Smoothed Analysis of Gaussian Elimination

by

Arvind Sankar

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2004

Author . . . .
Department of Mathematics
January 16, 2004

Certified by . . . .
Daniel A. Spielman
Associate Professor of Applied Mathematics
Thesis Supervisor

Accepted by . . . .
Rodolfo Ruben Rosales
Chairman, Applied Mathematics Committee

Accepted by . . . .
Pavel I. Etingof
Chairman, Department Committee on Graduate Students

Smoothed Analysis of Gaussian Elimination

by
Arvind Sankar

Submitted to the Department of Mathematics
on January 16, 2004, in partial fulfillment of the
requirements for the degree of
DOCTOR OF PHILOSOPHY

Abstract

We present a smoothed analysis of Gaussian elimination, both with partial pivoting and without pivoting. Let $\bar{A}$ be any matrix and let A be a slight random perturbation of $\bar{A}$. We prove that it is unlikely that A has large condition number. Using this result, we prove it is unlikely that A has large growth factor under Gaussian elimination without pivoting. By combining these results, we bound the smoothed precision needed to perform Gaussian elimination without pivoting. Our results improve the average-case analysis of Gaussian elimination without pivoting performed by Yeung and Chan (SIAM J. Matrix Anal. Appl., 1997).

We then extend the result on the growth factor to the case of partial pivoting, and present the first analysis of partial pivoting that gives a sub-exponential bound on the growth factor. In particular, we show that if the random perturbation is Gaussian with variance σ^2, then the growth factor is bounded by $(n/\sigma)^{\mathcal{O}(\log n)}$ with very high probability.

Thesis Supervisor: Daniel A. Spielman
Title: Associate Professor of Applied Mathematics

Credits

The work in this thesis is joint with my advisor, Dr. Daniel Spielman, and with Dr. Shang-Hua Teng of Boston University.

Acknowledgments

This thesis and the work that led to it would not have been possible without the constant support and encouragement of my advisor, Dan Spielman, and the many fruitful discussions I had with Dan and Shang-Hua Teng.

I thank the friends who have been an extended family to me through my years at the Institute. I would have been lost without them.

I thank the Institute, for challenging me.

I thank IIT Bombay, for a premier educational experience.

I thank my parents, without whom I would not be here today.

Contents

Chapter 1

Introduction

In this thesis, we present an analysis of the stability of Gaussian elimination, both without pivoting and with partial pivoting. The analysis is carried out under the smoothed model of complexity, as presented in [21]. We thus hope to explain the experimental observation that Gaussian elimination is stable in practice, even though it is extremely unstable in the worst case.

In the remainder of this chapter, we introduce the Gaussian elimination algorithm along with associated definitions of condition number and growth factors, and describe what is meant by a smoothed analysis of this algorithm.

1.1 Gaussian Elimination

Gaussian elimination is one of the simplest and perhaps the oldest numerical algorithm. It can be looked at in two slightly different but equivalent ways. One emphasizes the solution of the linear system of equations

$$A\mathbf{x} = \mathbf{b}$$

and the other the LU-factorization of the coefficient matrix

$$A = LU$$

into a lower triangular matrix with unit diagonal, and an upper triangular matrix.

The algorithm consists of choosing one of the equations and one of the variables, and using this equation to eliminate the variable from the remaining equations, thus giving a smaller system to which Gaussian elimination may be applied recursively. The choice of equation and variable is determined by which *pivoting rule* is being applied.

The simplest case is when no pivoting is done, when the first equation and first variable are chosen to be eliminated first.

The most commonly used pivoting rule is called partial pivoting, and it chooses the variables in order, but at each step to pick the equation that has the largest coefficient (in absolute value) of the variable to be eliminated. This leads to a matrix L in which all entries have absolute value at most 1.

A third pivoting rule is to choose the largest coefficient among the whole system, and eliminate using the variable and equation to which it corresponds. This is known as complete pivoting, and its worst case stability is provably better than that of partial pivoting. In spite of this, it is not commonly used as it requires twice as many floating point operations, and partial pivoting is usually stable enough.

It should be noted that the equation

$$A = LU$$

corresponds to no pivoting. For a general pivoting rule, the equation must be rewritten as

$$PAQ = LU$$

where P and Q are permutation matrices. Partial pivoting corresponds to $Q = I$ and $|L_{ij}| \leq 1$, while complete pivoting can be defined by

$$|L_{ij}| \leq 1 \text{ and } |U_{ij}| \leq |U_{ii}|$$

1.1.1 Error analysis: condition number and growth factors

Wilkinson [24] showed that the relative error when a linear system is solved using Gaussian elimination satisfies

$$\frac{\|\bar{x} - x\|}{\|x\|} \leq n^{\mathcal{O}(1)} \kappa(A) \rho_L(A) \rho_U(A) \epsilon$$

where $\kappa(A)$ is the *condition number* of A, $\rho_L(A)$ and $\rho_U(A)$ are the *growth factors*, ϵ is the machine precision, and the polynomial factor depends on the norms in which the condition numbers and growth factors are defined.

The condition number is an intrinsic property of the matrix, being defined as

$$\kappa(A) = \|A\| \left\|A^{-1}\right\|$$

and it measures how much the solution to the system $A\mathbf{x} = \mathbf{b}$ changes when there are slight changes in A or $\mathbf{b}$. Any technique to solve the system will incur this error.

The growth factors are a contribution to error that is specific to Gaussian elimination. They are defined as

$$\rho_L(A) = \|L\| \text{ and } \rho_U(A) = \frac{\|U\|}{\|A\|}$$

They measure how large intermediate entries become as Gaussian elimination is carried out. Partial pivoting eliminates the growth in L, since its entries remain bounded. However, $\rho_U(A)$ can grow exponentially with n in the worst case. In fact, Wilkinson showed that a tight bound on $\rho_U(A)$ with the $\mathtt{max}$-norm is 2^{n-1}. On the other hand, it is observed in practice that $\rho_U(A)$ is extremely well-behaved: for random matrices it grows sublinearly [23]. We will give a partial explanation for this behaviour in Chapter 3.

1.2 Smoothed analysis

Spielman and Teng [21], introduced the smoothed analysis of algorithms as a means of explaining the success of algorithms and heuristics that could not be well understood through traditional worst-case and average-case analyses. Smoothed analysis is a hybrid of worst-case and average-case analyses in which one measures the maximum over inputs of the expected value of a function on slight random perturbations of that input. For example, the smoothed complexity of an algorithm is the maximum over its inputs of the expected running time of the algorithm under slight perturbations of that input. If an algorithm has low smoothed complexity and its inputs are subject to noise, then it is unlikely that one will encounter an input on which the algorithm performs poorly. (See also the Smoothed Analysis Homepage [1])

Smoothed analysis is motivated by the existence of algorithms and heuristics that are known to work well in practice, but which are known to have poor worst-case performance. Average-case analysis was introduced in an attempt to explain the success of such heuristics. However, average-case analyses are often unsatisfying as the random inputs they consider may bare little resemblance to the inputs actually encountered in practice. Smoothed analysis attempts to overcome this objection by proving a bound that holds in every neighborhood of inputs.

1.3 Our results

For our analysis of Gaussian elimination, the model we use is that the input matrix A has additive Gaussian noise. In other words,

$$A = \bar{A} + \sigma G$$

where $\bar{A}$ is the "true" value of A, and σG represents noise. The matrix G is assumed to be composed of independent standard normal variables, that is, $G \sim \mathfrak{N}(0, I \otimes I)$.

We prove that perturbations of arbitrary matrices are unlikely to have large condition numbers or large growth factors under Gaussian elimination, both without pivoting in Chapter 2 and with partial pivoting in Chapter 3. In particular, we show that

$$\Pr[\kappa(A) \geq x] \leq \frac{9.4n\left(1+\sqrt{\log(x)/2n}\right)}{x\sigma}$$

$$\Pr[\rho_L(A) > x] \leq \sqrt{\frac{2}{\pi}}\frac{n^2}{x}\left(\frac{\sqrt{2}}{\sigma} + \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)$$

$$\Pr[\rho_U(A) > 1 + x] \leq \frac{1}{\sqrt{2\pi}}\frac{n(n+1)}{x\sigma}$$

for Gaussian elimination without pivoting.

For partial pivoting, we prove

$$\Pr_A[\rho_U(A) > x] \leq \left(\frac{1}{x}\left(\mathcal{O}\left(\frac{n(1+\sigma\sqrt{n})}{\sigma}\right)\right)^{12\log n}\right)^{-\frac{1}{21}\log n}$$

This is the first sub-exponential bound on the growth of partial pivoting, even in the average case. Hence we feel that the result is important, even though the bound of $(n/\sigma)^{\mathcal{O}(\log n)}$ it establishes on the growth remains far from experimental observations.

Chapter 2

Smoothed Analysis of Gaussian Elimination without Pivoting

2.1 Introduction

In this chapter, we consider the growth factor of Gaussian elimination when no pivoting is done. Since the matrix A has a Gaussian distribution, the event that a pivot is exactly zero (in which case the LU-factorization fails) occurs with probability zero. We will show that the growth factors ρ_U and ρ_L have tail distributions $\mathcal{O}(1/x)$, that is,

$$\Pr_A [\rho_{L,U}(A) > x] = \mathcal{O}\left(\frac{1}{x}\right)$$

We are able to show that the condition number of A has tail distribution

$$\Pr_A [\kappa(A) > x] = \mathcal{O}\left(\frac{\log x}{x}\right)$$

a slightly weaker bound.

The remaining sections are organized as follows: in Section 2.2, we bound the tail of the condition number. This section also contains the heart of the arguments we make, Theorem 2.2 on the distribution of the smallest singular value of a non-central Gaussian random matrix. We then bound the growth factor, in U and in L, and combine these three results to give a bound on the expected precision of Gaussian elimination without pivoting. We then extend the analysis to the case when A is symmetric, and certain entries are known to be zero, in Section 2.5.

2.2 Smoothed analysis of the condition number of a matrix

In his paper, "The probability that a numerical analysis problem is difficult", Demmel [7] proved that it is unlikely that a Gaussian random matrix centered at the origin has large condition number. Demmel's bounds on the condition number were improved by Edelman [10]. In this section, we present the smoothed analogue of this bound. That is, we show that for every matrix it is unlikely that a slight perturbation of that matrix has large condition number. For more information on the condition number of a matrix, we refer the reader to one of [12, 22, 8]. As bounds on the norm of a random matrix are standard, we focus on the norm of the inverse. Recall that $1/\left\|A^{-1}\right\| = \min_{x} \|Ax\| / \|x\|$.

The first step in the proof is to bound the probability that $\left\|A^{-1}v\right\|$ is small for a fixed unit vector v. This result is also used later (in Section 2.3.1) in studying the growth factor. Using this result and an averaging argument, we then bound the probability that $\left\|A^{-1}\right\|$ is large.

Lemma 2.1 (Projection of A^{-1}). Let $\bar{A}$ be an arbitrary square matrix in $\mathbb{R}^{n\times n}$, and A a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance σ^2. Let v be an arbitrary unit vector. Then

$$\Pr\left[\left\|A^{-1}v\right\| > x\right] < \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma}$$

Proof. First observe that by multiplying A by an orthogonal matrix, we may assume that $v = e_1$. In this case,

$$\left\|A^{-1}v\right\| = \left\|(A^{-1})_{:,1}\right\|,$$

the length of the first column of A^{-1}. The first column of A^{-1}, by the definition of the matrix inverse, is a vector orthogonal to $A_{2:n,:}$, *i.e.*, every row but the first. Also, it has inner product 1 with the first row. Hence its length is the reciprocal of the length of the projection of the first row onto the subspace orthogonal to the rest of the rows. This projection is a 1-dimensional Gaussian random variable of variance σ^2, and the probability that it is smaller than $1/x$ in absolute value is at most

$$\frac{1}{\sigma\sqrt{2\pi}}\int_{-1/x}^{1/x} e^{-t^2/2\sigma^2}\, dt \le \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma},$$

which completes the proof. □

Theorem 2.2 (Smallest singular value). Let $\bar{A}$ be an arbitrary square matrix in $\mathbb{R}^{n\times n}$, and A a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance σ^2. Then

$$\Pr\left[\left\|A^{-1}\right\| \ge x\right] \le 2.35\frac{\sqrt{n}}{x\sigma}$$

Proof. We apply Lemma 2.1 to a uniformly distributed random unit vector v and obtain

$$\Pr_{A,v}\left[\left\|A^{-1}v\right\| \ge x\right] \le \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma} \qquad (2.2.1)$$

Now let u be the unit vector such that $\left\|A^{-1}u\right\| = \left\|A^{-1}\right\|$ (this is unique with probability 1). From the inequality

$$\left\|A^{-1}v\right\| \ge \left\|A^{-1}\right\| \left|\langle u, v\rangle\right|,$$

we have that for any $c > 0$,

$$\Pr_{A,v}\left[\left\|A^{-1}v\right\| \geq x\sqrt{\frac{c}{n}}\right] \geq \Pr_{A,v}\left[\left\|A^{-1} \geq x\right\| \text{ and } |\langle u|v\rangle| \geq \sqrt{\frac{c}{n}}\right]$$
$$= \Pr_{A}\left[\left\|A^{-1} \geq x\right\|\right] \Pr_{A,v}\left[|\langle u|v\rangle| \geq \sqrt{\frac{c}{n}}\right].$$

So,

$$\begin{aligned}\Pr_{A}\left[\left\|A^{-1}\right\| \geq x\right] &\leq \frac{\Pr_{A,v}\left[\left\|A^{-1}v\right\| \geq x\sqrt{\frac{c}{n}}\right]}{\Pr_{A,v}\left[|\langle u|v\rangle| \geq \sqrt{\frac{c}{n}}\right]} \\ &\leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{n}}{x\sigma\sqrt{c}\Pr_{A,v}\left[|\langle u|v\rangle| \geq \sqrt{\frac{c}{n}}\right]} && \text{(by (2.2.1))} \\ &\leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{n}}{x\sigma\sqrt{c}\Pr_{g}\left[|g| \geq \sqrt{c}\right]}, && \text{(by Lemma A.5)}\end{aligned}$$

where g is a standard normal variable. Choosing $c = 0.57$, and evaluating the error function numerically, we get

$$\Pr_{A}\left[\left\|A^{-1}\right\| \geq x\right] \leq 2.35\frac{\sqrt{n}}{x\sigma}.$$

□

Theorem 2.3 (Condition number). Let $\bar{A}$ be an $n \times n$ matrix satisfying $\bar{A} \leq \sqrt{n}$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. Then,

$$\Pr\left[\kappa(A) \geq x\right] \leq \frac{9.4n\left(1 + \sqrt{\log(x)/2n}\right)}{x\sigma}.$$

Proof. As observed by Davidson and Szarek [6, Theorem II.11], one can apply inequality (1.4) of [17] to show that for all $k \geq 0$,

$$\Pr\left[\left\|\bar{A} - A\right\| \geq \sqrt{n} + k\right] \leq e^{-k^2/2}.$$

We rephrase this bound as

$$\Pr\left[\left\|\bar{A} - A\right\| \geq \sqrt{n} + \sqrt{2\log(1/\epsilon)}\right] \leq \epsilon,$$

for all $\epsilon \leq 1$. By assumption, $\left\|\bar{A}\right\| \leq \sqrt{n}$; so,

$$\Pr\left[\|A\| \geq 2\sqrt{n} + \sqrt{2\log(1/\epsilon)}\right] \leq \epsilon.$$

From the result of Theorem 2.2, we have

$$\Pr\left[\left\|A^{-1}\right\| \geq \frac{2.35\sqrt{n}}{\epsilon\sigma}\right] \leq \epsilon.$$

Combining these two bounds, we find

$$\Pr\left[\|A\|\left\|A^{-1}\right\| \geq \frac{4.7n + 2.35\sqrt{2n\log(1/\epsilon)}}{\epsilon\sigma}\right] \leq 2\epsilon.$$

We would like to express this probability in the form of $\Pr\left[\|A\|\left\|A^{-1}\right\| \geq x\right]$, for $x \geq 1$. By substituting

$$x = \frac{4.7n + 2.35\sqrt{2n\log(1/\epsilon)}}{\epsilon\sigma},$$

we observe that

$$2\epsilon = \frac{2\left(4.7n + 2.35\sqrt{2n\log(1/\epsilon)}\right)}{x\sigma} \leq \frac{9.4n\left(1 + \sqrt{\log(x)/2n}\right)}{x\sigma}$$

for

$$1 \leq \frac{9.4n\left(1 + \sqrt{\log(x)/2n}\right)}{\sigma}$$

which holds here, since $\sigma \leq 1$.

