CS 839 Probability and Learning in High Dimension Lecture 10 - 02/28/2022

Lecture 10: Max Norm and Nuclear Norm Relaxations for Matrix Completion

Lecturer: Yudong Chen *Scribe: Qilin Li*

In this lecture,[1] we will first review some matrix norms and introduce max norm. Based on max norm and nuclear norm respectively, we will present two convex relaxations for matrix completion. We will also compare their performances. Under the setting with non-uniform observations, the max-norm approach is more robust.

1 Matrix Norms

There are many matrix norms. Below we discuss some of them, which are useful in developing and analyzing convex relaxation methods for low-rank matrix estimation.

Caution: In the literature, sometimes a matrix norm has different names/notations, and the same notation/name is used for different norms. Make sure you know what exactly is the norm being mentioned.

Notation: Below X, Y are $d_1 \times d_2$ matrices unless otherwise specified. We use X_{ij} and X_{i-}, respectively, to denote the (i,j)-th entry and i-th row of X. Denote by $\langle X, Y \rangle := \sum_{ij} X_{ij} Y_{ij}$ the trace inner product between X and Y, and by $\sigma_i(X)$ the i-th singular value of X. For a positive integer d, let $[d] := \{1, 2, \dots, d\}$. Given a norm $\|\cdot\|_\triangle$, its dual norm is

$$\|X\|_\triangledown := \max_{Y : \|Y\|_\triangle \le 1} \langle X, Y \rangle .$$

By definition, the generalized Holder's inequality

$$\langle X, Y \rangle \le \|X\|_\triangledown \|Y\|_\triangle$$

holds for any X, Y and any dual norm pairs.

1.1 Vectorized Norms

These are vector norms applied to the vectorized version of a matrix.

- Frobenius norm: $\|X\|_F := \sqrt{\sum_{i,j} X_{ij}^2} = \sqrt{\sum_i \sigma_i(X)^2}$.
- Element-wise ℓ_1 norm: $\|X\|_1 := \sum_{i,j} |X_{ij}|$.
- Element-wise ℓ_∞ norm: $\|X\|_\infty := \max_{i,j} |X_{ij}|$.

[1] *Reading:*

- T. Tony Cai and Wen-Xin Zhou. Matrix completion via max-norm constrained optimization. Electronic Journal of Statistics, Vol. 10 (2016). `https://arxiv.org/abs/1303.0341` [Cai and Zhou, 2017]
- (Additional background) Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In Proceedings of the 18th Annual Conference on Learning Theory (COLT), 2005. `https://home.ttic.edu/~nati/Publications/SrebroShraibmanCOLT05.pdf` [Srebro and Shraibman, 2005]
- (Additional background) Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. (2004). Complexity measures of sign measures. Combinatorica 27, 439–463. `https://www.cs.huji.ac.il/~nati/PAPERS/complexity_matrices.pdf` [Linial et al., 2007]

Note that $\|\cdot\|_F$ is the dual to itself; $\|\cdot\|_1$ and $\|\cdot\|_\infty$ are dual to each other. We have the following inequalities

$$\begin{aligned}\frac{1}{\sqrt{d_1 d_2}} \|X\|_1 &\le \|X\|_F \le \|X\|_1 , \\ \|X\|_\infty &\le \|X\|_F \le \sqrt{d_1 d_2} \|X\|_\infty , \\ \|X\|_F^2 &= \langle X, X \rangle \le \|X\|_1 \|X\|_\infty .\end{aligned}$$

1.2 Schatten Norms

These are vector norms applied to the vector of the singular values of a matrix. For each $p \in [1, \infty]$, the Schatten-p norm is defined as$\|X\|_{S_p} := (\sum_i \sigma_i(X)^p)^{1/p}$.

- Nuclear/trace norm: $\|X\|_{S_1} = \|X\|_*$ =sum of singular values
- Spectral norm: $\|X\|_{S_\infty} = \|X\|_{op}$ =largest singular value
- Frobenius norm: $\|X\|_{S_2} = \|X\|_F$.

