5-2015

New Results on Randomized Matrix Computations

Jesse Lowell Wolf
Graduate Center, City University of New York

NEW RESULTS ON RANDOMIZED MATRIX COMPUTATIONS

by

JESSE WOLF

A dissertation submitted to the Graduate Faculty in Mathematics in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2015

This manuscript has been read and accepted for Graduate Faculty in Mathematics in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date

Chair of Examining Committee
Prof. Victor Y. Pan
The City University of New York

Date

Executive Officer
Prof. Linda Keen
The City University of New York

Prof. Victor Y. Pan
The City University of New York

Prof. Dana Draghicescu
The City University of New York

Prof. Roman Kossak
The City University of New York

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

NEW RESULTS ON RANDOMIZED MATRIX COMPUTATIONS

by **JESSE WOLF**

Adviser: Professor Victor Y. Pan

The aim of this thesis is to present new results in randomized matrix computations. Specifically, and ultimately, we show how to modify, or preprocess an ill conditioned matrix having small numerical nullity (co-rank) into a nonsingular well conditioned matrix. This has intrinsic theoretical interest and we show a sample application to accurate solutions of nonsingular and ill conditioned linear systems. We discuss both multiplicative and additive preprocessing; in fact the multiplicative case assists in the derivation of the additive case. In the additive case, we approximate a nonsingular ill conditioned matrix by a singular well conditioned matrix which is then preprocessed into a nonsingular well conditioned matrix, which will also assist in the aforementioned ultimate goal.

Preface

Given a nonsingular ill conditioned matrix A^*, we can readily produce a nonsingular well conditioned matrix C^* by applying additive preprocessing $A^* \Longrightarrow C^* = A^* + P$, with P to be defined in the sequel. Suppose that the numerical nullity of A^* is small in context, that is, only a small number of its singular values are much smaller than its norm. Such matrices make up a large and important subclass of nonsingular ill conditioned matrices, for which randomized additive preconditioning is supported by the following theorem. We prove it and suggest an extension to rectangular matrices in the sequel. In this thesis A^* is not the Hermitian conjugate of A; in fact, we initially assume that all components of A^* are real. In the sequel, when we allow A^* to contain complex entries, the Hermitian conjugate of A will be denoted A^H.

The following, to be restated in the Introduction and Chapter 1 and proven in Chapter 4, is our Key Result, where we use the concept of probabilistic order, to be defined in Section 3.1.

Theorem 0.0.1 *Suppose A^* is a nonsingular ill conditioned matrix with real components. Suppose $C^* \equiv A^* + UV^T$. Suppose further that $\|A^*\|_2 \approx \|UV^T\|_2$, where $U, V \in \mathcal{G}_{0,\sigma}^{n\times r}$, Gaussian random matrices with mean zero and standard deviation σ of order $n \times r$. Let r be the numerical nullity of A^*, denoted numnull(A^*). Then with Probability $= 1$, C^* is nonsingular and $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$.*

ACKNOWLEDGMENTS

To my mother Anita Wolf, father Jacob Wolf, sister Lisa Wolf - Bonilla, brother - in - law Jeffrey Bonilla and brother Joel Wolf, maternal grandparents Reinhilde and Frederick Lurch, paternal grandfather Joseph Wolf, Esther and Coleman Gregory, Esq., Minerva, Julius and Lisa Brodsky, for everything. To my non - mathematician but street - wise Philly friend, Rick Finestein, for his encouragement in my pursuit of the doctorate. To my senior consigliere, Elaine Grabel, and to my partner in no crime, Alan Rabunski, Esq., for the same. To Ulyana Bardyn, Esq. and Gary Meyerhoff, Esq. for pro bono support; to Mark Turken for the same. To Prof. Paul Chernoff for his generous support during my first year at Berkeley, and to Prof. David Gale, in memorium, for his equally generous support during my last year at Berkeley. To Professor Eli Passow for introducing me to higher mathematics. To Prof. Frederick P. Greenleaf for introducing me to more advanced mathematical techniques later in life. To my friend Prof. Chokri Cherif for invaluable career guidance. To my friend Prof. Lucio Prado for near infinite patience in teaching me Latex and many other things. And of course, to Distinguished Professor Victor Y. Pan, my advisor, both gentleman and scholar, and without whom this dissertation would not exist.

To my father Jacob Wolf.

Contents

INTRODUCTION

A standard problem in Linear Algebra is the following: Suppose A^* is a given non-singular ill-conditioned $n \times n$ matrix, x and b are vectors of dimension n. The vector b is given and one seeks the vector x: $A^*x = b$. Of course, standard techniques of Gaussian elimination can be applied. However, given the ill-conditioning of A^*, to be further described below, this entails numerical problems. In particular, define the ratio of [output error]/[input error] $= R$, a measure of inefficiency. We prove that there are non-standard techniques from applied linear algebra, which, with a high degree of probability, dramatically decrease R. That is, we utilize known probabilistic results combined with a new applied linear algebra algorithm to derive the solution x in a computationally efficient manner.

This probabilistic approach appears on average to reduce R by a factor on the order of 10^{-13}, in accordance with statistical tests in [**PQZ11**] and [**PQZ14**], by Victor Pan, et. al. Preconditioning tests from the latter paper are presented in the Appendix. Moreover, the results in this thesis are largely based on earlier results in the papers [**PGMQ**], [**PIMR10**], [**PQ10**], [**PQ12**], [**PQYa**], [**PQZ13**], and [**HMT11**].

