$\mathcal{X} = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$

Assume $\exists\, \theta_* \in \mathbb{R}^d$: pulling arm x_t @ time $t \in \mathbb{N}$

results in a reward $y_t = x_t^T \theta_* + \eta_t$

where $\mathbb{E}[\eta_t] = 0$, $\mathbb{E}[\exp(\lambda \eta_t)] \leq e^{\lambda^2/2}$.

$$\text{Regret} = \max_{x \in \mathcal{X}} \sum_{t=1}^{T} x^T \theta_* - x_t^T \theta_*$$

Suppose we choose $x_1, \dots, x_T \in \mathcal{X}$ and measure

corresponding rewards $y_1, \dots, y_T$ where $y_t = x_t^T \theta_* + \eta_t$

$$\hat{\theta} = \arg\min_{\theta \in \mathbb{R}^d} \sum_{t=1}^{T} (y_t - x_t^T \theta)^2$$

η_t IID

Recommend $\hat{x} = \arg\max_{x \in \mathcal{X}} x^T \hat{\theta}$.

What is $\mathbb{E}[(x_* - \hat{x})^T \theta_*]$ where $x_* = \arg\max_{x \in \mathcal{X}} x^T \theta_*$

$$\hat{\theta} = \arg\min_{\theta \in \mathbb{R}^d} \sum_{t=1}^{T} (y_t - x_t^T \theta)^2$$

$$0 = -\sum_t 2(y - x_t^T \theta) x_t \quad \Rightarrow \quad \left(\sum x_t x_t^T\right) \theta = \sum x_t y_t$$

$$\hat{\theta} = \left(\sum x_t x_t^T\right)^{-1} \sum_t x_t y_t$$

$$\hat{\theta} = \underset{\theta \in \mathbb{R}^d}{\operatorname{argmin}} \sum_{t=1}^{T} (y_t - x_t^T \theta)^2$$

$$= \left(\sum x_t x_t^T\right)^{-1} \sum_t x_t y_t$$

$$= \left(\sum_t x_t x_t^T\right)^{-1} \sum_t x_t (x_t^T \theta_* + z_t)$$

$$= \theta_* + \left(\sum_t x_t x_t^T\right)^{-1} \sum_t x_t z_t$$

$$= \theta_* + (X^T X)^{-1} X^T z$$

$$X = \begin{bmatrix} — x_1^T — \\ \vdots \\ — x_T^T — \end{bmatrix} \in \mathbb{R}^{T \times d} \qquad y = \begin{bmatrix} y_1 \\ \vdots \\ y_T \end{bmatrix} \in \mathbb{R}^T \qquad z = \begin{bmatrix} z_1 \\ \vdots \\ z_T \end{bmatrix} \in \mathbb{R}^T$$

Fix some $x \in \mathcal{X}$.

$$x^T(\hat{\theta} - \theta_*) = \underbrace{x^T (X^T X)^{-1} X^T}_{:= w^T} z$$

$$= w^T z$$

$$= \sum_{t=1}^{T} w_t z_t$$

$$\mathbb{E}[\exp(\lambda x^T(\hat{\theta} - \theta_*))] = \mathbb{E}\left[\exp\left(\lambda \sum_{t=1}^{T} w_t z_t\right)\right]$$

$$= \mathbb{E}\left[\prod_{t=1}^{T} \exp(\lambda w_t z_t)\right]$$

$$= \prod_{t=1}^{T} \mathbb{E}[\exp(\lambda w_t z_t)]$$

$$\le \sum_{t=1}^{T} \exp(\lambda^2 w_t^2/2)$$

$$= \exp(\lambda^2 \|w\|_2^2/2)$$

$$\|w\|_2^2 = x^T(X^TX)^{-1}x$$

$$\Rightarrow \quad \mathbb{P}\left((\hat\theta-\theta_*)^T x > \varepsilon\right) \le \exp\left(-\frac{\varepsilon^2}{2x^T(X^TX)^{-1}x}\right) = \delta$$

$$\Rightarrow \quad (\hat\theta-\theta_*)^T x \le \sqrt{2x^T(X^TX)^{-1}x \, \log(1/\delta)} \quad \text{w.p.} \ge 1-\delta.$$

$$(x_* - \hat x)^T\theta_* = \underbrace{(x_*-\hat x)^T\hat\theta}_{\le 0} + (x_*-\hat x)^T(\theta_*-\hat\theta)$$