Therefore, we conclude

$$\Pr\left[\|A\|\left\|A^{-1}\right\| \geq x\right] \leq \frac{9.4n\left(1 + \sqrt{\log(x)/2n}\right)}{x\sigma}.$$

□

We also conjecture that the $1 + \sqrt{\log(x)/2n}$ term should be unnecessary because those matrices for which $\|A\|$ is large are less likely to have $\left\|A^{-1}\right\|$ large as well.

Conjecture 1. *Let $\bar{A}$ be a $n \times n$ matrix satisfying $\|\bar{A}\|_{\max} \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. Then,*

$$\Pr\left[\kappa(A) \geq x\right] \leq \mathcal{O}\left(n/x\sigma\right).$$

2.3 Growth Factor of Gaussian Elimination without Pivoting

We now turn to proving a bound on the growth factor. With probability 1, none of the diagonal entries that occur during elimination will be 0. So, in the spirit of Yeung and Chan, we analyze the growth factor of Gaussian elimination without pivoting. When we specialize our smoothed analyses to the case $\bar{A} = 0$, we improve the bounds of Yeung and Chan by a factor of n. Our improved bound on ρ_U agrees with their experimental analyses.

2.3.1 Growth in U

We recall that

$$\rho_U(A) = \frac{\|U\|_\infty}{\|A\|_\infty} = \max_i \frac{\|U_{i,:}\|_1}{\|A\|_\infty},$$

and so we need to bound the ℓ_1-norm of each row of U. We denote the upper triangular segment of the kth row of U by $\mathbf{u} = U_{k,k:n}$, and observe that $\mathbf{u}$ can be obtained from the formula:

$$\mathbf{u} = \mathbf{a}^T - \mathbf{b}^T C^{-1} D \tag{2.3.1}$$

where

$$\mathbf{a}^T = A_{k,k:n} \qquad \mathbf{b}^T = A_{k,1:k-1} \qquad C = A_{1:k-1,1:k-1} \qquad D = A_{1:k-1,k:n}.$$

This expression for $\mathbf{u}$ follows immediately from

$$A_{1:k,:} = \begin{pmatrix} C & D \\ \mathbf{b}^T & \mathbf{a}^T \end{pmatrix} = \begin{pmatrix} L_{1:k-1,1:k-1} & 0 \\ L_{k,1:k-1} & 1 \end{pmatrix} \begin{pmatrix} U_{1:k-1,1:k-1} & U_{1:k-1,k:n} \\ 0 & \mathbf{u} \end{pmatrix}$$

In this section, we give two bounds on $\rho_U(A)$. The first will have a better dependence on σ, and second will have a better dependence on n. It is the later bound, Theorem 2.6, that agrees with the experiments of Yeung and Chan [25] when specialized to the average-case.

First bound

Theorem 2.4 (First bound on $\rho_U(A)$). Let $\bar{A}$ be an $n \times n$ matrix satisfying $\|\bar{A}\| \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. Then,

$$\Pr\left[\rho_U(A) > 1 + x\right] \leq \frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{x\sigma}.$$

Proof. From (2.3.1),

$$\begin{aligned} \|\mathbf{u}\|_1 = \left\|\mathbf{a}^T - \mathbf{b}^T C^{-1} D\right\|_1 &\leq \left\|\mathbf{a}^T\right\|_1 + \left\|\mathbf{b}^T C^{-1} D\right\|_1 \\ &\leq \left\|\mathbf{a}^T\right\|_1 + \left\|\mathbf{b}^T C^{-1}\right\|_1 \|D\|_\infty \qquad (\text{as } \|D\|_\infty = \left\|D^T\right\|_1) \\ &\leq \|A\|_\infty \left(1 + \left\|\mathbf{b}^T C^{-1}\right\|_1\right) \end{aligned} \tag{2.3.2}$$

We now bound the probability $\left\|\mathbf{b}^T C^{-1}\right\|_1$ is large. Now,

$$\left\|\mathbf{b}^T C^{-1}\right\|_1 \leq \sqrt{k-1} \left\|\mathbf{b}^T C^{-1}\right\|_2$$

Therefore,

$$\begin{aligned} \Pr_{\mathbf{b},C}\left[\left\|\mathbf{b}^T C^{-1}\right\|_1 > x\right] &\leq \Pr_{\mathbf{b},C}\left[\left\|\mathbf{b}^T C^{-1}\right\|_2 > x/\sqrt{k-1}\right] \\ &\leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{k-1}\sqrt{(k-1)\sigma^2 + 1}}{x\sigma} \leq \sqrt{\frac{2}{\pi}} \frac{k}{x\sigma}, \end{aligned}$$

where the second inequality follows from Lemma 2.5 below and the last inequality follows from the assumption $\sigma^2 \leq 1$.

We now apply a union bound over the n rows of U to obtain

$$\Pr\left[\rho_U(A) > 1 + x\right] \leq \sum_{k=2}^{n} \sqrt{\frac{2}{\pi}} \frac{k}{x\sigma} \leq \frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{x\sigma}.$$

□

Lemma 2.5. Let $\bar{C}$ be an arbitrary square matrix in $\mathbb{R}^{d\times d}$, and C be a random matrix of independent Gaussian variables of variance σ^2 centered at $\bar{C}$. Let $\bar{\mathbf{b}}$ be a vector in $\mathbb{R}^d$ such that $\left\|\bar{\mathbf{b}}\right\|_2 \leq 1$, and let $\mathbf{b}$ be a random Gaussian vector of variance σ^2 centered at $\bar{\mathbf{b}}$. Then

$$\Pr_{\mathbf{b},C}\left[\left\|\mathbf{b}^T C^{-1}\right\|_2 \geq x\right] \leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{\sigma^2 d + 1}}{x\sigma}$$

Proof. Let $\hat{\mathbf{b}}$ be the unit vector in the direction of $\mathbf{b}$. By applying Lemma 2.1, we obtain for all $\mathbf{b}$,

$$\Pr_C\left[\left\|\mathbf{b}^T C^{-1}\right\|_2 > x\right] = \Pr_C\left[\left\|\hat{\mathbf{b}}^T C^{-1}\right\|_2 > \frac{x}{\|\mathbf{b}\|_2}\right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \|\mathbf{b}\|_2 .$$

Therefore, we have

$$\Pr_{\mathbf{b},C}\left[\left\|\mathbf{b}^T C^{-1}\right\|_2 > x\right] = \mathbf{E}_{\mathbf{b}}\left[\Pr_C\left[\left\|\mathbf{b}^T C^{-1}\right\|_2 > x\right]\right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \mathbf{E}_{\mathbf{b}}\left[\|\mathbf{b}\|_2\right].$$

It is known [16, p. 277] that $\mathbf{E}_{\mathbf{b}}\left[\|\mathbf{b}\|_2^2\right] \leq \sigma^2 d + \left\|\bar{\mathbf{b}}\right\|^2$. As $\mathbf{E}[X] \leq \sqrt{\mathbf{E}[X^2]}$ for every positive random variable X, we have $\mathbf{E}_{\mathbf{b}}\left[\|\mathbf{b}\|_2\right] \leq \sqrt{\sigma^2 d + \left\|\bar{\mathbf{b}}\right\|^2} \leq \sqrt{\sigma^2 d + 1}$. □

Second Bound for $\rho_U(A)$

In this section, we establish an upper bound on $\rho_U(A)$ which dominates the bound in Theorem 2.4 for $\sigma \geq n^{-3/2}$.

If we specialize the parameters in this bound to $\bar{A} = 0$ and $\sigma^2 = 1$, we improve the average-case bound proved by Yeung and Chan [25] by a factor of n. Moreover, the resulting bound agrees with their experimental results.

Theorem 2.6 (Second bound on $\rho_U(A)$). Let $\bar{A}$ be an $n \times n$ matrix satisfying $\left\|\bar{A}\right\| \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. For $n \geq 2$,

$$\Pr\left[\rho_U(A) > 1 + x\right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{2}{3} n^{3/2} + \frac{n}{\sigma} + \frac{4}{3} \frac{\sqrt{n}}{\sigma^2}\right)$$

Proof. We will first consider the case $k \leq n-1$. By (2.3.1), we have

$$\|u\|_1 \leq \|a\|_1 + \left\|b^T C^{-1} D\right\|_1 \leq \|a\|_1 + \sqrt{k-1}\left\|b^T C^{-1} D\right\|_2.$$

Therefore, for all $k \leq n-1$,

$$\begin{aligned}\frac{\|u\|_1}{\|A\|_\infty} &\leq \frac{\|a\|_1 + \sqrt{k-1}\left\|b^T C^{-1} D\right\|_2}{\|A\|_\infty} \\ &\leq 1 + \frac{\sqrt{k-1}\left\|b^T C^{-1} D\right\|_2}{\|A\|_\infty} \\ &\leq 1 + \frac{\sqrt{k-1}\left\|b^T C^{-1} D\right\|_2}{\|A_{n,:}\|_1}\end{aligned}$$

We now observe that for fixed b and C, $(b^T C^{-1})D$ is a Gaussian random vector of variance $\left\|b^T C^{-1}\right\|_2^2 \sigma^2$ centered at $(b^T C^{-1})\bar{D}$, where $\bar{D}$ is the center of D. We have $\left\|\bar{D}\right\|_2 \leq \left\|\bar{A}\right\|_2 \leq 1$, by the assumptions of the theorem; so,

$$\left\|b^T C^{-1} \bar{D}\right\|_2 \leq \left\|b^T C^{-1}\right\|_2 \left\|\bar{D}\right\|_2 \leq \left\|b^T C^{-1}\right\|_2.$$

Thus, if we let $t = (b^T C^{-1} D)/\left\|b^T C^{-1}\right\|_2$, then for any fixed b and C, t is a Gaussian random vector of variance σ^2 centered at a point of norm at most 1. We also have

$$\Pr_{b,C,D}\left[\left\|b^T C^{-1} D\right\|_2 \geq x\right] = \Pr_{b,C,t}\left[\left\|b^T C^{-1}\right\|_2 \|t\|_2 \geq x\right].$$

It follows from Lemma 2.5 that

$$\Pr_{b,C}\left[\left\|b^T C^{-1}\right\|_2 \geq x\right] \leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{\sigma^2(k-1)+1}}{x\sigma}$$

Hence, we may apply Corollary A.10 to show

$$\Pr_{b,C,t}\left[\left\|b^T C^{-1}\right\|_2 \|t\|_2 \geq x\right] \leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{\sigma^2(k-1)+1}\sqrt{\sigma^2(n-k+1)+1}}{x\sigma}$$

Note that $A_{n,:}$ is a Gaussian random vector in $\mathbb{R}^n$ of variance σ^2. As $A_{n,:}$ is independent of b, C and D, we can again apply Lemma A.9 to show

$$\begin{aligned}\Pr\left[\frac{\sqrt{k-1}\left\|b^T C^{-1} D\right\|_2}{\|A_{n,:}\|_1} \geq x\right] &\leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{k-1}\sqrt{\sigma^2(k-1)+1}\sqrt{\sigma^2(n-k+1)+1}}{x\sigma}\times \\ &\quad \times \mathbf{E}\left[\frac{1}{\|A_{n,:}\|_1}\right] \\ &\leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{k-1}\left(1+\frac{n\sigma^2}{2}\right)}{x\sigma}\frac{2}{n\sigma},\end{aligned}$$

by Lemma A.4.

From the proof of Theorem 2.4, we have that for $k = n$

$$\Pr\left[\|\mathbf{u}\|_1 / \|A\|_\infty > 1 + x\right] \leq \sqrt{\frac{2}{\pi}} \frac{n}{x\sigma}. \tag{2.3.3}$$

Applying a union bound over the choices for k, we obtain

$$\begin{aligned}
\Pr\left[\rho_U(A) > 1 + x\right] &\leq \sum_{k=2}^{n-1} \sqrt{\frac{2}{\pi}} \frac{\sqrt{k-1}\left(1 + \frac{n\sigma^2}{2}\right)}{x\sigma} \frac{2}{n\sigma} + \sqrt{\frac{2}{\pi}} \frac{n}{x\sigma} \\
&\leq \sqrt{\frac{2}{\pi}} \frac{2}{3} \frac{\sqrt{n}\left(\frac{2}{\sigma^2} + n\right)}{x} + \sqrt{\frac{2}{\pi}} \frac{n}{x\sigma} \\
&= \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{2}{3} n^{3/2} + \frac{n}{\sigma} + \frac{4}{3} \frac{\sqrt{n}}{\sigma^2}\right)
\end{aligned}$$

□

2.3.2 Growth in L

Let L be the lower-triangular part of the LU-factorization of A. We have

$$L_{(k+1):n,k} = A^{(k-1)}_{(k+1):n,k} / A^{(k-1)}_{k,k},$$

where we let $A^{(k)}$ denote the matrix remaining after the first k columns have been eliminated. We will show that it is unlikely that $\left\|L_{(k+1):n,k}\right\|_\infty$ is large by proving that it is unlikely that $\left\|A^{(k-1)}_{(k+1):n,k}\right\|_\infty$ is large while $\left|A^{(k-1)}_{k,k}\right|$ is small.

Theorem 2.7 ($\rho_L(A)$)**.** Let $\bar{A}$ be an n-by-n matrix for which $\|A\| \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. Then,

$$\Pr\left[\rho_L(A) > x\right] \leq \sqrt{\frac{2}{\pi}} \frac{n^2}{x} \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2 \log n} + \frac{1}{\sqrt{2\pi} \log n}\right)$$

Proof. We have

$$\begin{aligned}
L_{(k+1):n,k} &= \frac{A^{(k-1)}_{(k+1):n,k}}{A^{(k-1)}_{k,k}} \\
&= \frac{A_{(k+1):n,k} - A_{(k+1):n,1:(k-1)} A^{-1}_{1:(k-1),1:(k-1)} A_{1:(k-1),k}}{A_{k,k} - A_{k,1:(k-1)} A^{-1}_{1:(k-1),1:(k-1)} A_{1:(k-1),k}} \\
&= \frac{A_{(k+1):n,k} - A_{(k+1):n,1:(k-1)} \boldsymbol{v}}{A_{k,k} - A_{k,1:(k-1)} \boldsymbol{v}}
\end{aligned}$$

where we let $\boldsymbol{v} = A^{-1}_{1:(k-1),1:(k-1)} A_{1:(k-1),k}$. Since $\|\bar{A}\| \leq 1$, and all the terms $A_{(k+1):n,k}$, $A_{(k+1):n,1:(k-1)}$, $A_{k,k}$, $A_{k,1:(k-1)}$ and $\boldsymbol{v}$ are independent, we can apply Lemma 2.8 to show that

$$\Pr\left[\left\|L_{(k+1):n,k}\right\|_\infty > x\right] \leq \sqrt{\frac{2}{\pi}}\frac{1}{x}\left(\frac{\sqrt{2}}{\sigma} + \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)$$

The theorem now follows by applying a union bound over the n choices for k and observing that $\|L\|_\infty$ is at most n times the largest entry in L. □

Lemma 2.8 (Vector Ratio). Let

- a be a Gaussian random variable of variance σ^2 with mean $\bar{a}$ of absolute value at most 1,
- $\mathbf{b}$ be a Gaussian random d-vector of variance σ^2 centered at a point $\bar{\mathbf{b}}$ of norm at most 1,
- $\boldsymbol{x}$ be a Gaussian random n-vector of variance σ^2 centered at a point of norm at most 1,
- Y be a Gaussian random n-by-d matrix of variance σ^2 centered at a matrix of norm at most 1, and
- let $\boldsymbol{v}$ be an arbitrary d-vector.

If a, $\mathbf{b}$, $\boldsymbol{x}$, and Y are independent and $\sigma^2 \leq 1$, then

$$\Pr\left[\frac{\|\boldsymbol{x} + Y\boldsymbol{v}\|_\infty}{|a + \mathbf{b}^T\boldsymbol{v}|} > x\right] \leq \sqrt{\frac{2}{\pi}}\frac{1}{x}\left(\frac{\sqrt{2}}{\sigma} + \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)$$

Proof. We begin by observing that $a + \mathbf{b}^T\boldsymbol{v}$ and each component of $\boldsymbol{x} + Y\boldsymbol{v}$ is a Gaussian random variable of variance $\sigma^2(1 + \|\boldsymbol{v}\|^2)$ whose mean has absolute value at most $1 + \|\boldsymbol{v}\|$, and that all these variables are independent.