For any conjugate pairs (p, q), i.e., $\frac{1}{p} + \frac{1}{q} = 1$, the norms $\|\cdot\|_{S_p}$ and $\|\cdot\|_{S_q}$ are dual to each other. In particular, the nuclear norm and spectral norm are duals.

1.3 Induced Operator Norms

For a pair of vector ℓ_a norm $\|\cdot\|_a$ and ℓ_b norm $\|\cdot\|_b$, the induced operator norm for a matrix X is defined as

$$\|X\|_{a \to b} := \max_{u \in \mathbb{R}^{d_2} : \|u\|_a \le 1} \|Xu\|_b .$$

- Spectral norm: $\|X\|_{2\to 2} = \|X\|_{op}$, a.k.a. THE operator norm. From its definition via induced norm, it is easy to verify the following sub-multiplicativity properties:

$$\begin{aligned}\|XY\|_{op} &\le \|X\|_{op} \|Y\|_{op} , \\ \|XY\|_F &\le \|X\|_{op} \|Y\|_F .\end{aligned}$$

- Max column sum: $\|X\|_{1\to 1} = \max_{1 \le j \le d_2} \sum_{i=1}^{d_1} |X_{ij}|$.
- Max row sum: $\|X\|_{\infty\to\infty} = \max_{1 \le i \le d_1} \sum_{j=1}^{d_2} |X_{ij}|$.
- Max row ℓ_2 norm: $\|X\|_{2\to\infty} = \max_{1 \le i \le d_1} \sqrt{\sum_{j=1}^{d_2} X_{ij}^2}$.
- Max entry: $\|X\|_{1\to\infty} = \|X\|_\infty = \max_{i,j} |X_{ij}|$.

Of particular importance to us is the ℓ_∞-to-ℓ_1 norm

$$\|X\|_{\infty\to 1} = \max_{\substack{u \in \{\pm 1\}^{d_1} \\ v \in \{\pm 1\}^{d_2}}} \left| \sum_{i,j} X_{ij} u_i v_j \right| = \max_{\substack{u \in \{\pm 1\}^{d_1} \\ v \in \{\pm 1\}^{d_2}}} \left| \langle X, uv^\top \rangle \right| .$$

Closely related is the so-called *cut norm*:

$$\|X\|_\square := \max_{\substack{I \subseteq [d_1] \\ J \subseteq [d_2]}} \left| \sum_{i \in I, j \in J} X_{ij} \right| .$$

These two norms are equivalent up to a *universal* constant:[2]

$$\|X\|_\square \le \|X\|_{\infty\to 1} \le 4 \|X\|_\square .$$

Note that exact computation of these two norms is at least as hard as MAX CUT.

[2]For a proof see Lemma 3.1 in Noga Alon, Assaf Naor: Approximating the cut-norm via Grothendieck's inequality. STOC 2004. [Alon and Naor, 2004]

1.4 Max norm and its dual

A less known matrix norm is the *max norm*:

$$\|X\|_{\max} := \min_{U,V:X=UV^\top} \|U\|_{2\to\infty} \|V\|_{2\to\infty}, \tag{1}$$

for which we recall that $\|U\|_{2\to\infty}$ is the maximum row ℓ_2 norm of U. (Not to confuse this max norm with the element-wise ℓ_∞ norm.) Compare (1) with an analogous variational formula for the nuclear norm:[3]

$$\|X\|_* = \min_{U,V:X=UV^\top} \|U\|_F \|V\|_F = \frac{1}{2} \min_{U,V:X=UV^\top} \left(\|U\|_F^2 + \|V\|_F^2 \right). \tag{2}$$

