

Smoothed Analysis of the Condition Numbers and Growth Factors of Matrices

Arvind Sankar *
Department of Mathematics
Massachusetts Institute of Technology

Daniel A. Spielman †
Department of Computer Science
Yale University

Shang-Hua Teng ‡
Department of Computer Science
Boston University and
Akamai Technologies Inc.

December 8, 2005

Abstract

Let $\bar{\mathbf{A}}$ be an arbitrary matrix and let \mathbf{A} be a slight random perturbation of $\bar{\mathbf{A}}$. We prove that it is unlikely that \mathbf{A} has large condition number. Using this result, we prove it is unlikely that \mathbf{A} has large growth factor under Gaussian elimination without pivoting. By combining these results, we show that the smoothed precision necessary to solve $\mathbf{Ax} = \mathbf{b}$, for any \mathbf{b} , using Gaussian elimination without pivoting is logarithmic. Moreover, when $\bar{\mathbf{A}}$ is an all-zero square matrix, our results significantly improve the average-case analysis of Gaussian elimination without pivoting performed by Yeung and Chan (SIAM J. Matrix Anal. Appl., 1997).

*Partially supported by NSF grant CCR-0112487

†Partially supported by an Alfred P. Sloan Foundation Fellowship, and NSF grants CCR-0112487 and CCR-0324914

‡Partially supported by an Alfred P. Sloan Foundation Fellowship, NSF grant CCR-9972532, and NSF grants CCR-0311430 and ITR CCR-0325630.

1 Introduction

Spielman and Teng [ST04], introduced the smoothed analysis of algorithms to explain 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 measure of the performance of an algorithm 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 [Smo])

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.

In this paper, we prove that perturbations of arbitrary matrices are unlikely to have large condition numbers or large growth factors under Gaussian Elimination without pivoting. As a consequence, we conclude that the smoothed precision necessary for Gaussian elimination is logarithmic. We obtain similar results for perturbations that affect only the non-zero and diagonal entries of symmetric matrices. We hope that these results will be a first step toward a smoothed analysis of Gaussian elimination with partial pivoting—an algorithm that is widely used in practice but known to have poor worst-case performance.

In the rest of this section, we recall the definitions of the condition numbers and growth factors of matrices, and review prior work on their average-case analysis. In Section 3, we perform a smoothed analysis of the condition number of a matrix. In Section 4, we use the results of Section 3 to obtain a smoothed analysis of the growth factors of Gaussian elimination without pivoting. In Section 5, we combine these results to obtain a smoothed bound on the precision needed by Gaussian elimination without pivoting. Definitions of zero-preserving perturbations and our results on perturbations that only affect the non-zero and diagonal entries of symmetric matrices appear in Section 6. In the conclusion section, we explain how our results may be extended to larger families of perturbations, present some counter-examples, and suggest future directions for research. Other conjectures and open questions appear in the body of the paper.

The analysis in this paper requires many results from probability. Where reasonable, these have been deferred to the appendix.

1.1 Condition numbers and growth factors

We use the standard notation for the 1, 2 and ∞ -norms of matrices and column vectors, and define

$$\|\mathbf{A}\|_{\max} = \max_{i,j} |\mathbf{A}_{i,j}|.$$

Definition 1.1 (Condition Number). For a square matrix \mathbf{A} , the condition number of \mathbf{A} is defined by

$$\kappa(\mathbf{A}) = \|\mathbf{A}\|_2 \left\| \mathbf{A}^{-1} \right\|_2.$$

The condition number measures how much the solution to a system $\mathbf{Ax} = \mathbf{b}$ changes as one makes slight changes to \mathbf{A} and \mathbf{b} . A consequence is that if one solves the linear system using fewer than $\log(\kappa(\mathbf{A}))$ bits of precision, one is likely to obtain a result far from a solution. For more information on the condition number of a matrix, we refer the reader to one of [GL83, TB97, Dem97].

The simplest and most often implemented method of solving linear systems is Gaussian elimination. Natural implementations of Gaussian elimination use $\mathcal{O}(n^3)$ arithmetic operations to solve a system of n linear equations in n variables. If the coefficients of these equations are specified using b bits, in the worst case it suffices to perform the elimination using $\mathcal{O}(bn)$ bits of precision [GLS91]. This high precision may be necessary because the elimination may produce large intermediate entries [TB97]. However, in practice one usually obtains accurate answers using much less precision. In fact, it is rare to find an implementation of Gaussian elimination that uses anything more than double precision, and high-precision solvers are rarely used or needed in practice [TB97, TS90] (for example, LAPACK uses 64 bits [ABB⁺99]). One of the main results of this paper is that $\mathcal{O}(b + \log n)$ bits of precision usually suffice for Gaussian elimination in the smoothed analysis framework.

Since Wilkinson’s seminal work [Wil61], it has been understood that it suffices to carry out Gaussian elimination with $b + \log_2(5n\kappa(\mathbf{A}) \|\mathbf{L}\|_{\infty} \|\mathbf{U}\|_{\infty} / \|\mathbf{A}\|_{\infty} + 3)$ bits of accuracy to obtain a solution that is accurate to b bits. In this formula, \mathbf{L} and \mathbf{U} are the LU-decomposition of \mathbf{A} ; that is, \mathbf{U} is the upper-triangular matrix and \mathbf{L} is the lower-triangular matrix with 1s on the diagonal for which $\mathbf{A} = \mathbf{LU}$.

1.2 Prior work

The average-case behaviors of the condition numbers and growth factors of matrices have been studied both analytically and experimentally. In his paper, “The probability that a numerical analysis problem is difficult”, Demmel [Dem88] 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 [Ede88].

Average-case analysis of growth factors began with the experimental work of Trefethen and Schreiber [TS90], who found that Gaussian random matrices rarely have large growth factors under partial or full pivoting.

Definition 1.2 (Gaussian Matrix). A matrix \mathbf{G} is a *Gaussian random matrix* of variance σ^2 if each entry of \mathbf{G} is an independent univariate Gaussian variable with mean 0 and standard deviation σ .

Yeung and Chan [YC97] study the growth factors of Gaussian elimination without pivoting on Gaussian random matrices of variance 1. They define $\rho_{\mathbf{U}}$ and $\rho_{\mathbf{L}}$ by

$$\begin{aligned}\rho_{\mathbf{U}}(\mathbf{A}) &= \|\mathbf{U}\|_{\infty} / \|\mathbf{A}\|_{\infty}, \text{ and} \\ \rho_{\mathbf{L}}(\mathbf{A}) &= \|\mathbf{L}\|_{\infty},\end{aligned}$$

where $\mathbf{A} = \mathbf{LU}$ is the LU-factorization of \mathbf{A} obtained without pivoting. They prove

Theorem 1.3 (Yeung-Chan). There exist constants $c > 0$ and $0 < b < 1$ such that if \mathbf{G} is an $n \times n$ Gaussian random matrix of variance 1 and $\mathbf{G} = \mathbf{LU}$ is the LU-factorization of \mathbf{G} , then

$$\begin{aligned}\Pr[\rho_{\mathbf{L}}(\mathbf{G}) > x] &\leq \frac{cn^3}{x}, \text{ and} \\ \Pr[\rho_{\mathbf{U}}(\mathbf{G}) > x] &\leq \min\left(\frac{cn^{7/2}}{x}, \frac{1}{n}\right) + \frac{cn^{5/2}}{x} + b^n.\end{aligned}$$

As it is generally believed that partial pivoting is better than no pivoting, their result provides some intuition for the experimental results of Trefethen and Schreiber demonstrating that random matrices rarely have large growth factors under partial pivoting. However, we note that it is difficult to make this intuition rigorous as there are matrices \mathbf{A} for which no pivoting has $\|\mathbf{L}\|_{\max}\|\mathbf{U}\|_{\max}/\|\mathbf{A}\|_{\max} = 2$ while partial pivoting has growth factor 2^{n-1} . (See also [Hig90])

The running times of many numerical algorithms depend on the condition numbers of their inputs. For example, the number of iterations taken by the method of conjugate gradients can be bounded in terms of the square root of the condition number. Similarly, the running times of interior-point methods can be bounded in terms of condition numbers [Ren95]. Blum [Blu89] suggested that a complexity theory of numerical algorithms should be parameterized by the condition number of an input in addition to the input size. Smale [Sma97] proposed a complexity theory of numerical algorithms in which one:

1. proves a bound on the running time of an algorithm solving a problem in terms of its condition number, and then
2. proves that it is unlikely that a random problem instance has large condition number.

This program is analogous to the average-case complexity of Theoretical Computer Science.

1.3 Our results

To better model the inputs that occur in practice, we propose replacing step 2 of Smale’s program with

- 2’. prove that for every input instance it is unlikely that a slight random perturbation of that instance has large condition number.

That is, we propose to bound the smoothed value of the condition number. Our first result in this program is presented in Section 3, where we improve upon Demmel’s [Dem88] and Edelman’s [Ede88] average-case results to show that a slight Gaussian perturbation of an arbitrary matrix is unlikely to have large condition number.

Definition 1.4 (Gaussian Perturbation). Let $\bar{\mathbf{A}}$ be an arbitrary $n \times n$ matrix. The matrix \mathbf{A} is a *Gaussian perturbation* of $\bar{\mathbf{A}}$ of variance σ^2 if \mathbf{A} can be written as $\mathbf{A} = \bar{\mathbf{A}} + \mathbf{G}$, where \mathbf{G} is a Gaussian random matrix of variance σ^2 . We also refer to \mathbf{A} as a Gaussian matrix of variance σ^2 centered at $\bar{\mathbf{A}}$.

In our smoothed analysis of the condition number, we consider an arbitrary $n \times n$ matrix $\bar{\mathbf{A}}$ of norm at most \sqrt{n} , and we bound the probability that $\kappa(\bar{\mathbf{A}} + \mathbf{G})$, the condition number of its Gaussian perturbation, is large, where \mathbf{G} is a Gaussian random matrix of variance $\sigma^2 \leq 1$. We bound this probability in terms of σ and n . In contrast with the average-case analysis of Demmel and Edelman, our analysis can be interpreted as demonstrating that if there is a little bit of imprecision or noise in the entries of a matrix, then it is unlikely it is ill-conditioned. On the other hand, Edelman [Ede92] writes of random matrices:

What is a mistake is to psychologically link a random matrix with the intuitive notion of a “typical” matrix or the vague concept of “any old matrix.”

The reader might also be interested in recent work on the smoothed analysis of the condition numbers of linear programs [BD02, DST02, ST03].

In Section 4, we use results from Section 3 to perform a smoothed analysis of the growth factors of Gaussian elimination without pivoting. If one specializes our results to perturbations of an all-zero square matrix, then one obtains a bound on $\rho_{\mathbf{U}}$ that improves the bound obtained by Yeung and Chan by a factor of n and which agrees with their experimental observations. The result obtained for $\rho_{\mathbf{L}}$ also improves the bound of Yeung and Chan [YC97] by a factor of n . However, while Yeung and Chan compute the density functions of the distribution of the elements in \mathbf{L} and \mathbf{U} , such precise estimates are not immediately available in our model. As a result, the techniques we develop are applicable to a wide variety of models of perturbations beyond the Gaussian. For example, one could use our techniques to obtain results of a similar nature if \mathbf{G} were a matrix of random variables chosen uniformly in $[-1, 1]$. We comment further upon this in the conclusions section of the paper.