By Lemma A.3,

$$\mathbf{E}\left[\|\boldsymbol{x} + Y\boldsymbol{v}\|_\infty\right] \leq 1 + \|\boldsymbol{v}\| + \left(\sigma\sqrt{(1 + \|\boldsymbol{v}\|^2)}\right)\left(\sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right).$$

On the other hand, Lemma A.2 implies

$$\Pr\left[\frac{1}{|a + \mathbf{b}^T\boldsymbol{v}|} > x\right] \leq \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma\sqrt{1 + \|\boldsymbol{v}\|^2}}. \tag{2.3.4}$$

Thus, we can apply Corollary A.9 to show

$$\Pr\left[\frac{\|x + Yv\|_\infty}{|a + b^T v|} > x\right] \leq \sqrt{\frac{2}{\pi}} \frac{1 + \|v\| + \left(\sigma\sqrt{1+\|v\|^2}\right)\left(\sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)}{x\sigma\sqrt{1+\|v\|^2}}$$

$$= \sqrt{\frac{2}{\pi}}\frac{1}{x}\left(\frac{1+\|v\|}{\sigma\sqrt{1+\|v\|^2}} + \left(\sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)\right)$$

$$\leq \sqrt{\frac{2}{\pi}}\frac{1}{x}\left(\frac{\sqrt{2}}{\sigma} + \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}\right)$$

□

2.4 Smoothed Analysis of Gaussian Elimination

We now combine the results from the previous sections to bound the smoothed precision needed to obtain b-bit answers using Gaussian elimination without pivoting.

Theorem 2.9 (Smoothed precision of Gaussian elimination). For $n > e^4$, let $\bar{A}$ be an n-by-n matrix for which $\|A\| \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1/4$. Then, the expected number of bits of precision necessary to solve $Ax = \mathbf{b}$ to b bits of accuracy using Gaussian elimination without pivoting is at most

$$b + \frac{7}{2}\log_2 n + 3\log_2\left(\frac{1}{\sigma}\right) + \log(1 + 2\sqrt{n}\sigma) + \log_2\sqrt{\log n} + \frac{1}{\log n} + 5.04$$

Proof. By Wilkinson's theorem, we need the machine precision, ϵ_{mach}, to satisfy

$$5 \cdot 2^b n \rho_L(A)\rho_U(A)\kappa(A)\epsilon_{mach} \leq 1 \implies$$
$$2.33 + b + \log_2 n + \log_2(\rho_L(A)) + \log_2(\rho_U(A)) + \log_2(\kappa(A)) \leq \log_2(1/\epsilon_{mach})$$

We will apply Lemma A.11 to bound these log-terms. For any matrix $\bar{A}$ satisfying $\|\bar{A}\| \leq 1$, Theorem 2.4 implies

$$\mathbf{E}\left[\log_2 \rho_U(A)\right] \leq 2\log_2 n + \log_2\left(\frac{1}{\sigma}\right) + 0.12,$$

and Theorem 2.7 implies

$$\mathbf{E}\left[\log_2 \rho_L(A)\right] \leq 2\log_2 n + \log_2\left(\frac{1}{\sigma} + \sqrt{\log n}\left(1 + \frac{1}{2\log n}\right)\right) + 1.62$$

using $\sigma \leq \frac{1}{2}$ and $n > e^4$,

$$\leq 2\log_2 n + \log_2\left(\frac{1}{\sigma}\right) + \log_2\sqrt{\log n} + \frac{1}{\log n} + 1.62$$

Theorem 2.2 implies

$$\mathbf{E}\left[\log_2 \|A^{-1}\|\right] \leq \frac{1}{2}\log_2 n + \log_2\left(\frac{1}{\sigma}\right) + 2.68,$$

and,

$$\mathbf{E}\left[\log_2(\|A\|)\right] \leq \log_2(1 + 2\sqrt{n}\sigma)$$

follows from the well-known fact that the expectation of $\|A - \bar{A}\|$ is at most $2\sqrt{n}\sigma$ (*c.f.*, [19]) and that $\mathbf{E}[\log(X)] \leq \log \mathbf{E}[X]$ for every positive random variable X. Thus, the expected number of digits of precision needed is at most

$$b + \frac{7}{2}\log_2 n + 3\log_2\left(\frac{1}{\sigma}\right) + \log(1 + 2\sqrt{n}\sigma) + \log_2\sqrt{\log n} + \frac{1}{\log n} + 5.04$$

□

The following conjecture would further improve the coefficient of $\log(1/\sigma)$.

Conjecture 2. Let $\bar{A}$ be a n-by-n matrix for which $\|A\| \leq 1$, and let A be a matrix of independent Gaussian random variables centered at $\bar{A}$, each of variance $\sigma^2 \leq 1$. Then

$$\Pr\left[\rho_L(A)\rho_U(A)\kappa(A) > x\right] \leq \frac{n^{c_1}\log^{c_2}(x)}{x\sigma},$$

for some constants c_1 and c_2.

2.5 Symmetric matrices

Many matrices that occur in practice are symmetric and sparse. Moreover, many matrix algorithms take advantage of this structure. Thus, it is natural to study the smoothed analysis of algorithms under perturbations that respect symmetry and non-zero structure. In this section, we study the condition numbers and growth factors of Gaussian elimination without pivoting of symmetric matrices under perturbations that only alter their diagonal and non-zero entries.

Definition 2.10 (Zero-preserving perturbations). Let $\bar{T}$ be a matrix. We define a *zero-preserving perturbation of* $\bar{T}$ *of variance* σ^2 to be the matrix T obtained by adding independent Gaussian random variables of mean 0 and variance σ^2 to the non-zero entries of $\bar{T}$.

In the lemmas and theorems of this section, when we express a symmetric matrix A as $T + D + T^T$, we mean that T is lower-triangular with zeros on the diagonal and D is a diagonal matrix. By making a zero-preserving perturbation to $\bar{T}$, we preserve the symmetry of the matrix. The main results of this section are that the smoothed condition number and growth factors of symmetric matrices under zero-preserving perturbations to T and diagonal perturbations to D have distributions similar those proved in Sections 2.2 and 2.3 for dense matrices under dense perturbations.

2.5.1 Bounding the condition number

We begin by recalling that the singular values and vectors of symmetric matrices are the eigenvalues and eigenvectors.

Lemma 2.11. Let $\bar{A} = \bar{T} + \bar{D} + \bar{T}^T$ be an arbitrary n-by-n symmetric matrix. Let T be a zero-preserving perturbation of $\bar{T}$ of variance σ^2, let G_D be a diagonal matrix of Gaussian random variables of variance σ^2 and mean 0, and let $D = \bar{D} + G_D$. Then, for $A = T + D + T^T$,

$$\Pr\left[\left\|A^{-1}\right\| \geq x\right] \leq \sqrt{\frac{2}{\pi}} n^{3/2}/x\sigma$$

Proof.

$$\Pr\left[\left\|(T + D + T^T)^{-1}\right\| \geq x\right] \leq \max_{T} \Pr_{G_D}\left[\left\|((T + \bar{D} + T^T) + G_D)^{-1}\right\| \geq x\right]$$

The proof now follow from Lemma 2.12, taking $T + \bar{D} + T^T$ as the base matrix. □

Lemma 2.12. Let $\bar{A}$ be an arbitrary n-by-n symmetric matrix, let G_D be a diagonal matrix of Gaussian random variables of variance σ^2 and mean 0, and let $A = \bar{A} + G_D$. Then,

$$\Pr\left[\left\|A^{-1}\right\| \geq x\right] \leq \sqrt{\frac{2}{\pi}} n^{3/2}/x\sigma.$$

Proof. Let $x_1, \ldots, x_n$ be the diagonal entries of G_D, and let

$$g = \frac{1}{n}\sum_{i=1}^{n} x_i, \text{ and}$$
$$y_i = x_i - g$$

Then,

$$\Pr_{y_1,\ldots,y_n,g}\left[\left\|(\bar{A} + G_D)^{-1}\right\| \geq x\right] = \Pr_{y_1,\ldots,y_n,g}\left[\left\|(\bar{A} + \mathrm{diag}(y_1, \ldots, y_n) + gI)^{-1}\right\| \geq x\right]$$
$$\leq \max_{y_1,\ldots,y_n} \Pr_{g}\left[\left\|(\bar{A} + \mathrm{diag}(y_1, \ldots, y_n) + gI)^{-1}\right\| \geq x\right].$$

The proof now follows from Proposition 2.13 and Lemma 2.14. □

Proposition 2.13. Let $x_1, \ldots, x_n$ be independent Gaussian random variables of variance σ^2 with means $a_1, \ldots, a_n$, respectively. Let

$$g = \frac{1}{n}\sum_{i=1}^{n} x_i, \text{ and}$$
$$y_i = x_i - g$$

Then, g is a Gaussian random variable of variance σ^2/n with mean $(1/n)\sum a_i$, independent of $y_1, \ldots, y_n$.

Lemma 2.14. Let $\bar{A}$ be an arbitrary n-by-n symmetric matrix, and let g be a Gaussian random variable of mean 0 and variance σ^2/n. Let $A = \bar{A} + gI$. Then,

$$\Pr\left[\left\|A^{-1}\right\| \geq x\right] \leq \sqrt{\frac{2}{\pi}} n^{3/2}/x\sigma.$$

Proof. Let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of $\bar{A}$. Then,

$$\left\|(\bar{A} + gI)^{-1}\right\|^{-1} = \min_i |\lambda_i + g|.$$

By Lemma A.2,

$$\Pr\left[|\lambda_i - g| < \epsilon\right] < \sqrt{\frac{2}{\pi}} \sqrt{n}\epsilon/\sigma; \text{ so,}$$

$$\Pr\left[\min_i |\lambda_i - g| < \epsilon\right] < \sqrt{\frac{2}{\pi}} n^{3/2}\epsilon/\sigma$$

□

As in Section 2.2, we can now prove:

Theorem 2.15 (Condition number of symmetric matrices). Let $\bar{A} = \bar{T} + \bar{D} + \bar{T}^T$ be an arbitrary n-by-n symmetric matrix satisfying $\left\|\bar{A}\right\| \leq 1$. Let $\sigma^2 \leq 1$, let T be a zero-preserving perturbation of $\bar{T}$ of variance σ^2, let G_D be a diagonal matrix of Gaussian random variables of variance σ^2 and mean 0, and let $D = \bar{D} + G_D$. Then, for $A = T + D + T^T$,

$$\Pr\left[\kappa(A) \geq x\right] \leq 4\sqrt{\frac{2}{\pi}} \frac{n^2}{x\sigma}\left(1 + \sqrt{\log(x)/2n}\right)$$

Proof. As in the proof of Theorem 2.3, we can apply the techniques used in the proof of [6, Theorem II.7], to show

$$\Pr\left[\left\|\bar{A} - A\right\| \geq \sqrt{d} + k\right] < e^{-k^2/2}.$$

The rest of the proof follows the outline of the proof of Theorem 2.3, using Lemma 2.11 instead of Theorem 2.2. □

2.5.2 Bounding entries in U

In this section, we will prove:

Theorem 2.16 ($\rho_U(A)$ of symmetric matrices). Let $\bar{A} = \bar{T} + \bar{D} + \bar{T}^T$ be an arbitrary n-by-n symmetric matrix satisfying $\bar{A} \leq 1$. Let $\sigma^2 \leq 1$, let T be a zero-preserving perturbation of $\bar{T}$ of variance σ^2, let G_D be a diagonal matrix of Gaussian random variables of variance σ^2 and mean 0, and let $D = \bar{D} + G_D$. Then, for $A = T + D + T^T$,

$$\Pr\left[\rho_U(A) > 1 + x\right] \leq \frac{2}{7}\sqrt{\frac{2}{\pi}} \frac{n^{7/2}}{x\sigma}$$

Proof. We proceed as in the proof of Theorem 2.4, where we derived (2.3.2)

$$\begin{aligned}\frac{\|U_{k,k:n}\|_1}{\|A\|_\infty} &\leq 1 + \left\|A_{k,1:k-1}A^{-1}_{1:k-1,1:k-1}\right\|_1 \\ &\leq 1 + \sqrt{k-1}\left\|A_{k,1:k-1}A^{-1}_{1:k-1,1:k-1}\right\|_2 \\ &\leq 1 + \sqrt{k-1}\left\|A_{k,1:k-1}\right\|_2\left\|A^{-1}_{1:k-1,1:k-1}\right\|_2\end{aligned}$$

Hence

$$\begin{aligned}\Pr\left[\frac{\|U_{k,k:n}\|_1}{\|A\|_\infty} > 1+x\right] &\leq \Pr\left[\|A_{k,1:k-1}\|_2\left\|A^{-1}_{1:k-1,1:k-1}\right\|_2 > \frac{x}{\sqrt{k-1}}\right] \\ &\leq \mathbf{E}\left[\|A_{k,1:k-1}\|_2\right]\sqrt{\frac{2}{\pi}}\frac{(k-1)^2}{x\sigma}, \text{ by Lemmas 2.11 and A.9,} \\ &\leq \sqrt{1+j\sigma^2}\sqrt{\frac{2}{\pi}}\frac{(k-1)^2}{x\sigma}\end{aligned}$$

where j is the number of non-zeros in $A_{k,1:k-1}$,

$$\leq \sqrt{\frac{2}{\pi}}\frac{\sqrt{k}(k-1)^2}{x\sigma}.$$

Applying a union bound over k,

$$\begin{aligned}\Pr\left[\rho_U(A) > x\right] &\leq \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma}\sum_{k=2}^{n}\sqrt{k}(k-1)^2 \\ &\leq \frac{2}{7}\sqrt{\frac{2}{\pi}}\frac{n^{7/2}}{x\sigma}.\end{aligned}$$

□

2.5.3 Bounding entries in L

As in Section 2.3.2, we derive a bound on the growth factor of L. As before, we will show that it is unlikely that $A^{(k-1)}_{j,k}$ is large while $A^{(k-1)}_{k,k}$ is small. However, our techniques must differ from those used in Section 2.3.2, as the proof in that section made critical use of the independence of $A_{k,1:(k-1)}$ and $A_{1:(k-1),k}$.

Theorem 2.17 ($\rho_L(A)$ of symmetric matrices). Let $\sigma^2 \leq 1$. Let $\bar{A} = \bar{T} + \bar{D} + \bar{T}^T$ be an arbitrary n-by-n symmetric matrix satisfying $\|\bar{A}\| \leq 1$. Let T be a zero-preserving perturbation of $\bar{T}$ of variance σ^2, let G_D be a diagonal matrix of Gaussian random variables of variance $\sigma^2 \leq 1$ and mean 0, and let $D = \bar{D} + G_D$. Then, for $A = T + D + T^T$,

$$\Pr\left[\rho_L(A) > x\right] \leq \frac{3.2n^4}{x\sigma^2}\log^{3/2}\left(e\sqrt{\pi/2}\,x\sigma^2\right).$$

Proof. Using Lemma 2.18, we obtain for all k

$$\Pr[\exists j > k : |L_{j,k}| > x] \le \Pr\left[\left\|L_{(k+1):n,k}\right\| > x\right] \le \frac{3.2n^2}{x\sigma^2}\log^{3/2}\left(e\sqrt{\pi/2}x\sigma^2\right).$$

Applying a union bound over the choices for k, we then have

$$\Pr[\exists j, k : |L_{j,k}| > x] \le \frac{3.2n^3}{x\sigma^2}\log^{3/2}\left(e\sqrt{\pi/2}x\sigma^2\right).$$

The result now follows from the fact that $\|L\|_\infty$ is at most n times the largest entry in L. □

Lemma 2.18. **Under the conditions of Theorem 2.17,**

$$\Pr\left[\left\|L_{(k+1):n,k}\right\| > x\right] \le \frac{3.2n^2}{x\sigma^2}\log^{3/2}\left(e\sqrt{\pi/2}x\sigma^2\right).$$

Proof. We recall that

$$L_{k+1:n,k} = \frac{A_{k+1:n,k} - A_{k+1:n,1:k-1}A^{-1}_{1:k-1,1:k-1}A_{1:k-1,k}}{A_{k,k} - A_{k,1:k-1}A^{-1}_{1:k-1,1:k-1}A_{1:k-1,k}}$$

Because of the symmetry of A, $A_{k,1:k-1}$ is the same as $A_{1:k-1,k}$, so we can no longer use the proof that worked in Section 2.3.2. Instead we will bound the tails of the numerator and denominator separately.

