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

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 $x_1, \ldots, x_T \in \mathcal{X}$ and measure corresponding rewards $y_1, \ldots, y_T$ 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 $x_* = \arg\max_{x \in \mathcal{X}} x^T \theta_*$

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

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

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

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

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

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

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

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

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

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

$$\mathbb{E}[(\hat{\theta} - \theta_*)(\hat{\theta} - \theta_*)^T] = X^T X$$

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

$$= w^T z$$

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

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

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

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

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

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

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

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

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

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

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

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

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

$= d$ if λ is G-optimal

$$X^T X = v_1 v_1^T \sigma_1 + v_2 v_2^T \sigma_2$$

$$(X^T X)^{-1} = v_1 v_1^T \sigma_1^{-1} + v_2 v_2^T \sigma_2^{-1}$$

$$x^T (X^T X)^{-1} x = (x^T v_1)^2 \sigma_1^{-1} + (x^T v_2)^2 \sigma_2^{-1}$$

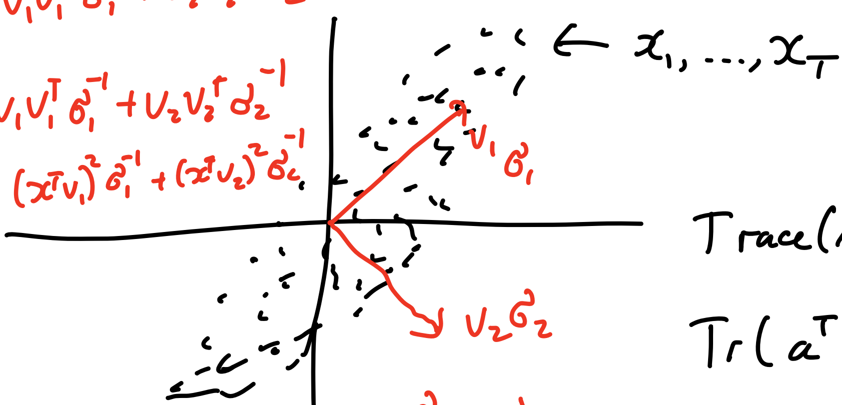

$\sigma_1 \gg \sigma_2$

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

$$\text{Tr}(a^T B a) = \text{Tr}(B a a^T)$$

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

$$\mathbb{E}\left[\|\hat{\theta} - \theta_\star\|_2^2\right] = \mathbb{E}\left[\|(X^TX)^{-1}X^T\xi\|_2^2\right]$$

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

$$\mathbb{E}[\xi_i \xi_j] \qquad = \mathbb{E}\left[\mathrm{Trace}\left((X^TX)^{-1}(X^TX)^{-1}X^T\xi\xi^T X\right)\right]$$

$$= \mathbb{E}[\xi_i^2]\,\mathbb{1}\{i=j\} \qquad = \mathrm{Trace}\left((X^TX)^{-1}(X^TX)^{-1}X^T\,\mathbb{E}[\xi\xi^T]\,X\right)$$

$$\leq \mathbb{1}\{i=j\} \qquad = \mathrm{Trace}\left((X^TX)^{-1}\right)$$

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

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

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

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

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

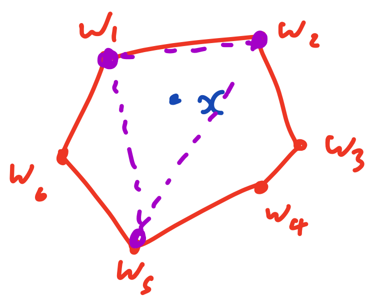

$$x = \sum_{i=1}^6 p_i w_i \quad \text{for some } p \in \Delta_6$$

$$= p_1' w_1 + p_2' w_2 + p_5' w_5$$

Caratheodory

Any $x \in \text{conv-hull}(w_1, \ldots, w_n) \subset \mathbb{R}^d$ can be represented as convex

Combo of $d+1$ ...

$A(\lambda) = \sum_{x \in \mathcal{X}} \lambda_x x x^T$

- **A-optimality**: minimize $f_A(\lambda) = \mathrm{Tr}(A_\lambda^{-1})$ minimizes $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2]$
- **E-optimality**: minimize $f_E(\lambda) = \max_{u:\|u\|\le 1} u^\top A_\lambda^{-1} u$ minimizes $\max_{u:\|u\|\le 1} \mathbb{E}[(\langle u, \widehat{\theta} - \theta\rangle)^2]$
- **D-optimality**: maximize $g_D(\lambda) = \log(|A_\lambda|)$ maximizes the entropy of distribution. Also, if $\mathcal{E}_\lambda = \{x : x^\top A_\lambda^{-1} x \le d\}$ 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*$(\lambda^*) = (d+1)d/2$

$\frac{\partial}{\partial x} \log|A(\lambda)|$

$= x^T A^{-1}(\lambda) x$

$\|x\|_A^2 = x^T A x$ so $\|x\|_2^2 = \|x\|_I^2$

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 $1-\delta$*

$X^T X = A(\lceil \tau\lambda \rceil)$

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

Pukelsheim

Pronzato ...

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

$\sigma(z) = \frac{e^z}{e^z + 1}$

$= \frac{1}{\sqrt{2\pi}} e^{-(\langle x, \theta_*\rangle - y)^2/2}$

In general for likelihood model $\ell_x(y)$ define

Fisher Info matrix as $\mathbb{E}_Y\left[\sum_x \nabla_\theta \log \ell_x(Y_x) \nabla_\theta \log \ell_x(Y_x)^T\right] = I(\theta)$

For linear regression $y \sim N(\langle x, \theta\rangle, 1)$ $\quad I(\theta) = \sum_x x x^T$

For logistic regression $y \sim \mathrm{Bernoulli}(\sigma(\langle x, \theta\rangle))$ $\quad I(\theta) = \sum_x x x^T \sigma(\langle x, \theta\rangle)(1 - \sigma(\langle x, \theta\rangle))$

$$\|x'\|^2_{\left(\sum_{x\in\mathcal{X}} xx^T \lceil \lambda_x \varsigma \rceil\right)^{-1}} = x'^T \left(\sum_{x\in\mathcal{X}} xx^T \lceil \lambda_x \varsigma \rceil \right)^{-1} x'$$

$$\leq x'^T \left(\sum_{x\in\mathcal{X}} xx^T \lambda_x \varsigma \right)^{-1} x'$$

$$= \frac{1}{\varsigma} \; x'^T \left(\sum_{x\in\mathcal{X}} xx^T \lambda_x \right)^{-1} x'$$

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

$$= \frac{1}{\varsigma} \|x'\|^2_{A(\lambda)^{-1}}$$

$$\mathbb{P}(A) = \sum_b \mathbb{P}(A \cap b)$$

$$\sum_{\ell=1}^{\infty} \frac{1}{\ell^2} = \frac{\pi^2}{6} \leq 2$$

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: $\mathcal{X}_\ell$

On $\mathcal{V} = \mathcal{X}_\ell$

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

$$\mathbb{P}\Big(\bigcap_{x \in \mathcal{X}} \bigcap_{\ell \ge 1}^{\infty} \mathcal{E}_{x,\ell}(\mathcal{X}_\ell)\Big) \ge 1 - \delta$$

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