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

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 $x_* = \arg\max_{x \in \mathcal{X}} x^T \theta_*$

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

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

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

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

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

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

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

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

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

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

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

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

$$= w^T z$$

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

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

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

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

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

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

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

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

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

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

$$\leq x_*^T (\theta_* - \hat\theta) - \hat{x}^T (\theta_* - \hat\theta)$$

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

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

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

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

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

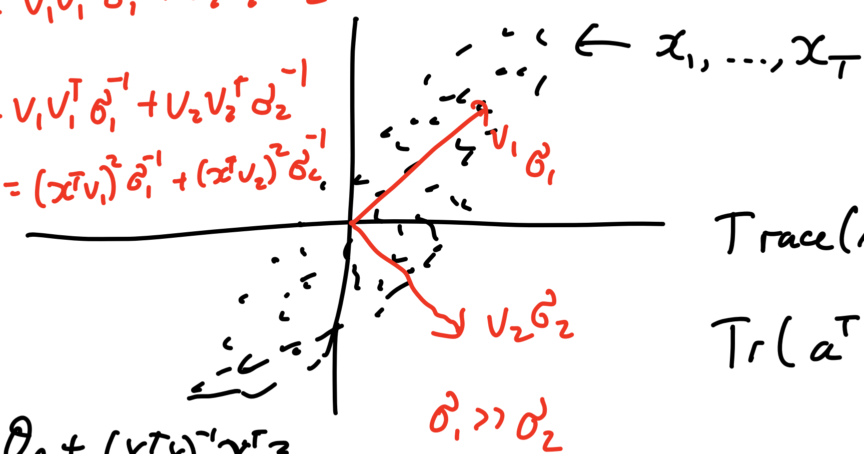

$$\text{Trace}(ABC) = \text{Trace}(BCA) = \text{Tr}(CAB)$$

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

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

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

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

$$\mathbb{E}[\eta_i \eta_j] = \mathbb{E}[\eta_i^2] \mathbb{1}\{i=j\} \leq \mathbb{1}\{i=j\}$$

$$= \mathbb{E}[\mathrm{Trace}((X^T X)^{-1} (X^T X)^{-1} X^T \eta \eta^T X)]$$

$$= \mathrm{Trace}((X^T X)^{-1} (X^T X)^{-1} X^T \mathbb{E}[\eta \eta^T] X)$$

$$= \mathrm{Trace}((X^T X)^{-1})$$

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

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

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

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

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

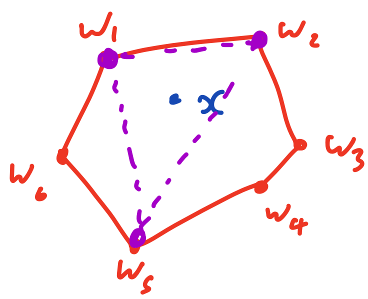

$$x = \sum_{i=1}^{6} p_i w_i \quad \text{for some } p \in \Delta_6$$

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

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 pts

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

- **A-optimality**: minimize $f_A(\lambda) = \mathrm{Tr}(A_\lambda^{-1})$ minimizes $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2]$
- **E-optimality**: minimize $f_E(\lambda) = \max_{u:\|u\|\leq 1} u^\top A_\lambda^{-1} u$ 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*$(\lambda^*) = (d+1)d/2$

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

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

$$\|x\|_A^2 = x^T A x \quad \text{so} \quad \|x\|_2^2 = \|x\|_I^2$$

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 $1 - \delta$*

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

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

Pudelsheim Pronzato...

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

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

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

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 $y \sim \mathcal{N}(\langle x, \theta \rangle, 1)$ $\quad I(\theta) = \sum_x x x^T$

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

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

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

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

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

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

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

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

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 $\underbrace{\lceil \widehat{\lambda}_{\ell,x} \tau_\ell \rceil}$ 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 $\mathcal{V} = \mathcal{X}_\ell$

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

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

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

$$\le \sum_{x \in \mathcal{V}} \underbrace{\mathbb{P}(\mathcal{E}_{x,\ell}(\mathcal{V})^c)} \le \frac{\delta}{2\ell^2 |\mathcal{X}|}$$

$x_* = \arg\max_x \langle x, \theta_* \rangle$ then WTS $x_* \in \mathcal{X}_\ell$ for all ℓ

$$\Rightarrow \quad \max_{x' \in \mathcal{X}_\ell} \langle x' - x_*, \hat{\theta}_\ell \rangle = \max_{x' \in \mathcal{X}_\ell} \underbrace{\langle x' - x_*, \theta_* \rangle}_{\le 0} + \underbrace{\langle x' - x_*, \hat{\theta} - \theta_* \rangle}_{\le 2\varepsilon_\ell}$$

$$\le 2\varepsilon_\ell$$

$\underline{\max_{x \in \mathcal{X}_\ell} \langle x_* - x, \theta_* \rangle \le 8\varepsilon_\ell}$ by same argument as MAB case.

