CS 58500 – Theoretical Computer Science Toolkit

Lecture 15 (03/31)

Matrix Analysis and Concentration

https://ruizhezhang.com/course_spring_2026.html

Today’s Lecture

- Matrix Analysis
- Matrix Concentration
- Application: Matrix Spectral Sparsification

Matrix Analysis

- The real $m \times n$ matrices, denoted $\mathbb{R}^{m\times n}$, are naturally identified with a vector space of dimension mn
- We define the Frobenius inner product of $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{m\times n}$ by

$$\langle \boldsymbol{A}, \boldsymbol{B} \rangle \coloneqq \sum_{i\in[m], j\in[n]} A_{ij} B_{ij} = \operatorname{tr}[\boldsymbol{A}^\top \boldsymbol{B}]$$

- The induced norm in this vector space is the Frobenius norm $\|\boldsymbol{A}\|_F = \sqrt{\langle \boldsymbol{A}, \boldsymbol{A} \rangle} = \sqrt{\operatorname{tr}[\boldsymbol{A}^\top \boldsymbol{A}]}$
- We are particularly interested in a subspace $\mathbb{S}^{d\times d} \subset \mathbb{R}^{d\times d}$, the real symmetric matrices. The dimension of $\mathbb{S}^{d\times d}$ is $d + \binom{d}{2} = d(d+1)/2$

Matrix Analysis

Theorem (Spectral theorem**).** Let $\boldsymbol{M} \in \mathbb{R}^{d\times d}$. Then $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ if and only if there is unitary $\boldsymbol{U} \in \mathbb{R}^{d\times d}$ ($\boldsymbol{U}^\top = \boldsymbol{U}^{-1}$) and diagonal $\boldsymbol{\Lambda} = \mathbf{Diag}(\boldsymbol{\lambda}) \in \mathbb{R}^{d\times d}$ such that $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^\top$, i.e., the columns of $\boldsymbol{U}$ are the real eigenvectors and $\lambda_1 \geq \cdots \geq \lambda_d$ are the real eigenvalues of $\boldsymbol{M}$

- If $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$ share a set of eigenvectors, i.e., simultaneously diagonalizable, then $\boldsymbol{M}$ and $\boldsymbol{N}$ commute: $[\boldsymbol{M}, \boldsymbol{N}] \coloneqq \boldsymbol{M}\boldsymbol{N} - \boldsymbol{N}\boldsymbol{M} = 0$
- The trace of $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ equals to the sum of its eigenvalues: $\mathrm{tr}[\boldsymbol{M}] = \mathrm{tr}[\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^\top] = \mathrm{tr}[\boldsymbol{U}^\top\boldsymbol{U}\boldsymbol{\Lambda}] = \mathrm{tr}[\boldsymbol{\Lambda}]$

Matrix Analysis: Positive Semi-definiteness

Let $\mathbb{S}_{\succcurlyeq 0}^{d\times d} \subseteq \mathbb{S}^{d\times d}$ denote the set of positive semi-definite matrices:

$$\boldsymbol{M} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d} \quad \Leftrightarrow \quad \lambda_i \geq 0 \;\; \forall i \in [d] \quad \Leftrightarrow \quad \boldsymbol{v}^\top \boldsymbol{M} \boldsymbol{v} \geq 0 \;\; \forall \boldsymbol{v} \in \mathbb{R}^d$$

- $\boldsymbol{M} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ if and only if there exists $\boldsymbol{A} \in \mathbb{R}^{n\times d}$ such that $\boldsymbol{M} = \boldsymbol{A}^\top \boldsymbol{A}$ for some $n > 0$
- We can similarly define the positive definite matrices $\mathbb{S}_{\succ 0}^{d\times d}$

Lemma. If $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$, then we have $\langle \boldsymbol{M}, \boldsymbol{N} \rangle \geq 0$

Proof.

- Let $\boldsymbol{M} = \sum_i \lambda_i \boldsymbol{u}_i \boldsymbol{u}_i^\top$ and $\boldsymbol{N} = \sum_i \gamma_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$ be the eigendecompositions with $\lambda_i, \gamma_i \geq 0$

$$\langle M, N \rangle = \sum_{i,j} \lambda_i \gamma_j \cdot \operatorname{tr}\left[\boldsymbol{u}_i \boldsymbol{u}_i^\top \boldsymbol{v}_j \boldsymbol{v}_j^\top\right] = \sum_{i,j} \lambda_i \gamma_j \cdot \left\langle \boldsymbol{u}_i, \boldsymbol{v}_j \right\rangle^2 \geq 0$$

■

Matrix Analysis: Positive Semi-definiteness

The positive semi-definiteness induces a natural ordering for $\mathbb{S}^{d\times d}$ (called the Löwner order):

$$\boldsymbol{M} \succcurlyeq \boldsymbol{N} \iff \boldsymbol{M}-\boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d} \iff \boldsymbol{v}^\top(\boldsymbol{M}-\boldsymbol{N})\boldsymbol{v} \geq 0 \ \ \forall \boldsymbol{v} \in \mathbb{R}^d$$

- If $\boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ and $\boldsymbol{M} \succcurlyeq \boldsymbol{M}'$, then $\langle \boldsymbol{M}, \boldsymbol{N}\rangle \geq \langle \boldsymbol{M}', \boldsymbol{N}\rangle$
- Conjugation rule: Let $\boldsymbol{M} \succcurlyeq \boldsymbol{N}$ and $\boldsymbol{A}$ be any matrix of the same dimension. Then
$$\boldsymbol{A}\boldsymbol{M}\boldsymbol{A}^\top \succcurlyeq \boldsymbol{A}\boldsymbol{N}\boldsymbol{A}^\top$$
  - For any $\boldsymbol{v} \in \mathbb{R}^d$, $\boldsymbol{v}^\top\boldsymbol{A}(\boldsymbol{M}-\boldsymbol{N})\boldsymbol{A}^\top\boldsymbol{v} = (\boldsymbol{A}^\top\boldsymbol{v})^\top(\boldsymbol{M}-\boldsymbol{N})(\boldsymbol{A}^\top\boldsymbol{v}) \geq 0$
- For $\boldsymbol{M} \in \mathbb{S}_{\succ 0}^{d\times d}$ and $\boldsymbol{N} \in \mathbb{S}^{d\times d}$, $\boldsymbol{M} \succcurlyeq \boldsymbol{N} \iff \boldsymbol{M}^{-1/2}\boldsymbol{N}\boldsymbol{M}^{-1/2} \preccurlyeq \boldsymbol{I}$
- For $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succ 0}^{d\times d}$, $\boldsymbol{M} \succcurlyeq \boldsymbol{N} \iff \boldsymbol{N}^{-1} \succcurlyeq \boldsymbol{M}^{-1}$
- $\boldsymbol{M} \succcurlyeq \boldsymbol{N}$ does NOT imply $\boldsymbol{M}^2 \succcurlyeq \boldsymbol{N}^2$!

Matrix Analysis: Schur complement

Schur complement

- Suppose $\boldsymbol{M}$ is partitioned into 2×2 blocks:

$$\boldsymbol{M} = \begin{bmatrix} \boldsymbol{M}_{11} & \boldsymbol{M}_{12} \\ \boldsymbol{M}_{21} & \boldsymbol{M}_{22} \end{bmatrix}$$

- If $\boldsymbol{M}_{11}$ and $\boldsymbol{M}_{22}$ are invertible, then the Schur complement onto the top-left block is

$$\mathrm{SC}_1(\boldsymbol{M}) \coloneqq \boldsymbol{M}_{11} - \boldsymbol{M}_{12}\boldsymbol{M}_{22}^{-1}\boldsymbol{M}_{21}$$

  and the Schur complement onto the bottom-right block is

$$\mathrm{SC}_2(\boldsymbol{M}) \coloneqq \boldsymbol{M}_{22} - \boldsymbol{M}_{21}\boldsymbol{M}_{11}^{-1}\boldsymbol{M}_{12}$$

- $\boldsymbol{M}$ is positive definite if and only if $\mathrm{SC}_1(\boldsymbol{M})$ and $\mathrm{SC}_2(\boldsymbol{M})$ are positive definte

Matrix Analysis: Schur complement

Why Schur complement?

