CS 6220: Data-Sparse Matrix Computations
Lecture 2 (01/30/2020)

Scribe: Vasilis Charisopoulos (vc333@cornell.edu)

February 19, 2020

Theory for the randomized SVD

Recall some notation from the previous lecture: given $A \in \mathbb{R}^{m \times n}$, we defined

$$Y = A\Omega, \quad \Omega \in \mathbb{R}^{n \times (k+p)}, \quad (\Omega)_{ij} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1). \tag{1}$$

and denoted P_Y for the spectral projector to the range of Y. Moreover, recall that we denoted

$$\Omega_1 := V_1^\top \Omega, \; \Omega_2 := V_2^\top \Omega, \tag{2}$$

where V_1, V_2 are the right singular vectors of A. Notice that because V_1, V_2 are orthogonal to each other, Ω_1 and Ω_2 are independent; moreover, since V_1, V_2 have orthonormal columns, Ω_1 and Ω_2 have i.i.d. $\mathcal{N}(0,1)$ elements as well.

We previously showed the following Lemma:

Lemma 1. *Assuming that Ω has full row rank, it holds that*

$$\|(I - P_Y)A\|_2^2 \leq \|\Sigma_2\|_2^2 + \|\Sigma_2 \Omega_2 \Omega_1^+\|_2^2, \tag{3}$$

*where X^+ denotes the **pseudoinverse** of the matrix X.*

To move forward, we need to understand $\Omega_2\Omega_1^+$ when Ω has i.i.d. standard normal entries. In that case, $V_1^\top \Omega$ and $V_2^\top \Omega$ also consist of i.i.d. Gaussian elements, by properties of orthogonal transforms of Gaussians. This would also hold for any choice of so-called **isotropic** distribution; for example, we could have sampled the columns of Ω from Rademacher random vectors. See (Vershynin, 2018, Chapter 3) for a discussion about isotropic random vectors.

We now introduce some technical results necessary for the remainder of the proof:

Proposition 2. *Fix S, T and draw a sample G with $G_{ij} \overset{i.i.d}{\sim} \mathcal{N}(0,1)$. Then:*

$$\left(\mathbb{E}\left[\|SGT\|_F^2\right]\right)^{1/2} = \|S\|_F \|T\|_F, \tag{4}$$

$$\mathbb{E}\left[\|SGT\|_2\right] \leq \|S\|_2 \|T\|_F + \|S\|_F \|T\|_2. \tag{5}$$

Theorem 3. *Let $G \in \mathbb{R}^{k \times (k+p)}$ have i.i.d. $\mathcal{N}(0,1)$ entries. Then its pseudoinverse satisfies:*

$$\left(\mathbb{E}\left[\|G^+\|_F^2\right]\right)^{1/2} = \sqrt{\frac{k}{p-1}} \tag{6}$$

$$\mathbb{E}\left[\|G^+\|_2\right] \leq \frac{e\sqrt{k+p}}{p} \tag{7}$$

The following theorem about Lipschitz concentration is standard (Vershynin, 2018, Chapter 5):

Theorem 4. *Let $h : \mathbb{R}^{m_1 \times m_2} \to \mathbb{R}$ satisfy $|h(X) - h(Y)| \leq L \|X - Y\|_F$, $\forall X, Y$. Then, if G is a Gaussian random matrix, we have*

$$\mathbb{P}\left(h(G) \geq \mathbb{E}\left[h(G)\right] + Lt\right) \leq \exp\left\{-\frac{t^2}{2}\right\}. \tag{8}$$

Finally, we need a tail bound for the Frobenius and spectral norms of the pseudoinverse:

Proposition 5. *For $G \in \mathbb{R}^{k \times (k+p)}$ with $G_{ij} \overset{i.i.d}{\sim} \mathcal{N}(0,1)$ and $p \geq 4$, we have*

$$\mathbb{P}\left(\|G^+\|_F \geq \sqrt{\frac{3k}{p+1}} t\right) \leq t^{-p}, \ \forall t \geq 1 \tag{9}$$

$$\mathbb{P}\left(\|G^+\|_2 \geq \frac{e\sqrt{k+p}}{p+1} t\right) \leq t^{-(p+1)}, \ \forall t \geq 1 \tag{10}$$

