ORIE 7790 High Dimensional Probability and Statistics	Lecture 9 - Feb. 18, 2020

Lecture 9: Matrix Concentration inequalities

Lecturer: Yudong Chen *Scribe: Vasilis Charisopoulos*

In this lecture, we state the Matrix Bernstein inequality and sketch a few interesting applications. References:

- The proof for the Matrix Bernstein inequality can be found in Vershynin's book [2, Chapter 5.4]. Also see Chapter 6.6 therein.
- An exposition of the matrix Chernoff method can be found in Tropp's paper [1], along with bounds extending beyond the case of rectangular bounded matrices.
- Also related: Wainwright's book [3, Chapter 6.4]

1 Matrix concentration inequalities

The general idea: write a random matrix X as the sum of "simple" random matrices $\sum_i X^{(i)}$.

Theorem 1 (Matrix Bernstein inequality). *Suppose that $X^{(1)}, \ldots, X^{(n)} \in \mathbb{R}^{m_1 \times m_2}$ are independent, zero-mean random matrices with*

$$\|X^{(i)}\|_{\text{op}} \leq b \quad a.s., \quad \max\left\{\left\|\sum_i \mathbb{E}\left(X^{(i)\top} X^{(i)}\right)\right\|_{\text{op}}, \left\|\sum_i \mathbb{E}\left(X^{(i)} X^{(i)\top}\right)\right\|_{\text{op}}\right\} \leq \sigma^2.$$

Then we have

$$\mathbb{P}\left(\left\|\sum_i X^{(i)}\right\|_{\text{op}} \geq t\right) \leq (m_1 + m_2) \exp\left(-c \min\left\{\frac{t^2}{\sigma^2}, \frac{t}{b}\right\}\right). \tag{1}$$

Remark. In the 1-dimensional case, the quantity σ^2 reduces to the sum of variances of each element.

The proof proceeds quite naturally by mimicking that of the scalar Bernstein inequality, with one important difference: in the scalar case, we have $\mathbb{E}\left(e^{X+Y}\right) = \mathbb{E}\left(e^X\right)\mathbb{E}\left(e^Y\right)$. This is no longer true generically in the matrix world, because matrices are not commutative in general. However, we can use the Golden-Thompson inequality instead:

$$\mathsf{tr}(e^{X+Y}) \leq \mathsf{tr}(e^X \cdot e^Y)$$

Alternatively, one may use the Lieb's theorem, as is done in [1]

2 Applications

Example 1 (Matrices with independent entries). Suppose $Y \in \mathbb{R}^{n \times n}$ with $Y_{ij} = \varepsilon_{ij} \overset{\text{i.i.d.}}{\sim} \text{Unif}\{-1, 1\}$. We write Y as

$$Y = \sum_{i,j} \varepsilon_{ij} e_i e_j^\top \equiv \sum_{i,j} Y^{(i,j)}.$$

To apply Theorem 1, we need to check the conditions. Verifying the first condition is easy since $\|Y^{(i,j)}\|_{\text{op}} = 1$, implying $b = 1$. On the other hand

$$\mathbb{E}\left((Y^{(i,j)})^\top Y^{(i,j)}\right) = \mathbb{E}\left(\varepsilon_{i,j}^2 e_j e_i^\top e_i e_j^\top\right) = e_j e_j^\top \Rightarrow \sum_{i,j} \mathbb{E}\left((Y^{(i,j)})^\top Y^{(i,j)}\right) = nI_n \tag{2}$$

The same argument applies to the other sum, hence $\sigma^2 = n$. By Matrix Bernstein with $t := C\sqrt{n\log n}$, we have

$$\mathbb{P}\left(\|Y\|_{\text{op}} \geq C\sqrt{n\log n}\right) \leq 2n\exp\left(-c\min\left\{\frac{\cancel{n}\log n}{\cancel{n}}, \frac{\sqrt{n\log n}}{1}\right\}\right) = 2n\exp\left(-c\log n\right) = 2n^{-c'} \tag{3}$$

for some constant $c' > 0$. Therefore, with probability at least $1 - 2n^{-c'}$, we get $\|Y\|_{\text{op}} \lesssim \sqrt{n\log n}$. Notice that we are loose by a factor of $\sqrt{\log n}$ compared to the bounds derived for Gaussian random matrices.