1. Schur complement gives a condition to check positive definiteness (or positive semi-definiteness) using sub-matrices, which is very useful in analyzing SDP relaxation and Sum-of-Squares hierarchy
   An example can be found in Lecture 10 of my CS 593 *Algorithms for Data Science* (Fall 2025)

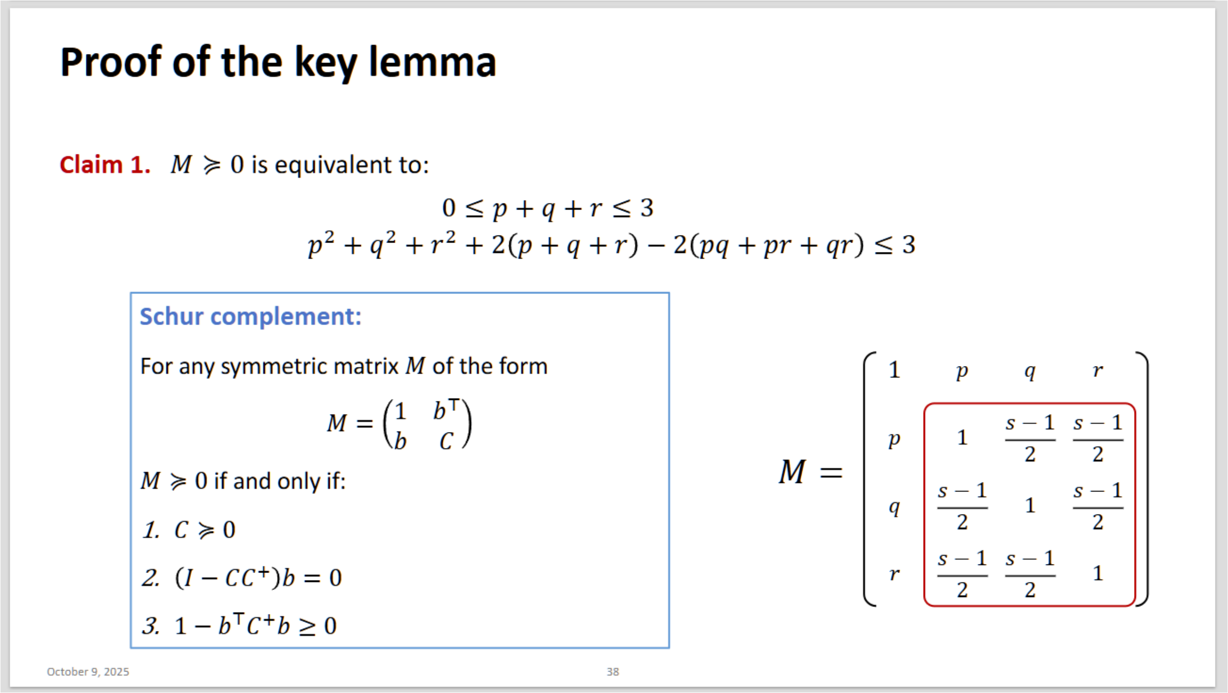

Proof of the key lemma

Claim 1. $M \succcurlyeq 0$ is equivalent to:

$$0 \le p + q + r \le 3$$
$$p^2 + q^2 + r^2 + 2(p + q + r) - 2(pq + pr + qr) \le 3$$

Schur complement:

For any symmetric matrix M of the form

$$M = \begin{pmatrix} 1 & b^\top \\ b & C \end{pmatrix}$$

$M \succcurlyeq 0$ if and only if:

1. $C \succcurlyeq 0$
2. $(I - CC^+)b = 0$
3. $1 - b^\top C^+ b \ge 0$

$$M = \begin{pmatrix} 1 & p & q & r \\ p & 1 & \frac{s-1}{2} & \frac{s-1}{2} \\ q & \frac{s-1}{2} & 1 & \frac{s-1}{2} \\ r & \frac{s-1}{2} & \frac{s-1}{2} & 1 \end{pmatrix}$$

October 9, 2025 38

Matrix Analysis: Schur complement

Why Schur complement?

2. Cholesky decomposition can be computed via iterative Schur complements:

- → We want to compute $\boldsymbol{A} = \boldsymbol{L}\boldsymbol{L}^\top$

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & \boldsymbol{a}_{21}^\top \\ \boldsymbol{a}_{21} & \boldsymbol{A}_{22} \end{bmatrix} = \begin{bmatrix} l_{11} & 0 \\ \boldsymbol{l}_{21} & \boldsymbol{L}_{22} \end{bmatrix} \begin{bmatrix} l_{11} & \boldsymbol{l}_{21}^\top \\ 0 & \boldsymbol{L}_{22}^\top \end{bmatrix}$$

- → We can solve the matrix equation:

$$a_{11} = l_{11}^2, \qquad \boldsymbol{a}_{21} = l_{11}\boldsymbol{l}_{21}, \qquad \boldsymbol{A}_{22} = \boldsymbol{l}_{21}\boldsymbol{l}_{21}^\top + \boldsymbol{L}_{22}\boldsymbol{L}_{22}^\top = \boldsymbol{L}_{22}\boldsymbol{L}_{22}^\top + \boldsymbol{a}_{21}a_{11}^{-1}\boldsymbol{a}_{21}^\top$$

- → Thus, we have $\boldsymbol{L}_{22}\boldsymbol{L}_{22}^\top = \boldsymbol{A}_{22} - \boldsymbol{a}_{21}a_{11}^{-1}\boldsymbol{a}_{21}^\top = \mathrm{SC}_2(\boldsymbol{A})$
- → That is, $\boldsymbol{L}_{22}$ is the Cholesky decomposition of $\mathrm{SC}_2(\boldsymbol{A}) \in \mathbb{R}^{(d-1)\times(d-1)}$
- → Particularly useful when $\boldsymbol{A}$ is the Laplacian of a graph (Kyng-Lee-Peng-Sachdeva-Spielman ’16, Kyng-Sachdeva ’16)

Matrix Analysis: Matrix Functions

For a function $f: \mathbb{R} \to \mathbb{R}$ and $\boldsymbol{M} \in \mathbb{S}^{d\times d}$, we define the matrix function $f(\boldsymbol{M}) = \boldsymbol{U} f(\boldsymbol{\Lambda}) \boldsymbol{U}^\top$

$$f(\boldsymbol{M}) = \boldsymbol{U} \begin{bmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_d) \end{bmatrix} \boldsymbol{U}^\top \quad \text{where} \quad \boldsymbol{M} = \boldsymbol{U} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_d \end{bmatrix} \boldsymbol{U}^\top$$

- Examples: $\exp(\boldsymbol{M})$, $\sin(\boldsymbol{M})$, $(\boldsymbol{M} - x\boldsymbol{I})^{-1}$, etc.
- **Spectral mapping theorem:** Let $f: I \to \mathbb{R}$ with $I \subset \mathbb{R}$ and $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ be a matrix whose eigenvalues are contained in I. Then, if λ is an eigenvalue of $\boldsymbol{M}$, then $f(\lambda)$ is an eigenvalue of $f(\boldsymbol{M})$
- **Transfer rule:** For an $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ have all its eigenvalues lying in $I \subset R$, and for two functions $f, g: I \to \mathbb{R}$ such that $f(x) \leq g(x)$ for all $x \in I$. Then, $f(\boldsymbol{M}) \preccurlyeq g(\boldsymbol{M})$

Matrix Analysis: Eigenvalue Inequalities

Theorem (Cauchy Interlacing Theorem). For any $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ and $\boldsymbol{V} \in \mathbb{R}^{d\times r}$ with $\boldsymbol{V}^\top \boldsymbol{V} = \boldsymbol{I}_r$,

$$\lambda_{d-r+k}(\boldsymbol{M}) \leq \lambda_k(\boldsymbol{V}^\top \boldsymbol{M}\boldsymbol{V}) \leq \lambda_k(\boldsymbol{M}) \quad \forall k \in [r]$$

Corollary. For any $\boldsymbol{M} \in \mathbb{S}^{d\times d}$ and $\boldsymbol{N} \in \mathbb{S}^{r\times r}$ a principal submatrix of $\boldsymbol{M}$,

$$\lambda_{d-r+k}(\boldsymbol{M}) \leq \lambda_k(\boldsymbol{N}) \leq \lambda_k(\boldsymbol{M}) \quad \forall k \in [r]$$

- If $r = d - 1$, then the eigenvalues of $\boldsymbol{N}$ interlace the eigenvalues of $\boldsymbol{M}$

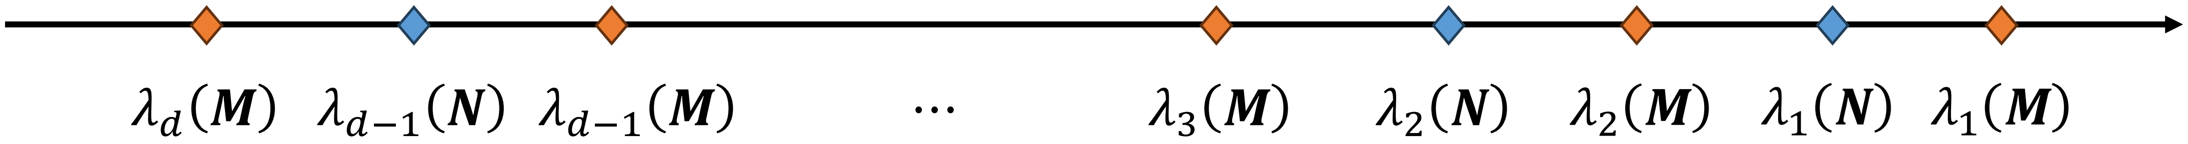

Matrix Analysis: Eigenvalue Inequalities

To prove the Cauchy Interlacing Theorem, we need the variational characterization of eigenvalues:

$$\lambda_k(\boldsymbol{M}) = \max_{\substack{\boldsymbol{U}\in\mathbb{R}^{d\times k} \\ \boldsymbol{U}^\top\boldsymbol{U}=\boldsymbol{I}}} \min_{\substack{\boldsymbol{v}\in\operatorname{span}(\boldsymbol{U}) \\ \|\boldsymbol{v}\|_2=1}} \boldsymbol{v}^\top \boldsymbol{M}\boldsymbol{v} = \min_{\substack{\boldsymbol{U}\in\mathbb{R}^{d\times(d-k+1)} \\ \boldsymbol{U}^\top\boldsymbol{U}=\boldsymbol{I}}} \max_{\substack{\boldsymbol{v}\in\operatorname{span}(\boldsymbol{U}) \\ \|\boldsymbol{v}\|_2=1}} \boldsymbol{v}^\top \boldsymbol{M}\boldsymbol{v}$$