$$R_T = \sum_{t=1}^{T} \langle x_* - x_t, \theta_* \rangle \qquad \Delta_x = \langle x_* - x, \theta_* \rangle$$

$$= \nu T + \sum_{\ell} \sum_{\substack{x \in \mathcal{X}_\ell \\ \Delta_x > \nu}} \underbrace{\langle x_* - x, \theta_* \rangle}_{\le 8\varepsilon_\ell} \lceil \lambda_{\ell,x} \zeta \rceil$$

$$\le \nu T + \sum_{\ell} \sum_{\substack{x \in \mathcal{X}_\ell \\ \Delta_x > \nu}} 8\varepsilon_\ell \lceil \lambda_{\ell,x} \zeta \rceil$$

$$\le \nu T + \sum_{\ell} \Big(8\varepsilon_\ell d^2 + \sum_{\substack{x \in \mathcal{X}_\ell \\ \Delta_x > \nu}} 8\varepsilon_\ell \lambda_{\ell,x} \zeta \Big)$$

$$= \nu T + \sum_{\ell} 8\varepsilon_\ell d^2 + \sum_{\substack{x \in \mathcal{X}_\ell \\ \Delta_x > \nu}} \sum_{\ell=1}^{\log_2 (\Delta_x \vee \varepsilon_\ell)^{-1}} 8\varepsilon_\ell \lambda_{\ell,x} \zeta$$

$$\lesssim \nu T + d^2 \log(|\mathcal{X}|/\delta) + d(\Delta \vee \nu)^{-1} \log(|\mathcal{X}|/\delta)$$

$$\nu = 0 \qquad R_T \le d^2 \log(|\mathcal{X}|/\delta) + \frac{d}{\Delta} \log(|\mathcal{X}|/\delta)$$

min wrt $\nu \implies R_T \le d^2 \log(|\mathcal{X}|/\delta) + \sqrt{dT \log(|\mathcal{X}|/\delta)}.$

Linear Bandits - Pure exploration / Best-arm identification

Motivation: If we choose $x_1, \ldots, x_3$ according to λ (i.e. after rounding via Caratheodory) and

$y_t = \langle x_t, \theta_* \rangle + \eta_t$ w/ $\eta_t \sim \mathcal{N}(0,1)$ then

$$\hat{\theta} = \theta_* + \left(\sum_t x_t x_t^T \right)^{-1} \sum_t x_t \eta_t .$$

If $x_* = \underset{x \in \mathcal{X}}{\text{argmax}} \langle x, \theta_* \rangle$ and $\hat{x} = \underset{x \in \mathcal{X}}{\text{argmax}} \langle x, \hat{\theta} \rangle$

we have that $\hat{x} = x_*$ if $\langle x_* - x, \hat{\theta} \rangle > 0 \quad \forall x \in \mathcal{X} \setminus x_*$.

$$\langle x_* - x, \hat{\theta} \rangle = \langle x_* - x, \theta_* \rangle + \underbrace{\langle x_* - x, \hat{\theta} - \theta_* \rangle}_{}$$

for any $v \in \mathbb{R}^d$ $\quad |v^T(\hat{\theta} - \theta_*)| \leq \|v\|_{(\sum x_t x_t^T)^{-1}} \sqrt{2 \log(2/\delta)}$

$$\sum_{t=1}^{3} x_t x_t^T \simeq 3 \cdot A(\lambda)$$

$$\geq - \|x_* - x\|_{(\sum x_t x_t^T)^{-1}} \sqrt{2 \log |\mathcal{X}| / \delta}$$

$$\Longrightarrow \quad \frac{\|x_* - x\|^2_{(\sum_t x_t x_t^T)^{-1}}}{\langle x_* - x, \theta_* \rangle^2} \leq \frac{1}{c \log |\mathcal{X}| / \delta} \quad \forall x$$

$$\underbrace{\min_{\lambda \in \Delta_{\mathcal{X}}} \max_{x \in \mathcal{X} \setminus x_*} \frac{\|x_* - x\|^2_{A(\lambda)^{-1}}}{\langle x_* - x, \theta_* \rangle^2}}_{=: \rho_*} \cdot c \log(|\mathcal{X}|/\delta) \leq \tau$$