$$\le x_*^T(\theta_*-\hat\theta) - \hat x^T(\theta_*-\hat\theta)$$

$$\le 2\cdot\sqrt{\max_x x^T(X^TX)^{-1}x \cdot 2\log(2n/\delta)}$$

$$\approx 2\cdot \sqrt{\underbrace{\max_x x^T A(\lambda)^{-1}x}_{=d} \cdot \frac{2\log(2n/\delta)}{T}}$$

$$X^TX = v_1v_1^T\sigma_1 + v_2v_2^T\sigma_2$$

$$(X^TX)^{-1} = v_1v_1^T\sigma_1^{-1} + v_2v_2^T\sigma_2^{-1}$$

$$x^T(X^TX)^{-1}x = (x^Tv_1)^2\sigma_1^{-1} + (x^Tv_2)^2\sigma_2^{-1}$$

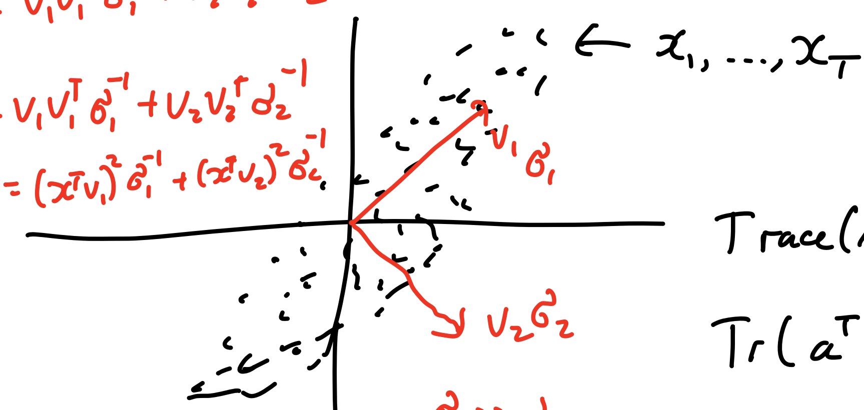

$\sigma_1 >> \sigma_2$

$$\mathrm{Trace}(ABC) = \mathrm{Trace}(BCA) = \mathrm{Tr}(CAB)$$

$$\mathrm{Tr}(a^TBa) = \mathrm{Tr}(Baa^T)$$

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

$$\mathbb{E}[\|\hat{\theta} - \theta_*\|_2^2] = \mathbb{E}[\|(X^TX)^{-1}X^T z\|_2^2]$$

$$= \mathbb{E}[z^T X (X^TX)^{-1}(X^TX)^{-1}X^T z]$$

$$\mathbb{E}[z_i z_j]$$

$$= \mathbb{E}[z_i^2]\mathbb{1}\{i=j\}$$

$$\leq \mathbb{1}\{i=j\}$$

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

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

$$= \text{Trace}\left((X^TX)^{-1}\right)$$

$$\exists\, \lambda \in \Delta_{\mathcal{X}} = \{\lambda_x \in \mathbb{R}_+^{|\mathcal{X}|} : \sum_{x \in \mathcal{X}} \lambda_x = 1\}$$

$$X^TX = \sum_{t=1}^{T} x_t x_t^T$$

$$= T \cdot \underline{\sum_{x \in \mathcal{X}} \lambda_x x x^T}$$

$$=: T \cdot \underline{A(\lambda)}$$

where $\lambda_x = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\{x_t = x\}$

- **A-optimality**: minimize $f_A(\lambda) = \mathrm{Tr}(A_\lambda^{-1})$ minimizes $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2]$

- **E-optimality**: minimize $f_E(\lambda) = \max_{u:\|u\|\leq 1} u^\top A_\lambda^{-1} u$ minimizes $\max_{u:\|u\|\leq 1} \mathbb{E}[(\langle u, \widehat{\theta} - \theta\rangle)^2]$

- **D-optimality**: maximize $g_D(\lambda) = \log(|A_\lambda|)$ maximizes the entropy of distribution. Also, if $\mathcal{E}_\lambda = \{x : x^\top A_\lambda^{-1} x \leq d\}$ then D-optimality is the minimum volume ellipsoid that contains $\mathcal{X}$.