Proof of $\lambda_k(\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V}) \leq \lambda_k(\boldsymbol{M})$:

- Let $\boldsymbol{U} \in \mathbb{R}^{d\times k}$ be the maximizer for $\lambda_k(\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V})$. Then $\lambda_k(\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V}) = \min_{\substack{\boldsymbol{v}\in\operatorname{span}(\boldsymbol{U}) \\ \|\boldsymbol{v}\|_2=1}} \boldsymbol{v}^\top\boldsymbol{M}\boldsymbol{v}$
- We have $\boldsymbol{v} = \boldsymbol{U}\boldsymbol{w}$ for some $\boldsymbol{w} \in \mathbb{R}^k$ and $1 = \|\boldsymbol{v}\|_2^2 = \boldsymbol{w}^\top\boldsymbol{U}^\top\boldsymbol{U}\boldsymbol{w} = \boldsymbol{w}^\top\boldsymbol{w} = \|\boldsymbol{w}\|_2^2$
- Thus, $\lambda_k(\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V}) = \min_{\substack{\boldsymbol{w}\in\mathbb{R}^k \\ \|\boldsymbol{w}\|_2=1}} \boldsymbol{w}^\top\boldsymbol{U}^\top\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V}\boldsymbol{U}\boldsymbol{w} = \min_{\substack{\boldsymbol{w}\in\mathbb{R}^k \\ \|\boldsymbol{w}\|_2=1}} \boldsymbol{w}^\top(\boldsymbol{V}\boldsymbol{U})^\top\boldsymbol{M}(\boldsymbol{V}\boldsymbol{U})\boldsymbol{w}$
- Since $(\boldsymbol{V}\boldsymbol{U})^\top(\boldsymbol{V}\boldsymbol{U}) = \boldsymbol{U}^\top\boldsymbol{V}^\top\boldsymbol{V}\boldsymbol{U} = \boldsymbol{I}$,

$$\lambda_k(\boldsymbol{V}^\top\boldsymbol{M}\boldsymbol{V}) \leq \max_{\substack{\boldsymbol{N}\in\mathbb{R}^{d\times k} \\ \boldsymbol{N}^\top\boldsymbol{N}=\boldsymbol{I}}} \min_{\substack{\boldsymbol{w}\in\mathbb{R}^k \\ \|\boldsymbol{w}\|_2=1}} \boldsymbol{w}^\top \boldsymbol{N}^\top\boldsymbol{M}\boldsymbol{N}\boldsymbol{w} = \max_{\substack{\boldsymbol{N}\in\mathbb{R}^{d\times k} \\ \boldsymbol{N}^\top\boldsymbol{N}=\boldsymbol{I}}} \min_{\substack{\boldsymbol{v}\in\operatorname{span}(\boldsymbol{N}) \\ \|\boldsymbol{v}\|_2=1}} \boldsymbol{v}^\top \boldsymbol{M}\boldsymbol{v} = \lambda_k(\boldsymbol{M})$$

$\blacksquare$

Matrix Analysis: Eigenvalue Inequalities

Theorem (Weyl's inequality). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$.

$$\lambda_k(\boldsymbol{M}) + \lambda_d(\boldsymbol{N}) \le \lambda_k(\boldsymbol{M}+\boldsymbol{N}) \le \lambda_k(\boldsymbol{M}) + \lambda_1(\boldsymbol{N}) \quad \forall k \in [d]$$

- Useful in perturbation analysis: let $\boldsymbol{A} \in \mathbb{S}^{d\times d}$ be the target matrix and $E \in \mathbb{S}^{d\times d}$ be an arbitrary error. Then, $\lambda_{\max}(\boldsymbol{A}+\boldsymbol{E}) \le \lambda_{\max}(\boldsymbol{A}) + \lambda_{\max}(\boldsymbol{E})$ and $\lambda_{\min}(\boldsymbol{A}+\boldsymbol{E}) \ge \lambda_{\min}(\boldsymbol{A}) + \lambda_{\min}(\boldsymbol{E})$

Lemma. Let $f: I \to \mathbb{R}$ be monotone nondecreasing, and let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_I^{d\times d}$. Then if $\boldsymbol{M} \preccurlyeq \boldsymbol{N}$, we have $\operatorname{tr}[f(\boldsymbol{M})] \le \operatorname{tr}[f(\boldsymbol{N})]$

Proof.

- We can use the variational characterization of eigenvalues to show that

$$\boldsymbol{M} \preccurlyeq \boldsymbol{N} \quad \Leftrightarrow \quad \lambda_k(\boldsymbol{M}) \le \lambda_k(\boldsymbol{N}) \quad \forall k \in [d]$$

- Thus, $\operatorname{tr}[f(\boldsymbol{M})] = \sum_{k\in[d]} f\big(\lambda_k(\boldsymbol{M})\big) \le \sum_{k\in[d]} f\big(\lambda_k(\boldsymbol{N})\big) = \operatorname{tr}[f(\boldsymbol{N})]$

■

Matrix Analysis: Eigenvalue Inequalities

Theorem (Weyl's inequality). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$.

$$\lambda_k(\boldsymbol{M}) + \lambda_d(\boldsymbol{N}) \le \lambda_k(\boldsymbol{M}+\boldsymbol{N}) \le \lambda_k(\boldsymbol{M}) + \lambda_1(\boldsymbol{N}) \quad \forall k \in [d]$$

- Useful in perturbation analysis: let $\boldsymbol{A} \in \mathbb{S}^{d\times d}$ be the target matrix and $E \in \mathbb{S}^{d\times d}$ be an arbitrary error. Then, $\lambda_{\max}(\boldsymbol{A}+\boldsymbol{E}) \le \lambda_{\max}(\boldsymbol{A}) + \lambda_{\max}(\boldsymbol{E})$ and $\lambda_{\min}(\boldsymbol{A}+\boldsymbol{E}) \ge \lambda_{\min}(\boldsymbol{A}) + \lambda_{\min}(\boldsymbol{E})$

Lemma. Let $f: I \to \mathbb{R}$ be monotone nondecreasing, and let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_I^{d\times d}$. Then if $\boldsymbol{M} \preccurlyeq \boldsymbol{N}$, we have $\operatorname{tr}[f(\boldsymbol{M})] \le \operatorname{tr}[f(\boldsymbol{N})]$

- It is NOT true in general that $f(\boldsymbol{M}) \preccurlyeq f(\boldsymbol{N})$
- Counterexample: $f(x) = x^2$, $\boldsymbol{M} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$, $\boldsymbol{N} = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$

Matrix Analysis: Operator Inequalities

Let $f: I \to \mathbb{R}$ and $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$ whose eigenvalues are contained in I

- If $f(\boldsymbol{M}) \preccurlyeq f(\boldsymbol{N})$ holds for any $\boldsymbol{M} \preccurlyeq \boldsymbol{N}$, then we say f is operator monotone
- If $f(t\boldsymbol{M} + (1-t)\boldsymbol{N}) \preccurlyeq tf(\boldsymbol{M}) + (1-t)f(\boldsymbol{N})$ for any $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$ and $t \in (0,1)$, then we say f is operator convex

Theorem (Löwner-Heinz). Let $f: \mathbb{R}_{\geq 0} \to \mathbb{R}$.