First, a bit of background: Suppose M is an $n \times n$ matrix; it is well known that M possesses a so-called Singular Value Decomposition [SVD]: $M = S_M \Sigma_M T_M^T$, a matrix decomposition in which S_M and T_M^T are orthogonal matrices and Σ_M is a diagonal matrix with diagonal entries σ_i, each respectively in the (i, i) position, in non-increasing order. σ_1 is the largest positive singular value and σ_ρ is the smallest positive singular value for $\rho = rank(M)$. So $\sigma_i = 0$ for $i > \rho$. There are a variety of ways to further characterize the σs explicitly; one way is to view them as the non-increasing square roots of the eigenvalues of MM^T. The ratio: σ_1/σ_ρ is defined as the condition number of M, $\kappa(M)$. M is ill-conditioned if $\kappa(M)$ is large. For the solution of a linear system of equations, as above, we have $\kappa(M) \approx R$, and so for ill-conditioned linear systems one must perform computations with a high precision to obtain meaningful output. Another way of viewing the σs and thus $\kappa(M)$: σ_1 is the 2-norm of M and σ_ρ is the reciprocal of the 2-norm of M^+, the generalized inverse of M. If M is non-singular, then $\rho = n$ and $M^+ = M^{-1}$.

Established Probabilistic Results we utilize: Suppose M is an element of the class of $n \times n$ Gaussian matrices, each component an independent identically distributed normal random variable of mean $\mu = 0$ and standard deviation σ. Then the cumulative distribution function [cdf] of the 2-norm of M^+ as a function of z is at least $1 - \sqrt{n}/\sigma z$. This is Theorem 3.4 in [**SST06**], Sankar, Spielman, and Teng, Smoothed Analysis of the Condition Numbers and Growth Factors of Matrices.

We use this theorem, which is sufficient for our purpose, but there are earlier and stronger although less explicit estimates in Edelman in [**E88**], Eigenvalues and Condition Numbers of Random Matrices.

Furthermore, the cdf of the 2-norm of M as a function of z is at least $1 - \exp(-(z - 2\sigma\sqrt{n})^2/(2\sigma^2))$. This is a consequence of a related Theorem in [**DS01**], Davidson and Szarek, Handbook on the Geometry of Banach Spaces.

General Approach of Thesis:

We employ a technique that, by following [**PIMR10**] and [**PQZ14**], we call Additive Preprocessing, more carefully defined below. We generate a sequence of five additional matrices: $A^* \xrightarrow{1} A \xrightarrow{2} C \xrightarrow{3} C^* \xrightarrow{4} C^{*-1} \xrightarrow{5} A^{*-1}$ in which the matrix A^* having a large condition number is transformed into matrices A, C, C^* and C^{*-1} having reasonably bounded condition numbers. The solution x of the original system is finally obtained as $A^{*-1}b$. A^{*-1} has, by definition of condition number, the same large condition number as A^* itself.

The derivation of A^{*-1} from the matrix C^{*-1}, having reasonably bounded condition number via the classical Sherman-Morrison-Woodbury [SMW] Formula, also further described below, has the effect of substantially reducing numerical problems. Pre-conditioning allows us to limit highly accurate computations to $O(n^2)$ flops, versus Gaussian elimination, which requires highly accurate computations on the order of $O(n^3)$ flops.

Some details of the General Approach:

Suppose M, any $n \times n$ matrix, has singular values $\sigma_1, \sigma_2, ..., \sigma_{\rho num}$ which are large in context and the remaining singular values $\sigma_{\rho num+1}, ..., \sigma_n$ are small in context. ρnum is defined as the numerical rank of M, the cardinality of the set of large singular values. $n - \rho num \equiv numnull$ is defined as the numerical nullity of M, the cardinality of the set of small singular values.

In our background paragraph above, ρ is by construction the rank of that matrix M and $n - \rho$ is the nullity of that same matrix M.

Let us return to a consideration of A^*, our non-singular ill-conditioned matrix; so the rank of A^* is n and its numerical rank ρnum is less than n.

We will first transform A^* into A, a singular and well conditioned matrix, by zeroing out the small singular values of A^* to create A. Equivalently A may be achieved by subtracting a perturbation matrix E, of norm $\sigma_{\rho num+1}(A^*)$ and rank $numnull(A^*)$ from A^* to obtain A: $A = A^* - E$. Now A has, by construction, a nullity equal to: $null(A) \equiv r(A) \equiv r \equiv numnull(A^*)$.

Here is where the additive preconditioning comes in: We next transform A into a matrix C, which is nonsingular with probability one and well conditioned in a probabilistic sense. We define this in the treatise as well conditioned up to probabilistic order. Let U and V be Gaussian random matrices of order $n \times r$ with mean 0 and standard deviation σ, r as above.

Define $C = A + UV^T$. We show that $\kappa(C)$ and $\kappa(A)$ are of the same probabilistic order, also defined in the sequel. Essentially, what will happen is this: We will first discover a product of random variables which acts as an upper bound to the ratio $\frac{\kappa(C)}{\kappa(A)}$. While the *cdf* of the product or even any individual factor of that product is difficult or impossible to ascertain, it is possible via the probabilistic results cited above to provide deterministic lower bounds to the *cdf*s of each individual factor, respectively, and thus implicitly provide what we characterize in the sequel as Probabilistic Upper Bounds ($PUBs$) to each individual factor, respectively.

