6 Linear Bandits: Regret Minimization

This section is inspired by [Lattimore and Szepesvári, 2020].

Model: $\mathbb{E}[Y_{i,t}] = \langle \theta^*, x_i \rangle$

> **Input**: Finite set $\mathcal{X} \subset \mathbb{R}^d$, confidence level $\delta \in (0,1)$.
> Let $\widehat{\mathcal{X}}_1 \leftarrow \mathcal{X}, \ell \leftarrow 1$
> **while** $|\widehat{\mathcal{X}}_\ell| > 1$ **do**
> Let $\widehat{\lambda}_\ell \in \triangle_{\widehat{\mathcal{X}}_\ell}$ be a $\frac{d(d+1)}{2}$-sparse minimizer of $f(\lambda) = \max_{x \in \widehat{\mathcal{X}}_\ell} \|x\|^2_{(\sum_{x \in \widehat{\mathcal{X}}_\ell} \lambda_x x x^\top)^{-1}}$
> $\epsilon_\ell = 2^{-\ell}$, $\tau_\ell = 2d\epsilon_\ell^{-2} \log(4\ell^2|\mathcal{X}|/\delta)$
> Pull arm $x \in \mathcal{X}$ exactly $\lceil \widehat{\lambda}_{\ell,x} \tau_\ell \rceil$ times and construct the least squares estimator $\widehat{\theta}_\ell$ using only the observations of this round
> $\widehat{\mathcal{X}}_{\ell+1} \leftarrow \widehat{\mathcal{X}}_\ell \setminus \{x \in \widehat{\mathcal{X}}_\ell : \max_{x' \in \widehat{\mathcal{X}}_\ell} \langle x' - x, \widehat{\theta}_\ell \rangle > 2\epsilon_\ell\}$
> $\ell \leftarrow \ell + 1$
> **Output**: $\widehat{\mathcal{X}}_\ell$

Recall: From standard MAB, the elimination alg. relied on just controlling $\hat{\theta}_i - \theta^*_i = \langle \hat{\theta} - \theta^*, e_i \rangle \quad \forall i \in [n]$

We showed $R_T \leq \sqrt{|\mathcal{X}| T \log(T)}$

Define $x^* = \arg\max_{x \in \mathcal{X}} \langle x, \theta^* \rangle$

Lemma Assume $\max_{x \in \mathcal{X}} \langle x^* - x, \theta^* \rangle \leq 4$. W.p. $\geq 1-\delta$ $x^* \in \hat{\mathcal{X}}_\ell$ for all $\ell \geq 1$ and $\max_{x \in \mathcal{X}_\ell} \underbrace{\langle x^* - x, \theta^* \rangle}_{\Delta x} \leq 8\varepsilon_\ell$.

Proof Fix $\mathcal{V} \subset \mathcal{X}$ and $x \in \mathcal{V}$. Define event

$$\mathcal{E}_{x,\ell}(\mathcal{V}) = \{ |\langle x, \hat{\theta}_\ell - \theta^* \rangle| \leq \varepsilon_\ell \}$$

where it is understood that $\hat{\theta}_\ell$ uses measurements from G-optimal design on $\underline{\underline{\mathcal{V}}}$ (i.e. $\hat{\mathcal{X}}_\ell = \mathcal{V}$)

$$\mathbb{P}\left(\bigcup_{\ell=1}^{\infty} \bigcup_{x \in \hat{\mathcal{X}}_\ell} \{ \mathcal{E}^c_{x,\ell}(\hat{\mathcal{X}}_\ell) \} \right) \leq \sum_{\ell=1}^{\infty} \mathbb{P}\left(\bigcup_{x \in \hat{\mathcal{X}}_\ell} \{ \mathcal{E}^c_{x,\ell}(\hat{\mathcal{X}}_\ell) \} \right)$$