- If $f(x) = x^p$ for $p \in [-1, 0]$, f is operator convex and $-f$ is operator monotone
- If $f(x) = x^p$ for $p \in [0, 1]$, or $f(x) = \log x$, f is operator monotone and operator concave
- If $f(x) = x^p$ for $p \in [1, 2]$, or $f(x) = x \log x$, f is operator convex

Matrix Analysis: Trace Inequalities

Inequalities are sometimes tight when eigenspaces of matrices are aligned

- Rearrangement inequality: for $x_1 \geq \cdots \geq x_d, y_1 \geq \cdots \geq y_d$, for any permutation σ, it holds that
$$x_1 y_d + \cdots + x_d y_1 \leq x_1 y_{\sigma(1)} + \cdots + x_d y_{\sigma(d)} \leq x_1 y_1 + \cdots x_d y_d$$
- For $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$, if they commute, then $\langle \boldsymbol{M}, \boldsymbol{N} \rangle = \lambda_1(\boldsymbol{M})\lambda'_{\sigma(1)}(\boldsymbol{N}) + \cdots + \lambda_d(\boldsymbol{M})\lambda'_{\sigma(d)}(\boldsymbol{N})$ and
$$\lambda_1\lambda'_d + \cdots + \lambda_d\lambda'_1 \leq \langle \boldsymbol{M}, \boldsymbol{N} \rangle \leq \lambda_1\lambda'_1 + \cdots + \lambda_d\lambda'_d$$

Theorem (von Neumann trace inequality). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$ have eigendecompositions $\boldsymbol{M} = \boldsymbol{U} \cdot \mathbf{Diag}(\boldsymbol{\lambda}) \cdot \boldsymbol{U}^\top$ and $\boldsymbol{N} = \boldsymbol{V} \cdot \mathbf{Diag}(\boldsymbol{\lambda}') \cdot \boldsymbol{V}^\top$, and we assume that the $\{\lambda_i\}_{i\in[d]}$ and $\{\lambda'_i\}_{i\in[d]}$ are sorted in nonincreasing order. Then

$$\sum_{i\in[d]} \lambda_i \lambda'_{d+1-i} \leq \langle \boldsymbol{M}, \boldsymbol{N} \rangle \leq \sum_{i\in[d]} \lambda_i \lambda'_i$$

The upper and lower bounds respectively hold if $\boldsymbol{U} = \boldsymbol{V}$ and if columns of $\boldsymbol{U}, \boldsymbol{V}$ are reversed

Matrix Analysis: Trace Inequalities

- For any unitary-invariant norm $\|\cdot\|_{\mathrm{UI}}$ (i.e., $\|\boldsymbol{M}\|_{\mathrm{UI}} = \|\boldsymbol{U}\boldsymbol{M}\boldsymbol{V}^\top\|_{\mathrm{UI}}$ for any unitaries $\boldsymbol{U} \in \mathbb{R}^{m\times m}, \boldsymbol{V} \in \mathbb{R}^{n\times n}$), it must be a function only on the eigenvalues (or singular values) of the matrix

- von Neumann trace inequality shows that the maximizing arguments of $\langle \boldsymbol{M}, \boldsymbol{N}\rangle$ align eigenspaces. Thus, we can apply the scalar Cauchy-Schwarz to obtain a matrix Cauchy-Schwarz inequality:
$$\langle \boldsymbol{M}, \boldsymbol{N}\rangle \le \|\boldsymbol{M}\|_{\mathrm{UI}} \cdot \|\boldsymbol{N}\|_*$$
where $\|\boldsymbol{N}\|_* = \sup\{\langle \boldsymbol{Z}, \boldsymbol{N}\rangle : \|\boldsymbol{Z}\|_{\mathrm{UI}} \le 1\}$ is the dual norm of $\|\cdot\|_{\mathrm{UI}}$

- The Schatten-p norm $\|\cdot\|_p$ is unitary-invariant, and its dual norm is $\|\cdot\|_q$ where $p^{-1} + q^{-1} = 1$. And we obtain a matrix Hölder's inequality:
$$\langle \boldsymbol{M}, \boldsymbol{N}\rangle \le \|\boldsymbol{M}\|_p \cdot \|\boldsymbol{N}\|_q$$

Matrix Analysis: Trace Inequalities

Disentangling products of matrices into portions which sort the same matrix with itself

- Consider $\mathrm{tr}[\boldsymbol{MNMN}]$ for $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$
- Applying matrix Cauchy-Schwarz inequality with the self-dual norm $\|\cdot\|_F$ gives:

$$\mathrm{tr}[\boldsymbol{MNMN}] = \langle \boldsymbol{NM}, \boldsymbol{MN}\rangle \leq \|\boldsymbol{NM}\|_F \|\boldsymbol{MN}\|_F = \mathrm{tr}[\boldsymbol{MNNM}]^{1/2}\mathrm{tr}[\boldsymbol{NMMN}]^{1/2} = \mathrm{tr}[\boldsymbol{M}^2\boldsymbol{N}^2]$$

Lemma (Extended Lieb-Thirring Inequality, Allen-Zhu-Lee-Orecchia ’16**).** Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ and $\alpha \in (0, 1)$. Then

$$\mathrm{tr}[\boldsymbol{M}^{\alpha}\boldsymbol{N}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}] \leq \mathrm{tr}[\boldsymbol{M}\boldsymbol{N}^2]$$

Matrix Analysis: Trace Inequalities

Lemma (Extended Lieb-Thirring Inequality, Allen-Zhu-Lee-Orecchia '16**).** Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ and $\alpha \in (0, 1)$. Then

$$\operatorname{tr}[\boldsymbol{M}^{\alpha}\boldsymbol{N}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}] \le \operatorname{tr}[\boldsymbol{M}\boldsymbol{N}^2]$$

Proof sketch:

- Let $f(\alpha) \coloneqq \operatorname{tr}[\boldsymbol{M}^{\alpha}\boldsymbol{N}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}]$. We can see that f is symmetric about $1/2$ in $[0,1]$
- We claim that $f(\alpha) \le \big(f(0) + f(2\alpha)\big)/2$ for any $\alpha \in [0,1/2]$ (You should think why it implies the result)
- To prove this claim, we construct two 2×2 block matrices $\boldsymbol{A}, \boldsymbol{B}$:

$$\boldsymbol{A} \coloneqq \begin{bmatrix} \boldsymbol{N} & -\boldsymbol{N}^{1/2}\boldsymbol{M}^{\alpha}\boldsymbol{N}^{1/2} \\ -\boldsymbol{N}^{1/2}\boldsymbol{M}^{\alpha}\boldsymbol{N}^{1/2} & \boldsymbol{N}^{1/2}\boldsymbol{M}^{2\alpha}\boldsymbol{N}^{1/2} \end{bmatrix} = \begin{bmatrix} \boldsymbol{N}^{1/2} \\ -\boldsymbol{N}^{1/2}\boldsymbol{M}^{\alpha} \end{bmatrix} [\boldsymbol{N}^{1/2} \quad -\boldsymbol{M}^{\alpha}\boldsymbol{N}^{1/2}] \succcurlyeq 0$$

$$\boldsymbol{B} \coloneqq \begin{bmatrix} \boldsymbol{N}^{1/2}\boldsymbol{M}\boldsymbol{N}^{1/2} & \boldsymbol{N}^{1/2}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}^{1/2} \\ \boldsymbol{N}^{1/2}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}^{1/2} & \boldsymbol{N}^{1/2}\boldsymbol{M}^{1-2\alpha}\boldsymbol{N}^{1/2} \end{bmatrix} = \begin{bmatrix} \boldsymbol{N}^{1/2}\boldsymbol{M}^{1/2} \\ \boldsymbol{N}^{1/2}\boldsymbol{M}^{1/2-\alpha} \end{bmatrix} [\boldsymbol{M}^{1/2}\boldsymbol{N}^{1/2} \quad \boldsymbol{M}^{1/2-\alpha}\boldsymbol{N}^{1/2}] \succcurlyeq 0$$

- Thus, $\langle \boldsymbol{A}, \boldsymbol{B} \rangle \ge 0$, which implies that

$$\operatorname{tr}[\boldsymbol{A}^{\top}\boldsymbol{B}] = \operatorname{tr}\begin{bmatrix} ? & \\ & ? \end{bmatrix} = f(0) - 2f(\alpha) + f(2\alpha) \ge 0 \qquad \blacksquare$$

Matrix Analysis: Trace Inequalities

Lemma (Extended Lieb-Thirring Inequality, Allen-Zhu-Lee-Orecchia ’16). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ and $\alpha \in (0, 1)$. Then

$$\operatorname{tr}[\boldsymbol{M}^{\alpha}\boldsymbol{N}\boldsymbol{M}^{1-\alpha}\boldsymbol{N}] \le \operatorname{tr}[\boldsymbol{M}\boldsymbol{N}^2]$$

Theorem (Lieb-Thirring Inequality). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succcurlyeq 0}^{d\times d}$ and $p \ge 1$. Then

$$\operatorname{tr}[(\boldsymbol{M}\boldsymbol{N})^p] \le \operatorname{tr}[\boldsymbol{M}^p\boldsymbol{N}^p]$$

By using limiting argument, it implies many powerful conclusions, such the Golden-Thompson inequality:

Theorem (Golden-Thompson Inequality). Let $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}^{d\times d}$. Then

$$\operatorname{tr}[\exp(\boldsymbol{M}+\boldsymbol{N})] \le \operatorname{tr}[\exp(\boldsymbol{M})\exp(\boldsymbol{N})]$$

- Note that $\exp(\boldsymbol{M}+\boldsymbol{N}) \ne \exp(\boldsymbol{M})\exp(\boldsymbol{N})$ unless $\boldsymbol{M}, \boldsymbol{N}$ commute

Today’s Lecture

- Matrix Analysis
- Matrix Concentration
- Application: Matrix Spectral Sparsification

