

Stochastic Linear Bandits

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

for $t = 1, 2, \dots, T$

Player chooses $I_t \in [n] = \{1, \dots, n\}$

Nature reveals $\langle x_{I_t}, \theta^* \rangle + z_t$ ($\mathbb{E}[z] = 0$,
 z -sub-Gaussian)

$$x_* = \operatorname{argmax}_{x \in \mathcal{X}} \langle x, \theta^* \rangle$$

$$\text{Regret}(T) = \max_{x \in \mathcal{X}} \sum_{t=1}^T \langle x, \theta^* \rangle - \langle x_{I_t}, \theta^* \rangle$$

$$= \sum_{t=1}^T \langle x_*, \theta^* \rangle - \langle x_{I_t}, \theta^* \rangle$$

$$= \sum_{t=1}^T \langle x_* - x_{I_t}, \theta^* \rangle$$

$$= \sum_{x \in \mathcal{X}} \langle x_* - x, \theta^* \rangle T_x$$

$$= \sum_{x \in \mathcal{X}} \Delta_x T_x$$

$$T_x = \sum_{t=1}^T \mathbb{1}\{x_{I_t} = x\}$$

$$\Delta_x = \langle x_* - x, \theta^* \rangle$$

Least squares $X = \begin{bmatrix} -x_1^T \\ \vdots \\ -x_n^T \end{bmatrix} \in \mathbb{R}^{n \times d}$ $y = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} \in \mathbb{R}^n$

Suppose I am given a dataset $\{(x_i, y_i)\}_{i=1}^n$ $z = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix} \in \mathbb{R}^n$

$$y_i = \langle x_i, \theta^* \rangle + z_i, \quad z_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$$

MLE (maximum likelihood estimator) is the LS estimator:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^n (y_i - \langle x_i, \theta \rangle)^2 \quad z \sim \mathcal{N}(0, I)$$

$$= \underset{\theta}{\operatorname{argmin}} \|y - X\theta\|_2^2 \quad (y = X\theta + z)$$

$$= (X^T X)^{-1} X^T y$$

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

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

(Note $\mathbb{E}[z] = 0$
so $\mathbb{E}[\hat{\theta}] = \theta_*$)

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

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

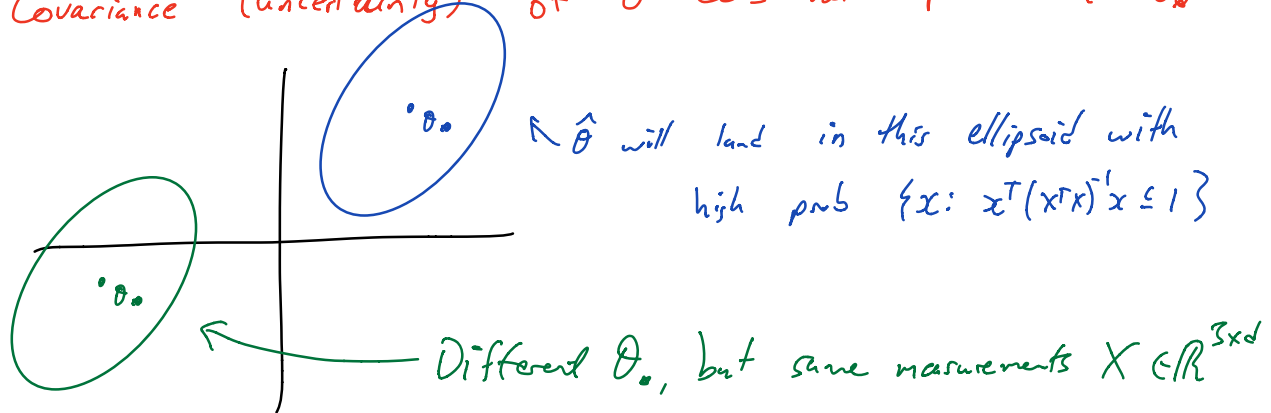
Any linear transformation
of Gaussian vector is
also a Gaussian vector

$$= (X^T X)^{-1} X^T X (X^T X)^{-1}$$

$$= (X^T X)^{-1}$$

$$\rightarrow \hat{\theta} \sim \mathcal{N}(\theta_*, (X^T X)^{-1}), \quad \hat{\theta} - \theta_* \sim \mathcal{N}(0, (X^T X)^{-1})$$

Covariance (uncertainty) of $\hat{\theta}$ does not depend on θ_*




Since the "shape" of uncertainty only depends on the measurement vectors $X = \begin{bmatrix} -x_1^+ \\ \vdots \\ -x_n^+ \end{bmatrix}$, if we have a desired "shape" we can choose X to hit it.

This observation \uparrow is the driving motivation for optimal linear experimental design.

To minimize regret we want to identify those arms $x \in X$ w/ $\langle x_p - x, \theta^* \rangle > 0$ ASAP to stop playing them and reducing regret.