The PUBs are in fact the respective arguments of the *cdf*s. One reason this construction is useful: Several of the deterministic lower bounds to the *cdf*s of the factors cited above will depend on r. In fact, as will be clearly seen later, they are differentiable functions of r with negative derivatives: smaller numerical nullity provides stronger (larger) lower bounds. It will also be seen that all of the deterministic lower bounds are also differentiable functions of σ in which the derivatives are positive: more randomness also provides stronger lower bounds. By strengthening deterministic lower bounds, one is, in effect, also strengthening the PUBs.

We then transform C into C^*: $C^* = C + E$. C^* is also nonsingular with Probability = 1, and is well conditioned since it is a small perturbation of a well conditioned matrix. [In actual fact, A^* is taken to C^* in one step; the creation of A and C is a theoretical construct useful for our analysis]. The intuition behind this step is to obtain $(C^*)^{-1}$ in the following step via Gaussian elimination from which $(A^*)^{-1}$ may be easily retrieved in the penultimate step via the SMW formula. The final step will yield the solution x, in which the number of high precision arithmetic computations will have been reduced from order n^3 to order n^2.

Invert C^* using Gaussian elimination. Note that, by construction, $C^* = A^* + UV^T$. Define $G = I_r - V^T(C^*)^{-1}U$, called the Schur Complement or the Gauss Transform. Then the SMW formula yields: $(A^*)^{-1} = (C^*)^{-1} + (C^*)^{-1}UG^{-1}V^T(C^*)^{-1}$.

Finally, $x = (A^*)^{-1}b$.

We state a Theorem below which provides the random variable product upper bound to the ratio: $\frac{\kappa(C)}{\kappa(A)}$. We note first, however, that once the Theorem is stated and proven in the sequel, the following comments apply: The established probabilistic results cited above will immediately imply deterministic lower bounds to the *cdf*s of each factor in that product of random variables. This means that each factor in our product has a PUB. This further implies that $\kappa(C)$ and $\kappa(A)$ are of the same probabilistic order. Moreover, $C^* \approx C$, $\kappa(C^*) \approx \kappa(C)$ and C is well conditioned.

And as just stated, $\kappa(C) \approx \kappa(A)$, up to probabilistic order. Also, by construction: $\kappa(A) \ll \kappa(A^*)$. We take this to mean that $\kappa(C^*) \ll \kappa(A^*)$,up to probabilistic order. This somewhat intuitive and philosophical approach to gauging the relative sizes of $\kappa(C^*)$ and $\kappa(A^*)$ appears to agree with empirical results. Again, statistical tests performed by Distinguished Professor Victor Pan's group at The City University of New York Graduate Center have shown that $\frac{\kappa(C^*)}{\kappa(A^*)}$ was consistently small for non - singular ill - conditioned matrices A^*.

In fact, probabilistically, $\kappa(C^*) \ll \kappa(A^*)$ for large numerical nullity r as well, but the computation of C^* becomes costly in this case.

The following presumes that we have scaled the matrix A so that the ratio $\frac{\|A^*\|_2}{||UV^T||_2}$ is close to 1 or is at least neither large nor small.

The aforementioned Theorem: Suppose $A, C, S, T \in \mathbb{R}^{n\times n}$ and $U, V \in \mathcal{G}_{\mu,\sigma}^{n\times r}$ for two positive integers r and n, $r = n - \rho(A)$, $0 < \rho(A) < n$. Define $\rho(A) \equiv \rho$.

Define $\sigma_j(A) \equiv \sigma_j$. $A = S\Sigma T^T$ is the SVD of the matrix A. So S and T are orthogonal matrices, $\Sigma = \mathrm{diag}(\sigma_j)_{j=1}^n$, $\sigma_1 \geq \sigma_2 \geq ... \geq \sigma_\rho$, and $\sigma_j = 0$ for all $j \in \rho + 1, \cdots, n$. Let $C \equiv A + UV^T$. Implicitly define the matrices $\bar{U}, \bar{V}, U_r, V_r$ by writing:

$$S^T U = \begin{pmatrix} \bar{U} \\ U_r \end{pmatrix},\ T^T V = \begin{pmatrix} \bar{V} \\ V_r \end{pmatrix},\ R_U = \begin{pmatrix} I_\rho & \bar{U} \\ O_{r,\rho} & U_r \end{pmatrix},\ R_V = \begin{pmatrix} I_\rho & \bar{V} \\ O_{r,\rho} & V_r \end{pmatrix}, \tag{0.0.2}$$

where U_r and V_r are $r \times r$ matrices.

Then: $\frac{\kappa(C)}{\kappa(A)} \leq (1+||U||)^2(1+||V||)^2 \ \max\{1,||U_r^{-1}||\} \ \max\{1,||V_r^{-1}||\}$.

This result and the preceding comments prove the Thesis' Key Result, Theorem 0.0.1: Suppose A^* is a nonsingular ill conditioned matrix with real components. Suppose $C^* \equiv A^* + UV^T$. Suppose further that $\|A^*\|_2 \approx ||UV^T||_2$, where $U, V \in \mathcal{G}_{0,\sigma}^{n\times r}$, Gaussian random matrices with mean zero and standard deviation σ of order $n \times r$. Let r be the numerical nullity of A^*, denoted $numnull(A^*)$. Then with Probability = 1, C^* is nonsingular and $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$.

Chapter 1

PRELIMINARIES

In this chapter, we present definitions, terminology, and collect certain facts that we will need in the sequel. A few constructions will be introduced later when appropriate. Here, some terminology, definitions, and theorems are new in the context of randomized matrix computations. Most are standard known results of Applied Linear Algebra.