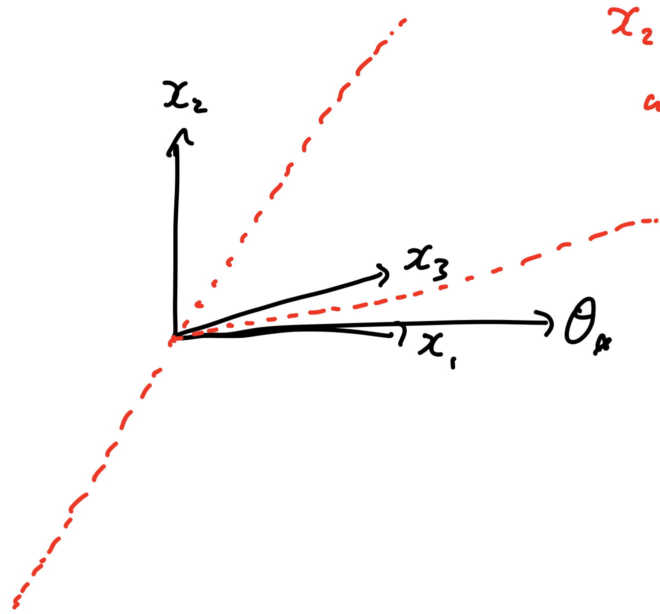

We need to estimate $\langle x_1 - x_3, \theta_* \rangle$ accurately. x_2 is most aligned w/ $x_1 - x_3$ and is thus optimally reduces # meas to identify x_1 as best.

Input: Finite set $\mathcal{X} \subset \mathbb{R}^d$, confidence level $\delta \in (0,1)$.
Let $\mathcal{X}_1 \leftarrow \mathcal{X}, t \leftarrow 1$
while $|\mathcal{X}_\ell| > 1$ **do**
  Let $\widehat{\lambda}_\ell \in \triangle_{\mathcal{X}}$ be a $\frac{d(d+1)}{2}$-sparse minimizer of $f(\lambda; \mathcal{X}_\ell)$ where
  $$f(\mathcal{V}) = \inf_{\lambda \in \mathcal{X}} f(\lambda; \mathcal{V}) = \inf_{\lambda \in \mathcal{X}} \max_{x, x' \in \mathcal{V}} \|x - x'\|^2_{(\sum_{x \in \mathcal{X}} \lambda_x x x^\top)^{-1}}$$
  Set $\epsilon_\ell = 2^{-\ell}$, $\tau_\ell = 2\epsilon_\ell^{-2} f(\mathcal{X}_\ell) \log(4\ell^2 |\mathcal{X}|/\delta)$
  Pull arm $x \in \mathcal{X}$ exactly $\lceil \tau_\ell \widehat{\lambda}_{\ell,x} \rceil$ times and construct $\widehat{\theta}_\ell$
  $\mathcal{X}_{\ell+1} \leftarrow \mathcal{X}_\ell \setminus \{x \in \mathcal{X}_\ell : \max_{x' \in \mathcal{X}_\ell} \langle x' - x, \widehat{\theta}_t \rangle > \epsilon_\ell\}$
  $t \leftarrow t + 1$
Output: $\mathcal{X}_{t+1}$

$$\Delta_x = \langle x_* - x, \theta_* \rangle \qquad \Delta = \min_{x \neq x_*} \Delta_x$$

$$\rho_* = \min_{\lambda \in \Delta_{\mathcal{X}}} \max_{x \in \mathcal{X} \setminus x_*} \frac{\|x - x_*\|^2_{A(\lambda)^{-1}}}{\Delta_x^2}$$

Theorem 5. *Assume that* $\max_{x \in \mathcal{X}} \langle x^\star - x, \theta^* \rangle \leq 2$. *Then with probability at least* $1 - \delta$, $x^\star$ *is returned from the algorithm at a time* τ *that satisfies*

$$\tau \leq c\rho^\star \log(\Delta^{-1})[\log(1/\delta) + \log(\log(\Delta^{-1})) + \log(|\mathcal{X}|)].$$

Furthermore, <u>any</u> algorithm requires $\mathbb{E}[\tau] \geq \rho_* \log(1/\delta)$.

Thompson Sampling / linear UCB

function class

Define prior $P_1 = \mathcal{N}(0, I_d)$, P_1 prior over $\mathcal{F}$

for $t = 1, 2, \ldots$

$\theta_t \sim P_t(\cdot)$ $\qquad f_t \sim P_t$

$x_t := \underset{x \in \mathcal{X}}{\operatorname{argmax}} \langle x, \theta_t \rangle$ $\qquad x_t = \underset{x \in \mathcal{X}}{\operatorname{argmax}} f_t(x)$

Observe $y_t = \langle x_t, \theta_* \rangle + \eta_t$

Model η_t as $\mathcal{N}(\langle x_t, \theta_* \rangle, 1)$

Update posterior $\underbrace{P_{t+1}(\theta \mid \{(x_s, y_s)\}_{s=1}^{t})}$ $\qquad$ Update $P_t \mapsto P_{t+1}$

$$\simeq \mathcal{N}\left(\hat{\theta}, \left(I + \sum_t x_t x_t^T\right)^{-1}\right)$$

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