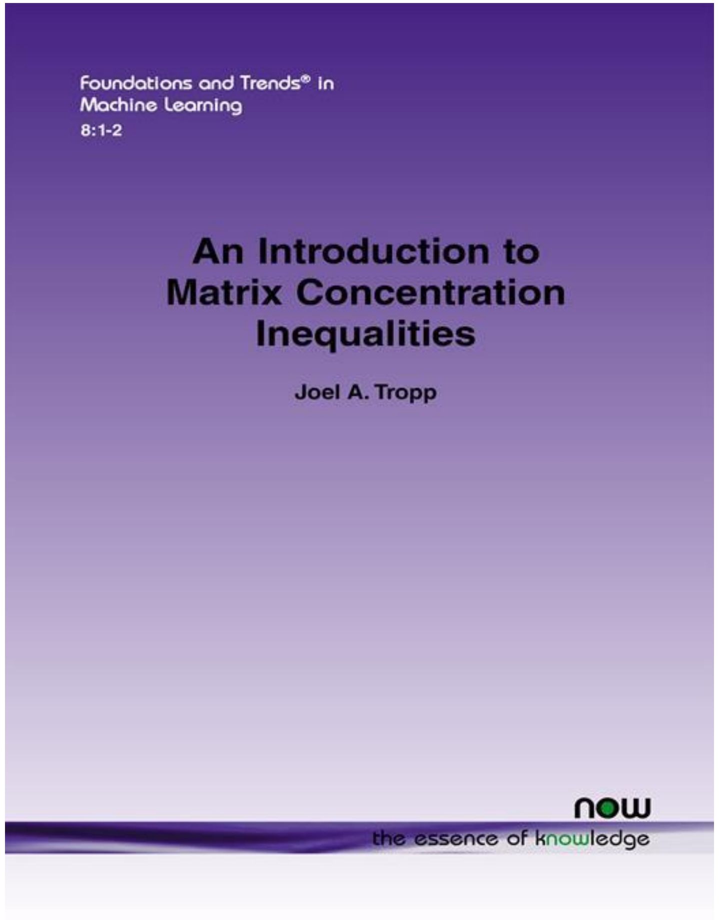
Foundations and Trends® in Machine Learning 8:1-2

An Introduction to Matrix Concentration Inequalities

Joel A. Tropp

now the essence of knowledge

Matrix Concentration

Theorem (Matrix Bernstein's inequality). Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices satisfying $\mathbb{E}[\boldsymbol{Z}_i] = \boldsymbol{0}_d$ and $\|\boldsymbol{Z}_i\| \le B$. Let $\boldsymbol{Z} := \sum_{i\in[n]} \boldsymbol{Z}_i$ and define the variance proxy $\nu := \left\|\sum_{i\in[n]} \mathbb{E}\left[\boldsymbol{Z}_i^2\right]\right\|$. Then,

$$\Pr[\|\boldsymbol{Z}\| \ge t] \le 2d \exp\left(-\frac{t^2/2}{\nu + Bt/3}\right) \qquad \forall t \ge 0$$

Theorem (Matrix Chernoff bound). Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices satisfying $\|\boldsymbol{Z}_i\| \le R$ with probability 1, and let $\boldsymbol{Z} := \sum_{i\in[n]} \boldsymbol{Z}_i$. Then,

$$\Pr[\lambda_1(\boldsymbol{Z}) \ge (1+\epsilon)\mu_{\max}] \le d \exp\left(-\frac{\epsilon^2 \mu_{\max}}{(2+\epsilon)R}\right) \quad \forall \epsilon \ge 0 \quad \text{where} \quad \mu_{\max} := \lambda_1(\mathbb{E}[\boldsymbol{Z}])$$

$$\Pr[\lambda_d(\boldsymbol{Z}) \le (1-\epsilon)\mu_{\min}] \le d \exp\left(-\frac{\epsilon^2 \mu_{\min}}{2R}\right) \quad \forall \epsilon \in (0,1) \quad \text{where} \quad \mu_{\min} := \lambda_d(\mathbb{E}[\boldsymbol{Z}])$$

Matrix Concentration: Matrix MGF

Lemma. Let $Z \in \mathbb{S}^{d\times d}$ be a random matrix. Then,

$$\Pr[\lambda_1(\boldsymbol{Z}) \geq t] \leq \inf_{\theta>0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

$$\Pr[\lambda_d(\boldsymbol{Z}) \leq t] \leq \inf_{\theta<0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

Proof.

- We only prove the upper tail:

$$\begin{aligned} \Pr[\lambda_1(\boldsymbol{Z}) \geq t] &= \Pr\big[\exp\big(\theta\lambda_1(\boldsymbol{Z})\big) \geq \exp(\theta t)\big] \\ &\leq \exp(-\theta t)\, \mathbb{E}\big[\exp\big(\theta\lambda_1(\boldsymbol{Z})\big)\big] \\ &= \exp(-\theta t)\, \mathbb{E}[\lambda_1(\exp(\theta \boldsymbol{Z}))] \\ &\leq \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big] \end{aligned}$$

■

Matrix Concentration: Matrix MGF

Lemma. Let $Z \in \mathbb{S}^{d \times d}$ be a random matrix. Then,

$$\Pr[\lambda_1(\boldsymbol{Z}) \geq t] \leq \inf_{\theta > 0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

$$\Pr[\lambda_d(\boldsymbol{Z}) \leq t] \leq \inf_{\theta < 0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

- If we follow the proof of Chernoff bound, the next step is to bound the matrix MGF:

$$\mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big] = \mathbb{E}\left[\mathrm{tr}\left[\exp\left(\theta \sum\nolimits_{i \in [n]} \boldsymbol{Z}_i\right)\right]\right]$$

- The key property for scalar MGF is the decomposition:

$$\mathbb{E}\left[\exp\left(\theta \sum\nolimits_{i \in [n]} X_i\right)\right] = \prod_{i \in [n]} \mathbb{E}[\exp(\theta X_i)]$$

- However, matrix MGF does not have such decomposition!

Matrix Concentration: Matrix MGF

Lemma. Let $Z \in \mathbb{S}^{d\times d}$ be a random matrix. Then,

$$\Pr[\lambda_1(\boldsymbol{Z}) \geq t] \leq \inf_{\theta>0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

$$\Pr[\lambda_d(\boldsymbol{Z}) \leq t] \leq \inf_{\theta<0} \exp(-\theta t)\, \mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big]$$

- If we follow the proof of Chernoff bound, the next step is to bound the matrix MGF:

$$\mathbb{E}\big[\mathrm{tr}[\exp(\theta \boldsymbol{Z})]\big] = \mathbb{E}\left[\mathrm{tr}\left[\exp\left(\theta \sum\nolimits_{i\in[n]} \boldsymbol{Z}_i\right)\right]\right]$$

- By Golden-Thompson inequality, $\mathrm{tr}[\exp(\boldsymbol{M}+\boldsymbol{N})] \leq \mathrm{tr}[\exp(\boldsymbol{M})\exp(\boldsymbol{N})]$. We may hope a weaker decomposition:

$$\mathrm{tr}\left[\exp\left(\theta \sum\nolimits_{i\in[n]} \boldsymbol{Z}_i\right)\right] \leq \mathrm{tr}\left[\prod_{i\in[n]} \exp(\theta \boldsymbol{Z}_i)\right]$$

- Unfortunately, this is still false (if $n \geq 3$)!

Matrix Concentration: Matrix MGF

The matrix Chernoff bound was first proved in a quantum information paper:

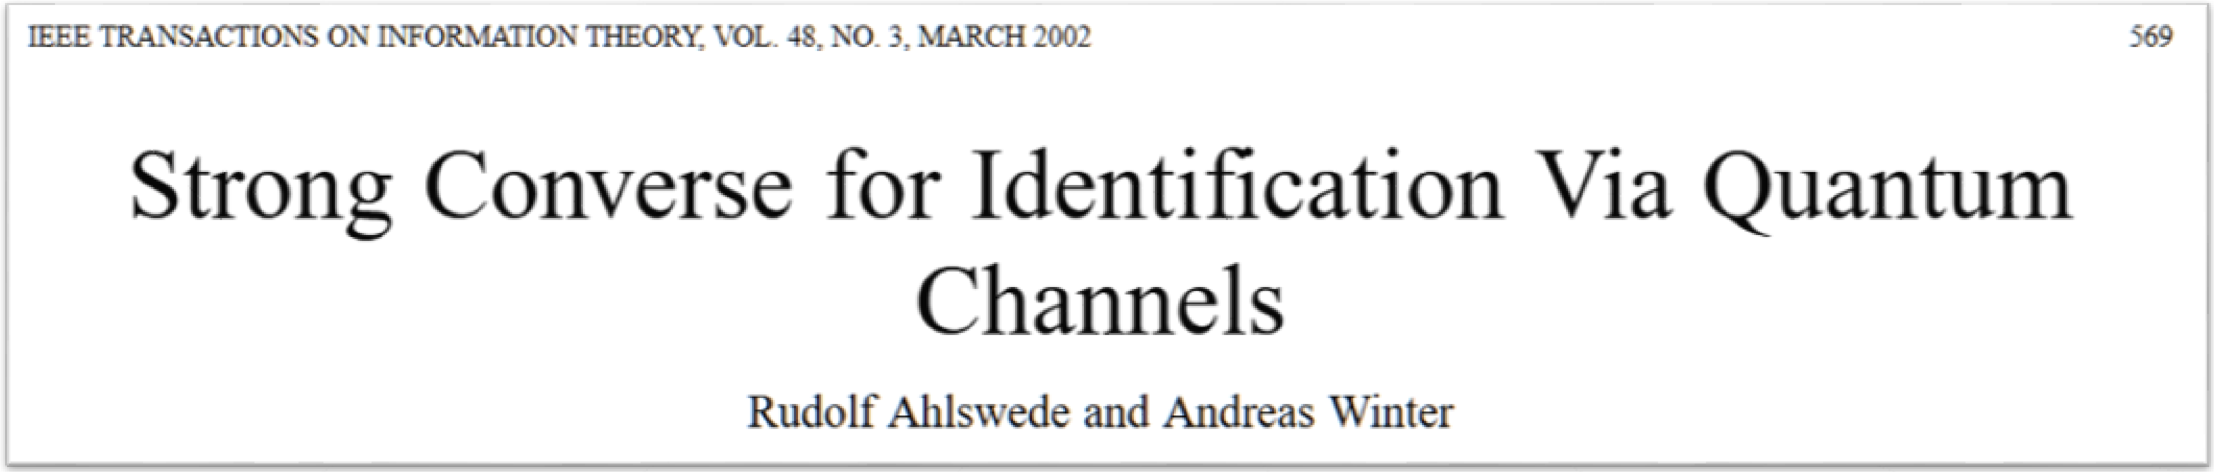

IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. 48, NO. 3, MARCH 2002 569