We don't know θ_* , but estimate gap w/ $\langle x_p - x, \hat{\theta} \rangle$

Idea: choose $X = (x_1, \dots, x_n)$ in order to minimize



Confidence intervals

$$\hat{\theta} - \theta_* \sim \mathcal{N}(0, (X^T X)^{-1})$$

For any vector $x \in \mathcal{X}$, $\langle x, \hat{\theta} - \theta_* \rangle$ is


$$\begin{aligned} \text{Gaussian w/ mean 0 and variance } x^T (X^T X)^{-1} x. \\ =: \|x\|_{(X^T X)^{-1}}^2 \end{aligned}$$

By sub-Gaussian tail-bound

$$\mathbb{P}(\langle x, \hat{\theta} - \theta_* \rangle \geq \|x\|_{(X^T X)^{-1}} \sqrt{2 \log(1/\delta)}) \leq \delta.$$

(see lecture notes, ch. 2)

Any $x \in \mathbb{R}^d$.



$$\begin{aligned} & \mathbb{P}(\exists u, \|u\|_2 = 1, \langle u, \hat{\theta} - \theta_* \rangle \geq \|u\|_{(X^T X)^{-1}} \sqrt{2 \log(m/\delta)}) \\ & \leq \mathbb{P}\left(\bigcup_{i=1}^m \left\{ \langle u_i, \hat{\theta} - \theta_* \rangle \geq \|u_i\|_{(X^T X)^{-1}} \sqrt{2 \log(m/\delta)} + \epsilon \right\}\right) \\ & \leq \delta \end{aligned}$$

$$P(\langle x, \hat{\theta} - \theta_* \rangle \geq \|x\|_{(X^T X)^{-1}} \sqrt{2 \log(1/\delta)}) \leq \delta. \quad (1)$$

$$\max_{x \in \mathcal{X}} \mathbb{E}[\langle x, \hat{\theta} - \theta_* \rangle^2] = \max_{x \in \mathcal{X}} \|x\|_{(X^T X)^{-1}}^2$$

Choose $X = (x_1, \dots, x_n)$
to minimize

For any set $(x_1, \dots, x_n) \in \mathcal{X}$

$$\exists \lambda \in \Delta_{\mathcal{X}} = \{z \in \mathbb{R}^{|\mathcal{X}|} : \sum_{x \in \mathcal{X}} z_x = 1, z_x \geq 0\}$$

$$\wedge / \quad X^T X = \sum_{t=1}^n x_t x_t^T = \mathfrak{S} \sum_{x \in \mathcal{X}} \lambda_x x x^T$$

\Rightarrow motivates relaxing to finding the

continuous / fractional design problem

$$\min_{\lambda \in \Delta_{\mathcal{X}}} \max_{x \in \mathcal{X}} \|x\|_{\left(\sum_{x \in \mathcal{X}} \lambda_x x x^T\right)^{-1}}^2 = \min_{\lambda} f(\lambda)$$

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

$$f(\lambda) = \max_{x \in \mathcal{X}} \|x\|_{A(\lambda)^{-1}}^2$$

$$g(\lambda) = \log \det(A(\lambda))$$

Lemma 4 (Kiefer-Wolfowitz (1960)). For any \mathcal{X} with $d = \dim(\text{span}(\mathcal{X}))$, there exists a $\lambda^* \in \Delta_{\mathcal{X}}$ that

- $\max_{\lambda} g_D(\lambda) = g_D(\lambda^*)$
- $\min_{\lambda} f_G(\lambda) = f_G(\lambda^*)$
- $f_G(\lambda^*) = d \frac{\log 2}{2}$
- $\text{support}(\lambda^*) = (d+1)d/2$

$$\lambda^* = \underset{\lambda}{\text{argmin}} f(\lambda)$$


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 $\hat{\theta}$. Then for each $x \in \mathcal{X}$ we have with probability at least $1 - \delta$

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

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

$$y_t \sim \text{Bernoulli}(\sigma(\langle x_t, \theta_* \rangle)), \quad \sigma(z) = \frac{1}{1 + e^{-z}}$$

$$\hat{\theta}_T \text{ MLE of } \{(x_t, y_t)\}_{t=1}^T$$



As $T \rightarrow \infty$ we have by CLT

$$\hat{\theta}_T \rightarrow \mathcal{N}\left(\theta_*, \left(\sum_{t=1}^T \sigma(\langle x_t, \theta_* \rangle) (1 - \sigma(\langle x_t, \theta_* \rangle)) x_t x_t^\top\right)^{-1}\right)$$

Input: Finite set $\mathcal{X} \subset \mathbb{R}^d$, confidence level $\delta \in (0, 1)$.

Let $\hat{\mathcal{X}}_1 \leftarrow \mathcal{X}, \ell \leftarrow 1$

while $|\hat{\mathcal{X}}_\ell| > 1$ **do**

Let $\hat{\lambda}_\ell \in \Delta_{\hat{\mathcal{X}}_\ell}$ be a $\frac{d(d+1)}{2}$ -sparse minimizer of $f(\lambda) = \max_{x \in \hat{\mathcal{X}}_\ell} \|x\|^2_{(\sum_{x \in \hat{\mathcal{X}}_\ell} \lambda_x x x^\top)^{-1}}$ } choosing
" design

$\epsilon_\ell = 2^{-\ell}, \tau_\ell = 2d\epsilon_\ell^{-2} \log(4\ell^2|\mathcal{X}|/\delta)$

Pull arm $x \in \mathcal{X}$ exactly $\lceil \hat{\lambda}_{\ell, x} \tau_\ell \rceil$ times and construct the least squares estimator $\hat{\theta}_\ell$ using only the observations of this round

$\hat{\mathcal{X}}_{\ell+1} \leftarrow \hat{\mathcal{X}}_\ell \setminus \{x \in \hat{\mathcal{X}}_\ell : \max_{x' \in \hat{\mathcal{X}}_\ell} \langle x' - x, \hat{\theta}_\ell \rangle > 2\epsilon_\ell\}$

$\ell \leftarrow \ell + 1$

Output: $\hat{\mathcal{X}}_\ell$

Lemma 5. Assume that $\max_{x \in \mathcal{X}} \langle x^* - x, \theta^* \rangle \leq 4$. With probability at least $1 - \delta$, we have $x^* \in \hat{\mathcal{X}}_\ell$ and $\max_{x \in \hat{\mathcal{X}}_\ell} \langle x^* - x, \theta^* \rangle \leq 8\epsilon_\ell$ for all $\ell \in \mathbb{N}$.