Consider the numerator first. Setting $\boldsymbol{v} = A^{-1}_{1:k-1,1:k-1}A_{1:k-1,k}$, the numerator can be written $A_{k+1:n,1:k}\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}$. We will now prove

$$\Pr_{A_{k+1:n,1:k},A_{1:k-1,1:k}}\left[\left\|A_{k+1:n,1:k}\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\|_\infty > x\right] \le \sqrt{\frac{2}{\pi}}\left(\frac{2n^2(1+\sigma\sqrt{2\log(x\sigma)})+n}{x\sigma}\right) \tag{2.5.1}$$

It suffices to prove this for all x for which the right-hand side is less than 1, so in particular it suffices to consider x for which

$$\frac{x}{1+\sigma\sqrt{2\log(x\sigma)}} \ge 1, \tag{2.5.2}$$

and $x\sigma \ge 2$. We divide this probability accordingly to a parameter c, which we will set so that $\frac{1-c}{c\sigma} = \sqrt{2\log(x\sigma)}$. We have

$$\Pr_{A_{k+1:n,1:k},A_{1:k-1,1:k}}\left[\left\|A_{k+1:n,1:k}\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\|_\infty > x\right]$$

$$\le \Pr_{A_{1:(k-1),1:k}}\left[\left\|\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\|_\infty > cx\right] \tag{2.5.3}$$

$$+ \Pr_{A_{k+1:n,1:k}}\left[\left\|A_{k+1:n,1:k}\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\|_\infty > \frac{1}{c}\left\|\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\| \,\middle|\, \left\|\begin{pmatrix}-\boldsymbol{v}^T\\1\end{pmatrix}\right\| < cx\right] \tag{2.5.4}$$

Once $\mathbf{v}$ is fixed, each component of $A_{k+1:n,1:k}\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}$ is a Gaussian random vector of variance

$$(1+\|\mathbf{v}\|^2)\sigma^2 \leq (1+\|\mathbf{v}\|)^2\sigma^2$$

and mean at most $\left\|\bar{A}_{k+1:n,1:k}\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}\right\| \leq \left\|\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}\right\|$. So,

$$\left\|A_{k+1:n,1:k}\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}\right\|_\infty > \frac{1}{c}\left\|\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}\right\|$$

implies some term in the numerator is more than $1/c-1$ standard deviations from its mean, and we can therefore apply Lemma A.1 and a union bound to derive

$$(2.5.4) \leq \sqrt{\frac{2}{\pi}}\frac{ne^{-\frac{1}{2}\left(\frac{1-c}{c\sigma}\right)^2}}{\frac{1-c}{c\sigma}} \leq \sqrt{\frac{2}{\pi}}\frac{n}{x\sigma\sqrt{2\log(x\sigma)}}.$$

To bound (2.5.4), we note that Lemma 2.11 and Corollary A.10 imply

$$\Pr_{A_{1:(k-1),1:k}}\left[\left\|A_{1:k-1,1:k-1}^{-1}A_{1:k-1,k}\right\| > y\right] \leq \sqrt{\frac{2}{\pi}}\frac{n^2}{y\sigma},$$

and so

$$\begin{aligned}\Pr_{A_{1:(k-1),1:k}}\left[\left\|\begin{pmatrix}-\mathbf{v}^T\\1\end{pmatrix}\right\| > cx\right] &\leq \Pr_{A_{1:(k-1),1:k}}\left[\left\|A_{1:k-1,1:k-1}^{-1}A_{1:k-1,k}\right\| > cx-1\right]\\ &\leq \sqrt{\frac{2}{\pi}}\frac{n^2}{(cx-1)\sigma}\\ &\leq \sqrt{\frac{2}{\pi}}\frac{2n^2(1+\sigma\sqrt{2\log(x\sigma)})}{x\sigma}, \text{ by (2.5.2).}\end{aligned}$$

So,

$$\begin{aligned}(2.5.1) &\leq \sqrt{\frac{2}{\pi}}\left(\frac{n}{x\sigma\sqrt{2\log(x\sigma)}} + \frac{2n^2(1+\sigma\sqrt{2\log(x\sigma)})}{x\sigma}\right)\\ &\leq \sqrt{\frac{2}{\pi}}\left(\frac{2n^2(1+\sigma\sqrt{2\log(x\sigma)})+n}{x\sigma}\right),\end{aligned}$$

by the assumption $x\sigma \geq 2$, which proves (2.5.1).

As for the denominator, we note that $A_{k,k}$ is a independent of all the other terms, and hence

$$\Pr\left[\left|A_{k,k} - A_{k,1:k-1}A_{1:k-1,1:k-1}^{-1}A_{1:k-1,k}\right| < 1/x\right] \leq \sqrt{\frac{2}{\pi}}\frac{1}{x\sigma},$$

by Lemma A.2.

Applying Corollary A.8 with

$$\alpha = \sqrt{\frac{2}{\pi}}\left(2n^2 + n\right) \qquad \beta = \frac{4n^2\sigma}{\sqrt{\pi}} \qquad \gamma = \sqrt{\frac{2}{\pi}}$$

to combine this inequality with (2.5.1), we derive the bound

$$\begin{aligned} \frac{2}{\pi x \sigma^2} & \left(2n^2 + n + \left(\left(2 + 4\sqrt{2}\sigma/3\right) n^2 + n\right) \log^{3/2}\left(\sqrt{\pi/2} x \sigma^2\right)\right) \\ & \leq \frac{2n^2}{\pi x \sigma^2}\left(3 + 4\sqrt{2}\sigma/3\right)\left(\log^{3/2}\left(\sqrt{\pi/2} x \sigma^2\right) + 1\right) \\ & \leq \frac{3.2n^2}{x\sigma^2} \log^{3/2}\left(e\sqrt{\pi/2} x \sigma^2\right), \end{aligned}$$

as $\sigma \leq 1$. □

Chapter 3

Smoothed Analysis of Gaussian Elimination with Partial Pivoting

3.1 Introduction

In this chapter, we consider the growth factor of Gaussian elimination with partial pivoting. With partial pivoting, the entries in L are necessarily bounded by 1 in absolute value, so the growth is confined to U. Thus

$$A = PLU$$

where P is a permutation matrix and $|L_{i,j}| \leq 1$. We will usually assume that A has been put into partial pivoting order, so that $P = I$. Recall that U is given by the equation

$$U_{k,:} = A_{k,:} - A_{k,1:k-1}A_{1:k-1,1:k-1}^{-1}A_{1:k-1,:}$$

and the growth factor ρ_U by

$$\rho_U = \frac{\max_k \|U_{k,:}\|_1}{\|A\|_\infty} \leq 1 + \left\|A_{k,1:k-1}A_{1:k-1,1:k-1}^{-1}\right\|_1$$

The remaining sections are organized as follows: in Sections 3.2 and 3.3, we establish a recursive formula for $A_{k,1:k-1}A_{1:k-1,1:k-1}^{-1}$. In Section 3.4 we give an outline of the probabilistic argument that will follow. This is followed by Section 3.5, which proves some technical results about Gaussian vectors and matrices. Sections 3.6 and 3.7 bound the tail distribution of two factors that appear in the recursive formula derived in Section 3.3, and the final Section 3.8 puts everything together to prove the main theorem, which we state below.

Theorem 3.1. *If $A \in \mathbb{R}^{n\times k}$ is a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$ with $\left\|\bar{A}\right\| \leq 1$, and $\rho_U(A)$ is the growth factor during Gaussian elimination with partial pivoting, then*

$$\Pr_A\left[\rho_U(A) > x\right] \leq \left(\frac{1}{x}\left(\mathcal{O}\left(\frac{n(1+\sigma\sqrt{n})}{\sigma}\right)\right)^{12\log k}\right)^{\frac{1}{21}\log k}$$

In the theorems and lemmas below, we will assume that all matrices are of full rank (in the probabilistic setting, this is true with probability 1).

3.2 Some algebraic results

First, we define some convenient notation.

Definition 3.2. Given a matrix $A \in \mathbb{R}^{n \times n}$, an index k and a submatrix $X = A_{k_1:k_2,l_1:l_2}$, define

$$\langle X \rangle_k = X - A_{k_1:k_2,1:k} A^{-1}_{1:k,1:k} A_{1:k,l_1:l_2}$$

and

$$\langle X \rangle = \langle X \rangle_{\min(k_1,l_1)-1}$$

Intuitively, $\langle X \rangle_k$ is what remains of X after k steps of elimination have been carried out.

Lemma 3.3 (Block-structured inverse). Consider an $n \times k$ real matrix with the block structure

$$\begin{bmatrix} A \\ X \end{bmatrix} = \begin{bmatrix} A_1 & C \\ R & A_2 \\ X_1 & X_2 \end{bmatrix}$$

where A, A_1 and A_2 are square matrices. Then

$$\begin{aligned} XA^{-1} &= \begin{bmatrix} X_1 & X_2 \end{bmatrix} \begin{bmatrix} A_1 & C \\ R & A_2 \end{bmatrix}^{-1} \\ &= \left[\left(X_1 - \langle X_2 \rangle \langle A_2 \rangle^{-1} R \right) A_1^{-1} \quad ; \quad \langle X_2 \rangle \langle A_2 \rangle^{-1} \right] \end{aligned}$$

Proof. Multiplying the RHS by A gives for the first component

$$X_1 - \langle X_2 \rangle \langle A_2 \rangle^{-1} R + \langle X_2 \rangle \langle A_2 \rangle^{-1} R = X_1$$

and for the second

$$\begin{aligned} X_1 A_1^{-1} C - \langle X_2 \rangle \langle A_2 \rangle^{-1} R A_1^{-1} C + \langle X_2 \rangle \langle A_2 \rangle^{-1} A_2 &= X_1 A_1^{-1} C + \langle X_2 \rangle \langle A_2 \rangle^{-1} \langle A_2 \rangle \\ &= X_2 \end{aligned}$$

□

Notice that XA^{-1} gives the coefficients when rows of X are expressed in a basis of the rows of A. According to the lemma, the coefficients of $\begin{bmatrix} R & A_2 \end{bmatrix}$ are given by $\langle X_2 \rangle \langle A_2 \rangle^{-1}$. Hence

$$X - \langle X_2 \rangle \langle A_2 \rangle^{-1} \begin{bmatrix} R & A_2 \end{bmatrix}$$

is the part of X that is spanned by $\begin{bmatrix} A_1 & C \end{bmatrix}$. Let us write this matrix as

$$X - \langle X_2 \rangle \langle A_2 \rangle^{-1} \begin{bmatrix} R & A_2 \end{bmatrix} = Y \begin{bmatrix} A_1 & C \end{bmatrix}$$

Then if S is any subset of columns of this matrix containing at least $|A_1|$ linearly independent columns, the coefficients Y of $\begin{bmatrix} A_1 & C \end{bmatrix}$ will be given by

$$\left(X_S - \langle X_2 \rangle \langle A_2 \rangle^{-1} \begin{bmatrix} R & A_2 \end{bmatrix}_S\right) \left(\begin{bmatrix} A_1 & C \end{bmatrix}_S\right)^{(r)}$$

where $M^{(r)}$ denotes any right-inverse of M, *i.e.*, $MM^{(r)} = I$.

We restate this as a corollary to Lemma 3.3.

Corollary 3.4. With the same notation as in Lemma 3.3,

$$\begin{aligned} \left(X_1 - \langle X_2 \rangle \langle A_2 \rangle^{-1} R\right) A_1^{-1} &= \left(X_S - \langle X_2 \rangle \langle A_2 \rangle^{-1} \begin{bmatrix} R & A_2 \end{bmatrix}_S\right) \left(\begin{bmatrix} A_1 & C \end{bmatrix}_S\right)^{(r)} \\ &= \begin{bmatrix} -\langle X_2 \rangle \langle A_2 \rangle^{-1} & I \end{bmatrix} \begin{bmatrix} R & A_2 \\ X_1 & X_2 \end{bmatrix}_S \left(\begin{bmatrix} A_1 & C \end{bmatrix}_S\right)^{(r)} \end{aligned}$$

for any subset S of the columns of A such that $\begin{bmatrix} A_1 & C \end{bmatrix}_S$ is right-invertible.

Corollary 3.5. Consider an $n \times m$ real matrix with the block structure

$$\begin{bmatrix} A_1 & C & Y_1 \\ R & A_2 & Y_2 \\ X_1 & X_2 & Z \end{bmatrix}$$

where A_1 and A_2 are square with dimensions k_1 and k_2 respectively, with $k_1+k_2 = k$. Let

$$A = \begin{bmatrix} A_1 & C \\ R & A_2 \end{bmatrix} \quad X = \begin{bmatrix} X_1 & X_2 \end{bmatrix} \quad Y = \begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix}$$

Then

$$\langle Z \rangle_k = \langle Z \rangle_{k_1} - \langle X_2 \rangle_{k_1} \langle A_2 \rangle_{k_1}^{-1} \langle Y_2 \rangle_{k_1}$$

Proof.

$$\begin{aligned} \langle Z \rangle_k &= Z - XA^{-1}Y \\ &= Z - \left(X_1 - \langle X_2 \rangle \langle A_2 \rangle^{-1} R\right) A_1^{-1} Y_1 + \langle X_2 \rangle \langle A_2 \rangle^{-1} Y_2 \\ &= (Z - X_1 A_1^{-1} Y_1) - \langle X_2 \rangle \langle A_2 \rangle^{-1} (Y_2 - RA_1^{-1} Y_1) \\ &= \langle Z \rangle_{k_1} - \langle X_2 \rangle \langle A_2 \rangle^{-1} \langle Y_2 \rangle \end{aligned}$$

□

Remark: Corollary 3.5 is the trivial fact that eliminating k_1 rows, then the next k_2 rows is the same as eliminating $k = k_1 + k_2$ rows.

Definition 3.6. For any matrix $A \in \mathbb{R}^{m \times n}$, $m \leq n$, define the *(right) pseudo-inverse* $A^\dagger \in \mathbb{R}^{n \times m}$ to be the matrix satisfying the conditions

$$AA^\dagger = I \text{ and } A^\dagger A \text{ is an orthogonal projection.}$$

Given the singular value decomposition

$$A = U\Sigma V^T$$

with $U \in O(m, m)$, $V \in O(n, m)$ and Σ an $m \times m$ diagonal matrix, $A^\dagger$ is given by

$$A^\dagger = V\Sigma^{-1}U^T$$

The pseudo-inverse is undefined if A does not have rank m.

If $m > n$, we similarly define a *left* pseudo-inverse.

Lemma 3.7. For matrices $X \in \mathbb{R}^{k\times n}$, $Y \in \mathbb{R}^{l\times n}$, $Z \in \mathbb{R}^{m\times n}$, such that $m \geq n \geq l$,

$$XY^{(r)} = (XZ^\dagger)(YZ^\dagger)^{(r)}$$

That is, X multiplied by any right-inverse of Y is equal to $XZ^\dagger$ multiplied by some right-inverse of $YZ^\dagger$ and *vice versa*.

Proof. Since $m \geq n$, we have $Z^\dagger Z = I$. Hence

$$XY^{(r)} = X\left(Z^\dagger Z\right)Y^{(r)} = \left(XZ^\dagger\right)\left(ZY^{(r)}\right)$$

Also $\left(YZ^\dagger\right)\left(ZY^{(r)}\right) = I$ which shows that $ZY^{(r)}$ is a right-inverse of $YZ^\dagger$.

Conversely, since

$$Y\left(Z^\dagger(YZ^\dagger)^{(r)}\right) = I$$

we may write $(XZ^\dagger)(YZ^\dagger)^{(r)} = X\left(Z^\dagger(YZ^\dagger)^{(r)}\right)$ as $XY^{(r)}$. □

3.3 Recursive bound

Now we will apply the results of the previous section to derive a recurrence relation.

Lemma 3.8. Let $A \in \mathbb{R}^{n\times k}$, and

$$1 \leq k_1 < k_2 < k$$

be two indices such that $k_2 \leq 2k_1$. Let S be a subset of $(k_1, k]$ such that

$$k_2 - k_1 \leq |S| \leq k_1$$

Define

$$\begin{aligned} A_i &= A_{(k_i,k],(k_i,k]} \\ X_i &= A_{(k,n],(k_i,k]} \\ B &= A_{(k_2,n],S} \\ C &= A_{(k_1,k_2],S} \\ Z &= A^{-1}_{(0,k_1],(0,k_1]}A_{(0,k_1],S} \end{aligned}$$

Then

$$\left[-\langle X_1\rangle\langle A_1\rangle^{-1} \quad I\right] = \left[-\langle X_2\rangle\langle A_2\rangle^{-1} \quad I\right]\left[-(\langle B\rangle_{k_1} Z^\dagger)(\langle C\rangle_{k_1} Z^\dagger)^\dagger \quad ; \quad I\right]$$

Proof. By Lemma 3.3 and Corollaries 3.4 and 3.5,

$$\langle X_1\rangle \langle A_1\rangle^{-1} = \left[\left[-\langle X_2\rangle \langle A_2\rangle^{-1} \quad I\right] \langle B\rangle_{k_1} \langle C\rangle_{k_1}^{(r)} \quad ; \quad -\langle X_2\rangle \langle A_2\rangle^{-1}\right]$$

and by Lemma 3.7,

$$= \left[\left[-\langle X_2\rangle \langle A_2\rangle^{-1} \quad I\right] \left(\langle B\rangle_{k_1} Z^\dagger\right)\left(\langle C\rangle_{k_1} Z^\dagger\right)^{(r)} \quad ; \quad -\langle X_2\rangle \langle A_2\rangle^{-1}\right]$$

□

The reason for choosing the matrix Z is that

$$\langle B\rangle_{k_1} Z^\dagger = BZ^\dagger - A_{(k_2,n],(0,k_1]}ZZ^\dagger$$

and $ZZ^\dagger$ is a projection matrix. This will prove useful in bounding the norm of $\langle B\rangle Z^\dagger$ (resp. $\left(\langle C\rangle Z^\dagger\right)^\dagger$) even though the norm of $\langle B\rangle$ (resp. $\langle C\rangle^\dagger$) is hard to bound directly.