From the variational formulae it is easy to obtain

$$\frac{1}{\sqrt{d_1 d_2}} \|X\|_* \le \|X\|_{\max} \le \|X\|_*.$$

The dual norm of the max norm is the so-called *Grothendieck norm*

$$\begin{aligned} \|X\|_G &:= \max_{U,V:\|U\|_{2\to\infty}=\|V\|_{2\to\infty}=1} \left|\left\langle X, UV^\top \right\rangle\right| \\ &= \max_{\substack{u_i\in\mathbb{R}^k:\|u_i\|_2=1, i\in[d_1] \\ v_j\in\mathbb{R}^k:\|v_j\|_2=1, j\in[d_2]}} \left| \sum_{i,j} X_{ij} \langle u_i, v_j \rangle \right|, \end{aligned}$$

where $k := \min\{d_1, d_2\}$. Compare the above definition to a similar variational formula for the spectral norm (the dual of the nuclear norm):

$$\|X\|_{op} = \max_{U,V:\|U\|_F=\|V\|_F=1} \left|\left\langle X, UV^\top \right\rangle\right|.$$

From the variational formulae it is easy to obtain

$$\|X\|_{op} \le \|X\|_G \le \sqrt{d_1 d_2} \|X\|_{op}.$$

By the Grothendieck's inequality, we have

$$\|X\|_{\infty\to 1} \le \|X\|_G \le K_G \|X\|_{\infty\to 1},$$

where $K_G \le 1.783$ is the Grothendieck's constant. Therefore, $\|\cdot\|_{\infty\to 1}$ can be viewed as an approximate dual to the max norm. This fact plays an important role in the analysis of max-norm-based convex relaxation.

We have the following analogous sets of inequalities for the nuclear norm and max norm:

$$\|X\|_F \le \|X\|_* \le \sqrt{\operatorname{rank}(X)} \|X\|_F, \tag{3}$$

$$\|X\|_\infty \le \|X\|_{\max} \le \sqrt{\operatorname{rank}(X)} \|X\|_\infty. \tag{4}$$

Equation (3) is well-known. Equation (4) is from [Linial et al., 2007].

1.5 Max norm, nuclear norm and rank

Both the max and nuclear norms are convex surrogate for the rank. From [Srebro and Shraibman, 2005]: "Whereas bounding the rank corresponds to constraining the *dimensionality* of each row of U and V in a factorization $X = UV^\top$, bounding the trace-norm and max-norm corresponds to constraining the *norms* of rows of U and V (average row-norm for the trace-norm, and maximal row-norm for the max-norm)."

[3]We note in passing that the RHS of (2) is related to the regularization effect of weight decay applied to a linear neural network.

The unit balls in these two norms are related to the convex hulls of certain sets of rank-1 matrices.

$$\{X : \|X\|_* \leq 1\} = \text{conv}\left\{uv^\top : \|u\|_2 = \|v\|_2 = 1\right\},$$
$$\text{conv}\mathcal{M}_\pm \subseteq \{X : \|X\|_{\max} \leq 1\} \subseteq K_G \cdot \text{conv}\mathcal{M}_\pm,$$

where $\mathcal{M}_\pm := \left\{M \in \{\pm 1\}^{d_1 \times d_2} : \text{rank}(M) = 1\right\}$ denotes the set of rank-1 sign matrices. The first equation can be proved by SVD.[4]. For the second equation above on the max-norm unit ball, see equation (2.3) in [Cai and Zhou, 2017] and the discussion therein.

Let $\mathcal{C}$ be a given convex set. The *convex envelop* of a (possibly nonconvex) function $f : \mathcal{C} \longrightarrow \mathbb{R}$ is defined at each point of $\mathcal{C}$ as the supremum of all convex function that lie under that function, i.e.,

$$(\text{conv}f)(x) := \sup_g \{g(x) : g \text{ is convex and } g \leq f \text{ over } \mathcal{C}\}, \text{for any } x \in \mathcal{C}.$$

$\text{conv}f$ is convex because the pointwise supremum preserves convexity. Hence, we can say that $\text{conv}f$ is the largest convex function lying under f.

By definition, we have that $\text{rank}(X) \cdot \|X\|_{op} \geq \|X\|_*$, so the nuclear norm is a convex lower bound of the rank function on the unit ball in the operator norm. In fact, it can be shown that this is the tightest convex lower bound, i.e., the convex envelope of $\text{rank}(X)$ on the set $\left\{X : \|X\|_{op} \leq 1\right\}$ is the nuclear norm $\|X\|_*$. (See Theorem 2.2 in [Recht et al., 2010].)