The less effect a perturbation has, the more meaningful the results of smoothed analysis are. As many matrices encountered in practice are sparse or have structure, it would be best to consider perturbations that respect their sparsity pattern or structure. Our first result in this direction appears in Section 6, in which we consider the condition numbers and growth factors of perturbations of symmetric matrices that only alter their non-zero and diagonal elements. We prove results similar to those proved for dense perturbations of arbitrary matrices.

2 Notation and Mathematical Preliminaries

We use bold lower-case Roman letters such as \mathbf{x} , \mathbf{a} , \mathbf{b}_j to denote vectors in $\mathbb{R}^?$. Whenever a vector, say $\mathbf{a} \in \mathbb{R}^n$ is present, its components will be denoted by lower-case Roman letters with subscripts, such as a_1, \dots, a_n . Matrices are denoted by bold upper-case Roman letters such as \mathbf{A} and scalars are denoted by lower-case roman letters. Indicator random variables and random event variables are denoted by upper-case Roman letters. Random variables taking real values are denoted by upper-case Roman letters, except when they are components of a random vector or matrix.

The probability of an event A is written $\Pr[A]$, and the expectation of a variable X is written $\mathbb{E}[X]$. The indicator random variable for an event A is written $[A]$.

We write \ln to denote the natural logarithm, base e , and explicitly write the base for all other logarithms.

For integers $a \leq b$, we let $a : b$ denote the set of integers $\{x : a \leq x \leq b\}$. For a matrix \mathbf{A} we let $\mathbf{A}_{a:b,c:d}$ denote the submatrix of \mathbf{A} indexed by rows in $a : b$ and columns in $c : d$.

We will bound many probabilities by applying the following proposition.

Proposition 2.1 (Minimum \leq Average \leq Maximum). Let $\mu(X, Y)$ be a non-negative integrable function, and let X and Y be random variables distributed according to $\mu(X, Y)$. If $A(X, Y)$ is an event and $F(X, Y)$ is a function, then

$$\begin{aligned} \min_X \Pr_Y [A(X, Y)] &\leq \Pr_{X,Y} [A(X, Y)] \leq \max_X \Pr_Y [A(X, Y)], \text{ and} \\ \min_X \mathbb{E}_Y [F(X, Y)] &\leq \mathbb{E}_{X,Y} [F(X, Y)] \leq \max_X \mathbb{E}_Y [F(X, Y)], \end{aligned}$$

where in the left-hand and right-hand terms, Y is distributed according to the induced distribution on $\mu(X, Y)$.

We recall that a matrix \mathbf{Q} is an orthonormal matrix if its inverse is equal to its transpose, that is, $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}$. In Section 3 we will use the following proposition.

Proposition 2.2 (Orthonormal Transformation of Gaussian). Let $\bar{\mathbf{A}}$ be a matrix in $\mathbb{R}^{n \times n}$ and \mathbf{Q} be an orthonormal matrix in $\mathbb{R}^{n \times n}$. If \mathbf{A} is a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance σ^2 , then $\mathbf{Q}\mathbf{A}$ is a Gaussian perturbation of $\mathbf{Q}\bar{\mathbf{A}}$ of variance σ^2 .

We will also use the following extension of Proposition 2.17 of [ST04].

Proposition 2.3 (Gaussian Measure of Halfspaces). Let \mathbf{t} be any unit vector in \mathbb{R}^n and r be any real. Let $\bar{\mathbf{b}}$ be a vector in \mathbb{R}^n and \mathbf{b} be a Gaussian perturbation of $\bar{\mathbf{b}}$ of variance σ^2 . Then

$$\Pr_{\mathbf{b}} \left[\left| \mathbf{t}^\top \mathbf{b} \right| \leq r \right] \leq \frac{1}{\sqrt{2\pi}\sigma} \int_{t=-r}^{t=r} e^{-t^2/2\sigma^2} dt.$$

In this paper we will use the following properties of matrix norms and vector norms.

Proposition 2.4 (Product). For any pair of matrices \mathbf{A} and \mathbf{B} such that \mathbf{AB} is defined, and for every $1 \leq p \leq \infty$,

$$\|\mathbf{AB}\|_p \leq \|\mathbf{A}\|_p \|\mathbf{B}\|_p.$$

Proposition 2.5 (Vector Norms). For any column vector \mathbf{a} in \mathbb{R}^n , $\|\mathbf{a}\|_1 / \sqrt{n} \leq \|\mathbf{a}\|_2 \leq \|\mathbf{a}\|_1$.

Proposition 2.6 (2-norm). For any matrix \mathbf{A} ,

$$\|\mathbf{A}\|_2 = \left\| \mathbf{A}^\top \right\|_2,$$

as both are equal to the largest eigenvalue of $\sqrt{\mathbf{A}^\top \mathbf{A}}$.

Proposition 2.7 ($\|\mathbf{A}\|_\infty$: the maximum absolute row sum norm). For every matrix \mathbf{A} ,

$$\|\mathbf{A}\|_\infty = \max_i \left\| \mathbf{a}_i^\top \right\|_1,$$

where $\mathbf{a}_1, \dots, \mathbf{a}_n$ are the rows of \mathbf{A} . Thus, for any submatrix \mathbf{D} of \mathbf{A} ,

$$\|\mathbf{D}\|_\infty \leq \|\mathbf{A}\|_\infty.$$

Proposition 2.8 ($\|\mathbf{A}\|_1$: the maximum absolute column sum norm). For every matrix \mathbf{A} ,

$$\|\mathbf{A}\|_1 = \max_i \|\mathbf{a}_i\|_1,$$

where $\mathbf{a}_1, \dots, \mathbf{a}_n$ are the columns of \mathbf{A} . Thus

$$\|\mathbf{A}\|_1 = \left\| \mathbf{A}^\top \right\|_\infty.$$

3 Smoothed analysis of the condition number of a matrix

In this section, we will prove the following theorem which shows that for every matrix it is unlikely that a slight perturbation of that matrix has large condition number.

Theorem 3.1 (Smoothed Analysis of Condition number). Let $\bar{\mathbf{A}}$ be an $n \times n$ matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq \sqrt{n}$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Then, $\forall x \geq 1$,

$$\Pr[\kappa(\mathbf{A}) \geq x] \leq \frac{14.1n \left(1 + \sqrt{2 \ln(x)/9n}\right)}{x\sigma}.$$

As bounds on the norm of a random matrix are standard, we focus on the norm of the inverse. Recall that $1/\|\mathbf{A}^{-1}\|_2 = \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x}\|_2 / \|\mathbf{x}\|_2$.

The first step in the proof is to bound the probability that $\|\mathbf{A}^{-1}\mathbf{v}\|_2$ is small for a fixed unit vector \mathbf{v} . This result is also used later (in Section 4.1) in studying the growth factor. Using this result and an averaging argument, we then bound the probability that $\|\mathbf{A}^{-1}\|_2$ is large.

Lemma 3.2 (Projection of \mathbf{A}^{-1}). Let $\bar{\mathbf{A}}$ be an arbitrary square matrix in $\mathbb{R}^{n \times n}$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance σ^2 . Let \mathbf{v} be an arbitrary unit vector. Then

$$\Pr\left[\|\mathbf{A}^{-1}\mathbf{v}\|_2 > x\right] < \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma}$$

Proof. Let \mathbf{Q} be an orthonormal matrix such that $\mathbf{Q}^T \mathbf{e}_1 = \mathbf{v}$. Let $\bar{\mathbf{B}} = \mathbf{Q}\bar{\mathbf{A}}$ and $\mathbf{B} = \mathbf{Q}\mathbf{A}$. By Proposition 2.2, \mathbf{B} is a Gaussian perturbation of $\bar{\mathbf{B}}$ of variance σ^2 . We have

$$\|\mathbf{A}^{-1}\mathbf{v}\|_2 = \|\mathbf{A}^{-1}\mathbf{Q}^T \mathbf{e}_1\|_2 = \|(\mathbf{Q}\mathbf{A})^{-1} \mathbf{e}_1\|_2 = \|\mathbf{B}^{-1} \mathbf{e}_1\|_2.$$

Thus, to prove the lemma it is sufficient to show

$$\Pr_{\mathbf{B}}\left[\|\mathbf{B}^{-1} \mathbf{e}_1\|_2 > x\right] < \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma}.$$

We observe that

$$\|\mathbf{B}^{-1} \mathbf{e}_1\|_2 = \|(\mathbf{B}^{-1})_{:,1}\|_2,$$

the length of the first column of \mathbf{B}^{-1} . The first column of \mathbf{B}^{-1} , by the definition of the matrix inverse, is the vector that is orthogonal to every row of \mathbf{B} but the first and that has inner product 1 with the first row of \mathbf{B} . Hence its length is the reciprocal of the length of the projection of the first row of \mathbf{B} onto the subspace orthogonal to the rest of the rows.

Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ be the rows of \mathbf{B} and $\bar{\mathbf{b}}_1, \dots, \bar{\mathbf{b}}_n$ be the rows of $\bar{\mathbf{B}}$. Note that \mathbf{b}_i is a Gaussian perturbation of $\bar{\mathbf{b}}_i$ of variance σ^2 . Let \mathbf{t} be the unit vector that is orthogonal to the span of $\mathbf{b}_2, \dots, \mathbf{b}_n$. Then

$$\|(\mathbf{B}^{-1})_{:,1}\|_2 = \left| \frac{1}{\mathbf{t}^T \mathbf{b}_1} \right|.$$

Thus,

$$\begin{aligned} \Pr_{\mathbf{B}} \left[\left\| \mathbf{B}^{-1} \mathbf{v} \right\|_2 > x \right] &= \Pr_{\mathbf{b}_1, \dots, \mathbf{b}_n} \left[\left| \frac{1}{\mathbf{t}^\top \mathbf{b}_1} \right| > x \right] \\ &\leq \max_{\mathbf{b}_2, \dots, \mathbf{b}_n} \Pr_{\mathbf{b}_1} \left[\left| \mathbf{t}^\top \mathbf{b}_1 \right| < 1/x \right] \\ &< \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma}, \end{aligned}$$

where the first inequality follows from Proposition 2.1 and the second inequality follows from Lemma A.2. \square

Theorem 3.3 (Smallest singular value). Let $\bar{\mathbf{A}}$ be an arbitrary square matrix in $\mathbb{R}^{n \times n}$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance σ^2 . Then

$$\Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq 2.35 \frac{\sqrt{n}}{x\sigma}$$

Proof. Let \mathbf{v} be a uniformly distributed random unit vector in \mathbb{R}^n . It follows from Lemma 3.2 that

$$\Pr_{\mathbf{A}, \mathbf{v}} \left[\left\| \mathbf{A}^{-1} \mathbf{v} \right\|_2 \geq x \right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \quad (3.1)$$