$$= \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \mathbb{P}\left(\bigcup_{x \in \mathcal{V}} \{ \mathcal{E}_{x,\ell}^c(\mathcal{V}) \}, \ \mathcal{V} = \hat{\mathcal{X}}_\ell \right)$$

$$= \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \underbrace{\mathbb{P}\left(\bigcup_{x \in \mathcal{V}} \{ \mathcal{E}_{x,\ell}^c(\mathcal{V}) \} \,\middle|\, \mathcal{V} = \hat{\mathcal{X}}_\ell \right)}_{= \mathbb{P}\left(\bigcup_{x \in \mathcal{V}} \{ \mathcal{E}^c \} \right)} \mathbb{P}(\mathcal{V} = \hat{\mathcal{X}}_\ell)$$

$$\leq \frac{\delta |\mathcal{V}|}{2\ell^2 |\mathcal{X}|} \leq \frac{\delta}{2\ell^2}$$

$$\leq \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \frac{\delta}{2\ell^2} \mathbb{P}(\mathcal{V} = \hat{\mathcal{X}}_\ell)$$

$$= \sum_{\ell=1}^{\infty} \frac{\delta}{2\ell^2} \leq \delta$$

$$\Rightarrow \quad \mathbb{P}\left(\bigcap_{\ell=1}^{\infty} \bigcap_{x \in \hat{\mathcal{X}}_\ell} \{ |\langle x, \hat{\theta}_\ell - \theta^* \rangle| \leq \varepsilon_\ell \} \right) \geq 1 - \delta.$$

Show $x^* \in \hat{\mathcal{X}}_\ell$ for all ℓ. $x^* \in \hat{\mathcal{X}}_1$ so assume in $\hat{\mathcal{X}}_\ell \ni x^*$

$$\langle x - x^*, \hat{\theta}_\ell \rangle = \langle x - x^*, \hat{\theta}_\ell - \theta^* \rangle + \langle x - x^*, \theta^* \rangle$$

$$= \underbrace{\langle x, \hat{\theta}_\ell - \theta^* \rangle}_{\leq \varepsilon_\ell} - \underbrace{\langle x^*, \hat{\theta}_\ell - \theta^* \rangle}_{\geq -\varepsilon_\ell} + \underbrace{\langle x - x^*, \theta^* \rangle}_{\leq 0}$$

$$\leq 2\varepsilon_\ell.$$

$\Delta_x = \langle x^* - x, \theta^* \rangle$. Fix $\nu \geq 0$

$$\sum_{x \in \mathcal{X} \setminus x^*} \Delta_x T_x = \sum_{\substack{x \in \mathcal{X} \setminus x^* \\ \Delta_x \leq \nu}} \Delta_x T_x + \sum_{\substack{x \in \mathcal{X} \setminus x^* \\ \Delta_x > \nu}} \Delta_x T_x$$

From above Lemma we have $\max_{x \in \mathcal{X}_\ell} \langle x^* - x, \theta^* \rangle \leq 8\varepsilon_\ell = 8 \cdot 2^{-\ell}$

$$\leq \nu T + \sum_{\ell=1}^{\infty} \sum_{\substack{x \in \mathcal{X} \setminus x^* \\ \Delta_x > \nu}} \Delta_x \underbrace{\lceil \tau_\ell \hat{\lambda}_{\ell,x} \rceil}_{\neq 0 \text{ if } T_x^{(\ell)} \neq 0}$$

