Matrix concentration inequalities

Recall from Lecture 2 the following version of Bernstein's ineq.:

Thm (Bernstein): If $X_1, \dots, X_n$ are independent, $\mathbb{E}X_i = 0$, $\mathbb{E}X_i^2 \le \sigma^2$, $X_i \le b$ a.s., then

$$\mathbb{P}\left[\sum_{i=1}^n X_i \ge t\right] \le e^{-\frac{t^2/2}{n\sigma^2 + bt/3}}$$

We will prove the following extension to matrices:

Thm (matrix Bernstein): If $X_1, \dots, X_n \in \mathbb{R}^{d\times d}$ are independent symmetric matrices, $\mathbb{E}X_i = 0$, $\lambda_{max}(X_i) \le b$ a.s., then

$$\mathbb{P}\left[\lambda_{max}\left(\sum_{i=1}^n X_i\right) \ge t\right] \le d \cdot e^{-\frac{t^2/2}{v + bt/3}}$$

where $v = \|\sum_{i=1}^n \mathbb{E}X_i^2\|_{op}$.

Remark: ① v is the matrix variance, representing "maximal variance of $S = \sum_{i=1}^n X_i$" in any direction:

$$v = \|\sum_{i=1}^n \mathbb{E}X_i^2\|_{op} = \|\mathbb{E}S^2\|_{op}$$

$$= \sup_{u : \|u\|_2 = 1} u^T(\mathbb{E}S^2)u = \sup_{u : \|u\|_2 = 1} \mathbb{E}\|Su\|_2^2.$$

② This shows $\lambda_{max}\left(\sum_{i=1}^n X_i\right) = O_p\left(\sqrt{v \log d} + b \log d\right)$.

The log d factors may not be sharp, but this bound is good enough for many applications.

③ There are versions that relax $\lambda_{max}(X_i) \le b$ to moment-type assumptions, see Tropp '12 "User-friendly tail bounds"

Example (covariance estimation): Let $X_1, \dots, X_n$ be iid, $\mathbb{E} X_i = 0$,

$\mathbb{E} X_i X_i^T = \Sigma$, $\quad \|X_i\|_2^2 \le K^2 \cdot \mathbb{E}\|X_i\|_2^2$ a.s. for some $K \ge 1$.

(E.g. $X_i = \Sigma^{1/2} z_i$, $\quad z_i \sim \mathrm{Unif}\{\pm\sqrt{d}\, e_1, \dots, \pm\sqrt{d}\, e_d\}$

$$\|X_i\|_2^2 = z_i^T \Sigma z_i \le d \cdot \max_{i=1}^{d} \Sigma_{ii}, \quad \mathbb{E}\|X_i\|_2^2 = \mathrm{Tr}\,\Sigma \ge d \cdot \min_{i=1}^{d} \Sigma_{ii}$$

so this holds w/ $K^2 = \max_i \Sigma_{ii} / \min_i \Sigma_{ii}$)

Let $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^{n} X_i X_i$ be sample covariance estimator.

$$\|\hat{\Sigma} - \Sigma\|_{op} = \Big\| \frac{1}{n} \sum_{i=1}^{n} \underbrace{(X_i X_i^T - \Sigma)}_{:= Y_i} \Big\|_{op}$$

- $\mathbb{E} Y_i = 0$
- $\|Y_i\|_{op} \le \|X_i\|_2^2 + \|\Sigma\|_{op} \le K^2 \mathrm{Tr}\,\Sigma + \|\Sigma\|_{op} \le 2K^2 \mathrm{Tr}\,\Sigma$.
- $\mathbb{E} Y_i^2 = \mathbb{E}(X_i X_i^T)^2 - \Sigma^2 \preceq \mathbb{E}(X_i X_i^T)^2$

$$= \mathbb{E}\|X_i\|_2^2 X_i X_i^T \preceq K^2 \mathrm{Tr}\,\Sigma \cdot \Sigma$$

$$\Rightarrow v := \Big\|\sum_{i=1}^{n} \mathbb{E} Y_i^2\Big\|_{op} \le K^2 n \cdot \mathrm{Tr}\,\Sigma \cdot \|\Sigma\|_{op}.$$