We are now in good shape to prove the desired statement. In fact, we will prove something stronger, as shown below:

Theorem 6. *For Y defined as in* (1), *we have*

$$\begin{aligned}\mathbb{E}\left[\|(I - P_Y)A\|_2\right] &\leq \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j>k}\sigma_j^2\right)^{1/2} \\ &= \left(1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{k+p}}{p}\sqrt{sr(\Sigma_2)}\right)\sigma_{k+1},\end{aligned} \tag{11}$$

*where $sr(B)$ denotes the **stable rank** of B, defined as*

$$sr(B) := \frac{\|B\|_F^2}{\|B\|_2^2} = \sum_{i=1}^{n}\left(\frac{\sigma_i}{\sigma_1}\right)^2 \tag{12}$$

Proof. Starting from our deterministic bound, we apply the expectation operator and obtain

$$\mathbb{E}\left[\|(I - P_Y)A\|_2\right] \leq \mathbb{E}\left[\left(\sigma_{k+1}^2 + \|\Sigma_2\Omega_2\Omega_1^+\|_2^2\right)^{1/2}\right] \qquad \leq \sigma_{k+1} + \mathbb{E}\left[\|\Sigma_2\Omega_2\Omega_1^+\|_2^2\right] \tag{13}$$

where (13) is simply $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$. We proceed by analyzing the second term above. We use the tower property of expectation and independence of Ω_1 and Ω_2 to write

$$\mathbb{E}\left[\|\Sigma_2\Omega_2\Omega_1^+\|_2\right] = \mathbb{E}_{\Omega_1}\left[\mathbb{E}\left[\|\Sigma_2\Omega_2\Omega_1^+\|_2 \mid \Omega_1\right]\right] \overset{(\sharp)}{\leq} \|\Sigma_2\|_2\,\mathbb{E}\left[\|\Omega_1^+\|_F\right] + \|\Sigma_2\|_F\,\mathbb{E}\left[\|\Omega_1^+\|_2\right] \tag{14}$$

where ($\sharp$) follows from Proposition 2. Applying Theorem 3 to upper bound the expectation yields the claim. □

Remark 1. The stable rank of a matrix is intimately related to the concept of *statistical dimension* or *stable dimension*. In particular, the algebraic dimension of a mathematical object may change abruptly even under small perturbations (e.g. adding gaussian noise εZ to a low rank matrix), but the stable dimension changes smoothly as a function of the perturbation. An excellent treatment is available in (Vershynin, 2018, Chapter 7).

We can also prove a corresponding high probability statement.

Theorem 7. *For all $t, u \geq 1$, we have*

$$\|(I - P_Y)A\|_2 \leq \left[1 + \sqrt{\frac{3k}{p+1}}t + t\frac{e\sqrt{k+p}}{p+1}\sqrt{\mathsf{sr}(\Sigma_2)}\right] \cdot \sigma_{k+1} + ut\frac{e\sqrt{k+p}}{p+1}\sigma_{k+1} \tag{15}$$

with probability at least $1 - 2t^{-p} - e^{-\frac{u^2}{2}}$. In particular, setting $t = p$ and $u = \sqrt{2p\log p}$ gives probability of success at least $1 - 3p^{-p}$, as stated in the previous lecture.