Example 2 (Matrix Completion)**.** Suppose we are given $Y^\star \in \mathbb{R}^{n\times n}$ with $\text{rank}(Y^\star) = r$ and $|Y^\star_{ij}| \leq 1, \forall i,j$. Observe:

$$Y_{ij} = \begin{cases} Y^\star_{ij}, & \text{w.p } p \\ 0, & \text{w.p. } 1-p \end{cases} \tag{4}$$

Moreover, assume that $p \ll 1$. It's easy to check that $\mathbb{E}\left(Y_{ij}\right) = pY^\star_{ij}$. This is a special case of the matrix estimation problem, so we'll take

$$\hat{Y} := \underset{Z:\text{rank}(Z)\leq r}{\arg\min} \left\|\frac{1}{p}Y - Z\right\|_F \tag{5}$$

To bound the error of this estimator (in Frobenius norm), we can write

$$\frac{1}{n^2}\left\|\hat{Y} - Y^\star\right\|_F^2 \lesssim \frac{r}{n^2}\left\|\frac{1}{p}Y - \hat{Y}^\star\right\|_{\text{op}}^2. \tag{6}$$

This reduces to bounding the operator norm of random matrix with independent, zero mean entries. Define

$$Z_{ij} \triangleq \frac{1}{p}Y_{ij} - Y^\star_{ij} \implies |Z_{ij}| \leq \max\left\{\left|\frac{Y^\star_{ij}}{p} - Y^\star_{ij}\right|, |Y^\star_{ij}|\right\} \leq \frac{1}{p}$$

Thus Z_{ij} is subgaussian with parameter $\frac{1}{p^2}$, so using the results of last week we obtain $\|Z\|_{\text{op}} \lesssim \sqrt{n}\cdot\frac{1}{p} \implies \frac{1}{n^2}\left\|\hat{Y} - Y^\star\right\|_F^2 \leq \frac{r}{np^2}$. Let us try the Matrix Bernstein inequality:

Lemma 1 (Homework)**.** *For the matrices defined above, we have*

$$\left\|\frac{1}{p}Y - Y^\star\right\|_{\text{op}} \lesssim \sqrt{\frac{n\log n}{p}} + \frac{\log n}{p}$$

with probability at least $1 - 2n^{-c}$.

Using Lemma 1, we have reduced the dependence to p from $\frac{1}{p}$ to $\frac{1}{\sqrt{p}}$, which is desirable since we are interested in $p \to 0$. Simplifying we get

$$\frac{1}{n^2}\left\|\hat{Y} - Y^\star\right\|_F^2 \lesssim \frac{r\log n}{np} + \frac{r\log^2 n}{n^2p^2}.$$

Observe that we can set p as low as $\frac{r\log n}{n\varepsilon^2}$, for some $\varepsilon \in (0,1]$ and still satisfy $\frac{1}{n^2}\left\|\hat{Y} - Y^\star\right\|_F^2 \lesssim \varepsilon^2$.

Example 3 (Preference matrix completion)**.** This problem can be posed as a ranking problem from pairwise comparisons. The setup follows:

- suppose we have n teams with unknown ranking (assume such a ranking exists).
- if team i is better than team j, we have

$$\mathbb{P}\left(i \text{ beats } k\right) \geq \mathbb{P}\left(j \text{ beats } k\right), \quad \forall k.$$

- a match is played between teams i and j with probability p. Let the probability of i beating j be $Y^\star_{ij}$.
- we observe

$$Y_{ij} = \begin{cases} 1, & \text{w.p.} \quad pY^\star_{ij}, \\ 0, & \text{w.p.} \quad p(1 - Y^\star_{ij}) \\ 0, & \text{w.p.} \quad (1-p) \end{cases} \tag{7}$$

  The first case in (7) corresponds to i beating j, the second case is j beating i, and the third case happens when no game was played.

For simplicity, let us assume that everything is independent across (i, j) (i.e., when teams are matched, they play two games independently — one of them is Y_{ij}, and the other is Y_{ji}). The goal is again to recover $Y^\star \in [0,1]^{n\times n}$ given the observation Y. The important difference between this example and Example 2 is that $Y^\star$ here is **not** exactly low rank.

Estimator: we still set $\hat{Y}$ equal to the best rank-r approximation of $\frac{1}{p}Y$, as in Eq. (5). The following Claim can guide us in setting the target rank r:

Claim 1. *The matrix $Y^\star$ is approximately low-rank with $r = \sqrt{pn}$, i.e., there exists Z with rank$(Z) \le r$ such that*

$$\begin{cases} \|Z - Y^\star\|_F^2 & \le \frac{n^2}{r} \\ |Z_{ij}| & \le 1, \quad \forall (i,j). \end{cases} \tag{8}$$

