Linear Bandits

Least squares review

Suppose I measure $x_1, \dots, x_T \in \mathbb{R}^d$ then observe

$\forall t \quad y_t = \langle x_t, \theta_* \rangle + \eta_t$ where η_t is mean-zero σ^2-sub-Gaussian

$$\hat{\theta} = \arg\min_\theta \sum_{t=1}^{T} (y_t - \langle x_t, \theta \rangle)^2 \qquad X = \begin{bmatrix} - x_1^T - \\ \vdots \\ - x_T^T - \end{bmatrix} \in \mathbb{R}^{T \times d}$$

$$= \arg\min_\theta \| y - X\theta \|_2^2 \qquad y = \begin{bmatrix} y_1 \\ \vdots \\ y_T \end{bmatrix} \in \mathbb{R}^T$$

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

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

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

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

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

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

$$\preceq \sigma^2 (X^T X)^{-1} X^T X (X^T X)^{-1}$$

$$= \sigma^2 (X^T X)^{-1}$$

Let $z \in \mathbb{R}^d$ be some new point.

$$\mathbb{E}[\langle z, \hat{\theta} \rangle] = \langle z, \theta_* \rangle \qquad \mathbb{V}(\langle z, \hat{\theta} \rangle) = \mathbb{E}[\langle z, \hat{\theta} - \theta_* \rangle^2]$$

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

For any PSD $A \in \mathbb{R}^{d \times d}$

define $\|x\|_A^2 = x^T A x$.