By matrix Bernstein (applied to Y_i and $-Y_i$),

$$\|\hat{\Sigma} - \Sigma\|_{op} = O_p\left(\sqrt{\frac{K^2 \log d}{n}} \cdot \sqrt{\operatorname{Tr}\Sigma \cdot \|\Sigma\|_{op}} + \frac{K^2 \log d}{n} \cdot \operatorname{Tr}\Sigma\right)$$

$$= O_p\left(\sqrt{\frac{K^2 r \log d}{n}} + \frac{K^2 r \log d}{n}\right) \cdot \|\Sigma\|_{op},$$

$r = \operatorname{Tr}\Sigma / \|\Sigma\|_{op}$ is the stable rank

This recovers result of Lecture 12 up to $K^2 \log d$ factor, without a subgaussian assumption for X_i.

Lieb's concavity theorem

For $X = U \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix} U^T \in \mathbb{R}^{n \times n}$, define $f(X)$ by functional calculus: $f(X) = U \begin{pmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_n) \end{pmatrix} U^T$.

Idea: Apply moment generating function approach to matrices.

$$\mathbb{P}\left[\lambda_{max}\left(\sum_{i=1}^n X_i\right) \geq t\right] \leq e^{-\lambda t}\, \mathbb{E}\, e^{\lambda \cdot \lambda_{max}\left(\sum_{i=1}^n X_i\right)} \quad (\forall \lambda \geq 0)$$

$$= e^{-\lambda t}\, \mathbb{E}\, \lambda_{max}\left(e^{\lambda \sum_{i=1}^n X_i}\right)$$

$$\leq e^{-\lambda t} \operatorname{Tr} \mathbb{E}\, e^{\lambda \sum_{i=1}^n X_i}$$

Issue: For matrices, $\operatorname{Tr} e^{\lambda \sum_{i=1}^n X_i} \neq \operatorname{Tr} e^{\lambda X_1} \times \cdots \times e^{\lambda X_n}$

Thm (Lieb): For any $H \in \mathbb{R}^{d \times d}$ symmetric,

$$A \mapsto \operatorname{Tr} e^{H + \log A}$$

is concave over $A \succ 0$.

Thus
$$\mathbb{E} \operatorname{Tr} e^{\lambda \sum_{i=1}^{n} X_i} = \mathbb{E} \operatorname{Tr} e^{\lambda \sum_{i=1}^{n-1} X_i + \log e^{\lambda X_n}}$$
$$\overset{\text{(Jensen)}}{\leq} \mathbb{E} \operatorname{Tr} e^{\lambda \sum_{i=1}^{n-1} X_i + \log \mathbb{E} e^{\lambda X_n}}$$
$$\leq \dots \leq \operatorname{Tr} e^{\sum_{i=1}^{n} \log \mathbb{E} e^{\lambda X_i}}$$

This will allow us to prove matrix Bernstein.

To prove Lieb's theorem:

Def: Let $A, H \in \mathbb{R}^{d \times d}$ be symmetric, $A, H \succ 0$. The <u>matrix relative entropy</u> is

$$D(A \| H) = \operatorname{Tr}\left[A(\log A - \log H) - (A - H)\right]$$

Remark: The vector analogue, for $a, h \in \mathbb{R}^d$ entrywise positive, is

$$D(a \| h) = \sum_{i=1}^{d} a_i \log \frac{a_i}{h_i} - (a_i - h_i).$$ If $\sum_i a_i = \sum_i h_i = 1$,

this is KL-divergence between two discrete distributions.

One may check that:

- $D(a \| h) \geq 0$ (non-negativity)
- $D(\lambda a + (1-\lambda) a' \| \lambda h + (1-\lambda) h') \leq \lambda D(a \| a') + (1-\lambda) D(h \| h')$

  $\forall \lambda \in [0,1]$ (convexity in (a,h)).

We aim to show the analogous properties for matrix relative entropy.