Let $\|\cdot\|_*$ denote the nuclear/trace norm of a matrix (i.e., the sum of its singular values). Since $\hat{Y}$ is the closest rank-r approximation to $\frac{1}{p}Y$, we can show that

$$\cancel{\left\| Z - \frac{1}{p}Y \right\|_F^2} \ge \left\| \frac{1}{p}Y - \hat{Y} \right\|_F^2 = \cancel{\left\| \frac{1}{p}Y - Z \right\|_F^2} + \left\| Z - \hat{Y} \right\|_F^2 + 2\left\langle \frac{1}{p}Y - Z, Z - \hat{Y} \right\rangle \tag{9}$$

$$\Rightarrow \left\| Z - \hat{Y} \right\|_F^2 \le 2\left\langle \frac{1}{p}Y - Z, \hat{Y} - Z \right\rangle = 2\left\langle \frac{1}{p}Y - Y^\star, \hat{Y} - Z \right\rangle + 2\left\langle Y^\star - Z, \hat{Y} - Z \right\rangle \tag{10}$$

$$\le 2\left\| \frac{1}{p}Y - Y^\star \right\|_{\text{op}} \left\| \hat{Y} - Z \right\|_* + 2\left\| \hat{Y} - Z \right\|_F \cdot \left\| Y^\star - Z \right\|_F, \tag{11}$$

where the penultimate inequality follows by adding and subtracting $Y^\star$ in the inner product, and the last inequality is the trace Hölder inequality + Cauchy-Schwarz inequality. Since the matrix $\hat{Y} - Y^\star$ has rank at most $2r$, we have that

$$\left\| \hat{Y} - Z \right\|_* \le \sqrt{2r} \left\| \hat{Y} - Z \right\|_F \Longrightarrow \left\| Z - \hat{Y} \right\|_F \lesssim \sqrt{r} \left\| \frac{1}{p}Y - Y^\star \right\|_{\text{op}} + \left\| Y^\star - Z \right\|_F.$$

From Lemma 1, we can readily bound the first term, while the second term is upper bounded by $\frac{n}{\sqrt{r}}$ using Claim 1 above. Now, we rewrite

$$\left\| \hat{Y} - Y^\star \right\|_F \le \left\| Z - Y^\star \right\|_F + \left\| \hat{Y} - Z \right\|_F \lesssim \frac{2n}{\sqrt{r}} + \sqrt{\frac{rn \log n}{p}}.$$

Plugging in $r = \sqrt{pn}$ and dividing by n^2 yields the error bound

$$\frac{1}{n^2} \left\| \hat{Y} - Y^\star \right\|_F^2 \lesssim \frac{\log n}{\sqrt{pn}}.$$

In particular, the error goes to zero when $p = \omega\left(\frac{\log^2 n}{n}\right)$.

Proof of Claim 1. Let us introduce some notation. Define

$$S_i := \sum_{j=1}^{n} Y^\star_{ij}, \quad i = 1, \ldots, n,$$
$$\mathcal{T}_\ell := \left\{ i \;\middle|\; S_i \in \left[\frac{n(\ell-1)}{r}, \frac{n\ell}{r}\right) \right\}, \quad \ell = 1, \ldots, r.$$

and let $k(\ell) :=$ first element in $\mathcal{T}_\ell$, which we will treat as the "representative" element. For all $i \in \mathcal{T}_\ell$, define the i^{th} row of our candidate low-rank matrix Z as

$$Z_{i,:} = Y^\star_{k(\ell),:}$$

which is a "discretization" of the rows of $Y^\star$. Then $Z \in [0,1]^{n\times n}$ has at most r distinct rows, and $\operatorname{rank}(Z) \le r$. Moreover, for each $i \in \mathcal{T}_\ell$:

- if team i is better than team $k(\ell)$, we can write

$$\begin{aligned}
\sum_{j=1}^{n}(Y^\star_{ij} - Z_{ij})^2 &= \sum_{j=1}^{n}(Y^\star_{ij} - Y^\star_{k(\ell)j})^2 \\
&\overset{(\flat)}{\le} \sum_{j=1}^{n}\left|Y^\star_{ij} - Y^\star_{k(\ell)j}\right| \\
&\overset{(\sharp)}{=} \sum_{j=1}^{n}(Y^\star_{ij} - Y^\star_{k(\ell)j}) = S_i - S_{k(\ell)} \\
&\overset{(\natural)}{\le} \frac{n}{r},
\end{aligned}$$

  where $(\flat)$ is due to the fact that each element in the sum is in $[0,1]$ and $(\sharp)$ is due to the fact that $Y^\star_{ij} > Y^\star_{k(\ell)j}$ if team i beats team $k(\ell)$. The last inequality, $(\natural)$, is due to the fact that the difference between row sums of elements in $\mathcal{T}_\ell$ is upper bounded by the width of that interval, $\frac{n}{r}$.

- if $k(\ell)$ beats team i, then the same argument applies leading to the difference $S_{k(\ell)} - S_i$ which is also upper bounded by $\frac{n}{r}$.

Combining, we obtain

$$\|Y^\star - Z\|_F^2 = \sum_{i=1}^{n}\left(\sum_{j=1}^{n}(Y^\star_{ij} - Z_{ij})^2\right) \le n \cdot \frac{n}{r} = \frac{n^2}{r}.$$

□

References

[1] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

[2] Vershynin, Roman. *High Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2018.

[3] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2019.