1.1 Basic Definitions

It is well known that random matrices tend to be well conditioned [**D88**], [**E88**], [**ES05**], [**CD05**], [**SST06**], [**B11**], and we employ this property to advance some fundamental matrix computations. Moreover, we provide an explicit definition of what we mean by "tend to be".

We prove that our techniques of randomized preprocessing are likely, i.e. up to probabilistic order, to be defined precisely in Chapter 3, precondition a large and important class of ill conditioned matrices. Some of our techniques and results can be of independent interest, e.g., our extensions of the Sherman–Morrison–Woodbury formula. Hereafter flop stands for arithmetic operation, and expected and likely mean up to probabilistic order, defined precisely in Section 3.1. $\sigma_j(A)$ denotes the jth largest singular value of an $n \times n$ matrix A, and the ratio $\kappa(A) = \sigma_1(A)/\sigma_\rho(A)$ for $\rho = \text{rank}(A)$ denotes its condition number. $\kappa(A) = ||A||\ ||A^{-1}||$ if $\rho = n$, that is if A is a nonsingular matrix. In general, $\kappa(A) = ||A||\ ||A^+||$ where A^+ is the Moore-Penrose generalized pseudo - inverse of A. If this number is large in context, then the matrix A is *ill conditioned*, otherwise *well conditioned*. For matrix inversion and solving linear systems of equations the condition number represents the output magnification of the errors of input perturbation.

$$\kappa(A) \approx \frac{||\text{OUTPUT ERROR}||}{||\text{INPUT ERROR}||} \tag{1.1.1}$$

1.1.1 Randomized Preconditioning: Statement of the Key Result

The following is a restatement of Theorem 0.0.1 and is our Key Result, to be proven in Chapter 4.

Theorem 1.1.2 *Suppose A^* is a nonsingular ill conditioned matrix with real components. Suppose $C^* \equiv A^* + UV^T$. Suppose further that $\|A^*\|_2 \approx ||UV^T||_2$, where $U, V \in \mathcal{G}_{0,\sigma}^{n\times r}$, Gaussian random matrices with mean zero and standard deviation σ of order $n \times r$. Let r be the numerical nullity of A^*, denoted numnull(A^*). Then with Probability $= 1$, C^* is nonsingular and $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$.*

Following are some definitions and basic results on matrix and related computations.

1.2 Other Basic Results

We here assume computations in the field $\mathbb{R}$ of real numbers, and comment on the extension to the field $\mathbb{C}$ of complex numbers in Chapter 4.

1.2.1 Matrix Computations

A^T is the transpose of a matrix A. A^H is its Hermitian transpose, that is, its complex conjugate transpose. A matrix A is symmetric if $A = A^T$. A real matrix Q is called *orthogonal* if $Q^TQ = I$ or $QQ^T = I$. I the identity matrix.

More generally, over the complex field $\mathbb{C}$ a matrix U is called *unitary* if $U^HU = I$ or $UU^H = I$.

1.2.2 Matrix Spaces

$\mathcal{R}(A)$ denotes the range of an $m \times n$ matrix A, that is the linear space $\{\mathbf{z} : \mathbf{z} = A\mathbf{x}\}$ generated by its columns. $\mathcal{N}(A)$ denotes its null space $\{\mathbf{v} : A\mathbf{v} = \mathbf{0}\}$, $\text{rank}(A) = \dim \mathcal{R}(A)$ its rank, and $\text{nul}(A) = \dim \mathcal{N}(A) = n - \text{rank}(A)$ its right nullity or just *nullity*, whereas $\text{nul}\, A^T = m - \text{rank}(A)$ is the left nullity of A, equal to $\text{nul}\, A$ if and only if $m = n$. $\mathbf{v}$ is a null vector of A if $A\mathbf{v} = \mathbf{0}$.

A matrix B having full column rank is said to be a *matrix basis* for its range $\mathcal{R}(B)$. Suppose a matrix B has full column rank and satisfies $\mathcal{R}(B) = \mathcal{N}(A)$ for a matrix A. Then we call B a *null matrix basis* or a *nmb* for this matrix A and write $B = \text{nmb}(A)$. $\mathcal{N}(A^T)$ is the left null space of a matrix A, and similarly the map $A \to A^T$ defines left null vectors, left nmbs, and the left nullity of a matrix A; in particular an $m \times n$ matrix of a rank ρ has the left nullity $m - \rho$.

1.2.3 Norms and SVD

$||A||_h$ is the h-norm of a matrix $A = (a_{i,j})_{i,j=1}^{m,n}$. We write $||A|| = ||A||_2 \equiv L_2$-norm , and $||\mathbf{v}|| = \sqrt{\mathbf{v}^T\mathbf{v}} = ||\mathbf{v}||_2$ and recall from [**GL96**] that

$$\max_{i,j=1}^{m,n}|a_{i,j}| \le ||A|| = ||A^T|| \le \sqrt{mn}\, \max_{i,j=1}^{m,n}|a_{i,j}|,$$

$$\frac{1}{\sqrt{m}}||A||_1 \le ||A|| \le \sqrt{n}||A||_1, \quad ||A||_1 = ||A^T||_\infty, \quad ||A||^2 \le ||A||_1||A||_\infty, \tag{1.2.1}$$

$$||AB||_h \le ||A||_h||B||_h \text{ for } h = 1, 2, \infty \text{ and any matrix product } AB. \tag{1.2.2}$$

Henceforth $||A|| \equiv ||A||_2$ as we employ the $||\ ||_2$ in the sequel.