Now we extend the idea of the preceding lemma to a whole sequence of indices k_i.

Theorem 3.9. Let $A \in \mathbb{R}^{n\times k}$, and

$$0 = k_0 < k_1 < k_2 < \ldots < k_r < k < n$$

be a sequence of indices with the property

$$k_{i+1} \leq 2k_i \text{ for } 1 \leq i < r.$$

Let S_i be a subset of $(k_i, k]$ such that

$$k_{i+1} - k_i \leq |S_i| \leq k_i \text{ for } 1 \leq i < r$$

and $k_1 \leq |S_0|$. Define

$$\begin{aligned} A_i &= A_{(k_i,k],(k_i,k]} \\ X_i &= A_{(k,n],(k_i,k]} \\ B_i &= A_{(k_{i+1},n],S_i} \\ C_i &= A_{(k_i,k_{i+1}],S_i} \\ Z_i &= A^{-1}_{(0,k_i],(0,k_i]}A_{(0,k_i],S_i} \end{aligned}$$

We define Z_i only for $i \geq 1$. Then

$$\left[-X_0A_0^{-1} \quad I\right] = \left[-\langle X_r\rangle \langle A_r\rangle^{-1} \quad I\right] \prod_{i=r-1}^{1} \left[-\left(\langle B_i\rangle_{k_i} Z_i^\dagger\right)\left(\langle C_i\rangle_{k_i} Z_i^\dagger\right)^\dagger \quad ; \quad I\right] \times$$
$$\times \left[-B_0C_0^\dagger \quad I\right]$$

Note that the index i in the product counts down.

Proof. Similar to Lemma 3.8 (except we do not apply Lemma 3.7 to rewrite $B_0C_0^{(r)}$), we have

$$\left[-X_0A_0^{-1} \quad I\right] = \left[-\langle X_1\rangle \langle A_1\rangle^{-1} \quad I\right] \left[-B_0C_0^\dagger \quad I\right]$$

The rest is immediate from Lemma 3.8 and induction. □

3.4 Outline of Argument

The idea now is to choose a suitable sequence of k_i's in Theorem 3.9 and bound the tails of each factor individually. In particular, we will want to limit r to $\mathcal{O}(\log k)$.

In the next section, we will prove some probabilistic results about Gaussian matrices and vectors, and use these to estimate the tail distributions of $\langle B_i \rangle Z_i^{\dagger}$ and $(\langle C_i \rangle Z_i^{\dagger})^{\dagger}$.

Once we have a handle on these two, it is a matter of nailing down appropriate values for the k_i and pulling everything together via a union bound over the r factors in Theorem 3.9. The very first factor in it will be bounded by ensuring that $k - k_r = \mathcal{O}(\log k)$, so that in the worst case, it is still only polynomial in k.

3.5 Some probabilistic results

3.5.1 The smallest singular value of a scaled Gaussian matrix

We will need to investigate the smallest singular value of a Gaussian random matrix that has been multiplied by a constant matrix. (For example, Z_i has this form.)

Theorem 3.10. If $\Sigma \in \mathbb{R}^{n\times n}$ is a matrix with singular values

$$\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_n$$

and $X \in \mathbb{R}^{n\times k}$ is a random matrix distributed as $\mathfrak{N}\left(\bar{X}, I_n \otimes I_k\right)$, then

$$\begin{aligned}\Pr_X\left[\left\|(\Sigma X)^{\dagger}\right\| > x\right] &\leq \frac{7(k/2)^{(n-k+1)/2}}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)} \prod_{i=1}^{n-k+1} \frac{1}{x\sigma_i} \\ &\leq \frac{7(k/2)^{(n-k+1)/2}}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)} \left(\frac{1}{x\sigma_1}\right)^{n-k+1}\end{aligned}$$

Proof. First, we will estimate

$$\Pr_X\left[\left\|\mathbf{u}^T(\Sigma X)^{\dagger}\right\| > x\right]$$

for a unit k-vector $\mathbf{u}$. Notice that rotating $\mathbf{u}$, *i.e.*, replacing $\mathbf{u}$ by $H\mathbf{u}$ where H is a square orthogonal matrix, is equivalent to replacing X by XH, since

$$(H\mathbf{u})^T(\Sigma X)^{\dagger} = \mathbf{u}^T H^T (\Sigma X)^{\dagger} = \mathbf{u}^T(\Sigma X H)^{\dagger}$$

Since this only changes the mean of X, and our bounds will be independent of this mean, we will assume $\mathbf{u} = \boldsymbol{e}_1$. In this case, we want the probability that the norm of the first row of $(\Sigma X)^{\dagger}$ is greater than x. The first row of $(\Sigma X)^{\dagger}$ is the vector that is the relative orthogonal complement of the column span of $\Sigma X_{:,2:k}$ in the column span of ΣX, and has inner product 1 with the first column of ΣX. Hence we are looking for a bound on the probability that the component of the first column of ΣX in the orthogonal complement of the span of the remaining columns has norm less than $1/x$. Let $V \in \mathbb{R}^{n\times n-k+1}$ be an orthonormal basis for this vector space. Then we need to bound

$$\Pr\left[\left\|V^T\Sigma X_{:,1}\right\| < 1/x\right]$$

If $\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_{n-k+1}$ are the singular values of $V^T\Sigma$, then this probability is the same as the probability that the component of $X_{:,1}$ in the row span of $V^T\Sigma$ is contained in an ellipsoid with semi-axes $1/x\tilde{\sigma}_i$, and hence is bounded by the volume of the ellipsoid, times the maximal density of the Gaussian, $(2\pi)^{-(n-k+1)/2}$. Thus

$$\begin{aligned}\Pr_X\left[\left\|\mathbf{u}^T(\Sigma X)^\dagger\right\| > x\right] &< \frac{2\pi^{(n-k+1)/2}\prod(1/x\tilde{\sigma}_i)}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)}\frac{1}{(2\pi)^{(n-k+1)/2}} \\ &= \frac{2^{-(n-k-1)/2}\prod(1/x\tilde{\sigma}_i)}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)} \\ &\leq \frac{2^{-(n-k-1)/2}\prod(1/x\sigma_i)}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)}\end{aligned}$$

since $\tilde{\sigma}_i \geq \sigma_i$. Now let $\mathbf{u}$ be a random unit vector, and $\boldsymbol{\nu}$ be the right singular vector corresponding to the smallest singular value of ΣX (which is the same as the left singular vector corresponding to the largest singular value of $(\Sigma X)^\dagger$). Since

$$\left\|\mathbf{u}^T(\Sigma X)^\dagger\right\| \geq \left\|(\Sigma X)^\dagger\right\| |\langle\mathbf{u},\boldsymbol{\nu}\rangle|$$

we have

$$\Pr_{X,\mathbf{u}}\left[\left\|\mathbf{u}^T(\Sigma X)^\dagger\right\| > x/\sqrt{k}\right] \geq \Pr_X\left[\left\|(\Sigma X)^\dagger\right\| > x\right]\cdot\Pr_{X,\mathbf{u}}\left[|\langle\mathbf{u},\boldsymbol{\nu}\rangle| > 1/\sqrt{k}\right]$$

or

$$\begin{aligned}\Pr_X\left[\left\|(\Sigma X)^\dagger\right\| > x\right] &< \frac{2^{-(n-k-1)/2}}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)}\prod_i\frac{\sqrt{k}}{x\sigma_i}\times \\ &\quad\times\frac{1}{\Pr_{X,\mathbf{u}}\left[|\langle\mathbf{u},\boldsymbol{\nu}\rangle| > 1/\sqrt{k}\right]} \\ &\leq \frac{7(k/2)^{(n-k+1)/2}}{(n-k+1)\Gamma\left(\frac{1}{2}(n-k+1)\right)}\prod_i\frac{1}{x\sigma_i}\end{aligned}$$

□

3.5.2 The moments of a Gaussian matrix

Theorem 3.11. If $X \in \mathbb{R}^{n\times n}$ is a random matrix distributed as $\mathfrak{N}\left(\bar{X}, \sigma^2 I_n \otimes I_n\right)$, with $\left\|\bar{X}\right\| \leq 1$ and $\sigma \leq 1/2$, then

$$\mathbf{E}_X\left[\|X\|^k\right] \leq 2^k\Gamma(k/2)(1+\sigma\sqrt{n})^k$$

Proof. First, note that

$$\|X\| \leq 1 + \sigma \|G\|$$

where $G \sim \mathfrak{N}(0, I_n \otimes I_n)$. Hence

$$\mathbf{E}_X\left[\|X\|^k\right] \leq \int_0^\infty (1+\sigma t)^k \, d\mu_{\|G\|}(t)$$

integrating by parts,

$$= 1 + k\sigma \int_0^\infty (1+\sigma t)^{k-1} \Pr_G [\|G\| > t] \, dt$$

applying the result of Theorem II.11 from [6] to bound the tail of $\|G\|$, we have for any $c > 0$,

$$\leq 1 + k\sigma \int_0^{c\sqrt{n}} (1+\sigma t)^{k-1} \, dt + k\sigma \int_{c\sqrt{n}}^\infty (1+\sigma t)^{k-1} e^{-\frac{1}{2}(t-2\sqrt{n})^2} \, dt$$
$$= (1 + c\sigma\sqrt{n})^k + k\sigma \int_{(c-2)\sqrt{n}}^\infty (1 + c\sigma\sqrt{n} + \sigma t)^{k-1} e^{-t^2/2} \, dt$$

setting $c = 2 + 1/\sigma\sqrt{n}$,

$$\leq (2 + 2\sigma\sqrt{n})^k + k\sigma(2 + 2\sigma\sqrt{n})^{k-1} \int_{1/\sigma}^\infty (\sigma t)^{k-1} e^{-t^2/2} \, dt$$
$$\leq (1 + \sigma\sqrt{n})^{k-1} \big(2^k(1+\sigma\sqrt{n}) + k\sigma 2^{k-1} \cdot \sigma^{k-1} 2^{k/2-1}\Gamma(k/2)\big)$$
$$\leq 2^k \Gamma(k/2)(1+\sigma\sqrt{n})^k$$

□

3.5.3 Partial pivoting polytope

Definition 3.12. Given a matrix $A \in \mathbb{R}^{n\times k}$, let B denote the result of ordering the rows of A in the partial pivoting order. Define the *partial pivoting polytope* of A to be the set of all points (row vectors) $\mathbf{x} \in \mathbb{R}^k$ such that

$$\left| x_i - \mathbf{x}_{1:i-1} B^{-1}_{1:i-1,1:i-1} B_{1:i-1,i} \right| \leq \left| B_{i,i} - B_{i,1:i-1} B^{-1}_{1:i-1,1:i-1} \right| \text{ for } 1 \leq i \leq k.$$

We will denote this polytope by the notation $\mathbf{PP}(A)$. Note that $B_{i,:} \in \mathbf{PP}(A)$ for all $i > k$, and that $\mathbf{PP}(A) = \mathbf{PP}(B) = \mathbf{PP}(B_{1:k,:})$.

Observe that $\mathbf{PP}(A)$ is symmetric about the origin. Define $r(A)$ to be the largest r such that the ball $B(0, r) \subseteq \mathbf{PP}(A)$. This can also be computed as the minimum distance r from the origin to one of the defining hyperplanes of the polytope, and is half the minimum width of the polytope.

Theorem 3.13. Let $A \in \mathbb{R}^{n \times k}$ be a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$. Then

$$\Pr_A [r(A) < r] \leq \binom{n}{k} \left(\sqrt{\frac{2}{\pi}} \frac{r}{\sigma}\right)^{n-k} \leq \left(\sqrt{\frac{2}{\pi}} \frac{nr}{\sigma}\right)^{n-k}$$

Proof. Letting S be the subset of rows that are chosen by partial pivoting, we have

$$\begin{aligned}
\Pr [r(A) < r] &= \sum_{S \in \binom{[n]}{k}} \Pr [r(A_{S,:}) < r \text{ and } \forall i \notin S,\ A_{i,:} \in \mathbf{PP}(A_{S,:})] \\
&\leq \sum_{S \in \binom{[n]}{k}} \Pr \left[\forall i \notin S,\ A_{i,:} \in \mathbf{PP}(A_{S,:}) \middle| r(A_{S,:}) < r\right]
\end{aligned}$$

and since $A_{i,:}$ is a Gaussian k-vector, and $A_{i,:} \in \mathbf{PP}(A_{S,:})$ implies that the component of $A_{i,:}$ along the normal to that defining hyperplane of $\mathbf{PP}(A_{S,:})$ which is at distance $r(A_{S,:})$ from the origin must be smaller in absolute value than $r(A_{S,:}) < r$,

$$\begin{aligned}
&\leq \sum_{S \in \binom{[n]}{k}} \left(\sqrt{\frac{2}{\pi}} \frac{r}{\sigma}\right)^{n-k} \\
&= \binom{n}{k} \left(\sqrt{\frac{2}{\pi}} \frac{r}{\sigma}\right)^{n-k}
\end{aligned}$$

□

3.5.4 Conditional Gaussian distribution

Lemma 3.14. If X_μ is a random variable with a density function of the form

$$\rho(x) e^{-\frac{1}{2}(x-\mu)^2}$$

with $\rho(x)$ a symmetric function of x, then for any $0 < r < R$,

$$\Pr_{X_\mu} \left[X_\mu \in [-r, r] \middle| X_\mu \in [-R, R]\right] \leq \Pr_{X_0} \left[X_0 \in [-r, r] \middle| X_0 \in [-R, R]\right]$$

Proof. The logarithmic derivative (wrt μ) of the conditional probability over X_μ is

$$\begin{aligned}
\frac{\partial}{\partial \mu} \ln \Pr_{X_\mu} \left[X_\mu \in [-r, r] \middle| X_\mu \in [-R, R]\right] &= \frac{\int_{-r}^{r} \rho(x)(x-\mu) e^{-\frac{1}{2}(x-\mu)^2}\, dx}{\int_{-r}^{r} \rho(x) e^{-\frac{1}{2}(x-\mu)^2}\, dx} \\
&\quad - \frac{\int_{-R}^{R} \rho(x)(x-\mu) e^{-\frac{1}{2}(x-\mu)^2}\, dx}{\int_{-R}^{R} \rho(x) e^{-\frac{1}{2}(x-\mu)^2}\, dx} \\
&= \mathbf{E}_{X_\mu} \left[X_\mu \middle| X_\mu \in [-r, r]\right] \\
&\quad - \mathbf{E}_{X_\mu} \left[X_\mu \middle| X_\mu \in [-R, R]\right]
\end{aligned}$$

Now, the logarithmic derivative of the first term on the RHS with respect to r is

$$\frac{re^{-\frac{1}{2}(r-\mu)^2} + re^{-\frac{1}{2}(-r-\mu)^2}}{\int_{-r}^{r} x\rho(x)e^{-\frac{1}{2}(x-\mu)^2}\,dx} - \frac{e^{-\frac{1}{2}(r-\mu)^2} + e^{-\frac{1}{2}(-r-\mu)^2}}{\int_{-r}^{r} \rho(x)e^{-\frac{1}{2}(x-\mu)^2}\,dx}$$

Its sign is hence the same as that of

$$\frac{r}{\mathbf{E}_{X_\mu}\left[X_\mu \middle| X_\mu \in [-r, r]\right]} - 1$$

Clearly, $-r < \mathbf{E}\left[X_\mu \middle| X_\mu \in [-r, r]\right] < r$, and the expectation will have the same sign as μ (this follows from the symmetry of ρ). Hence if $\mu > 0$, $\mathbf{E}\left[X_\mu \middle| X_\mu \in [-r, r]\right]$ is an increasing function of r, which implies that

$$\Pr_{X_\mu}\left[X_\mu \in [-r, r] \middle| X_\mu \in [-R, R]\right]$$

is a decreasing function of μ, and if $\mu < 0$, the reverse holds. Thus the conditional probability will be maximized at $\mu = 0$. □