Def: A function $f : \mathcal{I} \to \mathbb{R}$ on an interval $\mathcal{I} \subseteq \mathbb{R}$ is

$$\begin{cases} \text{operator monotone increasing} \\ \text{operator convex} \\ \text{trace monotone increasing} \end{cases}$$

if, for <u>all</u> $d \geq 1$ and $A, H \in \mathbb{R}^{d \times d}$ symmetric w/ all eigenvalues in $\mathcal{I}$,

$$\begin{cases} A \succeq H \text{ implies } f(A) \succeq f(H) \\ f(\lambda A + (1-\lambda) H) \preceq \lambda f(A) + (1-\lambda) f(H) \quad \forall \lambda \in [0,1] \\ A \preceq H \text{ implies } \operatorname{Tr} f(A) \leq \operatorname{Tr} f(H) \end{cases}$$

Trace monotonicity is the simplest condition:

Prop: If $f : \mathcal{I} \to \mathbb{R}$ is monotone increasing, then it is also trace monotone increasing.

Proof: Let $\lambda_1(A) \geq \dots \geq \lambda_d(A)$ be eigenvalues of A. If $A \preceq H$, by Courant-Fischer min-max theorem,

$$\lambda_i(A) = \max_{V:\, \dim(V)=i} \ \min_{u \in V:\, \|u\|_2 = 1} u^T A u \le \max_{V} \ \min_{u} u^T H u = \lambda_i(H).$$

Thus $\mathrm{Tr}\, f(A) = \sum_i f(\lambda_i(A)) \le \sum_i f(\lambda_i(H)) = \mathrm{Tr}\, f(H).$ ■

It is <u>not</u> true that f increasing / convex $\Rightarrow$ f operator increasing / convex.

E.g. $A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, $f(x) = x^2$ on $[0, \infty)$.

$A \preceq B$ but $A^2 \not\preceq B^2$

$A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, $B = \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}$, $f(x) = x^3$ on $[0, \infty)$.

$$\frac{A^3 + B^3}{2} \not\succeq \left(\frac{A+B}{2}\right)^3$$

Prop: (a) Fix any $u \ge 0$. $a \mapsto (a+u)^{-1}$ is operator monotone decreasing and operator convex on $(0, \infty)$.

(b) $a \mapsto \log a$ is operator monotone increasing and operator concave on $(0, \infty)$.

Proof: (a) Suppose $H \succeq A \succ 0$. Then $H + uI \succeq A + uI$.

Note if $A \succeq B$ then $M^T A M \succeq M^T B M$ for any M.

$$\Rightarrow I \succeq (H+uI)^{-1/2}(A+uI)(H+uI)^{-1/2}$$

$$\Rightarrow I \precsim \left[(H+uI)^{-1/2}(A+uI)(H+uI)^{-1/2}\right]^{-1} = (H+uI)^{1/2}(A+uI)^{-1}(H+uI)^{1/2}$$

$\Rightarrow (A+uI)^{-1} \succeq (H+uI)^{-1}$. This shows monotonicity.

Consider any $A, H \succeq 0$, $\lambda \in [0,1]$.

Note: If $A \succ 0$, then $\begin{pmatrix} A & B \\ B^T & C \end{pmatrix} \succeq 0$ if and only $C - B^T A^{-1} B \succeq 0$.

Then
$$0 \precsim \lambda \cdot \underbrace{\begin{pmatrix} A+uI & I \\ I & (A+uI)^{-1} \end{pmatrix}}_{\succeq 0} + (1-\lambda)\underbrace{\begin{pmatrix} H+uI & I \\ I & (H+uI)^{-1} \end{pmatrix}}_{\succeq 0}$$

$$= \begin{pmatrix} \lambda A + (1-\lambda)H + uI & I \\ I & \lambda(A+uI)^{-1} + (1-\lambda)(H+uI)^{-1} \end{pmatrix}$$

$$\Rightarrow \lambda(A+uI)^{-1} + (1-\lambda)(H+uI)^{-1} - (\lambda A + (1-\lambda)H + uI)^{-1} \succeq 0.$$

This shows convexity.