$$\leq \sigma^2 z^T (X^T X)^{-1} z.$$

$$= \sigma^2 \|z\|^2_{(X^T X)^{-1}}$$

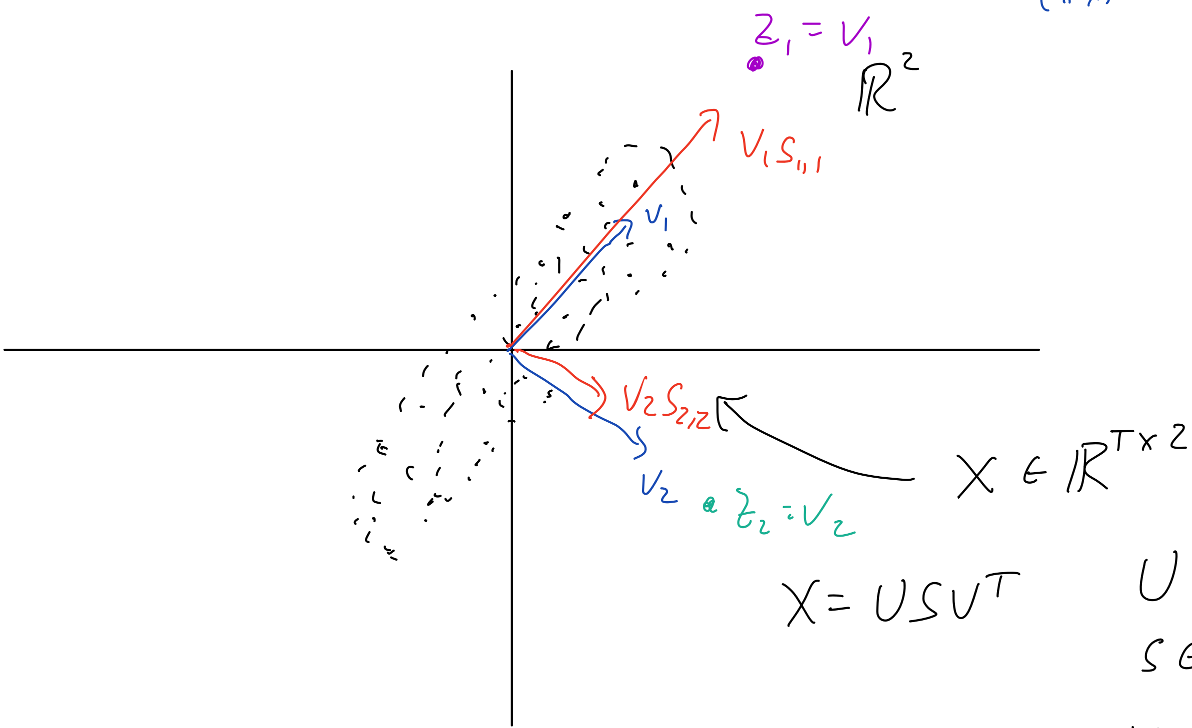

$X = USV^T$

$U \in \mathbb{R}^{T \times 2}$

$S \in \mathbb{R}^{2 \times 2}$

$V = \mathbb{R}^{2 \times 2}$

$$U = \begin{bmatrix} — u_1^T — \\ \vdots \\ — u_T — \end{bmatrix}$$

$U^T U = I_2$

$V^T V = I_2$

$$V = \begin{bmatrix} | & | \\ v_1 & v_2 \\ | & | \end{bmatrix}$$

$$x_i = u_i^T S V$$

$$= u_{i,1} s_{11} v_1 + u_{i,2} s_{22} v_2$$

$$(X^T X)^{-1} = (V S U^T U S V^T)^{-1}$$

$$= (V S^2 V^T)^{-1}$$

$$= V S^{-2} V^T$$

V is orthonormal

$V^{-1} = V^T$

If $z_1 = v_1$ then

$$z_1^T (X^T X)^{-1} z_1 = v_1^T V S^{-2} V^T v_1$$
$$= e_1^T S^{-2} e_1$$
$$= \frac{1}{S_{1,1}^2}$$

If $z_2 = v_2$

$$z_2^T (X^T X)^{-1} z_2 = \frac{1}{S_{2,2}^2}$$

How do I choose $x_1, \dots, x_T \in \mathcal{X}$ to minimize $\mathbb{E}[\langle z, \hat{\theta} - \theta_* \rangle^2]$?

Suppose I have a pool of potential measurements $\mathcal{X} \subset \mathbb{R}^d$.

For simplicity assume $|\mathcal{X}| = n$.

For any $X^T X = \sum_{i=1}^T x_i x_i^T$ $\exists \lambda \in \Delta_{\mathcal{X}} = \{ w_i \geq 0 : \sum_{i=1}^n w_i = 1 \}$ s.t. $X^T X = T \sum_{x \in \mathcal{X}} \lambda_x x x^T$.

$=: T \cdot A(\lambda)$

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

$$\mathbb{E}[\langle z, \hat{\theta} - \theta_* \rangle^2] \leq \sigma^2 z^T (X^T X)^{-1} z$$
$$= \sigma^2 \frac{z^T A(\lambda)^{-1} z}{T}.$$

Idea: Choose $\lambda \in \Delta_{\mathcal{X}}$ to minimize RHS! Before seeing any data then choose $x_1, \dots, x_T$ "in proportion" to $\lambda \in \Delta_{\mathcal{X}}$.

A-optimality

$$\sum_{i=1}^{d} \mathbb{E}[\langle e_i, \hat{\theta} - \theta_* \rangle^2] = \sum_{i=1}^{d} \mathbb{E}[(\hat{\theta} - \theta_*)^T e_i e_i^T (\hat{\theta} - \theta_*)]$$

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

$$= \mathbb{E}[\|\hat{\theta} - \theta_*\|_2^2]$$

$$\mathbb{E}[\|\hat{\theta} - \theta_*\|_2^2] = \sum_{i=1}^{d} \mathbb{E}[\langle e_i, \hat{\theta} - \theta_* \rangle^2]$$

$$[e_i]_j = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}$$

$$\leq \sum_{i=1}^{d} \sigma^2 \frac{e_i^T A(\lambda)^{-1} e_i}{T}$$

$$= \frac{\sigma^2}{T} \cdot \mathrm{Trace}(A(\lambda)^{-1}).$$

A-optimal design: $\underset{\lambda \in \Delta_{\mathcal{X}}}{\operatorname{argmin}} \ \mathrm{Trace}(A(\lambda)^{-1}).$

E-optimal

$$\max_{u : \|u\|_2 \leq 1} \mathbb{E}[\langle u, \hat{\theta} - \theta_* \rangle^2] \leq \frac{\sigma^2}{T} \max_{u : \|u\| \leq 1} u^T A(\lambda)^{-1} u$$

$$= \frac{\sigma^2}{T} \lambda_{max}(A(\lambda)^{-1})$$

$$= \frac{\sigma^2}{T} \cdot \frac{1}{\lambda_{min}(A(\lambda))}$$

(e, v_i) are eigen pairs

$Av_i = e_i v_i$

$$f_E(\lambda) = \lambda_{max}(A(\lambda)^{-1})$$

D-optimal

$$g_D(\lambda) = \log|A(\lambda)|$$

λ^* is D-optimal if $\lambda^* = \underset{\lambda}{\operatorname{argmax}} \ g_D(\lambda).$

D-optimal maximizes product of eigs of $A(\lambda)$.

$\mathcal{N}(\mu, \Sigma)$ has density $\frac{1}{(2\pi|\Sigma|)^{d/2}} \exp\left(-(x-\mu)^T \Sigma^{-1} (x-\mu)\right)$.

<u>G-optimal</u>

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

G-opt minimizes $\max_{x \in \mathcal{X}} \mathbb{E}[\langle x, \hat{\theta} - \theta_* \rangle^2]$.