1.6 The SDP formulation for max-norm relaxation and nuclear-norm relaxation

As shown in [Srebro et al., 2004], the max-norm of a $d_1 \times d_2$ matrix M can be computed via a semi-definite program (SDP):

$$\|M\|_{\max} = \min_{R, W_1, W_2} R \quad \text{s.t.} \begin{pmatrix} W_1 & M \\ M^\top & W_2 \end{pmatrix} \succeq 0, \ \text{diag}(W_1) \leq R, \ \text{diag}(W_2) \leq R.$$

Consequently, the following max-norm constrained problem

$$\begin{aligned} \min_M \quad & f(M) \\ \text{s.t.} \quad & \|M\|_{\max} \leq R. \end{aligned}$$

can be reformulated a problem with linear-inequality constraints:

$$\begin{aligned} \min_{M, W_1, W_2} \quad & f(M) \\ \text{s.t.} \quad & \begin{pmatrix} W_1 & M \\ M^\top & W_2 \end{pmatrix} \succeq 0, \ \text{diag}(W_1) \leq R, \ \text{diag}(W_2) \leq R. \end{aligned}$$

If f is a linear function or a convex quadratic function, then the above problem is a SDP problem.[5]

Similarly, as shown in [Recht et al., 2010, equaiton (2.6)], the nuclear-norm of M can be computed via a SDP:

$$\|M\|_* = \min_{R, W_1, W_2} R \quad \text{s.t.} \begin{pmatrix} W_1 & M \\ M^\top & W_2 \end{pmatrix} \succeq 0, \ \text{Tr}(W_1) + \text{Tr}(W_2) \leq 2R.$$

Consequently, the following nuclear-norm constrained problem

$$\begin{aligned} \min_M \quad & f(M) \\ \text{s.t.} \quad & \|M\|_* \leq R. \end{aligned}$$

[4]For more details, see `https://math.stackexchange.com/questions/3951902/convex-hull-of-rank-1-matrices-is-the\-nuclear-norm-unit-ball`

[5]For a proof, we can use the equivalence on Page 16 of [Freund, 2004]. `https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-251j-introduction-to-mathematical-programming-fall-2009/readings/MIT6_251JF09_SDP.pdf`

can be reformulated as a problem with linear-inequality constraints:

$$\begin{aligned} \min_{M,W_1,W_2} \quad & f(M) \\ \text{s.t.} \quad & \begin{pmatrix} W_1 & M \\ M^\top & W_2 \end{pmatrix} \succeq 0, \ \operatorname{Tr}(W_1) + \operatorname{Tr}(W_2) \leq 2R. \end{aligned}$$

Again, if f is a linear function or a convex quadratic function, then the above problem is an SDP.

Later we will see that the max norm has some nice properties, though it is a bit harder to compute/optimize than the nuclear norm (which is given by SVD).

2 Matrix Completion with Non-Uniform Observations

Let $M^* \in \mathbb{R}^{d\times d}$ be an unknown rank-r matrix satisfying $\|M^*\|_\infty \leq 1$. We observe a matrix $Y \in \mathbb{R}^{d\times d}$ of the form

$$Y_{ij} = \delta_{ij}\left(M^*_{ij} + E_{ij}\right) \qquad \text{for each } (i,j) \in [d]\times[d], \tag{5}$$

where

$$\delta_{ij} = \begin{cases} 1 & \text{with probability } p_{ij}, \\ 0 & \text{with probability } 1-p_{ij}, \end{cases}$$

is the indicator for observing the (i,j)-th entry, $E_{ij} \in [-\sigma,\sigma]$ is a zero-mean bounded random variable representing additive noise, and all the random variables $\delta_{ij}, E_{ij} : (i,j) \in [d]\times[d]$ are mutually independent. In words, each entry of M^* is observed with (non-uniform) probability p_{ij}, and the observed entries are contaminated by σ-bounded noise. This setting is a generalization of the problem in Lecture 1, which involves uniform, noiseless observations.