(b) Note $\displaystyle\int_0^\infty \left(\frac{1}{1+u} - \frac{1}{a+u}\right) du = \lim_{L\to\infty} \log(1+u)\Big|_0^L - \log(a+u)\Big|_0^L$

$$= \log a + \lim_{L\to\infty} \log \frac{1+L}{a+L} = \log a.$$

If $A \precsim H$, $\log A = \displaystyle\int_0^\infty \left[(1+u)^{-1} I - (A+uI)^{-1}\right] du$

$$\precsim \int_0^\infty \left[(1+u)^{-1} I - (H+uI)^{-1}\right] du = \log H$$

For any $A, H \succeq 0$, $\lambda \in [0,1]$,

$$\lambda \log A + (1-\lambda) \log H$$

$$= \int_0^\infty \left(\lambda \left[(1+u)^{-1} I - (A+uI)^{-1} \right] + (1-\lambda) \left[(1+u)^{-1} I - (H+uI)^{-1} \right] \right) du$$

$$= \int_0^\infty \left((1+u)^{-1} I - \left[\lambda (A+uI)^{-1} + (1-\lambda)(H+uI)^{-1} \right] \right) du$$

$$\preceq \int_0^\infty \left[(1+u)^{-1} I - (\lambda A + (1-\lambda) H + uI)^{-1} \right] du = \log(\lambda A + (1-\lambda) H). \blacksquare$$

Lemma (operator Jensen): Suppose $f: \mathcal{I} \to \mathbb{R}$ is operator convex, $A_1 \in \mathbb{R}^{d_1 \times d_1}$, $A_2 \in \mathbb{R}^{d_2 \times d_2}$ symmetric w/ eigenvalues in $\mathcal{I}$, $K_1 \in \mathbb{R}^{d_1 \times d}$, $K_2 \in \mathbb{R}^{d_2 \times d}$ s.t. $K_1^T K_1 + K_2^T K_2 = I$. Then

$$f(K_1^T A_1 K_1 + K_2^T A_2 K_2) \preceq K_1^T f(A_1) K_1 + K_2^T f(A_2) K_2.$$

Proof: Let $A = \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix}$, $U = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix}$. Note $\begin{pmatrix} K_1 \\ K_2 \end{pmatrix}$ has orthonormal columns, and pick L_1, L_2 s.t. $Q = \begin{pmatrix} K_1 & L_1 \\ K_2 & L_2 \end{pmatrix}$ is orthogonal.

Then
- $Q^T A Q = \begin{pmatrix} K_1^T A_1 K_1 + K_2^T A_2 K_2 & * \\ * & * \end{pmatrix}$
- $\frac{1}{2} \begin{pmatrix} T & B \\ B^T & M \end{pmatrix} + \frac{1}{2} U^T \begin{pmatrix} T & B \\ B^T & M \end{pmatrix} U = \begin{pmatrix} T & 0 \\ 0 & M \end{pmatrix} \quad \forall T, B, M.$

So $f(K_1^T A_1 K_1 + K_2^T A_2 K_2) = f([Q^T A Q]_{11})$

$$= f\left(\left[\tfrac{1}{2} Q^T A Q + \tfrac{1}{2} U^T Q^T A Q U\right]_{11}\right)$$

$$= \left[f\left(\tfrac{1}{2} Q^T A Q + \tfrac{1}{2} U^T Q^T A Q U\right)\right]_{11} \quad \text{b/c block diagonal}$$

By operator convexity,

$$f\left(\tfrac{1}{2} Q^T A Q + \tfrac{1}{2} U^T Q^T A Q U\right) \preceq \tfrac{1}{2} f(Q^T A Q) + \tfrac{1}{2} f(U^T Q^T A Q U)$$

$$= \tfrac{1}{2} Q^T f(A) Q + \tfrac{1}{2} U^T Q^T f(A) Q U \quad \text{b/c } Q, QU \text{ orthogonal}$$

$$\Rightarrow f(K_1^T A_1 K_1 + K_2^T A_2 K_2) \preceq \left[\tfrac{1}{2} Q^T f(A) Q + \tfrac{1}{2} U^T Q^T f(A) Q U\right]_{11}$$

$$= \left[Q^T f(A) Q\right]_{11} = K_1^T f(A_1) K_1 + K_2^T f(A_2) K_2. \quad \blacksquare$$

Lemma: $(A, H) \mapsto D(A \| H) = \mathrm{Tr}\left[A(\log A - \log H) - (A - H)\right]$

is non-negative and convex over $\{(A,H) : A, H \succ 0\}$.