Since \mathbf{A} is a Gaussian perturbation of $\bar{\mathbf{A}}$, with probability 1 there is a unique pair $(\mathbf{u}, -\mathbf{u})$ of unit vectors such that $\left\| \mathbf{A}^{-1} \mathbf{u} \right\|_2 = \left\| \mathbf{A}^{-1} \right\|_2$. From the inequality

$$\left\| \mathbf{A}^{-1} \mathbf{v} \right\|_2 \geq \left\| \mathbf{A}^{-1} \right\|_2 \left| \mathbf{u}^\top \mathbf{v} \right|,$$

we know that for every $c > 0$,

$$\begin{aligned} \Pr_{\mathbf{A}, \mathbf{v}} \left[\left\| \mathbf{A}^{-1} \mathbf{v} \right\|_2 \geq x\sqrt{c/n} \right] &\geq \Pr_{\mathbf{A}, \mathbf{v}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \text{ and } \left| \mathbf{u}^\top \mathbf{v} \right| \geq \sqrt{c/n} \right] \\ &= \Pr_{\mathbf{A}, \mathbf{v}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \Pr_{\mathbf{A}, \mathbf{v}} \left[\left| \mathbf{u}^\top \mathbf{v} \right| \geq \sqrt{c/n} \mid \left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \\ &= \Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \Pr_{\mathbf{A}, \mathbf{v}} \left[\left| \mathbf{u}^\top \mathbf{v} \right| \geq \sqrt{c/n} \mid \left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \\ &\geq \Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \min_{\mathbf{A}: \left\| \mathbf{A}^{-1} \right\|_2 \geq x} \Pr_{\mathbf{v}} \left[\left| \mathbf{u}^\top \mathbf{v} \right| \geq \sqrt{c/n} \right] \quad (\text{by Proposition 2.1}) \\ &\geq \Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \Pr_{\mathbf{G}} \left[\left| G \right| \geq \sqrt{c} \right], \quad (\text{by Lemma B.1}) \end{aligned}$$

where G is a Gaussian random variable with mean $\mathbf{0}$ and variance 1. To prove this last inequality, we first note that \mathbf{v} is a random unit vector and is independent from \mathbf{u} . Thus, in a basis of \mathbb{R}^n in which \mathbf{u} is the first vector, \mathbf{v} is a uniformly distributed random unit vector with the first coordinate equal to $\mathbf{u}^\top \mathbf{v}$, and so we may apply Lemma B.1 to bound $\Pr_{\mathbf{v}} \left[|\mathbf{u}^\top \mathbf{v}| \geq \sqrt{c/n} \right]$ from below by $\Pr_G [|G| \geq \sqrt{c}]$. So,

$$\begin{aligned} \Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] &\leq \frac{\Pr_{\mathbf{A}, \mathbf{v}} \left[\left\| \mathbf{A}^{-1} \mathbf{v} \right\|_2 \geq x \sqrt{c/n} \right]}{\Pr_G [|G| \geq \sqrt{c}]} \\ &\leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{n}}{x \sigma \sqrt{c} \Pr_G [|G| \geq \sqrt{c}]} \quad (\text{by (3.1)}). \end{aligned}$$

Because this inequality is true for every c , we will choose a value for c that almost maximizes $\sqrt{c} \Pr_G [|G| \geq \sqrt{c}]$ and which in turn almost minimizes the right hand side.

Choosing $c = 0.57$, and evaluating the error function numerically, we determine

$$\Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq 2.35 \frac{\sqrt{n}}{x \sigma}.$$

□

Note that Theorem 3.3 gives a smoothed analogue of the following bound of Edelman [Ede88] on Gaussian random matrices.

Theorem 3.4 (Edelman). Let $\mathbf{G} \in \mathbb{R}^{n \times n}$ be a Gaussian random matrix with variance σ^2 , then

$$\Pr_{\mathbf{G}} \left[\left\| \mathbf{G}^{-1} \right\|_2 \geq x \right] \leq \frac{\sqrt{n}}{x \sigma}.$$

As Gaussian random matrices can be viewed as Gaussian random perturbations of the $n \times n$ all-zero square matrix, Theorem 3.3 extends Edelman's theorem to Gaussian random perturbations of an arbitrary matrix. The constant 2.35 in Theorem 3.3 is bigger than Edelman's 1 for Gaussian random matrices. We conjecture that it is possible to reduce 2.35 in Theorem 3.3 to 1 as well.

Conjecture 1 (Smallest Singular Value). Let $\bar{\mathbf{A}}$ be an arbitrary square matrix in $\mathbb{R}^{n \times n}$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance σ^2 . Then

$$\Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq \frac{\sqrt{n}}{x \sigma}$$

We now apply Theorem 3.3 to prove Theorem 3.1.

Proof of Theorem 3.1. As observed by Davidson and Szarek [DS01, Theorem II.7], one can apply inequality (1.4) of [LT91] to show that for all $k \geq 0$,

$$\Pr_{\mathbf{A}} \left[\|\bar{\mathbf{A}} - \mathbf{A}\|_2 \geq \sigma (2\sqrt{n} + k) \right] \leq e^{-k^2/2}.$$

Replacing σ by its upper bound of 1 and setting $\epsilon = e^{-k^2/2}$, we obtain

$$\Pr_{\mathbf{A}} \left[\|\bar{\mathbf{A}} - \mathbf{A}\|_2 \geq 2\sqrt{n} + \sqrt{2\ln(1/\epsilon)} \right] \leq \epsilon,$$

for all $\epsilon \leq 1$. By assumption, $\|\bar{\mathbf{A}}\|_2 \leq \sqrt{n}$; so,

$$\Pr_{\mathbf{A}} \left[\|\mathbf{A}\|_2 \geq 3\sqrt{n} + \sqrt{2\ln(1/\epsilon)} \right] \leq \epsilon.$$

From the result of Theorem 3.3, we have

$$\Pr_{\mathbf{A}} \left[\|\mathbf{A}^{-1}\|_2 \geq \frac{2.35\sqrt{n}}{\epsilon\sigma} \right] \leq \epsilon.$$

Combining these two bounds, we find

$$\Pr_{\mathbf{A}} \left[\|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 \geq \frac{7.05n + 2.35\sqrt{2n\ln(1/\epsilon)}}{\epsilon\sigma} \right] \leq 2\epsilon.$$

So that we can express this probability in the form of $\Pr_{\mathbf{A}} [\|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 \geq x]$, for $x \geq 1$, we let

$$x = \frac{7.05n + 2.35\sqrt{2n\ln(1/\epsilon)}}{\epsilon\sigma}. \quad (3.2)$$

It follows Equation (3.2) and the assumption $\sigma \leq 1$ that $x\epsilon \geq 1$, implying $\ln(1/\epsilon) \leq \ln x$. From Equation (3.2), we derive

$$2\epsilon = \frac{2 \left(7.05n + 2.35\sqrt{2n\ln(1/\epsilon)} \right)}{x\sigma} \leq \frac{2 \left(7.05n + 2.35\sqrt{2n\ln x} \right)}{x\sigma} \leq \frac{14.1n \left(1 + \sqrt{2\ln(x)/9n} \right)}{x\sigma}.$$

Therefore, we conclude

$$\Pr \left[\|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 \geq x \right] \leq \frac{14.1n \left(1 + \sqrt{2\ln(x)/9n} \right)}{x\sigma}.$$

□

We conjecture that the $1 + \sqrt{2\ln(x)/9n}$ term should be unnecessary because those matrices for which $\|\mathbf{A}\|_2$ is large are less likely to have $\|\mathbf{A}^{-1}\|_2$ large as well.

4 Growth Factor of Gaussian Elimination without Pivoting

We now turn to proving a bound on the growth factor. We will consider a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ obtained from a Gaussian perturbation of variance σ^2 of an arbitrary matrix $\bar{\mathbf{A}}$ satisfying $\|\bar{\mathbf{A}}\|_2 \leq 1$. With probability 1, none of the diagonal entries that occur during elimination will be 0. So, in the spirit of Yeung and Chan [YC97], we analyze the growth factor of Gaussian elimination without pivoting. When we specialize our smoothed analyses to the case $\bar{\mathbf{A}} = 0$, we improve the bounds of Yeung and Chan (see Theorem 1.3) by a factor of n . Our improved bound on $\rho_{\mathbf{U}}$ agrees with their experimental analyses.

4.1 Growth in \mathbf{U}

We recall that

$$\rho_{\mathbf{U}}(\mathbf{A}) = \frac{\|\mathbf{U}\|_{\infty}}{\|\mathbf{A}\|_{\infty}}.$$

In this section, we give two bounds on $\rho_{\mathbf{U}}(\mathbf{A})$. The first will have a better dependence on σ , and second will have a better dependence on n . It is the later bound, Theorem 4.3, that agrees with the experiments of Yeung and Chan [YC97] when specialized to the average-case by setting $\bar{\mathbf{A}} = 0$ and $\sigma = 1$.

4.1.1 First bound

Theorem 4.1 (First bound on $\rho_{\mathbf{U}}(\mathbf{A})$). Let $\bar{\mathbf{A}}$ be an $n \times n$ matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Then,

$$\Pr[\rho_{\mathbf{U}}(\mathbf{A}) > 1 + x] < \frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{x\sigma}.$$

Proof. By Proposition 2.7.

$$\rho_{\mathbf{U}}(\mathbf{A}) = \frac{\|\mathbf{U}\|_{\infty}}{\|\mathbf{A}\|_{\infty}} = \max_i \frac{\|(\mathbf{U}_{i,:})^T\|_1}{\|\mathbf{A}\|_{\infty}}.$$

So, we need to bound the probability that the 1-norm of the vector defined by each row of \mathbf{U} is large and then apply a union bound to bound the overall probability.

Fix for now a k between 2 and n . We denote the upper triangular segment of the k th row of \mathbf{U} by $\mathbf{u}^T = \mathbf{U}_{k,k:n}$, and observe that \mathbf{u} can be obtained from the formula:

$$\mathbf{u}^T = \mathbf{a}^T - \mathbf{b}^T \mathbf{C}^{-1} \mathbf{D} \tag{4.1}$$

where

$$\mathbf{a}^T = \mathbf{A}_{k,k:n} \quad \mathbf{b}^T = \mathbf{A}_{k,1:k-1} \quad \mathbf{C} = \mathbf{A}_{1:k-1,1:k-1} \quad \mathbf{D} = \mathbf{A}_{1:k-1,k:n}.$$

This expression for \mathbf{u} follows immediately from

$$\mathbf{A}_{1:k,:} = \begin{pmatrix} \mathbf{C} & \mathbf{D} \\ \mathbf{b}^\top & \mathbf{a}^\top \end{pmatrix} = \begin{pmatrix} \mathbf{L}_{1:k-1,1:k-1} & 0 \\ \mathbf{L}_{k,1:k-1} & 1 \end{pmatrix} \begin{pmatrix} \mathbf{u}_{1:k-1,1:k-1} & \mathbf{u}_{1:k-1,k:n} \\ 0 & \mathbf{u}^\top \end{pmatrix}.$$

From (4.1), we derive

$$\begin{aligned} \|\mathbf{u}\|_1 &= \left\| \mathbf{a} - (\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D})^\top \right\|_1 \leq \|\mathbf{a}\|_1 + \left\| (\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D})^\top \right\|_1 \\ &\leq \|\mathbf{a}^\top\|_\infty + \left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_1 \|\mathbf{D}\|_\infty \quad \text{by Propositions 2.4 and 2.8} \\ &\leq \|\mathbf{A}\|_\infty \left(1 + \left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_1 \right) \quad \text{by Proposition 2.7} \end{aligned} \tag{4.2}$$

We now bound the probability $\left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_1$ is large. By Proposition 2.5,

$$\left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_1 \leq \sqrt{k-1} \left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_2.$$

Note that \mathbf{b} and \mathbf{C} are independent of each other. Therefore,

$$\begin{aligned} \Pr_{\mathbf{b}, \mathbf{C}} \left[\left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_1 > x \right] &\leq \Pr_{\mathbf{b}, \mathbf{C}} \left[\left\| (\mathbf{C}^\top)^{-1} \mathbf{b} \right\|_2 > x/\sqrt{k-1} \right] \\ &\leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{k-1} \sqrt{(k-1)\sigma^2 + 1}}{x\sigma} < \sqrt{\frac{2}{\pi}} \frac{k}{x\sigma}, \end{aligned} \tag{4.3}$$

where the second inequality follows from Lemma 4.2 below and the last inequality follows from the assumption $\sigma^2 \leq 1$.