Strong Converse for Identification Via Quantum Channels

Rudolf Ahlswede and Andreas Winter

For $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succ 0}^{d\times d}$, define the quantum KL divergence:

$$D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) := \operatorname{tr}[\boldsymbol{N}(\log \boldsymbol{N} - \log \boldsymbol{M}) + \boldsymbol{M} - \boldsymbol{N}]$$

Theorem. $D_{\mathcal{H}} : \mathbb{S}_{\succ 0}^{d\times d} \times \mathbb{S}_{\succ 0}^{d\times d} \to \mathbb{R}$ is a jointly convex function of its inputs:

$$D_{\mathcal{H}}\big(\lambda \boldsymbol{N} + (1-\lambda)\boldsymbol{N}' \big\| \lambda \boldsymbol{M} + (1-\lambda)\boldsymbol{M}'\big) \le \lambda D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) + (1-\lambda) D_{\mathcal{H}}\big(\boldsymbol{N}' \big\| \boldsymbol{M}'\big)$$

Matrix Concentration: Matrix MGF

For $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succ 0}^{d \times d}$, define the quantum KL divergence:

$$D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) \coloneqq \operatorname{tr}[\boldsymbol{N}(\log \boldsymbol{N} - \log \boldsymbol{M}) + \boldsymbol{M} - \boldsymbol{N}]$$

Theorem. $D_{\mathcal{H}} : \mathbb{S}_{\succ 0}^{d \times d} \times \mathbb{S}_{\succ 0}^{d \times d} \to \mathbb{R}$ is a jointly convex function of its inputs:

$$D_{\mathcal{H}}\left(\lambda \boldsymbol{N} + (1-\lambda)\boldsymbol{N}' \middle\| \lambda \boldsymbol{M} + (1-\lambda)\boldsymbol{M}'\right) \le \lambda D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) + (1-\lambda) D_{\mathcal{H}}\left(\boldsymbol{N}' \middle\| \boldsymbol{M}'\right)$$

- $D_{\mathcal{H}}$ is the Bregman divergence with respect to the convex function $\mathcal{H}(\boldsymbol{M}) = \operatorname{tr}[\boldsymbol{M} \log \boldsymbol{M}]$ (its convexity follows from Löwner-Heinz)

Theorem (Lieb's concavity theorem). For any $\boldsymbol{S} \in \mathbb{S}^{d \times d}$, the following is concave on $\mathbb{S}_{\succ 0}^{d \times d}$:

$$f(\boldsymbol{X}) \coloneqq \operatorname{tr}[\exp(\boldsymbol{S} + \log \boldsymbol{X})]$$

Matrix Concentration: Matrix MGF

For $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{\succ 0}^{d\times d}$, define the quantum KL divergence:

$$D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) \coloneqq \operatorname{tr}[\boldsymbol{N}(\log \boldsymbol{N} - \log \boldsymbol{M}) + \boldsymbol{M} - \boldsymbol{N}]$$

Theorem. $D_{\mathcal{H}} \colon \mathbb{S}_{\succ 0}^{d\times d} \times \mathbb{S}_{\succ 0}^{d\times d} \to \mathbb{R}$ is a jointly convex function of its inputs:

$$D_{\mathcal{H}}\left(\lambda \boldsymbol{N} + (1-\lambda)\boldsymbol{N}' \middle\| \lambda \boldsymbol{M} + (1-\lambda)\boldsymbol{M}'\right) \le \lambda D_{\mathcal{H}}(\boldsymbol{N} \| \boldsymbol{M}) + (1-\lambda) D_{\mathcal{H}}\left(\boldsymbol{N}' \middle\| \boldsymbol{M}'\right)$$

Theorem (Lieb's concavity theorem). For any $\boldsymbol{S} \in \mathbb{S}^{d\times d}$, the following is concave on $\mathbb{S}_{\succ 0}^{d\times d}$:

$$f(\boldsymbol{X}) \coloneqq \operatorname{tr}[\exp(\boldsymbol{S} + \log \boldsymbol{X})]$$

Corollary (Subadditivity of matrix MGF). Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices and $\theta \in \mathbb{R}$. Then

$$\mathbb{E}\left[\operatorname{tr}\left[\exp\left(\theta \sum_{i\in[n]} Z_i\right)\right]\right] \le \operatorname{tr}\left[\exp\left(\sum_{i\in[n]} \log \mathbb{E}[\exp(\theta \boldsymbol{Z}_i)]\right)\right]$$

Matrix Concentration: Matrix MGF

Corollary (Subadditivity of matrix MGF). Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices and $\theta \in \mathbb{R}$. Then

$$\mathbb{E}\left[\operatorname{tr}\left[\exp\left(\theta \sum\nolimits_{i\in[n]} Z_i\right)\right]\right] \le \operatorname{tr}\left[\exp\left(\sum\nolimits_{i\in[n]} \log \mathbb{E}[\exp(\theta \boldsymbol{Z}_i)]\right)\right]$$

Proof.

- Let $\boldsymbol{S}_k \coloneqq \sum_{i\in[k]} \boldsymbol{Z}_i$ for all $k \in [n]$

$$\mathbb{E}[\operatorname{tr}[\exp(\theta\boldsymbol{S}_n)]] = \mathbb{E}\left[\mathbb{E}[\operatorname{tr}[\exp(\theta(\boldsymbol{S}_{n-1} + \boldsymbol{Z}_n))]|\boldsymbol{S}_{n-1}]\right] = \mathbb{E}[\mathbb{E}[\operatorname{tr}[\exp(\theta\boldsymbol{S}_{n-1} + \log\exp(\theta\boldsymbol{Z}_n))]|\boldsymbol{S}_{n-1}]]$$

$$\text{(Jensen)} \le \mathbb{E}[\operatorname{tr}[\exp(\theta\boldsymbol{S}_{n-1} + \log\mathbb{E}[\exp(\theta\boldsymbol{Z}_n)])]]$$

$$\le \mathbb{E}[\operatorname{tr}[\exp(\theta\boldsymbol{S}_{n-2} + \log\mathbb{E}[\exp(\theta\boldsymbol{Z}_{n-1})] + \log\mathbb{E}[\exp(\theta\boldsymbol{Z}_n)])]]$$

$$\le \mathbb{E}[\operatorname{tr}[\exp(\log\mathbb{E}[\exp(\theta\boldsymbol{Z}_1)] + \cdots + \log\mathbb{E}[\exp(\theta\boldsymbol{Z}_n)])]]$$

■

Matrix Concentration: Matrix MGF

Corollary (Subadditivity of matrix MGF). Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices and $\theta \in \mathbb{R}$. Then

$$\mathbb{E}\left[\operatorname{tr}\left[\exp\left(\theta \sum_{i\in[n]} Z_i\right)\right]\right] \leq \operatorname{tr}\left[\exp\left(\sum_{i\in[n]} \log \mathbb{E}[\exp(\theta \boldsymbol{Z}_i)]\right)\right]$$

- Since $\exp(x)$ is monotone nondecreasing, we have

$$\operatorname{tr}[\exp(\boldsymbol{M})] \leq \operatorname{tr}[\exp(\boldsymbol{N})] \quad \forall \boldsymbol{M} \preccurlyeq \boldsymbol{N}$$

- Hence, we just need to upper bound $\log \mathbb{E}[\exp(\theta \boldsymbol{Z}_i)]$ for each $i \in [n]$

Matrix Concentration: Matrix Bernstein

Lemma. Let $\boldsymbol{Z} \in \mathbb{S}^{d\times d}$ be a random matrix satisfying $\mathbb{E}[\boldsymbol{Z}] = \boldsymbol{0}_d$ and $\|\boldsymbol{Z}\| \leq B$ with probability 1. Then, for all $0 < \theta < \frac{3}{B}$, we have $\log \mathbb{E}[\exp(\theta \boldsymbol{Z})] \preccurlyeq \frac{\theta^2/2}{1-B\theta/3} \mathbb{E}\left[\boldsymbol{Z}^2\right]$

Proof.