A matrix A is *normalized* if $||A|| = 1$. Define an *SVD* or *full SVD* of an $m \times n$ matrix A of a rank ρ as follows,

$$A = S_A \Sigma_A T_A^T. \tag{1.2.3}$$

Here $S_A S_A^T = S_A^T S_A = I_m$, $T_A T_A^T = T_A^T T_A = I_n$, $\Sigma_A = \mathrm{diag}(\widehat{\Sigma}_A, O_{m-\rho,n-\rho})$, $\widehat{\Sigma}_A = \mathrm{diag}(\sigma_j(A))_{j=1}^{\rho}$, $\sigma_j = \sigma_j(A) = \sigma_j(A^T)$ is the jth largest singular value of a matrix A for $j = 1, \dots, \rho$, and we write $\sigma_j = 0$ for $j > \rho$. These values have the minimax property

$$\sigma_j = \max_{\dim(\mathbb{S})=j} \min_{\mathbf{x}\in\mathbb{S},\ ||\mathbf{x}||=1} ||A\mathbf{x}||,\ j = 1, \dots, \rho, \tag{1.2.4}$$

where $\mathbb{S}$ denotes linear spaces [**GL96**]. Consequently, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_\rho$, $\sigma_\rho > 0$, $\sigma_1 = \max_{||\mathbf{x}||=1} ||A\mathbf{x}|| = ||A||$.

1.2.4 Generalized Inverses

$A^+ = T_A\, \mathrm{diag}(\widehat{\Sigma}_A^{-1}, O_{n-\rho,m-\rho}) S_A^T$ is the Moore–Penrose pseudo-inverse of an $m \times n$ matrix A, and

$$||A^+|| = 1/\sigma_\rho(A) \tag{1.2.5}$$

for a matrix A of a rank ρ. A^{+T} stands for $(A^+)^T = (A^T)^+$, and A^{-T} stands for $(A^{-1})^T = (A^T)^{-1}$. An $n \times m$ matrix $X = A^{(I)}$ is a left inverse of an $m \times n$ matrix A if $XA = I$ and is its right inverse if $AX = I$. A^+ is a left or right inverse $A^{(I)}$ if and only if a matrix A has full rank. $A^{(I)}$ is unique and is equal to A^{-1} if A is a nonsingular matrix.

1.2.5 SMW Formula

Theorem 1.2.6 *[**GL96**, page 50], [**S98**, Corollary 4.3.2]. Suppose that $U, V \in \mathbb{R}^{n\times r}$, the matrices $A \in \mathbb{R}^{n\times n}$ and $C = A + UV^T$ are nonsingular, and $0 < r < n$. Then the matrix $G = I_r - V^T C^{-1} U$ is nonsingular and we have the Sherman–Morrison–Woodbury (hereafter* SMW*) formula* $A^{-1} = C^{-1} + C^{-1} U G^{-1} V^T C^{-1}$.

1.2.6 Condition Number, Numerical Rank and Numerical Nullity

$\kappa(A) = \frac{\sigma_1(A)}{\sigma_\rho(A)} = ||A||\ ||A^+||$ is the condition number of an $m \times n$ matrix A of a rank ρ. Such a matrix is *ill conditioned* if $\sigma_1(A) \gg \sigma_\rho(A)$ and is *well conditioned* otherwise. See [**D83**], [**GL96**, Sections 2.3.2, 2.3.3, 3.5.4, 12.5], and [**S98**, Section 5.3] on the estimation of norms and condition numbers of nonsingular matrices.

An $m \times n$ matrix A has *numerical rank* $q \equiv \text{numrank}(A)$ *not exceeding* $\text{rank}(A) \equiv \rho(A) \equiv \rho$ *and has (right) numerical nullity* $r \equiv \text{numnull}(A) = n - q$, *if the ratios* $\sigma_j(A)/||A||$ *are small for* $j > q$ *but not for* $j \le q$.

If a rank deficient well conditioned $m \times n$ matrix A has a rank $\rho < l = \min\{m, n\}$, then almost all of its close neighbours have full rank l, and all of them have numerical rank $q = \rho$.

Alternatively, suppose an ill conditioned matrix A^* of full rank has a positive numerical rank $q < \rho \equiv l$, set to 0 all but the q largest singular values and denote the resulting matrix by $A = A^* - E$.

Then rank$(A) = q$, A is rank deficient and well conditioned, $||E|| = \sigma_{q+1}(A^*)$, and so A is a rank q approximation to the matrix A^*. We can obtain this approximation from the SVD of A^*.

1.3 Gaussian Random Matrices

1.3.1 Basic Results

Definition 1.3.1 *$F_\gamma(y) = Probability\{\gamma \le y\}$ for a real random variable γ is the cumulative distribution function (cdf) of γ evaluated at y. $F_{g(\mu,\sigma)}(y) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{y} \exp(-\frac{(x-\mu)^2}{2\sigma^2})dx$ for a Gaussian random variable $g(\mu, \sigma)$ with a mean μ and a positive variance σ^2, and so*

$$\mu - 4\sigma \le y \le \mu + 4\sigma \text{ with a probability near 1.} \tag{1.3.2}$$

Definition 1.3.3 *A matrix or a vector is a Gaussian or random matrix or vector with a mean μ and a positive variance σ^2 if it is filled with independent identically distributed Gaussian random variables, all having the mean μ and variance σ^2. $\mathcal{G}_{\mu,\sigma}^{m\times n}$ is the set of such Gaussian random $m \times n$ matrices (which are* standard *for $\mu = 0$ and $\sigma^2 = 1$).*

We do not use it, but provide the following definition for completion's sake:

Definition 1.3.4 *$\chi_{\mu,\sigma,n}(y)$ is the cdf of the norm $||\mathbf{v}|| = (\sum_{i=1}^{n} v_i^2)^{1/2}$ of a Gaussian random vector $\mathbf{v} = (v_i)_{i=1}^{n} \in \mathcal{G}_{\mu,\sigma}^{n\times 1}$. For $y \ge 0$ we have $\chi_{0,1,n}(y) =$*

$\frac{2}{2^{n/2}\Gamma(n/2)} \int_0^y x^{n-1} \exp(-x^2/2)dx$ *where* $\Gamma(h) = \int_0^\infty x^{h-1} \exp(-x)dx$, $\Gamma(n+1) = n!$ *for nonnegative integers* n.