We now apply a union bound over the choices of k to obtain

$$\Pr [\rho_{\mathbf{u}}(\mathbf{A}) > 1 + x] < \sum_{k=2}^n \sqrt{\frac{2}{\pi}} \frac{k}{x\sigma} \leq \frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{x\sigma}.$$

□

Lemma 4.2. Let $\bar{\mathbf{C}}$ be an arbitrary square matrix in $\mathbb{R}^{d \times d}$, and \mathbf{C} be a Gaussian perturbation of $\bar{\mathbf{C}}$ of variance σ^2 . Let $\bar{\mathbf{b}}$ be a column vector in \mathbb{R}^d such that $\|\bar{\mathbf{b}}\|_2 \leq 1$, and let \mathbf{b} be a Gaussian perturbation of $\bar{\mathbf{b}}$ of variance σ^2 . If \mathbf{b} and \mathbf{C} are independent of each other, then

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

Proof. Let $\hat{\mathbf{b}}$ be the unit vector in the direction of \mathbf{b} . By applying Lemma 3.2, we obtain for all \mathbf{b} ,

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

Let $\mu(\mathbf{b})$ denote the density according to which \mathbf{b} is distributed. Then, we have

$$\begin{aligned} \Pr_{\mathbf{b}, \mathbf{C}} \left[\left\| \mathbf{C}^{-1} \mathbf{b} \right\|_2 > x \right] &= \int_{\mathbf{b} \in \mathbb{R}^d} \Pr_{\mathbf{C}} \left[\left\| \mathbf{C}^{-1} \mathbf{b} \right\|_2 > x \right] \mu(\mathbf{b}) d\mathbf{b} \\ &\leq \int_{\mathbf{b} \in \mathbb{R}^d} \left(\sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \|\mathbf{b}\|_2 \right) \mu(\mathbf{b}) d\mathbf{b} \\ &= \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \mathbf{E}_{\mathbf{b}} [\|\mathbf{b}\|_2]. \end{aligned}$$

It is known [KJ82, p. 277] that $\mathbf{E}_{\mathbf{b}} \left[\|\mathbf{b}\|_2^2 \right] \leq \sigma^2 d + \|\bar{\mathbf{b}}\|_2^2$. As $\mathbf{E}[X] \leq \sqrt{\mathbf{E}[X^2]}$ for every positive random variable X , we have $\mathbf{E}_{\mathbf{b}} [\|\mathbf{b}\|_2] \leq \sqrt{\sigma^2 d + \|\bar{\mathbf{b}}\|_2^2} \leq \sqrt{\sigma^2 d + 1}$. □

4.1.2 Second Bound for $\rho_{\mathbf{U}}(\mathbf{A})$

In this section, we establish an upper bound on $\rho_{\mathbf{U}}(\mathbf{A})$ which dominates the bound in Theorem 4.1 for $\sigma \geq n^{-3/2}$.

If we specialize the parameters in this bound to $\bar{\mathbf{A}} = \mathbf{0}$ and $\sigma^2 = 1$, we improve the average-case bound proved by Yeung and Chan [YC97] (see Theorem 1.3) by a factor of n . Moreover, the resulting bound agrees with their experimental results.

Theorem 4.3 (Second bound on $\rho_{\mathbf{U}}(\mathbf{A})$). Let $\bar{\mathbf{A}}$ be an $n \times n$ matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. For $n \geq 2$,

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

Proof. As in the proof of Theorem 4.1, we will separately consider the k th row of \mathbf{U} for each $2 \leq k \leq n$. For any such k , define \mathbf{u} , \mathbf{a} , \mathbf{b} , \mathbf{C} and \mathbf{D} as in the proof of Theorem 4.1.

In the case when $k = n$, we may apply (4.3) in the proof of Theorem 4.1, to show

$$\Pr \left[\frac{\|\mathbf{u}\|_1}{\|\mathbf{A}\|_\infty} > 1 + x \right] \leq \sqrt{\frac{2}{\pi}} \frac{n}{x\sigma}. \quad (4.4)$$

We now turn to the case $k \leq n - 1$. By (4.1) and Proposition 2.5, we have

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

The last equation follows from Proposition 2.6. Therefore, for all $k \leq n-1$,

$$\begin{aligned} \frac{\|\mathbf{u}\|_1}{\|\mathbf{A}\|_\infty} &\leq \frac{\|\mathbf{a}\|_1 + \sqrt{k-1} \|\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D}\|_2}{\|\mathbf{A}\|_\infty} \leq 1 + \frac{\sqrt{k-1} \|\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D}\|_2}{\|\mathbf{A}\|_\infty} && \text{(by Proposition 2.7).} \\ &\leq 1 + \frac{\sqrt{k-1} \|\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D}\|_2}{\|(\mathbf{A}_{n,:})^\top\|_1} && \text{(also by Proposition 2.7).} \end{aligned}$$

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

$$\|\mathbf{b}^\top \mathbf{C}^{-1} \bar{\mathbf{D}}\|_2 \leq \|\mathbf{b}^\top \mathbf{C}^{-1}\|_2 \|\bar{\mathbf{D}}\|_2 \leq \|\mathbf{b}^\top \mathbf{C}^{-1}\|_2.$$

Thus, if we let $\mathbf{t}^\top = (\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D}) / \|\mathbf{b}^\top \mathbf{C}^{-1}\|_2$, then for every fixed \mathbf{b} and \mathbf{C} , \mathbf{t} is a Gaussian random column vector in \mathbb{R}^{n-k+1} of variance σ^2 centered at a vector of 2-norm at most 1. We also have

$$\Pr_{\mathbf{b}, \mathbf{C}, \mathbf{D}} \left[\|\mathbf{b}^\top \mathbf{C}^{-1} \mathbf{D}\|_2 \geq x \right] = \Pr_{\mathbf{b}, \mathbf{C}, \mathbf{t}} \left[\|\mathbf{b}^\top \mathbf{C}^{-1}\|_2 \|\mathbf{t}\|_2 \geq x \right]. \quad (4.5)$$

It follows from Lemma 4.2 that

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

Hence, we may apply Corollary C.5 to show

$$\begin{aligned} \Pr_{\mathbf{b}, \mathbf{C}, \mathbf{t}} \left[\|\mathbf{b}^\top \mathbf{C}^{-1}\|_2 \|\mathbf{t}\|_2 \geq x \right] &\leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{\sigma^2(k-1) + 1} \sqrt{\sigma^2(n-k+1) + 1}}{x\sigma} \\ &\leq \sqrt{\frac{2}{\pi}} \frac{\left(1 + \frac{n\sigma^2}{2}\right)}{x\sigma}. \end{aligned} \quad (4.6)$$

Note that $\mathbf{A}_{n,:}$ is a Gaussian perturbation of variance σ^2 of a row vector in \mathbb{R}^n . As $\mathbf{A}_{n,:}$ is independent of \mathbf{b} , \mathbf{C} and \mathbf{D} , we can apply (4.5), (4.6) and Lemma C.4 to show

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

by Lemma A.4.

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

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

where the second inequality follows from

$$\sum_{k=1}^{n-2} \sqrt{k} \leq \frac{2}{3} n^{3/2}.$$

□

4.2 Growth in \mathbf{L}

Let \mathbf{L} be the lower-triangular part of the LU-factorization of \mathbf{A} . We have

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

where we let $\mathbf{A}^{(k)}$ denote the matrix remaining after the first k columns have been eliminated. So, $\mathbf{A}^{(0)} = \mathbf{A}$.

Recall $\rho_{\mathbf{L}}(\mathbf{A}) = \|\mathbf{L}\|_{\infty}$, which is equal to the maximum absolute row sum of \mathbf{L} (Proposition 2.7). We will show that it is unlikely that $\|\mathbf{L}_{(k+1):n,k}\|_{\infty}$ is large by proving that it is unlikely that $\|\mathbf{A}_{(k+1):n,k}^{(k-1)}\|_{\infty}$ is large while $|\mathbf{A}_{k,k}^{(k-1)}|$ is small.

Theorem 4.4 ($\rho_{\mathbf{L}}(\mathbf{A})$). Let $\bar{\mathbf{A}}$ be an n -by- n matrix for which $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. If $n \geq 2$, then,

$$\Pr[\rho_{\mathbf{L}}(\mathbf{A}) > \varkappa] \leq \sqrt{\frac{2}{\pi}} \frac{n^2}{\varkappa} \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2 \ln n} + \frac{1}{\sqrt{2\pi \ln n}} \right)$$

Proof. For each k between 1 and $n - 1$, we have

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

where we let $\mathbf{v} = \mathbf{A}_{1:(k-1),1:(k-1)}^{-1} \mathbf{A}_{1:(k-1),k}$. Since $\|\bar{\mathbf{A}}\|_2 \leq 1$, and all the terms $\mathbf{A}_{(k+1):n,k}$, $\mathbf{A}_{(k+1):n,1:(k-1)}$, $\mathbf{A}_{k,k}$, $\mathbf{A}_{k,1:(k-1)}$ and \mathbf{v} are independent, we can apply Lemma 4.5 to show that

$$\begin{aligned}
\Pr \left[\|\mathbf{L}_{(k+1):n,k}\|_\infty > x \right] &\leq \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2 \ln(\max(n-k, 2))} + \frac{1}{\sqrt{2\pi} \ln(\max(n-k, 2))} \right) \\
&\leq \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2 \ln n} + \frac{1}{\sqrt{2\pi} \ln n} \right)
\end{aligned}$$

where the last inequality follows the facts that $\sqrt{2z} + \frac{1}{\sqrt{2\pi z}}$ is an increasing function when $z \geq \pi^{-1/3}$, and $\ln 2 \geq \pi^{-1/3}$.