- We can Taylor expand the matrix function $\exp(\theta\boldsymbol{Z})$:

$$\begin{aligned}\mathbb{E}[\exp(\theta\boldsymbol{Z})] &= \boldsymbol{I} + \theta\mathbb{E}[\boldsymbol{Z}] + \sum_{p\geq 2} \frac{\theta^p}{p!}\mathbb{E}[\boldsymbol{Z}^p] = \boldsymbol{I} + \sum_{p\geq 2} \frac{\theta^p}{p!}\mathbb{E}[\boldsymbol{Z}^p] \\ &\preccurlyeq \boldsymbol{I} + \sum_{p\geq 2} \frac{\theta^p B^{p-2}}{p!}\mathbb{E}[\boldsymbol{Z}^2] \preccurlyeq \boldsymbol{I} + \frac{\theta^2}{2}\mathbb{E}[\boldsymbol{Z}^2]\sum_{p\geq 0}\frac{\theta^p B^p}{3^p} \\ &= \boldsymbol{I} + \frac{\theta^2/2}{1-B\theta/3}\mathbb{E}[\boldsymbol{Z}^2]\end{aligned}$$

- Since $\log x$ is operator monotone and $\log(1+x) \leq x$, we have

$$\log\mathbb{E}[\exp(\theta\boldsymbol{Z})] \preccurlyeq \log\left(\boldsymbol{I} + \frac{\theta^2/2}{1-B\theta/3}\mathbb{E}[\boldsymbol{Z}^2]\right) \preccurlyeq \frac{\theta^2/2}{1-B\theta/3}\mathbb{E}[\boldsymbol{Z}^2]$$

■

Matrix Concentration: Matrix Bernstein

Theorem (Matrix Bernstein's inequality**).** Let $\{\boldsymbol{Z}_i\}_{i\in[n]} \in \mathbb{S}^{d\times d}$ be independent random matrices satisfying $\mathbb{E}[\boldsymbol{Z}_i] = \boldsymbol{0}_d$ and $\|\boldsymbol{Z}_i\| \le B$. Let $\boldsymbol{Z} := \sum_{i\in[n]} \boldsymbol{Z}_i$ and define the variance proxy $\nu := \left\|\sum_{i\in[n]} \mathbb{E}\left[\boldsymbol{Z}_i^2\right]\right\|$. Then,

$$\Pr[\|\boldsymbol{Z}\| \ge t] \le 2d \exp\left(-\frac{t^2/2}{\nu + Bt/3}\right) \qquad \forall t \ge 0$$

- Based on the MGF bound, the remaining proof is similar to the scalar case.
- It also implies an expectation bound: $\mathbb{E}[\|\boldsymbol{Z}\|] \le \sqrt{2\nu\log(d)} + \frac{1}{3}B\log(d)$
- This result can be generalized to rectangular matrices via Hermitian dilation: for $\boldsymbol{C} \in \mathbb{R}^{d_1\times d_2}$,

$$\mathfrak{H}(\boldsymbol{C}) := \begin{bmatrix} 0 & \boldsymbol{C} \\ \boldsymbol{C}^\top & 0 \end{bmatrix} \in \mathbb{S}^{(d_1+d_2)\times(d_1+d_2)}$$

- We have $\mathfrak{H}(\boldsymbol{C})^2 = \begin{bmatrix} \boldsymbol{C}\boldsymbol{C}^\top & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{C}^\top\boldsymbol{C} \end{bmatrix}$, which implies that $\lambda_1(\mathfrak{H}(\boldsymbol{C})^2) = \|\boldsymbol{C}\|$

Today’s Lecture

- Matrix Analysis
- Matrix Concentration
- Application: Matrix Spectral Sparsification

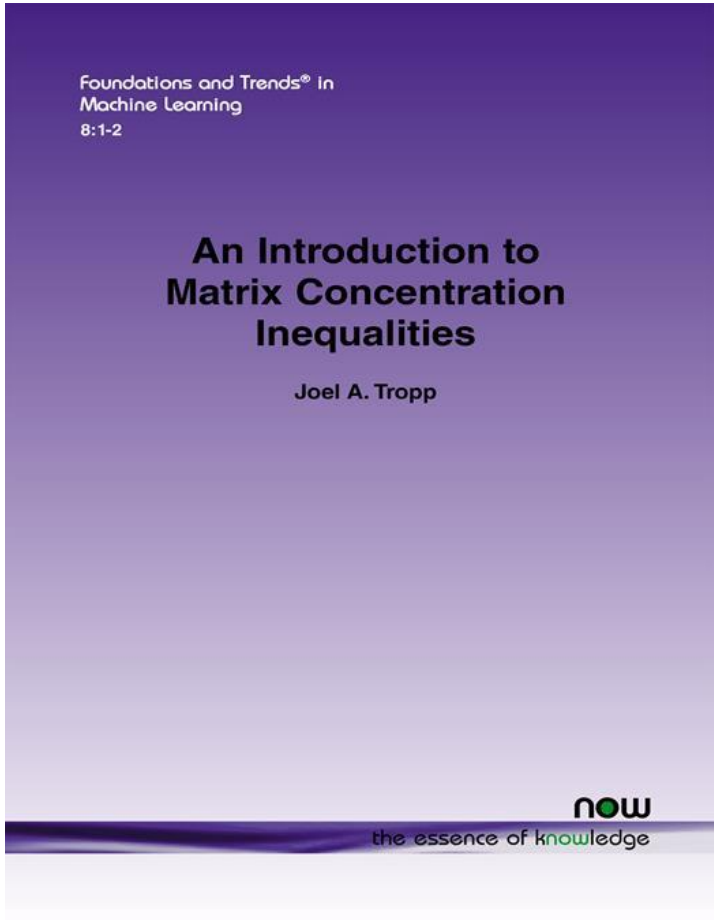

Foundations and Trends® in Machine Learning
8:1-2

An Introduction to Matrix Concentration Inequalities

Joel A. Tropp

now
the essence of knowledge

Randomized Sparsification of a Matrix

Matrix approximation by random sampling

- Let $\boldsymbol{B}$ be a target matrix represented as a sum of "simple" components:
$$\boldsymbol{B} = \boldsymbol{B}_1 + \boldsymbol{B}_2 + \cdots + \boldsymbol{B}_N$$
- We also need a sampling distribution $\boldsymbol{p} \in \mathbb{R}^N$
- We can construct an unbiased estimator of $\boldsymbol{B}$ by:
$$\boldsymbol{R} \coloneqq \frac{1}{p_i} \boldsymbol{B}_i \quad \text{with probability } \ p_i$$
- To improve the quality of approximation, we average n independent copies of $\boldsymbol{R}$:
$$\overline{\boldsymbol{R}}_n = \frac{1}{n} \sum\nolimits_{k=1}^{n} \boldsymbol{R}_k \quad \text{where each } \boldsymbol{R}_k \text{ is an independent copy of } \boldsymbol{R}$$
- If $n \ll N$ while $\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right]$ is small, then we obtain a good simple approximation of $\boldsymbol{B}$

Randomized Sparsification of a Matrix

Theorem (Matrix approximation by random sampling**).** Let $\boldsymbol{B} \in \mathbb{S}^{d\times d}$ be a fixed matrix. Construct a random matrix $\boldsymbol{R} \in \mathbb{S}^{d\times d}$ that satisfies $\mathbb{E}[\boldsymbol{R}] = \boldsymbol{B}$, $\|\boldsymbol{R}\| \le L$, and $m_2(\boldsymbol{R}) := \|\mathbb{E}[\boldsymbol{R}^2]\|$. Then, for the matrix sampling estimator $\overline{\boldsymbol{R}}_n = \sum_{k=1}^{n} \boldsymbol{R}_k$, we have

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right] \le \sqrt{\frac{2m_2(\boldsymbol{R})\log(d)}{n}} + \frac{2L\log(d)}{3n}$$

And for all $t \ge 0$,

$$\Pr\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\| \ge t\right] \le 2d\exp\left(\frac{-nt^2/2}{m_2(R) + 2Lt/3}\right)$$

Randomized Sparsification of a Matrix

Proof sketch.

- We'll apply matrix Bernstein's inequality to the centered, symmetric random matrix $\boldsymbol{Z}$:

$$\boldsymbol{Z} \coloneqq \overline{\boldsymbol{R}}_n - \boldsymbol{B} = \frac{1}{n}\sum_{k=1}^{n}(\boldsymbol{R}_k - \mathbb{E}[\boldsymbol{R}_k]) =: \sum_{k=1}^{n}\boldsymbol{Z}_k$$

- $\|\boldsymbol{Z}_k\| \le (1/n)(\|\boldsymbol{R}_k\| + \|\mathbb{E}[\boldsymbol{R}_k]\|) \le (1/n)(\|\boldsymbol{R}_k\| + \mathbb{E}[\|\boldsymbol{R}_k\|]) \le 2L/n$

- Variance proxy $\nu \coloneqq \left\|\sum_{i\in[n]}\mathbb{E}\left[\boldsymbol{Z}_i^2\right]\right\| = n\left\|\mathbb{E}[\boldsymbol{Z}_1^2]\right\|$

$$\boldsymbol{0} \preccurlyeq \mathbb{E}[\boldsymbol{Z}_1^2] = n^{-2}\mathbb{E}[(\boldsymbol{R} - \mathbb{E}[\boldsymbol{R}])^2] = n^{-2}(\mathbb{E}[\boldsymbol{R}^2] - \mathbb{E}[\boldsymbol{R}]^2) \preccurlyeq n^{-2}\mathbb{E}[\boldsymbol{R}^2]$$

- Thus, $\nu \le n \cdot n^{-2} m_2(\boldsymbol{R}) = m_2(\boldsymbol{R})/n$

- The theorem then follows from matrix Bernstein's inequality.