1.3.2 Nondegeneration of Gaussian Random Matrices

The total degree of a multivariate monomial is the sum of the degrees in all its variables. The total degree of a polynomial is the maximal total degree of its monomials. The following Lemma is well established.

Lemma 1.3.5 *For a set* Δ *of a cardinality* $|\Delta|$ *in any fixed ring let a polynomial in* m *variables have a total degree* d *and let it not vanish identically on this set. Then the polynomial vanishes in at most* $d/|\Delta|^{m-1}$ *points.*

We assume that a Gaussian random variable ranges over an infinite set Δ, the real line. Then the lemma implies that a nonzero polynomial vanishes with probability 0. Consequently a square Gaussian random matrix is nonsingular with probability 1 because its determinant is a polynomial in the entries. Likewise rectangular Gaussian random matrices have full rank with probability 1. Furthermore all entries of such a matrix A and of its adjoint $\operatorname{adj} A$ are subdeterminants and thus are nonzero with probability 1. Clearly this property of the adjoint also holds for the inverse $A^{-1} = \frac{\operatorname{adj} A}{\det A}$ if the matrix A is nonsingular. Hereafter, wherever this causes no confusion, we assume by default that Gaussian random matrices have full rank, and their inverses (if defined) have nonzero entries.

Similar properties with probability near 1 hold where a continuous random variable is sampled under the uniform probability distribution or a discrete random variable, from a finite set of a large cardinality, is sampled under an "equally likely" distribution.

1.3.3 Conditioning of Gaussian Random Matrices

Besides having full rank with probability 1, Gaussian random matrices are expected to be well conditioned [**D88**], [**E88**], [**ES05**], [**CD05**], [**B11**], and even the sum $M + P$ for $M \in \mathbb{R}^{m\times n}$ and $P \in \mathcal{G}_{\mu,\sigma}^{m\times n}$ is expected to be well conditioned unless the ratio $\sigma/||M||$ is small or large [**SST06**]. This notion of expectation is further clarified in the sequel; it is not necessarily the classical notion of expectation.

The following theorem yields a probabilistic lower bound $(PLB) \equiv y$, also to be defined carefully in the sequel, to the reciprocal of the norm of the generalized inverse of a matrix, by providing a deterministic upper bound to the cdf of that reciprocal. Intuitively, the probability that the reciprocal is small is also small.

Theorem 1.3.6 *Suppose* $M \in \mathcal{G}_{\mu,\sigma}^{m\times n}$, $l = \min\{m, n\}$, *and* $y \geq 0$. *Then* $F_{\frac{1}{\|M^+\|}}(y) \leq 2.35\ \sqrt{l}y/\sigma$, *as we restate in a manner to be employed in the sequel,* $F_{\|M^+\|}(y) \geq 1 - 2.35\ \sqrt{l}/\sigma y, y > 0$.

Proof. For $m = n$ this is a consequence of [**SST06**, Theorem 3.3]. Apply the fact that the jth singular value of a submatrix of a matrix is at most the jth singular value of the matrix and extend it to any pair $\{m, n\}$. ■

For small values of y the latter deterministic lower bound of Theorem 1.3.6 becomes negative, in which case the theorem becomes trivial.

We assume that $\mu = 0$ when we estimate the cdfs of $||A||$ and consequently of $\kappa(A)$ for Gaussian random matrices A, but we make no such assumption when we estimate the cdf of $1/||A^+||$.

Theorem 1.3.7 *[**DS01**, Theorem II.7]. Suppose $M \in \mathcal{G}_{0,\sigma}^{m\times n}$, $h = \max\{m, n\}$ and $z \geq 2\sigma\sqrt{h}$. Then $F_{||M||}(z) \geq 1 - \exp(-(z - 2\sigma\sqrt{h})^2/(2\sigma^2))$.*

Theorem 1.3.8 *[**SST06**, Theorem 3.1]. Suppose $0 < \sigma \leq 1$, $y \geq 1$, $M \in \mathcal{G}_{0,\sigma}^{n\times n}$. Then the matrix M has full rank with probability 1 and $F_{\kappa(M)}(y) \geq 1 - (14.1 + 4.7\sqrt{(2\ln y)/n})n/(y\sigma)$.*

Chapter 2

PRECONDITIONING OF RANDOMIZED MATRIX PRODUCTS

In the following, we provide a probabilistic upper bound (PUB) on the norm of the generalized inverse of the product of a fixed and random matrix. This will allow us in Chapter 3 to provide a probabilistic upper bound to the condition number of an additive pre-processed matrix, then to be defined.

2.1 Lemmas

Consider the following three Lemmas (and an interesting corollary to Lemma 1). The first two lemmas are standard results of Applied Linear Algebra. The third lemma is a consequence of Proposition 2.2 in [**SST06**].