Proof. The outline of the proof technique here is as follows: we have a "bad" event $\mathcal{B}$ that we want to control (i.e. $\|(I - P_Y)A\|_2$ being large), which depends on a deterministic quantity $\|\Sigma_2\|_2$ and a random quantity $\|\Sigma_2\Omega_2\Omega_1^+\|_2$. To control the latter, we can see that Ω_1^+ is "well-behaved" (call this event $\mathcal{E}_1$); then, for a convenient choice of parameters, it is easy to see that the contrary happens with very small probability. If $\mathcal{B}$ denotes the "bad" event where $\|\Sigma_2\Omega_2\Omega_1^+\|_2$ is not small, observe that

$$\mathbb{P}(\mathcal{B}) = \mathbb{P}(\mathcal{B} \cap \mathcal{E}) + \mathbb{P}(\mathcal{B} \cap \mathcal{E}^c) = \mathbb{P}(\mathcal{B} \mid \mathcal{E}) \cdot \underbrace{\mathbb{P}(\mathcal{E})}_{\leq 1} + \underbrace{\mathbb{P}(\mathcal{B} \mid \mathcal{E}^c)}_{\leq 1} \cdot \mathbb{P}(\mathcal{E}^c). \tag{16}$$

The remainder of the proof is devoted to control of the above.

- for any $t \geq 1$, we define the event $\mathcal{E}_t$ as

$$\mathcal{E}_t := \left\{ \|\Omega_1^+\|_2 \leq \frac{e\sqrt{k+p}}{p+1}t \text{ and } \|\Omega_1^+\|_F \leq \sqrt{\frac{3k}{p+1}}t \right\} \tag{17}$$

  We can then see that

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_t^c) &= \mathbb{P}\left(\left\{\|\Omega_1^+\|_2 \geq \frac{e\sqrt{k+p}}{p+1}t\right\} \cup \left\{\|\Omega_1^+\|_F \geq \sqrt{\frac{3k}{p+1}}t\right\}\right) &(18)\\
&\overset{(*)}{\leq} \mathbb{P}\left(\|\Omega_1^+\|_2 \geq \frac{e\sqrt{k+p}}{p+1}t\right) + \mathbb{P}\left(\|\Omega_1^+\|_F \geq \sqrt{\frac{3k}{p+1}}t\right) &(19)\\
&\overset{(\sharp)}{\leq} t^{-p} + t^{-(p+1)} \leq 2t^{-p}, &(20)
\end{aligned}$$

  where $(*)$ is simply the union bound, and $(\sharp)$ follows by Proposition 5.

- let $h(X) := \|\Sigma_2 X \Omega_1^+\|_2$. Then it follows that

$$|h(X) - h(Y)| \leq \|\Sigma_2\|_2 \|\Omega_1^+\|_2 \|X - Y\|_F,$$

  i.e. h is Lipschitz with $L := \|\Sigma_2\|_2 \|\Omega_1^+\|_2$. This is true since

$$\begin{aligned}
|h(X) - h(Y)| &= |\|\Sigma_2 X\Omega_1^+\|_2 - \|\Sigma_2 Y \Omega_1^+\|_2| \overset{(\sharp)}{\leq} \|\Sigma_2 X\Omega_1^+ - \Sigma_2 Y\Omega_1^+\|_2\\
&= \|\Sigma_2(X - Y)\Omega_1^+\|_2 \leq \|\Sigma_2\|_2 \|X - Y\|_2 \|\Omega_1^+\|_2\\
&\leq \|\Sigma_2\|_2 \|\Omega_1^+\|_2 \|X - Y\|_F,
\end{aligned}$$

  where $(\sharp)$ is just the (reverse) triangle inequality for norms, the penultimate inequality follows from the submultiplicative property of the spectral norm and the last inequality is $\|A\|_2 \leq \|A\|_F$.