Let $\Omega := \{(i,j) : \delta_{ij} = 1\}$ be the set of observed indices. Define the projection $\mathcal{P}_\Omega : \mathbb{R}^{d\times d} \to \mathbb{R}^{d\times d}$ via

$$(\mathcal{P}_\Omega(X))_{ij} = \begin{cases} X_{ij} & (i,j) \in \Omega, \\ 0 & (i,j) \notin \Omega, \end{cases}$$

which keeps the observed entries and zeros out the unobserved ones. The observation model (5) can be written succinctly as

$$Y = \mathcal{P}_\Omega(M^* + E).$$

We make note of the simple fact that

$$\langle \mathcal{P}_\Omega(A), \mathcal{P}_\Omega(B)\rangle = \langle A, \mathcal{P}_\Omega(B)\rangle = \langle \mathcal{P}_\Omega(A), B\rangle, \qquad \forall A, B \in \mathbb{R}^{d\times d}.$$

We develop estimators of M^* using convex relaxations. We begin by noting that in light of the inequalities (3) and (4), we have $\|M^*\|_* \leq \sqrt{r}\,\|M^*\|_F \leq \sqrt{r}d$ and $\|M^*\|_{\max} \leq \sqrt{r}\,\|M^*\|_\infty \leq \sqrt{r}$. These bounds motivate us to consider the nuclear-norm constrained convex program

$$\widehat{M}_{\text{nuc}} := \operatorname*{argmin}_{M:\|M\|_* \leq \sqrt{r}d} \|\mathcal{P}_\Omega(M-Y)\|_F^2. \tag{6}$$

and the max-norm constrained program

$$\widehat{M}_{\max} := \operatorname*{argmin}_{M:\|M\|_{\max} \leq \sqrt{r}} \|\mathcal{P}_\Omega(M-Y)\|_F^2. \tag{7}$$

Let $P = (p_{ij}) \in [0,1]^{d\times d}$ be the matrix of observation probabilities, and $p_{\min} := \min_{i,j} p_{ij}$ be the minimum probability. Recalling Section 1, we have $\|P\|_{1\to 1} = \max_j \sum_i p_{ij}$ (max column sum of probabilities), $\|P\|_{\infty\to\infty} = \max_i \sum_j p_{ij}$ (max row sum) and $\|P\|_1 = \sum_{i,j} p_{ij}$ (overall sum).

We establish the following.

Theorem 1. *Under the above setting, if $p_{\min} \geq \frac{1}{d}$, then with high probability we have*

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^*\right)\right\|_F^2 \lesssim \sigma\sqrt{rd^2 \cdot \max\left\{\|P\|_{1\to 1}, \|P\|_{\infty\to\infty}\right\}(\log d)^2}. \tag{8}$$

Theorem 2. *Under the above setting, if $p_{\min} \geq \frac{1}{d}$, then with high probability we have*

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\max} - M^*\right)\right\|_F^2 \lesssim \sigma\sqrt{rd\,\|P\|_1}. \tag{9}$$

Some remarks are in order.

- To compare the two theorems, we rewrite the nuclear-norm based bound (8) as

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^*\right)\right\|_F^2 \lesssim \sigma\sqrt{rd\,\|P\|_1(\log d)^2} \cdot \sqrt{\frac{\max\left\{\|P\|_{1\to 1}, \|P\|_{\infty\to\infty}\right\}/d}{\|P\|_1/d^2}}.$$

  The RHS has an additional factor $\gamma := \sqrt{\frac{\max\{\|P\|_{1\to 1}, \|P\|_{\infty\to\infty}\}/d}{\|P\|_1/d^2}}$ when compared with the max-norm based bound (9). Note that $\max\{\|P\|_{1\to 1}, \|P\|_{\infty\to\infty}\}/d$ is the average observation probability within the most-observed row/column, whereas $\|P\|_1/d^2$ is the average probability over the entire matrix. The more non-uniform the probabilities are, the larger γ is. Therefore, the max-norm approach is more robust against non-uniform observation probabilities. This advantage of max-norm has also been observed empirically.