Corollary 3.15. If X_μ is a normal random n-vector distributed as $\mathfrak{N}(\mu, I_n)$, and $0 < r < R$, then

$$\Pr_{X_\mu}\left[\|X_\mu\| < r \middle| \|X_\mu\| < R\right] \le \Pr_{X_0}\left[\|X_0\| < r \middle| \|X_0\| < R\right] \le \left(\frac{r}{R}\right)^n e^{\frac{1}{2}(R^2 - r^2)}$$

Proof.

$$\Pr_{X_\mu}\left[\|X_\mu\| < r \middle| \|X_\mu\| < R\right] = \frac{\int_{B(0,r)} e^{-\frac{1}{2}\|\mathbf{x}-\mu\|^2}\,d\mathbf{x}}{\int_{B(0,R)} e^{-\frac{1}{2}\|\mathbf{x}-\mu\|^2}\,d\mathbf{x}}$$

converting to polar coordinates,

$$= \frac{\int_{\|\mathbf{u}\|=1}\int_{-r}^{r} |x|^{n-1} e^{-\frac{1}{2}(x-\mu\cdot\mathbf{u})^2}\,dx\,d\mathbf{u}}{\int_{\|\mathbf{u}\|=1}\int_{-R}^{R} |x|^{n-1} e^{-\frac{1}{2}(x-\mu\cdot\mathbf{u})^2}\,dx\,d\mathbf{u}}$$

by Lemma 3.14,

$$\begin{aligned}
&\le \frac{\int_{-r}^{r} |x|^{n-1}\, e^{-\frac{1}{2}x^2}\,dx}{\int_{-R}^{R} |x|^{n-1}\, e^{-\frac{1}{2}x^2}\,dx} \\
&= \left(\frac{r}{R}\right)^n \frac{\int_{-R}^{R} |x|^{n-1}\, e^{-\frac{1}{2}\left(\frac{r}{R}x\right)^2}\,dx}{\int_{-R}^{R} |x|^{n-1}\, e^{-\frac{1}{2}x^2}\,dx} \\
&\le \left(\frac{r}{R}\right)^n \max_{x\in[-R,R]} e^{\frac{1}{2}\left(1-\left(\frac{r}{R}\right)^2\right)x^2} \\
&= \left(\frac{r}{R}\right)^n e^{\frac{1}{2}(R^2-r^2)}
\end{aligned}$$

□

Theorem 3.16. Let S be a closed, convex subset of $\mathbb{R}^n$. Let V be a k-dimensional subspace of $\mathbb{R}^n$, and assume that the k-dimensional ball of radius R centered at the origin $B_V(0,R)$ is contained in $S \cap V$. Let π_V denote the orthogonal projection map onto V. Then

$$\Pr_{z \sim \mathcal{N}(\mu, I_n)} \left[\|\pi_V(z)\| < \epsilon \middle| z \in S \right]$$
$$\leq \frac{1}{(1-\lambda)^{n-k}} \left(\frac{\epsilon}{\lambda R - (1-\lambda)\epsilon} \right)^k e^{\frac{1}{2}\left((\lambda R - (1-\lambda)\epsilon)^2 - \epsilon^2\right)} e^{\frac{\lambda \|\mu\|^2}{2(2-\lambda)}}$$

for all λ such that $2\epsilon/(R+\epsilon) < \lambda < 1$.

Proof.

$$\Pr_z \left[\|\pi_V(z)\| < \epsilon \middle| z \in S \right] = \frac{\int_{z \in S : \pi_V(z) \in B_V(0,\epsilon)} e^{-\frac{1}{2}\|z-\mu\|^2} \, dz}{\int_{z \in S} e^{-\frac{1}{2}\|z-\mu\|^2} \, dz}$$

in the denominator, substitute $z = T(z')$, where T acts as the identity on V and the contraction $1-\lambda$ on $V^\perp$, giving

$$= \frac{1}{(1-\lambda)^{n-k}} \frac{\int_{z \in S : \pi_V(z) \in B_V(0,\epsilon)} e^{-\frac{1}{2}\|z-\mu\|^2} \, dz}{\int_{T(z) \in S} e^{-\frac{1}{2}\|T(z)-\mu\|^2} \, dz}$$

let $z = x + y$, where $x = \pi_V(z)$, and similarly $\mu = \mu_V + \mu_{V^\perp}$,

$$= \frac{1}{(1-\lambda)^{n-k}} \frac{\int_{x+y \in S, x \in B_V(0,\epsilon)} e^{-\frac{1}{2}\|x-\mu_V\|^2} e^{-\frac{1}{2}\left\|y-\mu_{V^\perp}\right\|^2} \, dx \, dy}{\int_{x+(1-\lambda)y \in S} e^{-\frac{1}{2}\|x-\mu_V\|^2} e^{-\frac{1}{2}\left\|(1-\lambda)y-\mu_{V^\perp}\right\|^2} \, dx \, dy}$$

the ratio of the two y-integrals can be bounded by the maximum of the ratio of the integrands, giving

$$\leq \frac{1}{(1-\lambda)^{n-k}} \times$$
$$\times \max_y \frac{\int_{x : x+y \in S, x \in B_V(0,\epsilon)} e^{-\frac{1}{2}\|x-\mu_V\|^2} \, dx}{\int_{x : x+(1-\lambda)y \in S} e^{-\frac{1}{2}\|x-\mu_V\|^2} \, dx} e^{-\lambda y \cdot ((1-\lambda/2)y - \mu_{V^\perp})}$$

Let $S(y)$ denote the cross-section of S at y, *i.e.*, the set of x such that $x + y \in S$. By convexity,

$$(1-\lambda)S(y) + \lambda B_V(0,R) \subseteq S((1-\lambda)y)$$

Hence as long as $S(y) \cap B_V(0,\epsilon) \neq \emptyset$, we must have

$$B_V(0, \lambda R - (1-\lambda)\epsilon) \subseteq S((1-\lambda)y)$$

Thus applying Corollary 3.15 to the $\boldsymbol{x}$-integrals, we obtain

$$\begin{aligned}\Pr_{z}\left[\|\pi_V(z)\| < \epsilon \middle| z \in S\right] &\le \frac{1}{(1-\lambda)^{n-k}}\left(\frac{\epsilon}{\lambda R - (1-\lambda)\epsilon}\right)^k e^{\frac{1}{2}\left((\lambda R-(1-\lambda)\epsilon)^2-\epsilon^2\right)} \times \\ &\quad \times \max_{y} e^{-\lambda y\cdot((1-\lambda/2)y-\mu_{V^\perp})} \\ &\le \frac{1}{(1-\lambda)^{n-k}}\left(\frac{\epsilon}{\lambda R - (1-\lambda)\epsilon}\right)^k e^{\frac{1}{2}\left((\lambda R-(1-\lambda)\epsilon)^2-\epsilon^2\right)} e^{\frac{\lambda\left\|\mu_{V^\perp}\right\|^2}{2(2-\lambda)}}\end{aligned}$$

□

If we choose λ satisfying

$$\lambda \le \frac{1}{n} \quad \lambda \le \frac{1}{R}$$
$$\lambda \le \frac{1}{\|\mu_{V^\perp}\|^2} \quad \lambda \ge \frac{2\epsilon}{R+2\epsilon}$$

(assuming, of course, that ϵ is small enough that it is possible to satisfy these inequalities simultaneously), then we get

$$\frac{1}{(1-\lambda)^{n-k}} \le (1-1/n)^{-n} \le 4 \qquad \frac{\epsilon}{\lambda R-(1-\lambda)\epsilon} \le \frac{\epsilon}{\lambda R - \lambda R/2} = \frac{2\epsilon}{\lambda R}$$
$$e^{\frac{1}{2}\left((\lambda R-(1-\lambda)\epsilon)^2-\epsilon^2\right)} \le e^{\frac{1}{2}} \qquad e^{\frac{\lambda\left\|\mu_{V^\perp}\right\|^2}{2(2-\lambda)}} \le e^{\frac{1}{2}}$$

Hence for $k \ge 4$,

Corollary 3.17. For $\lambda = \min\left(1/n, 1/R, 1/\|\mu_{V^\perp}\|^2\right)$ and $\epsilon \le \lambda R/2(1-\lambda)$,

$$\Pr_{z\sim\mathfrak{N}(\mu, I_n)}\left[\|\pi_V(z)\| < \epsilon \middle| z \in S\right] \le \left(\frac{4\epsilon}{\lambda R}\right)^k$$

3.6 Bound on $\left\|\langle B\rangle Z^\dagger\right\|$

We will first bound the tail distribution of $BZ^\dagger$.

Lemma 3.18. Let A a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$, permuted into partial pivoting order. Let $0 < k_1 < k$, and $S \subseteq (k_1, k]$, with $s = |S|$. Define

$$B = A_{(k_1,n],S},\ A_1 = A_{(0,k_1],(0,k_1]} \text{ and } Z = A_1^{-1}A_{(0,k_1],S}$$

Then

$$\Pr\left[\left\|BZ^\dagger\right\| > x\right] \le \left(\frac{10k_1(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1}$$

Proof. Since $\left\|BZ^\dagger\right\| \leq \|B\| \left\|Z^\dagger\right\|$,

$$\Pr\left[\left\|BZ^\dagger\right\| > x\right] \leq \Pr\left[\|B\| \left\|Z^\dagger\right\| > x\right]$$

Conditioned on $A_{(0,n],(0,k_1]}$, $\|B\|$ and $\left\|Z^\dagger\right\|$ are independent, and we may apply Theorem 3.10 to obtain

$$\begin{aligned} \Pr\left[\|B\| \left\|Z^\dagger\right\| > x \middle| A_{(0,n],(0,k_1]}, \|B\|\right] & \\ & \leq \frac{7(s/2)^{(k_1-s+1)/2}}{(k_1-s+1)\Gamma\left(\frac{1}{2}(k_1-s+1)\right)} \left(\frac{\|A_1\| \|B\|}{x\sigma}\right)^{k_1-s+1} \end{aligned}$$

and so

$$\begin{aligned} \Pr\left[\left\|BZ^\dagger\right\| > x \middle| A_{(0,n],(0,k_1]}\right] \leq & \frac{7(s/2)^{(k_1-s+1)/2}}{(k_1-s+1)\Gamma\left(\frac{1}{2}(k_1-s+1)\right)} \times \\ & \times \frac{\mathbf{E}\left[\|A_1\|^{k_1-s+1} \|B\|^{k_1-s+1} \middle| A_{(0,n],(0,k_1]}\right]}{(x\sigma)^{k_1-s+1}} \end{aligned}$$

applying Theorem 3.11 twice,

$$\begin{aligned} \Pr\left[\left\|BZ^\dagger\right\| > x\right] & \leq \frac{7(s/2)^{(k_1-s+1)/2}}{(k_1-s+1)\Gamma\left(\frac{1}{2}(k_1-s+1)\right)} \frac{1}{(x\sigma)^{k_1-s+1}} \times \\ & \quad \times \left(2^{k_1-s+1}\Gamma\left(\tfrac{1}{2}(k_1-s+1)\right)(1+\sigma\sqrt{n})^{k_1-s+1}\right)^2 \\ & \leq 7\Gamma\left(\tfrac{1}{2}(k_1-s+1)\right)\left(\frac{2\sqrt{2}\sqrt{s}(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1} \\ & \leq \left(\frac{14\sqrt{2}\sqrt{s(k_1-s)}(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1} \\ & \leq \left(\frac{10k_1(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1} \end{aligned}$$

□

Lemma 3.19. Let A, B and Z be as in Lemma 3.8, but with A a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$ with $\sigma \leq 1$, permuted into partial pivoting order. Then

$$\Pr\left[\left\|\langle B\rangle_{k_1} Z^\dagger\right\| > x\right] \leq \left(\frac{30n(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1}$$

Proof. Since $\langle B\rangle_{k_1} = B - A_{(k_2,n],(0,k_1]}Z$, hence

$$\langle B\rangle_{k_1} Z^\dagger = BZ^\dagger - A_{(k_2,n],(0,k_1]}ZZ^\dagger$$

The norm of the second term is bounded by $\|A\|$, because $ZZ^\dagger$ is a projection. Hence for any $c \in [0,1]$, we have

$$\mathbf{Pr}\left[\left\|\langle B\rangle_{k_1} Z^\dagger\right\| > x\right] \leq \mathbf{Pr}\left[\left\|BZ^\dagger\right\| > cx\right] + \mathbf{Pr}\left[\|A\| > (1-c)x\right]$$

Now

$$\|A\| \leq \left\|\bar{A}\right\| + \sigma\|G\| \leq 1 + \sigma\|G\|$$

where $G \sim \mathfrak{N}(0, I_n \otimes I_k)$, and so

$$\|A\| > (1-c)x \implies \|G\| > \frac{(1-c)x-1}{\sigma}$$

and

$$\mathbf{Pr}\left[\|A\| > (1-c)x\right] \leq \mathbf{exp}\left(-\frac{1}{2}\left(\frac{(1-c)x-1}{\sigma} - 2\sqrt{n}\right)^2\right)$$

Choose c so that

$$\mathbf{exp}\left(-\frac{1}{2}\left(\frac{(1-c)x-1}{\sigma} - 2\sqrt{n}\right)^2\right) \leq \left(\frac{10k_1(1+\sigma\sqrt{n})^2}{cx\sigma}\right)^{k_1-s+1}$$

i.e.,

$$c \leq 1 - \frac{1}{x}\left(1 + \sigma\left(2(k_1-s+1)\ln\left(\frac{cx\sigma}{10k_1(1+\sigma\sqrt{n})^2}\right)\right)^{\frac{1}{2}} + 2\sigma\sqrt{n}\right)$$

This will be true if, say,

$$c = 1 - \frac{1 + 4\sigma\sqrt{n\ln x}}{x}$$

Suppose x is large enough that $c \geq 2/3$. Then we have

$$\mathbf{Pr}\left[\left\|\langle B\rangle_{k_1} Z^\dagger\right\| > x\right] \leq \left(\frac{30k_1(1+\sigma\sqrt{n})^2}{x\sigma}\right)^{k_1-s+1}$$

The statement of the lemma is vacuous unless

$$x \geq \frac{30n(1+\sigma\sqrt{n})^2}{\sigma}$$

which gives

$$\frac{\sqrt{\ln x}}{x} \leq x^{-\frac{3}{4}} \leq \left(\frac{\sigma}{30n}\right)^{\frac{3}{4}}$$

and so

$$\frac{1 + 4\sigma\sqrt{n\ln x}}{x} \leq \frac{1}{30n} + \frac{4}{30^{\frac{3}{4}}n^{\frac{1}{4}}} \leq \frac{1}{3}$$

□

3.7 Bound on $\left\|(\langle C\rangle Z^\dagger)^\dagger\right\|$

Lemma 3.20. Let A, C and Z be as in Lemma 3.8, but with A a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$, permuted into partial pivoting order. Then

$$\Pr\left[\left\|\langle C\rangle_{k_1} Z^\dagger\right\| > x\right] \le 10\left(\frac{6^4 n^{3.5+1/\alpha}(1+\sigma\sqrt{n})^4}{x\sigma^5}\right)^{\frac{1}{4}\min(k_1-s+1,\, s-(k_2-k_1)+1)}$$

where $1+\alpha = s/(k_2-k_1)$.

Proof. We will first bound the probability that $\left\|(\langle C\rangle Z^\dagger)^\dagger \mathbf{u}\right\|$ is large, where $\mathbf{u}$ is a fixed unit vector. Rotating $\mathbf{u}$ by an orthogonal matrix H is equivalent to rotating C and $A_{(k_1,k_2],(0,k_1]}$ by H^T, since

$$(\langle C\rangle Z^\dagger)^\dagger H\mathbf{u} = (H^T\langle C\rangle Z^\dagger)^\dagger \mathbf{u} = (H^T C - H^T A_{(k_1,k_2],(0,k_1]} Z)Z^\dagger)^\dagger \mathbf{u}$$

so we may assume as usual that $\mathbf{u} = \boldsymbol{e}_1$.