■

Randomized Sparsification of a Matrix

Dense data matrix → Low-rank matrix → Sparse matrix

We introduce a randomized sparsificaiton algorithm due to (Kundu-Drineas ’14)

- Let $B \in \mathbb{S}^{d\times d}$ be the target symmetric matrix, which can be expressed as a sum of its entries:

$$\boldsymbol{B} = \sum\nolimits_{i\le j} b_{ij}(\boldsymbol{E}_{ij} + \boldsymbol{E}_{ji}), \qquad \text{where } b_{ij} \coloneqq \boldsymbol{B}_{ij} \text{ for } i \neq j \text{ and } b_{ii} \coloneqq \frac{1}{2}\boldsymbol{B}_{ii}$$

- We introduce the sampling probabilities:

$$p_{ij} \coloneqq \left(\frac{|b_{ij}|^2}{\|\boldsymbol{B}\|_F^2} + \frac{|b_{ij}|}{\|\boldsymbol{B}\|_{\ell_1}}\right) \quad \forall 1 \le i < j \le d \quad \text{and} \quad p_{ii} \coloneqq \left(\frac{2|b_{ii}|^2}{\|\boldsymbol{B}\|_F^2} + \frac{|b_{ii}|}{\|\boldsymbol{B}\|_{\ell_1}}\right)$$

- Then, $\boldsymbol{R} \coloneqq p_{ij}^{-1} \cdot b_{ij}\left(\boldsymbol{E}_{ij} + \boldsymbol{E}_{ji}\right) \in \mathbb{S}^{d\times d}$ with probability p_{ij} is an unbiased estimator of $\boldsymbol{B}$

Randomized Sparsification of a Matrix

Theorem (Kundu-Drineas '14**).** Let $\overline{\boldsymbol{R}}_n \coloneqq \frac{1}{n}\sum_{k=1}^{n} \boldsymbol{R}_k$. Then, we have

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right] \leq \sqrt{\frac{4\|\boldsymbol{B}\|_F^2 d \log d}{n}} + \frac{4\|\boldsymbol{B}\|_{\ell_1} \log d}{3n}$$

Proof.

- To obtain this error bound, we'll apply Matrix approximation by random sampling theorem, which needs the uniform upper bound $\|\boldsymbol{R}\|$ and the second-moment $m_2(\boldsymbol{R})$
- We first note two lower bounds on the sampling probabilities p_{ij}:

$$p_{ij} \geq \frac{|b_{ij}|}{\|\boldsymbol{B}\|_{\ell_1}}, \qquad p_{ij} \geq \frac{|b_{ij}|^2}{\|\boldsymbol{B}\|_F^2} \quad \text{or} \quad p_{ii} \geq \frac{2|b_{ii}|^2}{\|\boldsymbol{B}\|_F^2}$$

- For the uniform upper bound, $\|\boldsymbol{R}\| \leq \max_{i \leq j} p_{ij}^{-1} \cdot |b_{ij}| \cdot \|\boldsymbol{E}_{ij} + \boldsymbol{E}_{ji}\| \leq 2\|\boldsymbol{B}\|_{\ell_1}$

Randomized Sparsification of a Matrix

- For the second-moment $m_2(\boldsymbol{R})$, we have

$$
\begin{aligned}
\mathbb{E}[\boldsymbol{R}^2] &= \sum_{i \le j} p_{ij}^{-2} \cdot b_{ij}^2 \left(\boldsymbol{E}_{ij} + \boldsymbol{E}_{ji}\right)^2 \cdot p_{ij} \\
&= \sum_{i<j} p_{ij}^{-1} \cdot b_{ij}^2 \left(\boldsymbol{E}_{ii} + \boldsymbol{E}_{jj}\right) + 4 \sum_i p_{ii}^{-1} \cdot b_{ii}^2\, \boldsymbol{E}_{ii} \\
&\preccurlyeq \|\boldsymbol{B}\|_F^2 \cdot \left(\sum_{i<j} \left(\boldsymbol{E}_{ii} + \boldsymbol{E}_{jj}\right) + 2 \sum_i \boldsymbol{E}_{ii} \right) \\
&= \|\boldsymbol{B}\|_F^2 \cdot \left((d-1) \sum_i \boldsymbol{E}_{ii} + 2 \sum_i \boldsymbol{E}_{ii} \right) = (d+1) \|\boldsymbol{B}\|_F^2 \cdot \boldsymbol{I}
\end{aligned}
$$

- By the matrix sampling theorem,

$$
\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right] \le \sqrt{\frac{2 m_2(\boldsymbol{R}) \log(d)}{n}} + \frac{2L \log(d)}{3n} = \sqrt{\frac{2\|\boldsymbol{B}\|_F^2 (d+1) \log d}{n}} + \frac{4\|\boldsymbol{B}\|_{\ell_1} \log d}{3n}
$$

■

Randomized Sparsification of a Matrix

Theorem (Kundu-Drineas ’14). Let $\overline{\boldsymbol{R}}_n \coloneqq \frac{1}{n}\sum_{k=1}^{n} \boldsymbol{R}_k$. Then, we have

$$\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right] \le \sqrt{\frac{4\|\boldsymbol{B}\|_F^2 d \log d}{n}} + \frac{4\|\boldsymbol{B}\|_{\ell_1} \log d}{3n}$$

- Note that $\|\boldsymbol{B}\|_{\ell_1} = \sum_{i,j}\left|B_{ij}\right| \le d\|\boldsymbol{B}\|_F$
- Thus, we obtain a relative-error bound:

$$\frac{\mathbb{E}\left[\left\|\overline{\boldsymbol{R}}_n - \boldsymbol{B}\right\|\right]}{\|\boldsymbol{B}\|} \le \boxed{\frac{\|\boldsymbol{B}\|_F}{\|\boldsymbol{B}\|}}\left(\sqrt{\frac{4d\log d}{n}} + \frac{4d\log d}{3n}\right) \le \epsilon$$

- We define the stable rank $\operatorname{srank}(\boldsymbol{B}) \coloneqq \frac{\|\boldsymbol{B}\|_F^2}{\|\boldsymbol{B}\|^2} \le \operatorname{rank}(B)$
- Thus, it suffices to take $n \ge \epsilon^{-2} \cdot \operatorname{srank}(\boldsymbol{B}) \cdot d\log d$ to obtain an ϵ-spectral sparsifier of $\boldsymbol{B}$

Randomized Sparsification of a Matrix

The more recent work (Braverman-Krauthgamer-Krishnan-Sapir ’21) improves the sparsity from

$$\tilde{O}(\epsilon^{-2}\mathrm{srank}(\boldsymbol{B})d)$$

to

$$\tilde{O}\left(\epsilon^{-2}\mathrm{srank}(\boldsymbol{B})\mathrm{ns}(\boldsymbol{B}) + \epsilon^{-1}\sqrt{\mathrm{srank}(\boldsymbol{B})\mathrm{ns}(\boldsymbol{B})d}\right)$$

where $\mathrm{ns}(\boldsymbol{a}) \coloneqq \min\left\{k \geq 0 : \|\boldsymbol{a}\|_1 \leq \sqrt{k}\|\boldsymbol{a}\|_2\right\}$ is the numerical sparsity of a vector $\boldsymbol{a}$ and $\mathrm{ns}(\boldsymbol{B})$ is defined as the maximum numerical sparsity of any of its rows and columns

More about Spectral Sparsification

There are many other forms of spectral sparsification

- The most important one is the graph sparsificaiton, where the target matrix is the Laplacian of a graph $\boldsymbol{L}_G \coloneqq \boldsymbol{D}_G - \boldsymbol{A}_G$ (Spielman-Teng ’04, Spielman-Teng ’11, Batson-Spielman-Srivastava ’14, Allen-Zhu-Liao-Orecchia ’15, Marcus-Spielman-Srivastava ’15, Lee-Sun ’17)
- Some “modern” spectral sparsification problems include:
  - Metric graphs and kernel graphs (Quanrud ’21)
  - Hypergraph sparsification (Lee ’23)
  - Sum of norms (Jambulapati-Lee-Liu-Sidford ’23)
  - Generalized linear models (Jambulapati-Lee-Liu-Sidford ’24)
  - Cayley graph (Khanna-Putterman-Sudan ’24, Hsieh-Lee-Mohanty-Putterman-Zhang ’25)
  - Sum of PSD matrices (Basu-Kothari-Liu-Meka ’26)