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

Lemma 4.5 (Vector Ratio). Let d and n be positive integers. Let \mathbf{a} , \mathbf{b} , \mathbf{x} , and \mathbf{Y} be Gaussian perturbations of $\bar{\mathbf{a}} \in \mathbb{R}^1$, $\bar{\mathbf{b}} \in \mathbb{R}^d$, $\bar{\mathbf{x}} \in \mathbb{R}^n$, and $\bar{\mathbf{Y}} \in \mathbb{R}^{n \times d}$, respectively, of variance σ^2 , such that $|\bar{\mathbf{a}}| \leq 1$, $\|\bar{\mathbf{b}}\|_2 \leq 1$, $\|\bar{\mathbf{x}}\|_2 \leq 1$, and $\|\bar{\mathbf{Y}}\|_2 \leq 1$. Let \mathbf{v} be an arbitrary vector in \mathbb{R}^d . If \mathbf{a} , \mathbf{b} , \mathbf{x} , and \mathbf{Y} are independent and $\sigma^2 \leq 1$, then

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

Proof. We begin by observing that $\mathbf{a} + \mathbf{b}^T \mathbf{v}$ and each component of $\mathbf{x} + \mathbf{Y}\mathbf{v}$ is a Gaussian random variable of variance $\sigma^2(1 + \|\mathbf{v}\|_2^2)$ whose mean has absolute value at most $1 + \|\mathbf{v}\|_2$, and that all these variables are independent. By Lemma A.3,

$$\mathbb{E}_{\mathbf{x}, \mathbf{Y}} [\|\mathbf{x} + \mathbf{Y}\mathbf{v}\|_\infty] \leq 1 + \|\mathbf{v}\|_2 + \left(\sigma \sqrt{(1 + \|\mathbf{v}\|_2^2)} \right) \left(\sqrt{2 \ln \max(n, 2)} + \frac{1}{\sqrt{2\pi} \ln \max(n, 2)} \right).$$

On the other hand, Lemma A.2 implies

$$\Pr_{\mathbf{a}, \mathbf{b}} \left[\frac{1}{|\mathbf{a} + \mathbf{b}^\top \mathbf{v}|} > x \right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x \sigma \sqrt{1 + \|\mathbf{v}\|_2^2}}. \quad (4.7)$$

Thus, we can apply Corollary C.4 to show

$$\begin{aligned} \Pr \left[\frac{\|\mathbf{x} + \mathbf{Y}\mathbf{v}\|_\infty}{|\mathbf{a} + \mathbf{b}^\top \mathbf{v}|} > x \right] &\leq \sqrt{\frac{2}{\pi}} \frac{1 + \|\mathbf{v}\|_2 + \left(\sigma \sqrt{1 + \|\mathbf{v}\|_2^2} \right) \left(\sqrt{2 \ln \max(n, 2)} + \frac{1}{\sqrt{2\pi \ln \max(n, 2)}} \right)}{x \sigma \sqrt{1 + \|\mathbf{v}\|_2^2}} \\ &= \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{1 + \|\mathbf{v}\|_2}{\sigma \sqrt{1 + \|\mathbf{v}\|_2^2}} + \frac{\left(\sigma \sqrt{1 + \|\mathbf{v}\|_2^2} \right) \left(\sqrt{2 \ln \max(n, 2)} + \frac{1}{\sqrt{2\pi \ln \max(n, 2)}} \right)}{\sigma \sqrt{1 + \|\mathbf{v}\|_2^2}} \right) \\ &\leq \sqrt{\frac{2}{\pi}} \frac{1}{x} \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2 \ln \max(n, 2)} + \frac{1}{\sqrt{2\pi \ln \max(n, 2)}} \right), \end{aligned}$$

where the last inequality follows from $(1+z)^2 \leq 2(1+z^2)$, $\forall z \geq 0$. \square

5 Smoothed Analysis of Gaussian Elimination

We now combine the results from the previous sections to bound the smoothed precision needed in the application of Gaussian elimination without pivoting to obtain solutions to linear systems accurate to b bits.

Theorem 5.1 (Smoothed precision of Gaussian elimination). For $n > e^4$, let $\bar{\mathbf{A}}$ be an n -by- n matrix for which $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1/4$. Then, the expected number of bits of precision necessary to solve $\mathbf{A}\mathbf{x} = \mathbf{b}$ to b bits of accuracy using Gaussian elimination without pivoting is at most

$$b + \frac{11}{2} \log_2 n + 3 \log_2 \left(\frac{1}{\sigma} \right) + \log_2(1 + 2\sqrt{n}\sigma) + \frac{1}{2} \log_2 \log_2 n + 6.83$$

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

$$\begin{aligned} 5 \cdot 2^b n \rho_{\mathbf{L}}(\mathbf{A}) \rho_{\mathbf{U}}(\mathbf{A}) \kappa(\mathbf{A}) \epsilon_{\text{mach}} &\leq 1 \quad \implies \\ 2.33 + b + \log_2 n + \log_2(\rho_{\mathbf{L}}(\mathbf{A})) + \max(0, \log_2(\rho_{\mathbf{U}}(\mathbf{A}))) + \log_2(\kappa(\mathbf{A})) &\leq \log_2(1/\epsilon_{\text{mach}}). \end{aligned}$$

We will apply Lemma C.6 to bound these log terms. Theorem 4.1 tells us that

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

To put this inequality into a form to which Lemma C.6 may be applied, we set

$$\mathbf{y} = x \left(1 + \frac{\sqrt{2\pi}\sigma}{n(n+1)} \right),$$

to obtain

$$\Pr [\rho_{\mathbf{U}}(\mathbf{A}) > \mathbf{y}] \leq \left(\frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{\sigma} + 1 \right) \frac{1}{\mathbf{y}}.$$

By Lemma C.6,

$$\begin{aligned} \mathbb{E} [\max(0, \log_2 \rho_{\mathbf{U}}(\mathbf{A}))] &\leq \log_2 \left(\frac{1}{\sqrt{2\pi}} \frac{n(n+1)}{\sigma} + 1 \right) + \log_2 e \\ &\leq \log_2 (n(n+1) + \sigma\sqrt{2\pi}) + \log_2 \left(\frac{1}{\sigma} \right) + \log_2 \left(\frac{e}{\sqrt{2\pi}} \right) \\ &\leq \log_2 (1.02n^2) + \log_2 \left(\frac{1}{\sigma} \right) + \log_2 \left(\frac{e}{\sqrt{2\pi}} \right) \\ &\leq 2\log_2 n + \log_2 \left(\frac{1}{\sigma} \right) + 0.15, \end{aligned}$$

where in the second-to-last inequality, we used the assumptions $n \geq e^4$ and $\sigma \leq 1/2$. In the last inequality, we numerically computed $\log_2(1.02e/\sqrt{2\pi}) < 0.15$.

Theorem 4.4 and Lemma C.6 imply

$$\begin{aligned} \mathbb{E} [\log_2 \rho_{\mathbf{L}}(\mathbf{A})] &\leq \log_2 \left(\sqrt{\frac{2}{\pi}} n^2 \left(\frac{\sqrt{2}}{\sigma} + \sqrt{2\ln n} + \frac{1}{\sqrt{2\pi\ln n}} \right) \right) + \log_2 e \\ &\leq 2\log_2 n + \log_2 \left(\frac{1}{\sigma} + \sqrt{\ln n} \left(1 + \frac{1}{2\sqrt{\pi\ln n}} \right) \right) + \log_2 \left(\frac{2e}{\sqrt{\pi}} \right) \\ &= 2\log_2 n + \log_2 \left(\frac{1}{\sigma} \right) + \log_2 \sqrt{\ln n} + \log_2 \left(\frac{1}{\sqrt{\ln n}} + \sigma \left(1 + \frac{1}{2\sqrt{\pi\ln n}} \right) \right) + \log_2 \left(\frac{2e}{\sqrt{\pi}} \right) \end{aligned}$$

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

$$\begin{aligned} &\leq 2\log_2 n + \log_2 \left(\frac{1}{\sigma} \right) + \frac{1}{2} \log_2 \log_2 n + \log_2 \left(1 + \frac{1}{16\sqrt{\pi}} \right) + \log_2 \left(\frac{2e}{\sqrt{\pi}} \right) \\ &\leq 2\log_2 n + \log_2 \left(\frac{1}{\sigma} \right) + \frac{1}{2} \log_2 \log_2 n + 1.67, \end{aligned}$$

as $\log_2(1+1/16\sqrt{\pi}) + \log_2(2e/\sqrt{\pi}) < 1.67$. Theorem 3.3 and Lemma C.6, along with the observation that $\log_2(2.35e) < 2.68$, imply

$$\mathbb{E} \left[\log_2 \left\| \mathbf{A}^{-1} \right\|_2 \right] \leq \frac{1}{2} \log_2 n + \log_2 \left(\frac{1}{\sigma} \right) + 2.68.$$

Finally,

$$\mathbf{E} [\log_2(\|\mathbf{A}\|_2)] \leq \log_2(1 + 2\sqrt{n}\sigma)$$

follows from the well-known facts that the expectation of $\|\mathbf{A} - \bar{\mathbf{A}}\|_2$ is at most $2\sqrt{n}\sigma$ (c.f., [Seg00]) and that $\mathbf{E} [\log_2(X)] \leq \log_2 \mathbf{E} [X]$ for every positive random variable X . Thus, the expected number of digits of precision needed is at most

$$b + \frac{11}{2} \log_2 n + 3 \log_2 \left(\frac{1}{\sigma} \right) + \log_2(1 + 2\sqrt{n}\sigma) + \frac{1}{2} \log_2 \log_2 n + 6.83.$$

□

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

Conjecture 2. Let $\bar{\mathbf{A}}$ be a n -by- n matrix for which $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Then

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

for some constants c_1 and c_2 .

6 Zero-preserving perturbations of symmetric matrices with diagonals

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

Throughout this section, when we express a symmetric matrix \mathbf{A} as $\mathbf{T} + \mathbf{D} + \mathbf{T}^T$, we mean that \mathbf{T} is lower-triangular with zeros on the diagonal and \mathbf{D} is a diagonal matrix. By making a zero-preserving perturbation to $\bar{\mathbf{T}}$, 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 \mathbf{T} and diagonal perturbations to \mathbf{D} have distributions similar those proved in Sections 3 and 4 for dense matrices under dense perturbations.