Consider $\ell : 8 \cdot 2^{-\ell} \leq \nu$

$\Rightarrow \ell \leq \log_2(8\nu^{-1})$

$$\leq \nu T + \sum_{\ell=1}^{\infty} 8\varepsilon_\ell \sum_{\substack{x \in \mathcal{X} \setminus x^* \\ \Delta_x > \nu}} (\tau_\ell \hat{\lambda}_{\ell,x} + 1) \mathbb{1}\{\hat{\lambda}_{\ell,x} > 0\}$$

$$= \nu T + \sum_{\ell=1}^{\log_2((\Delta \vee \nu)^{-1})} 8\varepsilon_\ell \left(\tau_\ell + \text{support}(\hat{\lambda}_\ell) \right)$$

$$\leq \nu T + d^2 + c \sum_{\ell=1}^{\log_2(8\Delta^{-1} \wedge 8\nu)} \varepsilon_\ell \frac{d \log(\ell^2 |\mathcal{X}| / \delta)}{\varepsilon_\ell^2}$$

$$\leq \nu T + d^2 + c\, d \log\left(|\mathcal{X}| \log(\Delta^{-1} \wedge \nu^{-1}) / \delta\right) \cdot \sum_{\ell}^{\log(\Delta^{-1} \wedge \nu^{-1})} \frac{1}{\varepsilon_\ell}$$

$$\leq \nu T + d^2 + c \frac{1}{(\nu \vee \Delta)^{-1}} d \log\left(|\mathcal{X}| \log(\Delta^{-1} \wedge \nu^{-1}) / \delta\right)$$

$\delta = 1/T$

$$\nu = 0 \implies R_T \le \frac{d \log(|\mathcal{X}|T)}{\Delta}$$

$$\nu \ge 0 \implies R_T \le c\sqrt{dT \log(|\mathcal{X}|T)}$$

Crux of above proof is that

$$\max_{x_i \in \hat{\mathcal{X}}_\ell} |\langle \hat{\theta}_\ell - \theta^*, x_i \rangle| \le \sqrt{\frac{2d \log(4\ell^2|\mathcal{X}|/\delta)}{\tau_\ell}} \le \varepsilon_\ell$$

<u>Misspecified model.</u>

Suppose that $y_t = \langle x_{I_t}, \theta^* \rangle + \varepsilon_t + \eta_t$

ε_t: $\mathcal{N}(0,1)$

η_t: Deterministic Misspecification of model $|\eta_\ell| \le h$ for some unknown h.

$$X = \begin{bmatrix} - x_{I_1}^T - \\ \vdots \\ - x_{I_\tau}^T - \end{bmatrix} \quad Y = \begin{bmatrix} y_{I_1} \\ \vdots \\ y_{I_T} \end{bmatrix}$$

$$\underline{Y} = X\theta^* + \varepsilon + \eta \qquad \varepsilon \sim \mathcal{N}(0, I_\tau)$$

$$\|\eta\|_\infty \le h$$

$$\hat{\theta}_\ell = (X^TX)^{-1}X^TY$$

$$= \theta^* + (X^TX)^{-1}X^T\varepsilon + (X^TX)^{-1}X^T\eta$$

$$|\langle \hat{\theta}_\ell - \theta^*, x \rangle| \leq |\langle (X^TX)^{-1}X^T\varepsilon, x\rangle| + |\langle (X^TX)^{-1}X^T\eta, x\rangle|$$

$$\leq \varepsilon_\ell + |\langle (X^TX)^{-1}X^T\eta, x\rangle|$$

↑ From analysis above

$$x^T(X^TX)^{-1}X^T\eta \leq \sup_{\|e\|_\infty \leq h} x^T(X^TX)^{-1}X^Te$$

$$= \sup_{\|e\|_\infty \leq h} \sum_{t=1}^{\zeta} x^T(X^TX)^{-1}x_{I_t}e_t$$

$$\leq h \cdot \sum_{t=1}^{\zeta} |x^T(X^TX)^{-1}x_{I_t}|$$

$$= \zeta_\ell h \cdot \sum_{x' \in \mathcal{X}} |x^T(X^TX)^{-1}x'| \frac{T_{x'}}{\zeta_\ell}$$

$\zeta_\ell := \sum_{x' \in \mathcal{X}} T_{x'}$

$$= \zeta_\ell \cdot h \cdot \sqrt{\left(\sum_{x' \in \mathcal{X}} \frac{T_{x'}}{\zeta_\ell} |x^T(X^TX)^{-1}x'|\right)^2}$$