The first column of $(\langle C\rangle Z^\dagger)^\dagger$ has length equal to the reciprocal of the component of the first row of $\langle C\rangle Z^\dagger$ orthogonal to the span of the remaining rows. Now

$$\langle C\rangle = C - A_{(k_1,k_2],(0,k_1]} Z$$

We will use a union bound over the possible choices for $A_{(k_1,k_2],(0,k_1]}$ from the $n-k_1$ rows that remain after the first k_1 rounds of partial pivoting. For each fixed choice of the subset of k_2-k_1 rows, the distribution of these rows is Gaussian conditioned on being contained in $\mathbf{PP}\left(A_{(0,k_1],(0,k_1]}\right)$. Hence, for any choice of R and M, we have

$$\begin{aligned}\Pr\left[\left\|(\langle C\rangle Z^\dagger)^\dagger \mathbf{u}\right\| > x\right] &\le \Pr\left[r(A_{(0,k_1],(0,k_1]}) < R\right] + \Pr\left[1+\left\|CZ^\dagger\right\| > M\right] \\ &+ \binom{n-k_1}{k_2-k_1}\Pr\left[\left\|(\langle C\rangle Z^\dagger)^\dagger \mathbf{u}\right\| > x \,\middle|\, r(A_{(0,k_1],(0,k_1]}) \ge R,\, 1+\left\|CZ^\dagger\right\| \le M\right]\end{aligned}$$

The first term on the RHS is bounded by Theorem 3.13,

$$\Pr\left[r(A_{(0,k_1],(0,k_1]}) < R\right] \le \left(\sqrt{\frac{2}{\pi}}\frac{nR}{\sigma}\right)^{n-k_1}$$

the second by Lemma 3.18,

$$\Pr\left[1+\left\|CZ^\dagger\right\| > M\right] \le \left(\frac{10k_1(1+\sigma\sqrt{n})^2}{(M-1)\sigma}\right)^{k_1-s+1} \le \left(\frac{20n(1+\sigma\sqrt{n})^2}{M\sigma}\right)^{k_1-s+1}$$

(where we assume $M \ge 2$) and the third by Corollary 3.17,

$$\begin{aligned}&\binom{n-k_1}{k_2-k_1}\Pr\left[\left\|(\langle C\rangle Z^\dagger)^\dagger \mathbf{u}\right\| > x \,\middle|\, r(A_{(0,k_1],(0,k_1]}) \ge R,\, 1+\left\|CZ^\dagger\right\| \le M\right] \\ &\qquad \le \binom{n-k_1}{k_2-k_1}\left(\frac{4}{x\lambda R}\right)^{s-(k_2-k_1)+1}\end{aligned}$$

where

$$\lambda = \min\left(\frac{1}{k_1}, \frac{\sigma}{R}, \frac{\sigma^2}{M^2}\right)$$

and $1 + \left\|CZ^\dagger\right\|$ upper bounds the $\|\mu\|$ that appears in Corollary 3.17. For the choice of parameters that we will make, it will turn out that $\lambda = \sigma^2/M^2$. Assume this for now and set $s = (1+\alpha)(k_2 - k_1)$, so that

$$\binom{n-k_1}{k_2-k_1} \Pr\left[\left\|(\langle C\rangle\, Z^\dagger)^\dagger \mathbf{u}\right\| > x \middle| r(A_{(0,k_1],(0,k_1]}) \geq R, 1 + \left\|CZ^\dagger\right\| \leq M\right] \leq \left(\frac{4M^2 n^{1/\alpha}}{x\sigma^2 R}\right)^{s-(k_2-k_1)+1}$$

Now choose R and M such that

$$\sqrt{\frac{2}{\pi}}\frac{nR}{\sigma} = \frac{20n(1+\sigma\sqrt{n})^2}{M\sigma} = \frac{4M^2 n^{1/\alpha}}{x\sigma^2 R}$$
$$= \left(\sqrt{\frac{2}{\pi}}\frac{nR}{\sigma}\left(\frac{20n(1+\sigma\sqrt{n})^2}{M\sigma}\right)^2 \frac{4M^2 n^{1/\alpha}}{x\sigma^2 R}\right)^{1/4} = (5.9\cdots)\frac{n^{(3+1/\alpha)/4}(1+\sigma\sqrt{n})}{x^{1/4}\sigma^{5/4}}$$

This gives

$$M \geq \frac{10}{3}(n^{1-1/\alpha}\sigma x)^{1/4}(1+\sigma\sqrt{n}) \geq \sigma\sqrt{n}$$
$$\frac{M^2}{\sigma R} \geq \frac{5n^{3/4(1-1/\alpha)}(1+\sigma\sqrt{n})x^{3/4}}{4\sigma^{1/4}} \geq 1$$

for $n^{1-1/\alpha}\sigma x \geq 1$, which is true whenever the bound in the statement of the lemma is non-trivial. So with this choice of R and M, in fact $\lambda = \sigma^2/M^2$.

Hence

$$\Pr\left[\left\|(\langle C\rangle\, Z^\dagger)^\dagger \mathbf{u}\right\| > x\right] \leq 3\left(\frac{6^4 n^{3+1/\alpha}(1+\sigma\sqrt{n})^4}{x\sigma^5}\right)^{\frac{1}{4}\min(k_1-s+1, s-(k_2-k_1)+1)}$$

since $n - k_1 > k - k_1 \geq s \geq s - (k_2 - k_1) + 1$. By a now familiar argument (see proof of Theorem 3.10, for example), this leads to

$$\Pr\left[\left\|(\langle C\rangle\, Z^\dagger)^\dagger\right\| > x\right] \leq 10\left(\frac{6^4 n^{3.5+1/\alpha}(1+\sigma\sqrt{n})^4}{x\sigma^5}\right)^{\frac{1}{4}\min(k_1-s+1, s-(k_2-k_1)+1)}$$

□

3.8 Choosing parameters

In Theorem 3.9, we will choose

$$\begin{aligned} k_1 &= 2k/3 \\ k - k_{i+1} &= \frac{2}{3}(k - k_i) \\ k_r &= \log k \\ s_i &= k - k_i \implies S_i = (k_i, k] \end{aligned}$$

This corresponds to having $\alpha = 2$ in Lemma 3.20. The number of factors r will be

$$\frac{\log(k/3 \log k)}{\log(3/2)} \leq 2 \log k$$

So for $1 \leq i \leq r-1$, we have by Lemma 3.19 and 3.20,

$$\begin{aligned} \mathbf{Pr}\left[\left\| \langle B_i \rangle_{k_i} Z_i^\dagger \right\| > x_1\right] &\leq \left(\frac{30n(1+\sigma\sqrt{n})^2}{x_1\sigma}\right)^{2k_i - k + 1} &\leq \left(\frac{30n(1+\sigma\sqrt{n})^2}{x_1\sigma}\right)^{k/3} \\ \mathbf{Pr}\left[\left\| \langle C_i \rangle_{k_i} Z_i^\dagger \right\| > x_2\right] &\leq 10\left(\frac{6n(1+\sigma\sqrt{n})}{x_2^{1/4}\sigma^{5/4}}\right)^{k-k_{i+1}+1} &\leq \left(\frac{10n(1+\sigma\sqrt{n})}{x_2^{1/4}\sigma^{5/4}}\right)^{k-k_{i+1}} \end{aligned}$$

for k large enough. Since $k - k_{i+1} \leq 2k/9$, we pick x_1 and x_2 such that

$$\begin{aligned} \left(\frac{30n(1+\sigma\sqrt{n})^2}{x_1\sigma}\right) &= \left(\frac{10n(1+\sigma\sqrt{n})}{x_2^{1/4}\sigma^{5/4}}\right)^{2/3} \\ &= \left(\frac{30 \cdot 10^4 n^5 (1+\sigma\sqrt{n})^6}{x_1 x_2 \sigma^6}\right)^{1/7} \leq \frac{7n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{(x_1x_2)^{1/7}\sigma^{6/7}} \end{aligned}$$

and obtain

$$\mathbf{Pr}\left[\left\| \left(\langle B_i \rangle_{k_i} Z_i^\dagger\right)\left(\langle C_i \rangle_{k_i} Z_i^\dagger\right)^\dagger \right\| > x\right] \leq 2\left(\frac{7n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{x^{1/7}\sigma^{6/7}}\right)^{\frac{3}{2}(k-k_{i+1})}$$

which implies

$$\mathbf{Pr}\left[1 + \left\| \left(\langle B_i \rangle_{k_i} Z_i^\dagger\right)\left(\langle C_i \rangle_{k_i} Z_i^\dagger\right)^\dagger \right\| > x_i\right] \leq \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{x_i^{1/7}\sigma^{6/7}}\right)^{k-k_i}$$

for $x_i \geq 2$. Choose the x_i so that

$$\begin{aligned}
\left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{x_1^{1/7}\sigma^{6/7}}\right)^{k-k_1} &= \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{x_2^{1/7}\sigma^{6/7}}\right)^{k-k_2} \\
&= \ldots = \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{x_{r-1}^{1/7}\sigma^{6/7}}\right)^{k-k_{r-1}} \\
&= \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{\sigma^{6/7}}\right)^{(r-1)/\sum_i \frac{1}{k-k_i}} \times \\
&\quad \times \left(\prod_i x_i\right)^{-1/7\sum_i \frac{1}{k-k_i}} \\
&\leq \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{\sigma^{6/7}}\right)^{\frac{2}{3}\log^2 k} \left(\prod x_i\right)^{-\frac{1}{21}\log k}
\end{aligned}$$

Thus,

$$\begin{aligned}
&\Pr\left[\left\|\prod_{i=r-1}^{1}\left[-\left(\langle B_i\rangle_{k_i} Z_i^{\dagger}\right)\left(\langle C_i\rangle_{k_i} Z_i^{\dagger}\right)^{\dagger} \quad ; \quad I\right]\right\| > x\right] \\
&\qquad\qquad \leq \left(\frac{10n^{5/7}(1+\sigma\sqrt{n})^{6/7}}{\sigma^{6/7}}\right)^{\frac{2}{3}\log^2 k} x^{-\frac{1}{21}\log k}
\end{aligned}$$

We also have, from [6] and Theorem 3.10,

$$\Pr\left[\|B_0\| > x\right] \leq \exp\left(-\frac{1}{2}\left(\frac{x-1}{\sigma} - 2\sqrt{n}\right)^2\right)$$

$$\Pr\left[\left\|C_0^{\dagger}\right\| > x\right] \leq \binom{n}{2k/3}\frac{7(k/3)^{k/6}}{(k/3)\Gamma(k/6)(x\sigma)^{k/3}} \leq \left(\frac{n^2}{x\sigma}\right)^{k/3}$$

and since A_r has only $\log k$ rows, we can use the worst-case growth of partial pivoting to obtain

$$\left\|\langle X_r\rangle\langle A_r\rangle^{-1}\right\| \leq k\log k$$

Putting everything together, we get

Theorem 3.21. **If $A \in \mathbb{R}^{n\times k}$ is a random matrix distributed as $\mathfrak{N}\left(\bar{A}, \sigma^2 I_n \otimes I_k\right)$ with $\left\|\bar{A}\right\| \leq 1$, and $\rho_U(A)$ is the growth factor during Gaussian elimination with partial pivoting, then**

$$\Pr_A\left[\rho_U(A) > x\right] \leq \left(\frac{1}{x}\left(\mathcal{O}\left(\frac{n(1+\sigma\sqrt{n})}{\sigma}\right)\right)^{12\log k}\right)^{\frac{1}{21}\log k}$$

Chapter 4

Conclusions and open problems

The most important contribution of this thesis is to establish a theoretical bound on the growth in Gaussian elimination with partial pivoting that is better than 2^{n-1}. The bound that we have managed to prove is, however, still much larger than the experimentally observed growth factors. In this chapter, we attempt to clarify the limitations of the method used.

4.1 Limitations of the proof

The argument presented in Chapter 3 is remarkable in that very little use is made of the improvement from partial pivoting that one expects and indeed, observes experimentally. Most of the proof is in fact devoted to showing that rearranging the rows does not significantly *worsen* the situation, as compared to not pivoting. We have used a union bound in Lemma 3.20 to prove this, and this is the technical reason why we require a logarithmic number of "stages" in the proof, and ultimately the reason why our bound is of order $(n/\sigma)^{\mathcal{O}(\log n)}$.

This technique thus does not take advantage of the fact that partial pivoting appears to significantly mitigate the effects of a large pivot. That is, if $A_{(0,k],(0,k]}$ has a small singular value, typically the next step chooses a row that removes this small singular value, so that $A_{(0,k+1],(0,k+1]}$ is much better conditioned. So the strength of our method of argument, that it manages to get by only "touching" the algorithm at a logarithmic number of places, is also its weakness, since it cannot take advantage of the systematic improvement that partial pivoting produces.

4.2 Improvements

The most direct way of improving the bound we have proved is to reduce the number of stages we use in Section 3.8. This in turn depends on improving the proof of Lemma 3.20 so that we can make larger steps between stages. This will cut the exponent in the bound for the growth factor, but ultimately cannot reduce it to a constant.

To get a polynomial bound on the growth factor it appears necessary to understand better the effect of partial pivoting on the distribution of the remaining rows after each step of elimination. This would appear to be the most fruitful area for future research.

Appendix A

Technical Results

A.1 Gaussian random variables

We recall that the probability density function of a d-dimensional Gaussian random vector with covariance matrix $\sigma^2 I_d$ and mean $\bar{\mu}$ is given by

$$\mathfrak{n}(\mu, \sigma^2 I_d)(\mathbf{x}) = \frac{1}{(2\pi\sigma^2)^{d/2}} e^{-\frac{1}{2\sigma^2}\mathbf{dist}(\mathbf{x},\mu)^2}$$

Lemma A.1. Let x be a standard normal variable. Then,

$$\Pr[x \geq k] \leq \frac{1}{\sqrt{2\pi}} \frac{e^{-\frac{1}{2}k^2}}{k}.$$

for all $k > 1$.

Proof. We have

$$\Pr[x \geq k] = \frac{1}{\sqrt{2\pi}} \int_k^\infty e^{-\frac{1}{2}x^2}\, dx$$

putting $t = \frac{1}{2}x^2$,

$$\begin{aligned} &\leq \frac{1}{\sqrt{2\pi}} \int_{\frac{1}{2}k^2}^\infty \frac{e^{-t}}{k}\, dt \\ &= \frac{1}{\sqrt{2\pi}} \frac{e^{-\frac{1}{2}k^2}}{k}. \end{aligned}$$

$\square$

Lemma A.2. Let $\mathbf{x}$ be a d-dimensional Gaussian random vector of variance σ^2 and let $\mathcal{H}$ be a hyperplane. Then,

$$\Pr[\mathbf{dist}(\mathbf{x}, \mathcal{H}) \leq \epsilon] \leq \sqrt{2/\pi}\epsilon/\sigma.$$

Lemma A.3. Let $g_1, \ldots, g_n$ be Gaussian random variables of mean 0 and variance 1. Then,

$$\mathrm{E}\left[\max_i |g_i|\right] \leq \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}.$$

Proof.

$$\begin{aligned}
\mathrm{E}\left[\max_i |g_i|\right] &= \int_{t=0}^{\infty} \Pr\left[\max_i |g_i| \geq t\right]\, dt \\
&\leq \int_{t=0}^{\sqrt{2\log n}} 1\, dt + \int_{\sqrt{2\log n}}^{\infty} n \Pr\left[|g_1| \geq t\right]\, dt
\end{aligned}$$

applying Lemma A.1,

$$\begin{aligned}
&\leq \sqrt{2\log n} + \int_{\sqrt{2\log n}}^{\infty} n \frac{2}{\sqrt{2\pi}} \frac{e^{-\frac{1}{2}t^2}}{t}\, dt \\
&\leq \sqrt{2\log n} + \frac{n}{\sqrt{\log n}} \int_{\sqrt{2\log n}}^{\infty} \frac{1}{\sqrt{\pi}} e^{-\frac{1}{2}t^2}\, dt \\
&\leq \sqrt{2\log n} + \frac{n}{\sqrt{\log n}} \frac{e^{-\frac{1}{2}(\sqrt{2\log n})^2}}{\sqrt{2\pi\log n}} \\
&= \sqrt{2\log n} + \frac{1}{\sqrt{2\pi}\log n}
\end{aligned}$$

□

Lemma A.4 (Expectation of reciprocal of the L1 norm of a Gaussian vector). Let $\mathbf{a}$ be an n-dimensional Gaussian random vector of variance σ^2, for $n \geq 2$. Then

$$\mathrm{E}\left[\frac{1}{\|\mathbf{a}\|_1}\right] \leq \frac{2}{n\sigma}$$

Proof. Let $\mathbf{a} = (a_1, \ldots, a_n)$. Without loss of generality, we assume $\sigma^2 = 1$. For general σ, we can simply scale the bound by the factor $1/\sigma$. It is also clear that the expectation of $1/\|\mathbf{a}\|_1$ is maximized if the mean of $\mathbf{a}$ is zero, so we will make this assumption.