6.1 Bounding the condition number

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

Lemma 6.2. Let $\bar{\mathbf{A}} = \bar{\mathbf{T}} + \bar{\mathbf{D}} + \bar{\mathbf{T}}^\top$ be an arbitrary n -by- n symmetric matrix. Let \mathbf{T} be a zero-preserving perturbation of $\bar{\mathbf{T}}$ of variance σ^2 , let \mathbf{G}_D be a diagonal matrix of independent Gaussian random variables of variance σ^2 and mean 0 that are independent of \mathbf{T} , and let $\mathbf{D} = \bar{\mathbf{D}} + \mathbf{G}_D$. Then, for $\mathbf{A} = \mathbf{T} + \mathbf{D} + \mathbf{T}^\top$,

$$\Pr \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq \sqrt{\frac{2}{\pi}} \frac{n^{3/2}}{x\sigma}.$$

Proof. By Proposition 2.1,

$$\Pr_{\mathbf{T}, \mathbf{G}_D} \left[\left\| (\mathbf{T} + \mathbf{D} + \mathbf{T}^\top)^{-1} \right\|_2 \geq x \right] \leq \max_{\mathbf{T}} \Pr_{\mathbf{G}_D} \left[\left\| ((\mathbf{T} + \bar{\mathbf{D}} + \mathbf{T}^\top) + \mathbf{G}_D)^{-1} \right\|_2 \geq x \right].$$

The proof now follows from Lemma 6.3, taking $\mathbf{T} + \bar{\mathbf{D}} + \mathbf{T}^\top$ as the base matrix. \square

Lemma 6.3. Let $\bar{\mathbf{A}}$ be an arbitrary n -by- n symmetric matrix, let \mathbf{G}_D be a diagonal matrix of independent Gaussian random variables of variance σ^2 and mean 0, and let $\mathbf{A} = \bar{\mathbf{A}} + \mathbf{G}_D$. Then,

$$\Pr \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq \sqrt{\frac{2}{\pi}} \frac{n^{3/2}}{x\sigma}.$$

Proof. Let x_1, \dots, x_n be the diagonal entries of \mathbf{G}_D , and let

$$\begin{aligned} g &= \frac{1}{n} \sum_{i=1}^n x_i, \text{ and} \\ y_i &= x_i - g. \end{aligned}$$

Then,

$$\begin{aligned} \Pr_{y_1, \dots, y_n, g} \left[\left\| (\bar{\mathbf{A}} + \mathbf{G}_D)^{-1} \right\|_2 \geq x \right] &= \Pr_{y_1, \dots, y_n, g} \left[\left\| (\bar{\mathbf{A}} + \text{diag}(y_1, \dots, y_n) + g\mathbf{I})^{-1} \right\|_2 \geq x \right] \\ &\leq \max_{y_1, \dots, y_n, g} \Pr \left[\left\| (\bar{\mathbf{A}} + \text{diag}(y_1, \dots, y_n) + g\mathbf{I})^{-1} \right\|_2 \geq x \right], \end{aligned}$$

where the last inequality follows from Proposition 2.1. The proof now follows from Proposition 6.4 and Lemma 6.5. \square

Proposition 6.4. Let X_1, \dots, X_n be independent Gaussian random variables of variance σ^2 with means a_1, \dots, a_n , respectively. Let

$$\begin{aligned} G &= \frac{1}{n} \sum_{i=1}^n X_i, \text{ and} \\ Y_i &= X_i - G. \end{aligned}$$

Then, G is a Gaussian random variable of variance σ^2/n with mean $(1/n) \sum a_i$, independent of Y_1, \dots, Y_n .

Lemma 6.5. Let $\bar{\mathbf{A}}$ 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 $\mathbf{A} = \bar{\mathbf{A}} + G\mathbf{I}$. Then,

$$\Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] \leq \sqrt{\frac{2}{\pi}} \frac{n^{3/2}}{x\sigma}.$$

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

$$\left\| (\bar{\mathbf{A}} + G\mathbf{I})^{-1} \right\|_2^{-1} = \min_i |\lambda_i + G|.$$

Thus,

$$\Pr_{\mathbf{A}} \left[\left\| \mathbf{A}^{-1} \right\|_2 \geq x \right] = \Pr_G \left[\min_i |\lambda_i - G| < \frac{1}{x} \right] \leq \sum_i \Pr_G \left[|\lambda_i - G| < \frac{1}{x} \right] \leq \sum_i \sqrt{\frac{2}{\pi}} \frac{\sqrt{n}}{x\sigma} \leq \sqrt{\frac{2}{\pi}} \frac{n^{3/2}}{x\sigma},$$

where the second-to-last inequality follows from Lemma A.2 for \mathbb{R}^1 . \square

As in Section 3, we can now prove:

Theorem 6.6 (Condition number of symmetric matrices). Let $\bar{\mathbf{A}} = \bar{\mathbf{T}} + \bar{\mathbf{D}} + \bar{\mathbf{T}}^\top$ be an arbitrary n -by- n symmetric matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq \sqrt{n}$. Let $\sigma^2 \leq 1$, let \mathbf{T} be a zero-preserving perturbation of $\bar{\mathbf{T}}$ of variance σ^2 , let \mathbf{G}_D be a diagonal matrix of independent Gaussian random variables of variance σ^2 and mean 0 that are independent of \mathbf{T} , and let $\mathbf{D} = \bar{\mathbf{D}} + \mathbf{G}_D$. Then, for $\mathbf{A} = \mathbf{T} + \mathbf{D} + \mathbf{T}^\top$,

$$\Pr [\kappa(\mathbf{A}) \geq x] \leq 6\sqrt{\frac{2}{\pi}} \frac{n^{7/2}}{x\sigma} \left(1 + \sqrt{2\ln(x)/9n} \right)$$

Proof. As in the proof of Theorem 3.1, we can apply the techniques used in the proof of [DS01, Theorem II.7], to show

$$\Pr \left[\left\| \bar{\mathbf{A}} - \mathbf{A} \right\|_2 \geq 2\sqrt{n} + k \right] < e^{-k^2/2}.$$

The rest of the proof follows the outline of the proof of Theorem 3.1, using Lemma 6.2 instead of Theorem 3.3. \square

6.2 Bounding entries in \mathbf{U}

In this section, we will prove:

Theorem 6.7 ($\rho_{\mathbf{U}}(\mathbf{A})$ of symmetric matrices). Let $\bar{\mathbf{A}} = \bar{\mathbf{T}} + \bar{\mathbf{D}} + \bar{\mathbf{T}}^{\top}$ be an arbitrary n -by- n symmetric matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq 1$. Let $\sigma^2 \leq 1$, let \mathbf{T} be a zero-preserving perturbation of $\bar{\mathbf{T}}$ of variance σ^2 , let $\mathbf{G}_{\mathbf{D}}$ be a diagonal matrix of independent Gaussian random variables of variance σ^2 and mean 0 that are independent of \mathbf{T} , and let $\mathbf{D} = \bar{\mathbf{D}} + \mathbf{G}_{\mathbf{D}}$. Then, for $\mathbf{A} = \mathbf{T} + \mathbf{D} + \mathbf{T}^{\top}$,

$$\Pr[\rho_{\mathbf{U}}(\mathbf{A}) > 1 + x] \leq \frac{2}{7} \sqrt{\frac{2}{\pi}} \frac{n^3}{x\sigma}$$

Proof. We proceed as in the proof of Theorem 4.1. For k between 2 and n , we define \mathbf{u} , \mathbf{a} , \mathbf{b} and \mathbf{C} as in the proof of Theorem 4.1. By (4.2)

$$\frac{\|\mathbf{u}\|_1}{\|\mathbf{A}\|_{\infty}} \leq 1 + \left\| \left(\mathbf{C}^{\top} \right)^{-1} \mathbf{b} \right\|_1 \leq 1 + \sqrt{k-1} \left\| \mathbf{b}^{\top} \mathbf{C}^{-1} \right\|_2 \leq 1 + \sqrt{k-1} \|\mathbf{b}\|_2 \left\| \mathbf{C}^{-1} \right\|_2.$$

Hence

$$\begin{aligned} \Pr \left[\frac{\|\mathbf{u}\|_1}{\|\mathbf{A}\|_{\infty}} > 1 + x \right] &\leq \Pr \left[\|\mathbf{b}\|_2 \left\| \mathbf{C}^{-1} \right\|_2 > \frac{x}{\sqrt{k-1}} \right] \\ &\leq \mathbf{E} [\|\mathbf{b}\|_2] \sqrt{\frac{2}{\pi}} \frac{(k-1)^2}{x\sigma}, && \text{by Lemmas 6.2 and C.4,} \\ &\leq \sqrt{1+j}\sigma^2 \sqrt{\frac{2}{\pi}} \frac{(k-1)^2}{x\sigma}, && \text{where } j \text{ is the number of non-zeros in } \mathbf{b}, \\ &\leq \sqrt{\frac{2}{\pi}} \frac{\sqrt{k}(k-1)^2}{x\sigma}. \end{aligned}$$

Applying a union bound over k ,

$$\Pr[\rho_{\mathbf{U}}(\mathbf{A}) > x] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma} \sum_{k=2}^n \sqrt{k}(k-1)^2 \leq \frac{2}{7} \sqrt{\frac{2}{\pi}} \frac{n^{7/2}}{x\sigma}.$$

□

6.3 Bounding entries in \mathbf{L}

As in Section 4.2, we derive a bound on the growth factor of \mathbf{L} . As before, we will show that it is unlikely that $\mathbf{A}_{j,k}^{(k-1)}$ is large while $\mathbf{A}_{k,k}^{(k-1)}$ is small. However, our techniques must differ from those used in Section 4.2, as the proof in that section made critical use of the independence of $\mathbf{A}_{k,1:(k-1)}$ and $\mathbf{A}_{1:(k-1),k}$.

Theorem 6.8 ($\rho_{\mathbf{L}}(\mathbf{A})$ of symmetric matrices). Let $\sigma^2 \leq 1$ and $n \geq 2$. Let $\bar{\mathbf{A}} = \bar{\mathbf{T}} + \bar{\mathbf{D}} + \bar{\mathbf{T}}^{\top}$ be an arbitrary n -by- n symmetric matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq 1$. Let \mathbf{T} be a zero-preserving perturbation of $\bar{\mathbf{T}}$ of variance σ^2 , let $\mathbf{G}_{\mathbf{D}}$ be a diagonal matrix of independent Gaussian random variables of

variance $\sigma^2 \leq 1$ and mean 0 that are independent of \mathbf{T} , and let $\mathbf{D} = \bar{\mathbf{D}} + \mathbf{G}_D$. Let $\mathbf{A} = \mathbf{T} + \mathbf{D} + \mathbf{T}^\top$. Then,

$$\forall x \geq \sqrt{\frac{2}{\pi}} \frac{1}{\sigma^2}, \quad \Pr [\rho_{\mathbf{L}}(\mathbf{A}) > x] \leq \frac{3.2n^4}{x\sigma^2} \ln^{3/2} \left(e \sqrt{\frac{\pi}{2}} x \sigma^2 \right).$$

Proof. Using Lemma 6.9, we obtain for all k

$$\Pr [\exists j > k : |\mathbf{L}_{j,k}| > x] \leq \Pr \left[\|\mathbf{L}_{(k+1):n,k}\|_2 > x \right] \leq \frac{3.2n^2}{x\sigma^2} \ln^{3/2} \left(e \sqrt{\frac{\pi}{2}} x \sigma^2 \right).$$

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

$$\Pr [\exists j, k : |\mathbf{L}_{j,k}| > x] \leq \frac{3.2n^3}{x\sigma^2} \ln^{3/2} \left(e \sqrt{\frac{\pi}{2}} x \sigma^2 \right).$$

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

Lemma 6.9. Under the conditions of Theorem 6.8,

$$\forall x \geq \sqrt{\frac{2}{\pi}} \frac{1}{\sigma^2}, \quad \Pr \left[\|\mathbf{L}_{(k+1):n,k}\|_2 > x \right] \leq \frac{3.2n^2}{x\sigma^2} \ln^{3/2} \left(e \sqrt{\frac{\pi}{2}} x \sigma^2 \right).$$

Proof. We recall that

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

Because of the symmetry of \mathbf{A} , $\mathbf{A}_{k,1:k-1}$ is the same as $\mathbf{A}_{1:k-1,k}$, so we can no longer use the proof technique that worked in Section 4.2. Instead, we will bound the tails of the numerator and denominator separately, exploiting the fact that only the denominator depends upon $\mathbf{A}_{k,k}$.