Proof: Let $A \otimes H = \begin{pmatrix} a_{11} H & \cdots & a_{1d} H \\ \vdots & & \\ a_{d1} H & \cdots & a_{dd} H \end{pmatrix} \in \mathbb{R}^{d^2 \times d^2}$ be the Kronecker product.

If $A = \sum_i \lambda_i u_i u_i^T$, $H = \sum_j \mu_j v_j v_j^T$ are the eigendecomps, then

$$A \otimes H = \sum_{ij} \lambda_i \mu_j (u_i \otimes v_j)(u_i \otimes v_j)^T \Rightarrow \log(A \otimes H) = \log A \otimes I + I \otimes \log H.$$

Set $f(x) = x - 1 - \log x$, nonnegative + operator convex on $(0, \infty)$.

Let $\mathcal{C} : \mathbb{R}^{d^2 \times d^2} \to \mathbb{R}$ be given by $\mathcal{C}(B) = \underbrace{\mathrm{vec}(I)}_{\in \mathbb{R}^{d^2}}^{T} B \, \mathrm{vec}(I)$.

This satisfies

- $B \succeq 0 \Rightarrow \mathcal{C}(B) \geq 0$
- $\mathcal{C}(A \otimes H) = \sum_{ij} (A \otimes H)_{(i,i),(j,j)} = \sum_{ij} a_{ij} h_{ij} = \mathrm{Tr}\, AH.$

$$\Rightarrow D(A \| H) = \mathcal{C}\Big(A \log A \otimes I - A \otimes \log H - (A \otimes I) + (I \otimes H)\Big)$$

$$= \mathcal{C}\Big((A \otimes I)\big[A^{-1} \otimes H - I \otimes I - \underbrace{\log(A^{-1} \otimes H)}_{= -\log A \otimes I + I \otimes \log H}\big]\Big)$$

$$= \mathcal{C}\Big((A \otimes I) \cdot f(A^{-1} \otimes H)\Big)$$

Note $(A \otimes I) f(A^{-1} \otimes H) = (A \otimes I)^{1/2} f\big((A \otimes I)^{-1/2} (I \otimes H)(A \otimes I)^{-1/2}\big)(A \otimes I)^{1/2}$

b/c $A \otimes I$, $A^{-1} \otimes H$ have same eigenvectors

$$\succeq 0 \Rightarrow D(A \| H) \geq 0.$$

Consider any $A_1, A_2, H_1, H_2 \succeq 0$, $\lambda \in [0,1]$.

Set $A = \lambda A_1 + (1-\lambda) A_2$, $H = \lambda H_1 + (1-\lambda) H_2$

$$K_1 = \sqrt{\lambda}\, A_1^{1/2} A^{-1/2}, \quad \sqrt{1-\lambda}\, A_2^{1/2} A^{-1/2} \Rightarrow K_1^T K_1 + K_2^T K_2 = I$$

$$\Rightarrow A^{1/2} f(A^{-1/2} H A^{-1/2}) A^{1/2}$$

$$= A^{1/2} f\big(\lambda \cdot A^{-1/2} H_1 A^{-1/2} + (1-\lambda) A^{-1/2} H_2 A^{-1/2}\big) \cdot A^{1/2}$$

$$= A^{1/2} f\big(K_1^T A_1^{-1/2} H_1 A_1^{-1/2} K_1 + K_2^T A_2^{-1/2} H_2 A_2^{-1/2} K_2\big) A^{1/2}$$

(operator Jensen)

$$\succsim A^{1/2}\left[K_1^T f(A_1^{-1/2} H_1 A_1^{-1/2}) K_1 + K_2^T f(A_2^{-1/2} H_2 A^{-1/2}) K_2\right] A^{1/2}$$

$$= \lambda A_1^{1/2} f(A_1^{-1/2} H_1 A_1^{-1/2}) A_1^{1/2} + (1-\lambda) A_2^{1/2} f(A_2^{-1/2} H_2 A_2^{-1/2}) A_2^{1/2}$$