Recall that the Laplace transform of a positive random variable X is defined by

$$\mathcal{L}[X](t) = \mathrm{E}_X\left[e^{-tX}\right]$$

and the expectation of the reciprocal of a random variable is simply the integral of its Laplace transform.

Let X be the absolute value of a standard normal random variable. The Laplace transform of X is given by

$$\begin{aligned}
\mathcal{L}[X](t) &= \sqrt{\frac{2}{\pi}}\int_0^\infty e^{-tx}e^{-\frac{1}{2}x^2}\,dx \\
&= \sqrt{\frac{2}{\pi}}e^{\frac{1}{2}t^2}\int_0^\infty e^{-\frac{1}{2}(x+t)^2}\,dx \\
&= \sqrt{\frac{2}{\pi}}e^{\frac{1}{2}t^2}\int_t^\infty e^{-\frac{1}{2}x^2}\,dx \\
&= e^{\frac{1}{2}t^2}\operatorname{erfc}\left(\frac{t}{\sqrt{2}}\right).
\end{aligned}$$

We now set a constant $c = 2.4$ and set α to satisfy

$$1 - \frac{\sqrt{c/\pi}}{\alpha} = e^{\frac{1}{2}(c/\pi)}\operatorname{erfc}\left(\frac{\sqrt{c/\pi}}{\sqrt{2}}\right).$$

As $e^{\frac{1}{2}t^2}\operatorname{erfc}\left(\frac{t}{\sqrt{2}}\right)$ is convex, we have the upper bound

$$e^{\frac{1}{2}t^2}\operatorname{erfc}\left(\frac{t}{\sqrt{2}}\right) \leq 1 - \frac{t}{\alpha},\ \text{for } 0 \leq t \leq \sqrt{c/\pi}.$$

For $t > \sqrt{c/\pi}$, we apply the upper bound

$$e^{\frac{1}{2}t^2}\operatorname{erfc}\left(\frac{t}{\sqrt{2}}\right) \leq \sqrt{\frac{2}{\pi}}\frac{1}{t}.$$

We now have

$$\begin{aligned}
\mathbf{E}\left[\frac{1}{\|\mathbf{a}\|_1}\right] &= \int_0^\infty \left(e^{\frac{1}{2}t^2}\operatorname{erfc}(t/\sqrt{2})\right)^n dt \\
&\leq \int_0^{\sqrt{c/\pi}} \left(1 - \frac{t}{\alpha}\right)^n dt + \int_{\sqrt{c/\pi}}^\infty \left(\sqrt{\frac{2}{\pi}}\frac{1}{t}\right)^n dt \\
&\leq \frac{\alpha}{n+1} + \sqrt{\frac{2}{\pi}}\frac{(2/c)^{(n-1)/2}}{n-1} \\
&\leq \frac{2}{n-1},
\end{aligned}$$

for $n \geq 2$. □

A.2 Random point on sphere

Lemma A.5. Let $u_1, \ldots, u_d$ be a unit vector chosen uniformly at random in $\mathbb{R}^d$. Then, for $c \leq 1$,

$$\Pr\left[|u_1| \geq \sqrt{\frac{c}{d}}\right] \geq \Pr\left[|g| \geq \sqrt{c}\right],$$

where g is a Gaussian random variable of variance 1 and mean 0.

Proof. We may obtain a random unit vector by choosing d Gaussian random variables of variance 1 and mean 0, $x_1, \ldots, x_d$, and setting

$$u_i = \frac{x_i}{\sqrt{x_1^2 + \cdots + x_d^2}}.$$

We have

$$\begin{aligned}
\Pr\left[u_1^2 \geq \frac{c}{d}\right] &= \Pr\left[\frac{x_1^2}{x_1^2 + \cdots + x_d^2} \geq \frac{c}{d}\right] \\
&= \Pr\left[\frac{(d-1)x_1^2}{x_2^2 \cdots + x_d^2} \geq \frac{(d-1)c}{d-c}\right] \\
&\geq \Pr\left[\frac{(d-1)x_1^2}{x_2^2 \cdots + x_d^2} \geq c\right], \text{ since } c \leq 1.
\end{aligned}$$

We now note that

$$t_d \stackrel{\text{def}}{=} \frac{\sqrt{(d-1)}x_1}{\sqrt{x_2^2 \cdots + x_d^2}}$$

is the random variable distributed according to the t-distribution with d degrees of freedom. The lemma now follows from the fact (*c.f.* [15, Chapter 28, Section 2] or [2, 26.7.5]) that, for $c > 0$,

$$\Pr\left[t_d > \sqrt{c}\right] \geq \Pr\left[g > \sqrt{c}\right],$$

and that the distributions of t_d and g are symmetric about the origin. □

A.3 Combination Lemma

Lemma A.6. Let A and B be two positive random variables. Assume

1. $\Pr[A \geq x] \leq f(x)$.

2. $\Pr[B \geq x | A] \leq g(x)$.

where g is monotonically decreasing and $\lim_{x \to \infty} g(x) = 0$. Then,

$$\Pr[AB \geq x] \leq \int_0^\infty f\left(\frac{x}{t}\right)(-g'(t))\, dt$$

Proof. Let μ_A denote the probability measure associated with A. We have

$$\begin{aligned}
\Pr[AB \geq x] &= \int_0^\infty \Pr_B[B \geq x/t | A]\; d\mu_A(t) \\
&\leq \int_0^\infty g\left(\frac{x}{t}\right)\; d\mu_A(t)
\end{aligned}$$

integrating by parts,

$$\begin{aligned}
&= \int_0^\infty \Pr[A \geq t] \frac{d}{dt} g\left(\frac{x}{t}\right) dt \\
&\leq \int_0^\infty f(t) \frac{d}{dt} g\left(\frac{x}{t}\right) dt \\
&= \int_0^\infty f\left(\frac{x}{t}\right)(-g'(t))\, dt
\end{aligned}$$

□

Corollary A.7 (linear-linear). Let A and B be two positive random variables. Assume

1. $\Pr[A \geq x] \leq \frac{\alpha}{x}$ and
2. $\Pr[B \geq x|A] \leq \frac{\beta}{x}$

for some $\alpha, \beta > 0$. Then,

$$\Pr[AB \geq x] \leq \frac{\alpha\beta}{x}\left(1 + \ln\left(\frac{x}{\alpha\beta}\right)\right)$$

Proof. As the probability of an event can be at most 1,

$$\begin{aligned}
\Pr[A \geq x] &\leq \min\left(\frac{\alpha}{x}, 1\right) \stackrel{\text{def}}{=} f(x), \text{ and} \\
\Pr[B \geq x] &\leq \min\left(\frac{\beta}{x}, 1\right) \stackrel{\text{def}}{=} g(x).
\end{aligned}$$

Applying Lemma A.6 while observing

- $g'(t) = 0$ for $t \in [0, \beta]$, and
- $f(x/t) = 1$ for $t \geq x/\alpha$,

we obtain

$$\begin{aligned}
\Pr[AB \geq x] &\leq \int_0^\beta \frac{\alpha t}{x} \cdot 0\, dt + \int_\beta^{x/\alpha} \frac{\alpha t}{x} \frac{\beta}{t^2}\, dt + \int_{x/\alpha}^\infty \frac{\beta}{t^2}\, dt \\
&= \frac{\alpha\beta}{x} \int_\beta^{x/\alpha} \frac{dt}{t} + \frac{\alpha\beta}{x} \\
&= \frac{\alpha\beta}{x}\left(1 + \ln\left(\frac{x}{\alpha\beta}\right)\right).
\end{aligned}$$

□

Corollary A.8. Let A and B be two positive random variables. Assume

1. $\Pr[A \geq x] \leq \min\left(1, \frac{\alpha+\beta\sqrt{\ln x\sigma}}{\sigma x}\right)$.

2. $\Pr[B \geq x | A] \leq \frac{\gamma}{x\sigma}$.

for some $\alpha \geq 1$ and $\beta, \gamma, \sigma > 0$. Then,

$$\Pr[AB \geq x] \leq \frac{\alpha\gamma}{x\sigma^2}\left(1 + \left(\frac{2\beta}{3\alpha} + 1\right)\ln^{3/2}\left(\frac{x\sigma^2}{\gamma}\right)\right).$$

Proof. Define f and g by

$$f(x) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{for } x \leq \frac{\alpha}{\sigma} \\ \frac{\alpha+\beta\sqrt{\ln x\sigma}}{x\sigma} & \text{for } x > \frac{\alpha}{\sigma} \end{cases}$$

$$g(x) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{for } x \leq \frac{\gamma}{\sigma} \\ \frac{\gamma}{x\sigma} & \text{for } x > \frac{\gamma}{\sigma} \end{cases}$$

Applying Lemma A.6 while observing

- $g'(t) = 0$ for $t \in \left[0, \frac{\gamma}{\sigma}\right]$, and
- $f(x/t) = 1$ for $t \geq x\sigma/\alpha$,

we obtain

$$\begin{aligned}
\Pr[AB \geq x] &\leq \int_{\gamma/\sigma}^{x\sigma/\alpha} \frac{\alpha + \beta\sqrt{\ln(x\sigma/t)}}{x\sigma/t}\frac{\gamma}{t^2\sigma}\,dt + \int_{x\sigma/\alpha}^{\infty}\frac{\gamma}{\sigma t^2}\,dt \\
&= \int_{\gamma/\sigma}^{x\sigma/\alpha} \frac{\alpha + \beta\sqrt{\ln(x\sigma/t)}}{x\sigma^2}\frac{\gamma}{t}\,dt + \frac{\alpha\gamma}{x\sigma^2}
\end{aligned}$$

(substituting $s = \sqrt{\ln(x\sigma/t)}$, $t = x\sigma e^{-s^2}$)

$$\begin{aligned}
&= \int_{\sqrt{\ln(x\sigma^2/\gamma)}}^{\sqrt{\ln\alpha}} \frac{\alpha + \beta s}{x\sigma^2}\frac{\gamma}{x\sigma e^{-s^2}}x\sigma(-2se^{-s^2})\,ds + \frac{\alpha\gamma}{x\sigma^2} \\
&= \frac{\gamma}{x\sigma^2}\int_{\sqrt{\ln\alpha}}^{\sqrt{\ln(x\sigma^2/\gamma)}} 2s(\alpha + \beta s)\,ds + \frac{\alpha\gamma}{x\sigma^2} \\
&= \frac{\alpha\gamma}{x\sigma^2}\left(1 + \ln\left(\frac{x\sigma^2}{\alpha\gamma}\right) + \frac{2\beta}{3\alpha}\left(\ln^{3/2}\left(\frac{x\sigma^2}{\gamma}\right) - \ln^{3/2}\alpha\right)\right) \\
&\leq \frac{\alpha\gamma}{x\sigma^2}\left(1 + \left(\frac{2\beta}{3\alpha} + 1\right)\ln^{3/2}\left(\frac{x\sigma^2}{\gamma}\right)\right),
\end{aligned}$$

as $\alpha \geq 1$. $\square$

Lemma A.9 (linear-bounded expectation). Let A, B and C be positive random variables such that

$$\Pr[A \geq x] \leq \frac{\alpha}{x},$$

for some $\alpha > 0$, and

$$\forall A, \ \Pr[B \geq x|A] \leq \Pr[C \geq x].$$

Then,

$$\Pr[AB \geq x] \leq \frac{\alpha}{x} \mathbf{E}[C].$$

Proof. Let $g(x)$ be the distribution function of C. By Lemma A.6, we have

$$\begin{aligned}
\Pr[AB \geq x] &\leq \int_0^\infty \left(\frac{\alpha t}{x}\right) (-(1-g)'(t)) \, dt \\
&= \frac{\alpha}{x} \int_0^\infty t(g'(t)) \, dt \\
&= \frac{\alpha}{x} \mathbf{E}[C].
\end{aligned}$$

□

Corollary A.10 (linear-chi). Let A a be positive random variable such that

1. $\Pr[A \geq x] \leq \frac{\alpha}{x}$.

for some $\alpha > 0$. For every A, let $\mathbf{b}$ be a d-dimensional Gaussian random vector (possibly depending upon A) of variance at most σ^2 centered at a point of norm at most k, and let $B = \|\mathbf{b}\|$. Then,

$$\Pr[AB \geq x] \leq \frac{\alpha\sqrt{\sigma^2 d + k^2}}{x}$$

Proof. As $\mathbf{E}[B] \leq \sqrt{\mathbf{E}[B^2]}$, and it is known [16, p. 277] that the expected value of B^2—the non-central χ^2-distribution with non-centrality parameter $\left\|\bar{\mathbf{b}}\right\|^2$—is $d + \left\|\bar{\mathbf{b}}\right\|^2$, the corollary follows from Lemma A.9. □

Lemma A.11 (Linear to log). Let A be a positive random variable. Assume

$$\Pr_A[A \geq x] \leq \frac{\alpha}{x},$$

for some $\alpha \geq 1$. Then,

$$\mathbf{E}_A[\log A] \leq \log \alpha + 1.$$

Proof.

$$\begin{aligned}
\mathbf{E}_A[\log A] &= \int_{x=0}^\infty \Pr_A[\log A \geq x] \, dx = \int_{x=0}^\infty \min(1, \frac{\alpha}{e^x}) dx \\
&= \int_{x=0}^{\log \alpha} dx + \int_{x=\log \alpha}^\infty \alpha e^{-x} dx = \log \alpha + 1.
\end{aligned}$$

□

Bibliography

[1] `http://math.mit.edu/~spielman/SmoothedAnalysis`.

[2] Milton Abramowitz and Irene A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, volume 55 of *Applied mathematics series*. U. S. Department of Commerce, Washington, DC, USA, 1964. Tenth printing, with corrections (December 1972).

[3] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide, Third Edition*. SIAM, Philadelphia, 1999.

[4] Avrim Blum and John Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *SODA '02*, pages 905–914, 2002.

[5] Lenore Blum. Lectures on a theory of computation and complexity over the reals (or an arbitrary ring). In Erica Jen, editor, *The Proceedings of the 1989 Complex Systems Summer School, Santa Fe, New Mexico*, volume 2, pages 1–47, June 1989.

[6] K. R. Davidson and S. J. Szarek. *Handbook on the Geometry of Banach spaces*, chapter Local operator theory, random matrices, and Banach spaces, pages 317–366. Elsevier Science, 2001.

[7] James Demmel. The probability that a numerical analysis problem is difficult. *Math. Comp.*, pages 499–480, 1988.

[8] James Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.

[9] John Dunagan, Daniel A. Spielman, and Shang-Hua Teng. Smoothed analysis of Renegar's condition number for linear programming. Available at `http://math.mit.edu/~spielman/SmoothedAnalysis`, 2002.

[10] Alan Edelman. Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.*, 9(4):543–560, 1988.

[11] Alan Edelman. Eigenvalue roulette and random test matrices. In Marc S. Moonen, Gene H. Golub, and Bart L. R. De Moor, editors, *Linear Algebra for Large Scale and Real-Time Applications*, NATO ASI Series, pages 365–368. 1992.

[12] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins Series in the Mathematical Sciences. The Johns Hopkins University Press and North Oxford Academic, Baltimore, MD, USA and Oxford, England, 1983.

[13] Martin Grotschel, Laszlo Lovasz, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1991.

[14] Nick Higham. How accurate is Gaussian elimination? In *Numerical Analysis 1989, Proceedings of the 13th Dundee Conference, volume 228 of Pitman Research Notes in Mathematics*, pages 137–154, 1990.

[15] N. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 2. Wiley-Interscience, 1995.

[16] Samuel Kotz and Norman L. Johnson, editors. *Encyclopedia of Statistical Sciences*, volume 6. John Wiley & Sons, 1982.

[17] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer-Verlag, 1991.

[18] J. Renegar. Incorporating condition measures into the complexity theory of linear programming. *SIAM J. Optim.*, 5(3):506–524, 1995.

[19] Yoav Seginer. The expected norm of random matrices. *Combinatorics, Probability and Computing*, 9:149–166, 2000.

[20] Steve Smale. Complexity theory and numerical analysis. *Acta Numerica*, pages 523–551, 1997.

[21] Daniel Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing (STOC '01)*, pages 296–305, 2001.

[22] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.

[23] Lloyd N. Trefethen and Robert S. Schreiber. Average-case stability of Gaussian elimination. *SIAM Journal on Matrix Analysis and Applications*, 11(3):335–360, 1990.

[24] J. H. Wilkinson. Error analysis of direct methods of matrix inversion. *J. Assoc. Comput. Mach.*, 8:261–330, 1961.

[25] Man-Chung Yeung and Tony F. Chan. Probabilistic analysis of Gaussian elimination without pivoting. *SIAM J. Matrix Anal. Appl.*, 18(2):499–517, 1997.