Consider the numerator first. Setting $\mathbf{v} = \mathbf{A}_{1:k-1,1:k-1}^{-1} \mathbf{A}_{1:k-1,k}$, the numerator can be written $\mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ \mathbf{1} \end{pmatrix}$. We will now prove that for all $x \geq 1/\sigma$,

$$\Pr_{\substack{\mathbf{A}_{k+1:n,1:k} \\ \mathbf{A}_{1:k-1,1:k}}} \left[\left\| \mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ \mathbf{1} \end{pmatrix} \right\|_\infty > x \right] \leq \sqrt{\frac{2}{\pi}} \left(\frac{2n^2(1 + \sigma\sqrt{2\ln(x\sigma)}) + n}{x\sigma} \right). \quad (6.1)$$

Let

$$c = \frac{1}{1 + \sigma\sqrt{2\ln(x\sigma)}}, \quad (6.2)$$

which implies $\frac{1-c}{c\sigma} = \sqrt{2\ln(x\sigma)}$. It suffices to prove (6.1) for all x for which the right-hand side is less than 1. Given that $x \geq 1/\sigma$, it suffices to consider x for which $cx \geq 2$ and $x\sigma \geq 2$.

We use the parameter c to divide the probability as follows:

$$\begin{aligned} & \Pr_{\substack{\mathbf{A}_{k+1:n,1:k} \\ \mathbf{A}_{1:k-1,1:k}}} \left[\left\| \mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_{\infty} > x \right] \\ & \leq \Pr_{\mathbf{A}_{1:(k-1),1:k}} \left[\left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2 > cx \right] \end{aligned} \quad (6.3)$$

$$+ \Pr_{\mathbf{A}_{k+1:n,1:k}} \left[\left\| \mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_{\infty} > \frac{1}{c} \left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2 \mid \left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2 \leq cx \right] \quad (6.4)$$

To evaluate (6.4), we note that once \mathbf{v} is fixed, each component of $\mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix}$ is a Gaussian random variable of variance $\left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2^2 \sigma^2$ and mean at most $\|\bar{\mathbf{A}}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix}\|_2 \leq \left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2$. So,

$$\left\| \mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_{\infty} > \frac{1}{c} \left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2$$

implies one of the Gaussian random variables differs from its mean by more than $(1/c - 1)/\sigma$ times its standard deviation, and we can therefore apply Lemma A.1 and a union bound to derive

$$(6.4) \leq \sqrt{\frac{2}{\pi}} \frac{n e^{-\frac{1}{2} \left(\frac{1-c}{c\sigma}\right)^2}}{\frac{1-c}{c\sigma}} = \sqrt{\frac{2}{\pi}} \frac{n}{x\sigma \sqrt{2 \ln(x\sigma)}}.$$

To bound (6.3), we note that Lemma 6.2 and Corollary C.5 imply

$$\Pr_{\mathbf{A}_{1:(k-1),1:k}} \left[\left\| \mathbf{A}_{1:k-1,1:k-1}^{-1} \mathbf{A}_{1:k-1,k} \right\|_2 > y \right] \leq \sqrt{\frac{2}{\pi}} \frac{n^2}{y\sigma},$$

and so

$$\begin{aligned} \Pr_{\mathbf{A}_{1:(k-1),1:k}} \left[\left\| \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_2 > cx \right] & \leq \Pr_{\mathbf{A}_{1:(k-1),1:k}} \left[\left\| \mathbf{A}_{1:k-1,1:k-1}^{-1} \mathbf{A}_{1:k-1,k} \right\|_2 > cx - 1 \right] \\ & \leq \sqrt{\frac{2}{\pi}} \frac{n^2}{(cx - 1)\sigma} \\ & = \sqrt{\frac{2}{\pi}} \frac{n^2}{cx\sigma(1 - 1/cx)} \\ & = \sqrt{\frac{2}{\pi}} \frac{n^2(1 + \sigma\sqrt{2 \ln(x\sigma)})}{x\sigma(1 - 1/cx)} \\ & \leq \sqrt{\frac{2}{\pi}} \frac{2n^2(1 + \sigma\sqrt{2 \ln(x\sigma)})}{x\sigma}, \text{ by } cx \geq 2. \end{aligned}$$

So,

$$\begin{aligned}
\Pr_{\substack{\mathbf{A}_{k+1:n,1:k} \\ \mathbf{A}_{1:k-1,1:k}}} \left[\left\| \mathbf{A}_{k+1:n,1:k} \begin{pmatrix} -\mathbf{v} \\ 1 \end{pmatrix} \right\|_{\infty} > x \right] &\leq \sqrt{\frac{2}{\pi}} \left(\frac{n}{x\sigma\sqrt{2\ln(x\sigma)}} + \frac{2n^2(1 + \sigma\sqrt{2\ln(x\sigma)})}{x\sigma} \right) \\
&\leq \sqrt{\frac{2}{\pi}} \left(\frac{2n^2(1 + \sigma\sqrt{2\ln(x\sigma)}) + n}{x\sigma} \right), \tag{6.5}
\end{aligned}$$

by the assumption $x\sigma \geq 2$, which proves (6.1).

As for the denominator, we note that $\mathbf{A}_{k,k}$ is independent of all other terms, and hence

$$\Pr \left[\left| \mathbf{A}_{k,k} - \mathbf{A}_{k,1:k-1} \mathbf{A}_{1:k-1,1:k-1}^{-1} \mathbf{A}_{1:k-1,k} \right| < 1/x \right] \leq \sqrt{\frac{2}{\pi}} \frac{1}{x\sigma}, \tag{6.6}$$

by Lemma A.2. Applying Corollary C.3 with

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

to combine (6.5) with (6.6), we derive the bound

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

as $\sigma \leq 1$. □

7 Conclusions and open problems

7.1 Generality of results

In this paper, we have presented bounds on the smoothed values of the condition number and growth factors assuming the input matrix is subjected to a slight Gaussian perturbation. We would like to point out here that our results can be extended to some other families of perturbations.

With the exception of the proof of Theorem 3.3, the only properties of Gaussian random vectors that we used in Sections 3 and 4 are

1. there is a constant c for which the probability that a Gaussian random vector has distance less than ϵ to a hyperplane is at most $c\epsilon$, and

2. it is exponentially unlikely that a Gaussian random vector lies far from its mean.

Moreover, a result similar to Theorem 3.3 but with an extra factor of d could be proved using just fact 1.

In fact, results of a character similar to ours would still hold if the second condition were reduced to a polynomial probability. Many other families of perturbations share these properties. For example, similar results would hold if we let $\mathbf{A} = \bar{\mathbf{A}} + \mathbf{U}$, where \mathbf{U} is a matrix of variables independently uniformly chosen in $[-\sigma, \sigma]$, or if $\mathbf{A} = \bar{\mathbf{A}} + \mathbf{S}$, where the columns of \mathbf{S} are chosen uniformly among those vectors of norm at most σ .

7.2 Counter-Examples

The results of sections 3 and 4 do not extend to zero-preserving perturbations for non-symmetric matrices. For example, the following matrix remains ill-conditioned under zero-preserving perturbations.

$$\begin{pmatrix} 1 & -2 & 0 & 0 & 0 \\ 0 & 1 & -2 & 0 & 0 \\ 0 & 0 & 1 & -2 & 0 \\ 0 & 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

A symmetric matrix that remains ill-conditioned under zero-preserving perturbations that do not alter the diagonal can be obtained by locating the above matrix in the upper-right quadrant, and its transpose in the lower-left quadrant:

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 & -2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -2 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The following matrix maintains large growth factor under zero-preserving perturbations, re-

ardless of whether partial pivoting or no pivoting is used.

$$\begin{array}{cccccc}
 1.1 & 0 & 0 & 0 & 0 & 1 \\
 -1 & 1.1 & 0 & 0 & 0 & 1 \\
 -1 & -1 & 1.1 & 0 & 0 & 1 \\
 -1 & -1 & -1 & 1.1 & 0 & 1 \\
 -1 & -1 & -1 & -1 & 1.1 & 1 \\
 -1 & -1 & -1 & -1 & -1 & 1
 \end{array}$$

These examples can be easily normalized to so that their 2-norms are equal to 1.

7.3 Open Problems

Questions that naturally follow from this work are:

- What is the probability that the perturbation of an arbitrary matrix has large growth factors under Gaussian elimination with partial pivoting?
- What is the probability that the perturbation of an arbitrary matrix has large growth factors under Gaussian elimination with complete pivoting?
- Can zero-preserving perturbations of symmetric matrices have large growth factors under partial pivoting or under complete pivoting?
- Can zero-preserving perturbations of arbitrary matrices have large growth factors under complete pivoting?

For the first question, we point out that experimental data of Trefethen and Bau [TB97, p. 168] suggest that the probability that the perturbation of an arbitrary matrix has large growth factor under partial pivoting may be exponentially smaller than without pivoting. This leads us to conjecture:

Conjecture 3. Let $\bar{\mathbf{A}}$ be an n -by- n matrix for which $\|\bar{\mathbf{A}}\|_2 \leq 1$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Let \mathbf{U} be the upper-triangular matrix obtained from the LU-factorization of \mathbf{A} with partial pivoting. There exist absolute constants k_1 , k_2 and α for which

$$\Pr [\|\mathbf{U}\|_{\max}/\|\mathbf{A}\|_{\max} > \kappa + 1] \leq n^{k_1} e^{-\alpha \kappa^{k_2} \sigma}$$

Finally, we ask whether similar analyses can be performed for other algorithms of Numerical Analysis. One might start by extending Smale's program by analyzing the smoothed values of other condition numbers.

7.4 Recent Progress

Since the announcement of our result, Wschebor [Wsc04] improved the smoothed bound on the condition number.

Theorem 7.1 (Wschebor). Let $\bar{\mathbf{A}}$ be an $n \times n$ matrix and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Then,

$$\Pr[\kappa(\mathbf{A}) \geq x] \leq \frac{n}{x} \left(\frac{1}{4\sqrt{2\pi n}} + 7 \left(5 + \frac{4\|\bar{\mathbf{A}}\|_2^2(1 + \log n)}{\sigma^2 n} \right)^{1/2} \right)$$

When $\|\bar{\mathbf{A}}\|_2 \leq \sqrt{n}$, his result implies

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

We conjecture

Conjecture 4. Let $\bar{\mathbf{A}}$ be an $n \times n$ matrix satisfying $\|\bar{\mathbf{A}}\|_2 \leq \sqrt{n}$, and let \mathbf{A} be a Gaussian perturbation of $\bar{\mathbf{A}}$ of variance $\sigma^2 \leq 1$. Then,

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

8 Acknowledgments

We thank Alan Edelman for suggesting the name “smoothed analysis”, for suggesting we examine growth factors, and for his continuing support of our efforts. We thank Juan Cuesta and Mario Wschebor for pointing out some mistakes in an early draft of this paper. We thank Felipe Cucker for bringing Wschebor’s paper [Wsc04] to our attention. Finally, we thank the referees for their extraordinary efforts and many helpful suggestions.