Apply this w/ $A \otimes I$, $I \otimes H$ in place of A, H:

$$(A \otimes I) f(A^{-1} \otimes H) \succsim \lambda (A_1 \otimes I) f(A_1^{-1} \otimes H_1) + (1-\lambda)(A_2 \otimes I) f(A_2^{-1} \otimes H_2)$$

$$\Rightarrow D(A \| H) = \mathcal{l}\left((A \otimes I) f(A^{-1} \otimes H)\right)$$

$$\leq \lambda D(A_1 \| H_1) + (1-\lambda) D(A_2 \| H_2).$$

■

Proof of Lieb's Thm: For any $M, T \succ 0$,

$$D(T \| M) = \operatorname{Tr}\left[T(\log T - \log M) - (T - M)\right] \geq 0$$

$$\Rightarrow \operatorname{Tr} M \geq \operatorname{Tr}\left[T \log M - T \log T + T\right]$$

Equality holds for $T = M$. So

$$\operatorname{Tr} M = \sup_{T \succ 0} \operatorname{Tr}\left[T \log M - T \log T + T\right]$$

$$\Rightarrow \operatorname{Tr} e^{H + \log A} = \sup_{T \succ 0} \operatorname{Tr}\left[T(H + \log A) - T \log T + T\right]$$

$$= \sup_{T \succ 0} \underbrace{\operatorname{Tr} TH + \operatorname{Tr} A - D(T \| A)}_{\text{concave in } (T, A) \text{ for any } H \succ 0}.$$

$\Rightarrow A \mapsto \operatorname{Tr} e^{H + \log A}$ concave.

■

Proof of matrix Bernstein

Let $S = \sum_{i=1}^n X_i$, $\mathbb{E}X_i = 0$, $\lambda_{max}(X_i) \le b$ a.s.

$$\mathbb{P}[\lambda_{max}(S) \ge t] \le e^{-\lambda t} \cdot \mathbb{E} e^{\lambda \cdot \lambda_{max}(S)} \qquad \forall \lambda \ge 0$$

$$\le e^{-\lambda t} \operatorname{Tr} \mathbb{E} e^{\lambda S}$$

$$\le e^{-\lambda t} \operatorname{Tr} e^{\sum_{i=1}^n \log \mathbb{E} e^{\lambda X_i}} \quad \text{by Lieb's Thm.}$$

Write $e^{\lambda X} = I + \lambda X + X f(X) X$, $\quad f(x) = \frac{e^{\lambda x} - x - 1}{x^2}$.

Can check that
- f is increasing
- $f(x) \le \frac{\lambda^2/2}{1 - \lambda x/3} \quad \forall x < \frac{3}{\lambda}$.

$$\Rightarrow e^{\lambda X_i} \preceq I + \lambda X_i + f(b) X_i^2 \preceq I + \lambda X_i + \underbrace{\frac{\lambda^2/2}{1 - \lambda b/3}}_{:= g(\lambda)} X_i^2 \quad \forall \lambda < 3/b.$$

$$\Rightarrow \mathbb{E} e^{\lambda X_i} \preceq I + g(\lambda) \mathbb{E} X_i^2$$

$$\Rightarrow \log \mathbb{E} e^{\lambda X_i} \preceq \log(I + g(\lambda)\mathbb{E}X_i^2) \preceq g(\lambda) \mathbb{E} X_i^2 \quad \text{by operator monotonicity of log}$$

By trace monotonicity of exp,

$$\mathbb{P}[\lambda_{max}(S) \ge t] \le e^{-\lambda t} \operatorname{Tr} e^{g(\lambda) \sum_{i=1}^n \mathbb{E}X_i^2} \le d e^{-\lambda t} e^{g(\lambda) \cdot v},$$

where $v = \| \sum_{i=1}^n \mathbb{E} X_i^2 \|_{op}$. Pick $\lambda = \frac{t}{v + bt/3}$ and substitute into $g(\lambda)$. ∎