- **G-optimality**: minimize $f_G(\lambda) = \max_{x\in\mathcal{X}} x^\top A_\lambda^{-1} x$ minimizes $\max_{x\in\mathcal{X}} \mathbb{E}[(\langle x, \widehat{\theta} - \theta^*\rangle)^2]$

Lemma 10 (Kiefer-Wolfowitz (1960)). *For any $\mathcal{X}$ with $d = dim(span(\mathcal{X}))$, there exists a $\lambda^* \in \triangle_{\mathcal{X}}$ that*

- $\max_\lambda g_D(\lambda) = g_D(\lambda^*)$

- $\min_\lambda f_G(\lambda) = f_G(\lambda^*)$

- $f_G(\lambda^*) = d$

- *support*$(\lambda^*) = (d+1)d/2$

Proposition 2. *If λ^* is the G-optimal design for $\mathcal{X}$ then if we pull arm $x \in \mathcal{X}$ exactly $\lceil \tau\lambda_x^* \rceil$ times for some $\tau > 0$ and compute the least squares estimator $\widehat{\theta}$. Then for each $x \in \mathcal{X}$ we have with probability at least $1 - \delta$*

$$\begin{aligned}
\langle x, \widehat{\theta} - \theta^*\rangle &\leq \|x\|_{(\sum_{x\in\mathcal{X}} \lceil \tau\lambda_x^* \rceil xx^\top)^{-1}} \sqrt{2\log(1/\delta)} \\
&\leq \frac{1}{\sqrt{\tau}} \|x\|_{(\sum_{x\in\mathcal{X}} \lambda_x^* xx^\top)^{-1}} \sqrt{2\log(1/\delta)} \\
&\leq \sqrt{\frac{2d\log(1/\delta)}{\tau}}
\end{aligned}$$

and we have taken at most $\tau + \frac{d(d+1)}{2}$ pulls. Thus, for any $\delta' \in (0,1)$ we have $\mathbb{P}(\bigcup_{x\in\mathcal{X}} \{|\langle x, \widehat{\theta} - \theta^*\rangle| > \sqrt{\frac{2d\log(2|\mathcal{X}|/\delta')}{\tau}}\}) \leq \delta'$.

Input: Finite set $\mathcal{X} \subset \mathbb{R}^d$, confidence level $\delta \in (0,1)$.
Let $\mathcal{X}_1 \leftarrow \mathcal{X}, \ell \leftarrow 1$
while $|\mathcal{X}_\ell| > 1$ **do**
Let $\widehat{\lambda}_\ell \in \triangle_{\mathcal{X}_\ell}$ be a $\frac{d(d+1)}{2}$-sparse minimizer of $f(\lambda) = \max_{x \in \mathcal{X}_\ell} \|x\|^2_{(\sum_{x \in \mathcal{X}_\ell} \lambda_x x x^\top)^{-1}}$
$\epsilon_\ell = 2^{-\ell}$, $\tau_\ell = 2d\epsilon_\ell^{-2} \log(4\ell^2|\mathcal{X}|/\delta)$
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
$\mathcal{X}_{\ell+1} \leftarrow \mathcal{X}_\ell \setminus \{x \in \mathcal{X}_\ell : \max_{x' \in \mathcal{X}_\ell} \langle x' - x, \widehat{\theta}_\ell \rangle > 2\epsilon_\ell\}$
$\ell \leftarrow \ell + 1$
Output: $\mathcal{X}_\ell$

$$\mathcal{E}_{x,\ell}(\mathcal{V}) = \{|\langle x, \widehat{\theta}_\ell(\mathcal{V}) - \theta^* \rangle| \le \epsilon_\ell\}$$

$$\begin{aligned}
\mathbb{P}\left(\bigcup_{\ell=1}^{\infty} \bigcup_{x \in \mathcal{X}_\ell} \{\mathcal{E}^c_{x,\ell}(\mathcal{X}_\ell)\}\right) &\le \sum_{\ell=1}^{\infty} \mathbb{P}\left(\bigcup_{x \in \mathcal{X}_\ell} \{\mathcal{E}^c_{x,\ell}(\mathcal{X}_\ell)\}\right) \\
&= \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \mathbb{P}\left(\bigcup_{x \in \mathcal{V}} \{\mathcal{E}^c_{x,\ell}(\mathcal{V})\}, \mathcal{X}_\ell = \mathcal{V}\right) \\
&= \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \mathbb{P}\left(\bigcup_{x \in \mathcal{V}} \{\mathcal{E}^c_{x,\ell}(\mathcal{V})\}\right) \mathbb{P}(\mathcal{X}_\ell = \mathcal{V}) \\
&\le \sum_{\ell=1}^{\infty} \sum_{\mathcal{V} \subseteq \mathcal{X}} \frac{\delta|\mathcal{V}|}{2\ell^2|\mathcal{X}|} \mathbb{P}(\mathcal{X}_\ell = \mathcal{V}) \le \delta
\end{aligned}$$