Lemma 2.1.1 *Let $M \in \mathbb{R}^{n\times n}$ and let $\Sigma = \mathrm{diag}(\sigma_i(M))_{i=1}^n$ represent the diagonal matrix of singular values of M within the Singular Value Decomposition of M. Let $G \in \mathbb{R}^{r\times n}$, and $H \in \mathbb{R}^{n\times r}$. [$(\sigma_i(X))$ continues to represent the ith singular value of any matrix X, here, and in the sequel]. Then: $\sigma_j(G\Sigma) \geq \sigma_j(G)\sigma_n(M)$, $\sigma_j(\Sigma H) \geq \sigma_j(H)\sigma_n(M)$ for all $j \in \{1, 2, \cdots, min\{r, n\}\}$. Moreover, if $\sigma_n > 0$ and $r < n$, then $\mathrm{rank}(G\Sigma) = \mathrm{rank}(G)$, $\mathrm{rank}(\Sigma H) = \mathrm{rank}(H)$.*

Corollary 2.1.2 *We have $\kappa(AB) \leq \kappa(A)\kappa(B)$ if A or B is a nonsingular matrix.*

Lemma 2.1.3 *Let $M \in \mathbb{R}^{m\times n}$. Let $S \in \mathbb{R}^{m\times m}$. Let $T \in \mathbb{R}^{n\times n}$. Let S and T be orthogonal matrices. Then $\sigma_j(SM) = \sigma_j(MT) = \sigma_j(M)$ for all $j \in \{1, 2, \cdots, min\{m, n\}\}$.*

Lemma 2.1.4 *[**SST06**, Proposition 2.2]. Suppose $H \in \mathcal{G}_{\mu,\sigma}^{m\times n}$. Let $S \in \mathbb{R}^{m\times m}$. Let $T \in \mathbb{R}^{n\times n}$; both S and T orthogonal. Then $SH \in \mathcal{G}_{\mu,\sigma}^{m\times n}$ and $HT \in \mathcal{G}_{\mu,\sigma}^{m\times n}$.*

2.2 Multiplicative Theorem

The following Theorem provides the claimed probabilistic upper bound (PUB), to be clearly defined in more general terms in Chapter 3, to the norm of the generalized inverse of the product of a fixed and random matrix. As in Chapter 1, define $(GM)^+$ to be the generalized inverse of the matrix product GM. Similarly, $(MH)^+$ is the generalized inverse of MH.

In the course of the following proof we will first consider the reciprocal of the norm of this generalized inverse. Given prior results in this mathematical area, it is natural to first compute a probabilistic lower bound (PLB) to this quantity. The (PLB) will also be addressed more formally in Chapter 3. This will then immediately provide the claimed (PUB) to the norm of the generalized inverse of the product matrix itself.

Theorem 2.2.1 *Suppose* $r \leq \min\{m, n\}$. $G \in \mathcal{G}_{\mu,\sigma}^{r\times m}$. $H \in \mathcal{G}_{\mu,\sigma}^{n\times r}$. $M \in \mathbb{R}^{m\times n}$. *Let* $\hat{r} = \min\{\rho(M), r\}$. *Then:*

$$F_{\|(MH)^+\|}(y) \geq 1 - \frac{2.35\sqrt{\hat{r}}}{\sigma_{\rho(M)}(M)\sigma y}$$

where y *is the claimed* PUB.

Proof. Since $H \in \mathcal{G}_{\mu,\sigma}^{n\times r}$, $n \geq r$, which implies that H has full rank $\rho(H) = r$ with Probability $= P = 1$. For any matrices A, B, $\rho(AB) \leq \min\{\rho(A), \rho(B)\}$. If this minimum is attained we define AB to be of maximal rank. Then MH has maximal rank $\hat{r} = \min\{\rho(M), r\}$ with $P = 1$. So:

$$\frac{1}{\|(MH)^+\|} = \sigma_{\hat{r}}(MH). \tag{2.2.2}$$

Let $M = S_M \Sigma_M T_M^T$ be the Singular Value Decomposition of M, where $\Sigma_M = \text{diag}(\hat{\Sigma}_M, 0)$, $\hat{\Sigma}_M = \text{diag}(\sigma_i(M))_{i=1}^{\rho(M)}$, and 0 is the zero matrix of order $(m - \rho(M)) \times (n - \rho(M))$. Then $MH = S_M \Sigma_M T_M^T H$. Note that $\sigma_j(S_M \Sigma_M T_M^T H) = \sigma_j(\Sigma_M T_M^T H)$ by virtue of Lemma 2.1.3 since S is orthogonal. So

$$\sigma_j(MH) = \sigma_j(\Sigma_M T_M^T H), \text{ for all } j \in \{1, 2, \cdots, \hat{r}\}. \tag{2.2.3}$$

Write

$$H_{\rho(M)} = (I \mid 0) T_M^T H, \tag{2.2.4}$$

where I is the $\rho(M) \times \rho(M)$ identity matrix and 0 is the $(\rho(M)) \times (n - \rho(M))$ zero matrix. In other words, this $\rho(M) \times r$ matrix $H_{\rho(M)}$ consists of the top $\rho(M)$ rows of $T_M^T H$. Now consider the matrix $\hat{\Sigma}_M H_{\rho(M)}$. This is also a $\rho(M) \times r$ matrix whose ith row is the ith row of $H_{\rho(M)}$ multiplied by $\sigma_i(M)$. Compare this with the matrix $\Sigma_M T_M^T H$, which is an $m \times r$ matrix whose top $\rho(M)$ rows are identical, respectively, to the $\rho(M)$ rows of $\hat{\Sigma}_M H_{\rho(M)}$ and whose bottom $(m - \rho(M))$ rows are zero, and obtain:

$$\sigma_j(\Sigma_M T_M^T H) = \sigma_j(\hat{\Sigma}_M H_{\rho(M)}), \text{ for all } j \in \{1, 2, \cdots, \hat{r}\}. \tag{2.2.5}$$

