boldsymbol{\theta}^*(3) \end{pmatrix} = \begin{pmatrix} \boxed{\boldsymbol{x}_{t1}(3)^\top} \\ \boldsymbol{x}_{t2}(3)^\top \end{pmatrix} \begin{pmatrix} \boldsymbol{\theta}^*(3) \end{pmatrix}$

Three 2-armed problems

$$\begin{pmatrix} \boldsymbol{x}_{t1}^\top \boldsymbol{\theta}^* \\ \boldsymbol{x}_{t2}^\top \boldsymbol{\theta}^* \\ \boxed{\boldsymbol{x}_{t3}^\top \boldsymbol{\theta}^*} \\ \boldsymbol{x}_{t4}^\top \boldsymbol{\theta}^* \\ \boldsymbol{x}_{t5}^\top \boldsymbol{\theta}^* \\ \boldsymbol{x}_{t6}^\top \boldsymbol{\theta}^* \\ \boldsymbol{x}_{t7}^\top \boldsymbol{\theta}^* \\ \boldsymbol{x}_{t8}^\top \boldsymbol{\theta}^* \end{pmatrix} = \begin{pmatrix} \boldsymbol{x}_{t1}(1)^\top & \boldsymbol{x}_{t1}(2)^\top & \boldsymbol{x}_{t1}(3)^\top \\ \boldsymbol{x}_{t2}(1)^\top & \boldsymbol{x}_{t1}(2)^\top & \boldsymbol{x}_{t1}(3)^\top \\ \boxed{\boldsymbol{x}_{t1}(1)^\top} & \boxed{\boldsymbol{x}_{t2}(2)^\top} & \boxed{\boldsymbol{x}_{t1}(3)^\top} \\ \boldsymbol{x}_{t2}(1)^\top & \boldsymbol{x}_{t2}(2)^\top & \boldsymbol{x}_{t1}(3)^\top \\ \boldsymbol{x}_{t1}(1)^\top & \boldsymbol{x}_{t1}(2)^\top & \boldsymbol{x}_{t2}(3)^\top \\ \boldsymbol{x}_{t2}(1)^\top & \boldsymbol{x}_{t1}(2)^\top & \boldsymbol{x}_{t2}(3)^\top \\ \boldsymbol{x}_{t1}(1)^\top & \boldsymbol{x}_{t2}(2)^\top & \boldsymbol{x}_{t2}(3)^\top \\ \boldsymbol{x}_{t2}(1)^\top & \boldsymbol{x}_{t2}(2)^\top & \boldsymbol{x}_{t2}(3)^\top \end{pmatrix} \begin{pmatrix} \boldsymbol{\theta}^*(1) \\ \boldsymbol{\theta}^*(2) \\ \boldsymbol{\theta}^*(3) \end{pmatrix}$$

A 2^3-armed problem

FIGURE: Selecting arms $(1, 2, 1)$ in (a) yields the same regret as selecting arm 3 in (b). Arm selection combinations for (a) are in one-to-one correspondence with arms for (b).

Applications

Model	Upper Bound from Literature	Lower Bound from Literature	Our Lower Bound
Without a margin condition	$\tilde{O}(\sqrt{Td})$	$\Omega(dT)$	$\Omega(\sqrt{dT\log K})$
With a margin condition	$O(d^3 \log T)$	$\Omega(d \log T)$	$\Omega(d\log(\frac{T\log K}{d}))$
With a batch constraint (stochastic context)	$O\left(\sqrt{dT}\left(\frac{T}{d^2}\right)^{\frac{1}{2(2^M-1)}}\right)$	$\Omega\left(\sqrt{dT}\left(\frac{T}{d^2}\right)^{\frac{1}{2(2^M-1)}}\right)$	$\Omega\left(\sqrt{dT}\left(\frac{T}{d^2}\right)^{\frac{1}{2(2^M-1)}}\right)$
Sparse	$O(\sqrt{s_0T\log(dT)})$	$\Omega(\sqrt{s_0T})$	$\Omega(\sqrt{s_0T\log K\log(d/s_0)})$
Jointly differentially private	$O(d^{3/4}\sqrt{T}/\sqrt{\varepsilon})$ for adversarial context $\tilde{O}(\sqrt{dT}+\frac{d\log(1/\delta)}{\varepsilon})$ can be achieved (absent from literature)		$\Omega(\sqrt{dT\log K}+\frac{d}{\varepsilon+\delta})$
Locally differentially private with sparsity			$\frac{\sqrt{dT\log d}}{\varepsilon}$

TABLE: Comparison

Thank you!