(Jensen's) $$\leq \zeta_\ell \cdot h \cdot \sqrt{\sum_{x' \in \mathcal{X}} \frac{T_{x'}}{\zeta_\ell} |x^T(X^TX)^{-1}x'|^2}$$

$$= \tau_\ell \cdot h \cdot \sqrt{\sum_{x' \in \mathcal{X}} \frac{T_{x'}}{\tau_\ell} \; x^T (X^T X)^{-1} x' x' (X^T X)^{-1} x}$$

$$= h \cdot \sqrt{\tau_\ell} \cdot \sqrt{x^T (X^T X)^{-1} \Big(\underbrace{\sum_{x'} T_{x'} x' x'^T}_{= X^T X} \Big) (X^T X)^{-1} x}$$

$$= h \cdot \sqrt{\tau_\ell} \cdot \sqrt{x^T (X^T X)^{-1} x}$$

$$\leq h \cdot \sqrt{\tau_\ell} \cdot \sqrt{\frac{d}{\tau_\ell}} \quad \text{by G-optimal design.}$$

$$|\langle \hat{\theta}_\ell - \theta^*, x \rangle| \leq |\langle (X^T X)^{-1} X^T \varepsilon, x \rangle| + |\langle (X^T X)^{-1} X^T \eta, x \rangle|$$

$$\leq \underbrace{\varepsilon_\ell + h\sqrt{d}}$$

Plug back in x8 to upper bound Δ_x in proof

$$\Longrightarrow R_T \leq c\sqrt{dT \log(|\mathcal{X}|T)} + c' h \sqrt{d} \cdot T$$

$\Longrightarrow$ Algorithm doesn't need to know h and can tolerate misspecification up to $h \approx \frac{1}{\sqrt{T}}$

What if $|\mathcal{X}|$ is really big?

What if $\mathcal{X} = \{x : \|x\|_2 \leq 1\}$

Need upper bound on this quantity.

$$\max_{x \in \hat{\mathcal{X}}_\ell} |\langle x, \hat{\theta}_\ell - \theta_* \rangle|.$$

Fix $x \in \hat{\mathcal{X}}_\ell$.

$$|\langle x, \hat{\theta}_\ell - \theta_* \rangle| = |\langle (X^T X)^{-1/2} x, (X^T X)^{1/2} (\hat{\theta}_\ell - \theta_*) \rangle|$$

$$\leq \|(X^T X)^{-1/2} x\|_2 \cdot \|(X^T X)^{1/2} (\hat{\theta} - \theta_*)\|_2$$

$$= \|x\|_{(X^T X)^{-1}} \cdot \|\hat{\theta} - \theta_*\|_{(X^T X)}$$

by G-optimal $\rightarrow$ $$\leq \sqrt{\frac{d}{T_\ell}} \cdot \underbrace{\|\hat{\theta} - \theta_*\|_{(X^T X)}}_{\text{Indepent of } |\mathcal{X}|}$$

$$\Longrightarrow \max_{x \in \hat{\mathcal{X}}_\ell} |\langle x, \hat{\theta}_\ell - \theta_* \rangle| \leq \sqrt{\frac{d}{T_\ell}} \cdot \|\hat{\theta} - \theta_*\|_{(X^T X)}$$

$$\hat{\theta} = \theta^* + (X^TX)^{-1}X^T\varepsilon$$

$$\|\hat{\theta} - \theta_*\|_{(X^TX)} = \sqrt{\varepsilon^T X (X^TX)^{-1}(X^TX)(X^TX)^{-1}X^T\varepsilon}$$