- Theorems 1 and 2 control $\left\|\mathcal{P}_\Omega\left(\widehat{M} - M^*\right)\right\|_F^2$, the error on the observed entries. Under mild assumptions, one can further show that w.h.p.,

$$\left\|\mathcal{P}_\Omega\left(\widehat{M} - M^*\right)\right\|_F^2 \gtrsim \sum_{i,j} p_{ij}\left(\widehat{M}_{ij} - M^*_{ij}\right)^2 \geq p_{\min}\left\|\widehat{M} - M^*\right\|_F^2. \tag{10}$$

  Combining (10) with the two theorems controls the error over the entire matrix.

- The proof of Theorems 1 and 2 actually does not rely on M^* being low-rank. They hold as long as any M^* has small nuclear norm or max norm.

- In the uniform setting $p_{ij} = p, \forall i,j$ and assuming (10) holds, the error bounds (8) and (9) become, after some algebra,

$$\frac{1}{d^2}\left\|\widehat{M} - M^*\right\|_F^2 \lesssim \sqrt{\frac{r\,\text{polylog}(d)}{pd}};$$

  note the square root on the RHS. This bound is called a *slow rate*; as mentioned, it is valid for any M^* with small nuclear/max norm. This rate is in fact minimax optimal for estimating low nuclear/max norm matrices. Compare with the bound for the spectral method from Lecture 1:

$$\frac{1}{d^2}\left\|\widehat{M}_{\text{spectral}} - M^*\right\|_F^2 \lesssim \frac{r\log d}{pd},$$

  which is called a *fast rate* but only valid under the stronger assumption of M^* being low-rank. This rate is minimax optimal for estimating low-rank matrices.

3 Proof of Theorem 1

Since $\widehat{M}_{\text{nuc}}$ is optimal and M^* is feasible to the program (6), we have the basic inequality

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - Y\right)\right\|_F^2 \leq \left\|\mathcal{P}_\Omega\left(M^* - Y\right)\right\|_F^2,$$

$$\Rightarrow \left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^* - E\right)\right\|_F^2 \leq \left\|\mathcal{P}_\Omega\left(E\right)\right\|_F^2,$$

where the last step follows from $Y = \mathcal{P}_\Omega(M^* + E)$. Expanding the squares and rearranging terms gives

$$\begin{aligned}\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^*\right)\right\|_F^2 &\leq 2\left\langle \mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^*\right), \mathcal{P}_\Omega(E)\right\rangle \\ &\lesssim \left|\left\langle \widehat{M}_{\text{nuc}}, \mathcal{P}_\Omega(E)\right\rangle\right| + \left|\langle M^*, \mathcal{P}_\Omega(E)\rangle\right|.\end{aligned}$$

The constraint in the program (6) ensures that $\left\|\widehat{M}_{\text{nuc}}\right\|_* \leq \sqrt{r}d$ and $\|M^*\|_* \leq \sqrt{r}d$. By duality between the nuclear and spectral norms, we have

$$\begin{aligned}\left|\left\langle \widehat{M}_{\text{nuc}}, \mathcal{P}_\Omega(E)\right\rangle\right| + \left|\langle M^*, \mathcal{P}_\Omega(E)\rangle\right| &\leq \left\|\widehat{M}_{\text{nuc}}\right\|_* \|\mathcal{P}_\Omega(E)\|_{op} + \|M^*\|_* \|\mathcal{P}_\Omega(E)\|_{op} \\ &\lesssim \sqrt{r}d \, \|\mathcal{P}_\Omega(E)\|_{op}.\end{aligned}$$

The (i,j)-th entry of $\mathcal{P}_\Omega(E)$ is zero-mean, σ-bounded with variance $\lesssim \sigma^2 p_{ij}$. (Exercise) By Matrix Bernstein we have w.h.p.

$$\begin{aligned}\|\mathcal{P}_\Omega(E)\|_{op} &\lesssim \sigma\sqrt{\max\left\{\|P\|_{1\to1}, \|P\|_{\infty\to\infty}\right\}\log d} + \sigma \log d \\ &\lesssim \sigma\sqrt{\max\left\{\|P\|_{1\to1}, \|P\|_{\infty\to\infty}\right\}(\log d)^2},\end{aligned}$$

where the last steps holds under the assumption $p_{\min} \geq \frac{1}{d}$. Combining pieces, we obtain

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\text{nuc}} - M^*\right)\right\|_F^2 \lesssim \sqrt{r}d \cdot \sigma\sqrt{\max\left\{\|P\|_{1\to1}, \|P\|_{\infty\to\infty}\right\}(\log d)^2}.$$

This completes the proof of Theorem 1.