We can now apply Theorem 4 *conditionally* on $\mathcal{E}_t$ for $h(X)$ defined as above; using Eq. (14)), we obtain:

$$\mathbb{P}\left(h(X) \geq \|\Sigma_2\|_2 \sqrt{\frac{3k}{p+1}}t + \|\Sigma_2\|_F \frac{e\sqrt{k+p}}{p+1}t + \|\Sigma_2\|_2 \frac{e\sqrt{k+p}}{p+1}tu \;\middle|\; \mathcal{E}_t\right) \tag{21}$$

$$\leq \mathbb{P}\left(h(X) \geq \mathbb{E}\left[h(X)\right] + \|\Sigma_2\|_2 \|\Omega_1^+\|_2 \, u \mid \mathcal{E}_t\right) \leq \exp\left\{-\frac{u^2}{2}\right\}, \tag{22}$$

where we used the fact that if $\beta \leq \alpha$, then $\mathbb{P}(X \geq \alpha) \leq \mathbb{P}(X \geq \beta)$. Combining with the bound for $\mathbb{P}(\mathcal{E}_t^c)$ gives probability of failure at most $2t^{-p} + \exp\left\{-\frac{u^2}{2}\right\}$, completing the proof after substituting the bound for $h(X)$ into the deterministic expression (3). □

Extensions to other random matrices

It is possible to extend the above analysis to other types of random matrices, including structured random matrices enabling quick application. At the cost of an extra $\log k$ term, one can use the Subsampled Random Fourier Transform (SRFT) from Woolfe et al. (2008), which is defined as

$$\Omega := \sqrt{\frac{n}{\ell}} DFR^\top, \tag{23}$$

with $D = \text{diag}(d_1, \ldots, d_n)$ and d_i are uniformly sampled from the unit complex circle. If we wish to remain in the real world, the Subsampled Random Hadamard Transform (SRHT) replaces D with random signs and F with the $n \times n$ Hadamard matrix H_n (assuming for simplicity that $n = 2^p$ for some p). Both F and H can be applied from the right in time $\mathcal{O}(mn \log n)$; the former is possible via the FFT, while the latter is via the Fast Walsh-Hadamard Transform. The analysis for the SRHT appears in Tropp (2011); see also Drineas and Mahoney (2018). The aforementioned structured random matrices find applications in fields extending beyond numerical linear algebra; the SRHT has been used as a design matrix in compressed sensing Do et al. (2011) and adjacent fields such as phase retrieval Duchi and Ruan (2019).

Here, we show how to adapt Theorem 7 to the SRHT sketching matrix. First, we need the following theorem from Tropp (2011):

Theorem 8 (SRHT guarantees). *Let r be the target rank, and draw an $n \times \ell$ SRHT matrix Ω, where ℓ satisfies*

$$n \geq \ell \gtrsim \sqrt{r \log n}.$$

*Then, for a **fixed** orthonormal matrix V, it holds that*

$$\mathbb{P}\left(\left\{\sigma_1(V^\top \Omega) \leq 1.48\right\} \cap \left\{\sigma_r(V^\top \Omega) \geq 0.40\right\}\right) \geq 1 - c_1 r^{-1}$$

for a universal constant $c_1 > 0$.

Then it suffices to note that for any A, it holds that $\|A^+\|_2 = \sigma_{\min}^{-1}(A)$, so $\sigma_r(V^\top \Omega) \geq 0.40$ implies that $\|\Omega_1^+\|_2 \leq \frac{1}{0.40}$. Similarly, $\|\Omega_2\|_2 \leq 1.48$, and combining with the deterministic bound from (3) yields the proof for SRHT matrices.

References

T. T. Do, L. Gan, N. H. Nguyen, and T. D. Tran. Fast and efficient compressive sensing using structurally random matrices. *IEEE Transactions on signal processing*, 60(1):139–154, 2011.

P. Drineas and M. W. Mahoney. Lectures on randomized numerical linear algebra. *The Mathematics of Data*, 25:1, 2018.

J. C. Duchi and F. Ruan. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *Information and Inference: A Journal of the IMA*, 8(3):471–529, 2019.

J. A. Tropp. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(01n02):115–126, 2011.

R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert. A fast randomized algorithm for the approximation of matrices. *Applied and Computational Harmonic Analysis*, 25(3):335 – 366, 2008. ISSN 1063-5203. doi: https://doi.org/10.1016/j.acha.2007.12.002. URL `http://www.sciencedirect.com/science/article/pii/S1063520307001364`.