$$= \sqrt{\varepsilon^T X (X^TX)^{-1} X^T\varepsilon}$$

Intuition: $\mathbb{E}\left[\|\hat{\theta} - \theta_*\|_{(X^TX)}\right] \leq \sqrt{\mathbb{E}\left[\|\hat{\theta} - \theta_*\|^2_{X^TX}\right]}$

$$= \sqrt{\mathbb{E}\left[\varepsilon^T X (X^TX)^{-1} X^T \varepsilon\right]}$$

$$= \sqrt{\mathbb{E}\left[\mathrm{Trace}\left((X^TX)^{-1}X^T\varepsilon\varepsilon^T X\right)\right]}$$

$$= \sqrt{\mathrm{Tr}\left(\underbrace{(X^TX)^{-1}X^TX}_{I_d}\right)}$$

$$= \sqrt{d}.$$

Formally show $\|\hat{\theta}-\theta_*\|_{X^TX}$ bounded by a cover. let $\varepsilon \geq 0$ and let $N = \left(\frac{3}{\varepsilon}\right)^d$

$$\{x_i\}_{i=1}^{N} : \|x_i\|_2 = 1 \quad \sup_{\|y\|_2=1} \min_i \|x_i - y\|_2 \leq \varepsilon.$$

$$\|\hat{\theta}-\theta_*\|_{X^TX} = \|(X^TX)^{1/2}(\hat{\theta}-\theta_*)\|_2$$

$$= \sup_{\|y\|_2=1} \langle y, (X^TX)^{1/2}(\hat{\theta}-\theta_*)\rangle$$

$$= \sup_{\|y\|_2=1} \min_i \langle y - x_i, (X^TX)^{1/2}(\hat{\theta}-\theta_*)\rangle + \langle x_i, (X^TX)^{1/2}(\hat{\theta}-\theta_*)\rangle$$

$$\leq \varepsilon \cdot \|\hat{\theta}-\theta_*\|_{X^TX} + \sqrt{2\log(4\ell^2 N/\delta)}$$

$$\|\hat{\theta}-\theta_*\|_{X^TX} \leq \frac{1}{1-\varepsilon}\sqrt{2\log\left(4\ell^2\left(\tfrac{3}{\varepsilon}\right)^d/\delta\right)}$$

Pick $\varepsilon = 1/2$

$$\leq c\sqrt{d + \log(4\ell^2/\delta)}$$

$$\Longrightarrow \max_{x \in \hat{\mathcal{X}}_\ell} |\langle x, \hat{\theta}_\ell - \theta_*\rangle| \leq \sqrt{\frac{d}{\tau_\ell}} \cdot \|\hat{\theta}-\theta_*\|_{(X^TX)}$$

$$\leq c\sqrt{\frac{d}{\tau_\ell}} \cdot \sqrt{d + \log(4\ell^2/\delta)}$$

$$= c\sqrt{\frac{d^2 + \log(4\ell^2/\delta)}{\tau_\ell}}$$

If we want RHS $\leq \varepsilon_\ell$ take

$$\tau_\ell = c\,\varepsilon_\ell^{-2} \cdot \left(d^2 + d\log(4\ell^2/\delta)\right)$$

$$R_T \leq \sqrt{d^2 T + d\,T\log(T)}$$

Holds for <u>any</u> $\mathcal{X}$

$$f(\lambda) = \max_{x\in\mathcal{X}} x^T A(\lambda)^{-1} x \qquad \text{(G-optimal)}$$

$$A(\lambda) := \sum_{x'} \lambda_{x'} x' x'^T \qquad \tilde{x} = \arg\max_{x\in\mathcal{X}} x^T A(\lambda)^{-1} x$$

$$\frac{\partial}{\partial x'} f(\lambda) = -\tilde{x}^T A(\lambda)^{-1} x' x'^T A(\lambda)^{-1} \tilde{x}$$

To find steepest comp. descent solve $\arg\max_{x'\in\mathcal{X}} |x'^T A(\lambda)^{-1}\tilde{x}|$