References

- [ABB⁺99] 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.
- [AS64] 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).

- [BD02] Avrim Blum and John Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *SODA '02*, pages 905–914, 2002.
- [Blu89] 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.
- [Dem88] James Demmel. The probability that a numerical analysis problem is difficult. *Math. Comp.*, pages 499–480, 1988.
- [Dem97] James Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [DS01] K. R. Davidson and S. J. Szarek. In W. B. Johnson and J. Lindenstrauss, editors, *Handbook on the Geometry of Banach spaces*, chapter Local operator theory, random matrices, and Banach spaces, pages 317–366. Elsevier Science, 2001.
- [DST02] 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.
- [Ede88] Alan Edelman. Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.*, 9(4):543–560, 1988.
- [Ede92] 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.
- [GL83] 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.
- [GLS91] Martin Grottschel, Laszlo Lovasz, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1991.
- [Hig90] 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.
- [JKB95] N. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 2. Wiley-Interscience, 1995.
- [KJ82] Samuel Kotz and Norman L. Johnson, editors. *Encyclopedia of Statistical Sciences*, volume 6. John Wiley & Sons, 1982.

- [LT91] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer-Verlag, 1991.
- [Ren95] J. Renegar. Incorporating condition measures into the complexity theory of linear programming. *SIAM J. Optim.*, 5(3):506–524, 1995.
- [Seg00] Yoav Seginer. The expected norm of random matrices. *Combinatorics, Probability and Computing*, 9:149–166, 2000.
- [Sma97] Steve Smale. Complexity theory and numerical analysis. *Acta Numerica*, pages 523–551, 1997.
- [Smo] <http://math.mit.edu/~spielman/SmoothedAnalysis>.
- [ST03] Daniel Spielman and Shang-Hua Teng. Smoothed analysis of termination of linear programming algorithms. *Mathematical Programming, Series B*, 97:375–404, 2003.
- [ST04] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, 2004.
- [TB97] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.
- [TS90] 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.
- [Wil61] J. H. Wilkinson. Error analysis of direct methods of matrix inversion. *J. Assoc. Comput. Mach.*, 8:261–330, 1961.
- [Wsc04] M. Wschebor. Smoothed analysis of $\kappa(\mathbf{a})$. *J. of Complexity*, 20(1):97–107, February 2004.
- [YC97] Man-Chung Yeung and Tony F. Chan. Probabilistic analysis of gaussian elimination without pivoting. *SIAM J. Matrix Anal. Appl.*, 18(2):499–517, 1997.

A Gaussian random variables

Lemma A.1. Let X be a univariate Gaussian random variable with mean 0 and standard deviation 1. Then for all $k \geq 1$,

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

Proof. We have

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

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

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

□

Lemma A.2. Let \mathbf{x} be a d -dimensional Gaussian random vector of variance σ^2 , let \mathbf{t} be a unit vector, and let λ be a real. Then,

$$\Pr\left[|\mathbf{t}^\top \mathbf{x} - \lambda| \leq \epsilon\right] \leq \sqrt{\frac{2}{\pi}} \frac{\epsilon}{\sigma}.$$

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

$$\mathbb{E}\left[\max_i |g_i|\right] \leq \sqrt{2 \ln(\max(n, 2))} + \frac{1}{\sqrt{2\pi} \ln(\max(n, 2))}.$$

Proof. For any $a \geq 1$,

$$\begin{aligned} \mathbb{E}\left[\max_i |g_i|\right] &= \int_{t=0}^\infty \Pr\left[\max_i |g_i| \geq t\right] dt \\ &\leq \int_{t=0}^a 1 dt + \int_a^\infty n \Pr[|g_1| \geq t] dt \\ &\leq a + \int_a^\infty n \frac{2}{\sqrt{2\pi}} \frac{e^{-\frac{1}{2}t^2}}{t} dt && \text{(applying Lemma A.1,)} \\ &= a + \frac{2n}{\sqrt{2\pi}} \int_a^\infty \frac{e^{-\frac{1}{2}t^2}}{t^2} d\left(\frac{1}{2}t^2\right) \\ &\leq a + \frac{2n}{\sqrt{2\pi}} \frac{1}{a^2} \int_a^\infty e^{-\frac{1}{2}t^2} d\left(\frac{1}{2}t^2\right) \\ &= a + \frac{2n}{\sqrt{2\pi}} \frac{1}{a^2} e^{-\frac{1}{2}a^2}. \end{aligned}$$

Setting $\mathbf{a} = \sqrt{2\ln(\max(n, 2))}$, which is greater than 1 for all $n \geq 1$, we obtain the following upper bound on the expectation:

$$\sqrt{2\ln(\max(n, 2))} + \frac{2n}{\sqrt{2\pi}} \frac{1}{2\ln(\max(n, 2))} \frac{1}{\max(n, 2)} \leq \sqrt{2\ln(\max(n, 2))} + \frac{1}{\sqrt{2\pi}\ln(\max(n, 2))}.$$

□

Lemma A.4 (Expectation of reciprocal of the 1-norm of a Gaussian vector). Let $\bar{\mathbf{a}}$ be an arbitrary column vector in \mathbb{R}^n for $n \geq 2$. Let \mathbf{a} be a Gaussian perturbation of $\bar{\mathbf{a}}$ of variance σ^2 . Then

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

Proof. Let $\mathbf{a} = (a_1, \dots, a_n)$. It is clear that the expectation of $1/\|\mathbf{a}\|_1$ is maximized if $\bar{\mathbf{a}} = \mathbf{0}$, so we will make this assumption. Without loss of generality, we also assume $\sigma^2 = 1$. For general σ , we can simply scale the bound by the factor $1/\sigma$.

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

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

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

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

Taking second derivatives, and applying the inequality (*c.f.* [AS64, 26.2.13])

$$\frac{1}{\sqrt{2\pi}} \int_t^\infty e^{-\frac{1}{2}x^2} dx \geq \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} \frac{1}{x + 1/x},$$

we find that $e^{\frac{1}{2}t^2} \operatorname{erfc} \left(\frac{t}{\sqrt{2}} \right)$ is convex.

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

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

Numerically, we find that $\alpha \approx 1.9857 < 2$.

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

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

which follows from Lemma A.1.

We now have

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

for $n \geq 2$. To verify this last equality, one can multiply through by $(n+1)(n-1)$ to obtain

$$\sqrt{\frac{2}{\pi}}(n+1)(2/c)^{(n-1)/2} \leq 4,$$

which one can verify by taking the derivative of the left-hand side to find the point where it is maximized, $n = (2 + \ln(5/6))/\ln(6/5)$. \square

B Random point on a sphere

Lemma B.1. Let $d \geq 2$ and let (u_1, \dots, u_d) be a unit vector chosen uniformly at random in \mathbb{R}^d . Then, for $c \leq 1$,

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

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

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

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

We have

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

We now note that

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

is a random variable distributed according to the t -distribution with $d-1$ degrees of freedom. The lemma now follows from the fact (*c.f.* [JKB95, Chapter 28, Section 2] or [AS64, 26.7.5]) that, for $c > 0$,

$$\Pr [t_d > \sqrt{c}] \geq \Pr [G > \sqrt{c}],$$

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

C Combination Lemmas

Lemma C.1. 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 $\lim_{x \rightarrow \infty} g(x) = 0$. 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

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

integrating by parts,

$$\begin{aligned} &= \int_0^\infty \Pr [A \geq s] \frac{d}{ds} g\left(\frac{x}{s}\right) ds \\ &\leq \int_0^\infty f(s) \frac{d}{ds} g\left(\frac{x}{s}\right) ds, \end{aligned}$$

setting $t = x/s$

$$= \int_0^\infty f\left(\frac{x}{t}\right) (-g'(t)) dt.$$

□

Corollary C.2 (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,

$$\Pr [AB \geq x] \leq \frac{\alpha\beta}{x} \left(1 + \max\left(0, \ln\left(\frac{x}{\alpha\beta}\right)\right) \right)$$

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

$$\Pr [A \geq x] \leq \min\left(\frac{\alpha}{x}, 1\right) \stackrel{\text{def}}{=} f(x), \text{ and}$$

$$\Pr [B \geq x] \leq \min\left(\frac{\beta}{x}, 1\right) \stackrel{\text{def}}{=} g(x).$$

Applying Lemma C.1 while observing

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

we obtain

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

where the **max** appears in case $x/\alpha < \beta$. □

Corollary C.3. Let A and B be two positive random variables. If

1. $\forall x \geq 1/\sigma$, $\Pr[A \geq x] \leq \min\left(1, \frac{\alpha + \beta\sqrt{\ln x\sigma}}{\sigma x}\right)$ and
2. $\Pr[B \geq x|A] \leq \frac{\gamma}{x\sigma}$

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

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

Proof. Define f and g by

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

Applying Lemma C.1 while observing

- $g'(t) = 0$ for $t \in [0, \frac{\gamma}{\sigma}]$, 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 $s = \sqrt{\ln(x\sigma/t)}$, $t = x\sigma e^{-s^2}$, which is defined as $x \geq \gamma/\sigma^2$,)

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

as $\alpha \geq 1$. □

Lemma C.4 (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,

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

Proof. Let $g(x)$ be the distribution function of C. By Lemma C.1, we have

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

□

Corollary C.5 (linear-chi). Let A be a positive random variable such that

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

for some $\alpha > 0$. Let \mathbf{b} be a d -dimensional Gaussian random vector (possibly depending upon A) of variance at most σ^2 centered at a vector of norm at most t , and let $B = \|\mathbf{b}\|_2$. Then,

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

Proof. As $\mathbb{E}[B] \leq \sqrt{\mathbb{E}[B^2]}$, and it is known [KJ82, p. 277] that the expected value of B^2 —the non-central χ^2 -distribution with non-centrality parameter $\|\bar{\mathbf{b}}\|_2^2$ —is $\sigma^2 d + \|\bar{\mathbf{b}}\|_2^2$, the corollary follows from Lemma C.4. \square

Lemma C.6 (Linear to log). Let A be a positive random variable. If there exists an $A_0 \geq 1$ and an $\alpha \geq 1$ such that for all $x \geq A_0$,

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

Then,

$$\mathbb{E}_A [\max(0, \ln A)] \leq \ln \max(A_0, \alpha) + 1.$$

Proof.

$$\begin{aligned} \mathbb{E}_A [\max(0, \ln A)] &= \int_{x=0}^{\infty} \Pr_A [\max(0, \ln A) \geq x] dx \\ &\leq \int_{x=0}^{\ln \max(A_0, \alpha)} 1 dx + \int_{x=\ln \max(A_0, \alpha)}^{\infty} \Pr_A [\ln A \geq x] dx \\ &\leq \int_{x=0}^{\ln \max(A_0, \alpha)} dx + \int_{x=\ln \max(A_0, \alpha)}^{\infty} \alpha e^{-x} dx \\ &\leq \ln \max(A_0, \alpha) + 1. \end{aligned}$$

\square