Proof of the Exercise:

Proof To apply Matrix Bernstein inequality, we write $\mathcal{P}_\Omega(E)$ as a summation of d^2 matrices.

$$\mathcal{P}_\Omega(E) = \sum_{i,j=1}^{d} (\delta_{ij}E_{ij}) \cdot I_{ij},$$

where I_{ij} is an indicator matrix with the (i,j)-th entry as 1 and all other entries as 0.

Next, we check the conditions of Matrix Bernstein inequality:

- Since $\delta_{ij}E_{ij}$ is the (i,j)-th entry of $\mathcal{P}_\Omega(E)$, we have $\mathbb{E}[(\delta_{ij}E_{ij}) \cdot I_{ij}] = 0, \ \forall\, i,\, j$.
- $\|(\delta_{ij}E_{ij}) \cdot I_{ij}\|_{op} \leq \|(\delta_{ij}E_{ij}) \cdot I_{ij}\|_F = \sqrt{\delta_{ij}E_{ij}^2} \leq \sigma$.
- First, we consider

$$\begin{aligned}&((\delta_{ij}E_{ij}) \cdot I_{ij}) \cdot ((\delta_{ij}E_{ij}) \cdot I_{ij})^\top = \delta_{ij}E_{ij}^2 \cdot I_{ii} \\ \Rightarrow &\sum_{i,j=1}^{d} ((\delta_{ij}E_{ij}) \cdot I_{ij}) \cdot ((\delta_{ij}E_{ij}) \cdot I_{ij})^\top = \operatorname{diag}\Big(\sum_{j=1}^{d} \delta_{1j}E_{1j}^2, \sum_{j=1}^{d} \delta_{2j}E_{2j}^2, \ldots, \sum_{j=1}^{d} \delta_{dj}E_{dj}^2\Big) \\ \Rightarrow &\left\|\mathbb{E}\Big[\sum_{i,j=1}^{d} ((\delta_{ij}E_{ij}) \cdot I_{ij}) \cdot ((\delta_{ij}E_{ij}) \cdot I_{ij})^\top\Big]\right\|_{op} \leq \max_i \sum_{j=1}^{d} \mathbb{E}[\delta_{ij}E_{ij}^2] \lesssim \max_i \sum_{j=1}^{d} \sigma^2 p_{ij} = \sigma^2 \|P\|_{\infty\to\infty}.\end{aligned}$$

  Symmetrically, we can prove

$$\left\|\mathbb{E}\Big[\sum_{i,j=1}^{d} ((\delta_{ij}E_{ij}) \cdot I_{ij})^\top \cdot ((\delta_{ij}E_{ij}) \cdot I_{ij})\Big]\right\|_{op} \lesssim \sigma^2 \|P\|_{1\to1}.$$

Therefore, by the user-friendly form of Matrix Bernstein inequality, we have w.h.p.,

$$\|\mathcal{P}_\Omega(E)\|_{op} \lesssim \sigma\sqrt{\max\left\{\|P\|_{1\to1}, \|P\|_{\infty\to\infty}\right\}\log d} + \sigma \log d.$$

□

4 Proof of Theorem 2

By the exactly the same argument as in the proof of Theorem 1, we obtain

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\max} - M^*\right)\right\|_F^2 \lesssim \left|\left\langle \widehat{M}_{\max}, \mathcal{P}_\Omega(E)\right\rangle\right| + \left|\left\langle M^*, \mathcal{P}_\Omega(E)\right\rangle\right|.$$

The constraint in the program (7) ensures that $\left\|\widehat{M}_{\max}\right\|_{\max} \le \sqrt{r}$ and $\|M^*\|_{\max} \le \sqrt{r}$. By duality between the max norm and Grothendieck norm, we have

$$\begin{aligned}\left|\left\langle \widehat{M}_{\max}, \mathcal{P}_\Omega(E)\right\rangle\right| + \left|\left\langle M^*, \mathcal{P}_\Omega(E)\right\rangle\right| &\le \left\|\widehat{M}_{\max}\right\|_{\max}\|\mathcal{P}_\Omega(E)\|_G + \|M^*\|_{\max}\|\mathcal{P}_\Omega(E)\|_G \\ &\lesssim \sqrt{r}\,\|\mathcal{P}_\Omega(E)\|_G \\ &\lesssim \sqrt{r}\,\|\mathcal{P}_\Omega(E)\|_{\infty\to1},\end{aligned}$$

where the last step follows from Grothendieck's inequality.

Define the shorthand $Z := \mathcal{P}_\Omega(E)$. Recall that $\|Z\|_{\infty\to1} := \max_{u,v\in\{\pm1\}^d}\left|\sum_{i,j} Z_{ij}u_i v_j\right|$. Fix an arbitrary pair $(u,v) \in \{\pm1\}^d \times \{\pm1\}^d$. The quantity $\sum_{i,j} Z_{ij}u_iv_j$ is the sum of d^2 independent, zero-mean, σ-bounded random variables, each of which has variance $\lesssim \sigma^2 p_{ij}$. The (scalar) Bernstein gives

$$\mathbb{P}\left\{\left|\sum_{i,j} Z_{ij}u_iv_j\right| \ge t\right\} \le 2\exp\left(-\frac{ct^2}{\sum_{i,j}\sigma^2 p_{ij} + \sigma t}\right) = 2\exp\left(-\frac{ct^2}{\sigma^2\|P\|_1 + \sigma t}\right),$$

where $c > 0$ is a universal constant. By a union bound over all 2^{2d} possible pairs of (u,v), we get

$$\mathbb{P}\left\{\|Z\|_{\infty\to1} \ge t\right\} \le 2^{2d}\cdot 2\exp\left(-\frac{ct^2}{\sigma^2\|P\|_1 + \sigma t}\right).$$

Taking $t = C\left(\sigma\sqrt{d\|P\|_1} + \sigma d\right)$ for a sufficiently large constant C, we obtain that with probability at least $1 - 2^{-d}$,

$$\|Z\|_{\infty\to1} \lesssim \sigma\sqrt{d\|P\|_1} + \sigma d \lesssim \sigma\sqrt{d\|P\|_1},$$

where the last steps holds under the assumption $p_{\min} \ge \frac{1}{d}$. Combining pieces, we obtain

$$\left\|\mathcal{P}_\Omega\left(\widehat{M}_{\max} - M^*\right)\right\|_F^2 \lesssim \sqrt{r}\cdot\sigma\sqrt{d\|P\|_1}.$$

References

[Alon and Naor, 2004] Alon, N. and Naor, A. (2004). Approximating the cut-norm via grothendieck's inequality. *SIAM Journal on Computing*.

[Cai and Zhou, 2017] Cai, T. T. and Zhou, W. X. (2017). Matrix completion via max-norm constrained optimization. *Electronic Journal of Statistics*, 10(1).

[Freund, 2004] Freund, R. M. (2004). Introduction to semidefinite programming (sdp). *Massachusetts Institute of Technology*, pages 8–11.

[Linial et al., 2007] Linial, N., Mendelson, S., Schechtman, G., and Shraibman, A. (2007). Complexity measures of sign matrices. *Combinatorica*.

[Recht et al., 2010] Recht, B., Fazel, M., and Parrilo, P. A. (2010). Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501.

[Srebro et al., 2004] Srebro, N., Rennie, J., and Jaakkola, T. (2004). Maximum-margin matrix factorization. In Saul, L., Weiss, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems*, volume 17. MIT Press.

[Srebro and Shraibman, 2005] Srebro, N. and Shraibman, A. (2005). Rank, trace-norm and max-norm. *Springer-Verlag*.