So (2.2.3) and (2.2.5) imply

$$\sigma_j(MH) = \sigma_j(\hat{\Sigma}_M H_{\rho(M)}), \text{ for all } j \in \{1, 2, \cdots, \hat{r}\}. \tag{2.2.6}$$

In Lemma 2.1.1 we replace (Σ, H) with$(\hat{\Sigma}_M, H_{\rho(M)})$. n in Lemma 2.1.1 is $\rho(M)$ here. r in Lemma 2.1.1 is also r here. Let a specific j in Lemma 2.1.1 be $\hat{r}$ here. Then:

$$\sigma_{\hat{r}}(\hat{\Sigma}_M H_{\rho(M)}) \geq \sigma_{\hat{r}}(H_{\rho(M)})\sigma_{\rho(M)}(M). \tag{2.2.7}$$

So (2.2.6) and (2.2.7) imply

$$\sigma_{\hat{r}}(MH) \geq \sigma_{\hat{r}}(H_{\rho(M)})\sigma_{\rho(M)}(M). \tag{2.2.8}$$

Further note: $T_M^T H$, the right factor of $H_{\rho(M)}$, is an element of $\mathcal{G}_{\mu,\sigma}^{n\times r}$ because T_M and its transpose are orthogonal, $H \in \mathcal{G}_{\mu,\sigma}^{n\times r}$ by hypothesis and (2.1.4) holds. Left multiplication of $T_M^T H$ by $(I \mid 0)$, of order $\rho(M) \times n$, then generates $H_{\rho(M)}$ as a Gaussian random matrix of order $\rho(M) \times r$, that is, $H_{\rho(M)} \in \mathcal{G}_{\mu,\sigma}^{\rho(M)\times r}$. Theorem 1.3.6 from Chapter 1 applied to $H_{\rho(M)}$ yields:

$$F_{\frac{1}{\|(H_{\rho(M)})^+\|}}(x) \leq \frac{2.35\sqrt{\hat{r}}x}{\sigma} \tag{2.2.9}$$

or

$$F_{\sigma_{\hat{r}}(H_{\rho(M)})}(x) \leq \frac{2.35\sqrt{\hat{r}}x}{\sigma}. \tag{2.2.10}$$

We also rewrite (2.2.8) as

$$\frac{\sigma_{\hat{r}}(MH)}{\sigma_{\rho(M)}(M)} \geq \sigma_{\hat{r}}(H_{\rho(M)}). \tag{2.2.11}$$

If $X_1 \geq X_2$, where X_1, X_2 are Gaussian random variables, we take as a tautology that this implies that $F_{X_1} \leq F_{X_2}$, where F_{X_i} is the cumulative distribution function of X_i. Then

$$F_{\frac{\sigma_{\hat{r}}(MH)}{\sigma_{\rho(M)}(M)}}(x) \leq F_{\sigma_{\hat{r}}}(H_{\rho(M)})(x) \tag{2.2.12}$$

Then (2.2.10) and (2.2.12) imply

$$F_{\sigma_{\hat{r}}(MH)}(x\sigma_{\rho(M)}(M)) \leq F_{\sigma_{\hat{r}}(H_{\rho(M)})}(x) \leq \frac{2.35\sqrt{\hat{r}}x}{\sigma} \tag{2.2.13}$$

or upon replacing x with $\frac{x}{\sigma_{\rho(M)}(M)}$, we have,

$$F_{\sigma_{\hat{r}}(MH)}(x) \leq \frac{2.35\sqrt{\hat{r}}x}{\sigma_{\rho(M)}(M)\sigma} \tag{2.2.14}$$

x is the claimed PLB to $\sigma_{\hat{r}}$, which is the reciprocal of $\|(MH)^+\|$. But for any cdf F_X, and any non-negative random variable X, any probabilistic lower bound $PLB = x$ and deterministic upper bound f of F_X , to say: $F_X(x) \leq f(x)$, upon letting $y = \frac{1}{x}$, is equivalent to $F_{\frac{1}{X}}(y) \geq 1 - f(\frac{1}{y})$. So (2.2.14) is equivalent to

$$F_{\frac{1}{\sigma_{\hat{r}}(MH)}}(y) \geq 1 - \frac{2.35\sqrt{\hat{r}}}{\sigma_{\rho(M)}(M)\sigma y} \tag{2.2.15}$$

or

$$F_{\|(MH)^+\|}(y) \geq 1 - \frac{2.35\sqrt{\hat{r}}}{\sigma_{\rho(M)}(M)\sigma y}. \quad \blacksquare \tag{2.2.16}$$

Corollary 2.2.17 *If* $y = \|M^+\|$, *then* $F_{\|(MH)^+\|}(\|M^+\|) \geq 1 - \frac{2.35\sqrt{\hat{r}}}{\sigma}$.

Proof. It follows from: $\|M^+\| = \frac{1}{\sigma_{\rho(M)}(M)}$. ■

Remark: A similar result holds for $F_{\|(GM)^+\|}(y)$.

2.3 Bounds: Introduction

The concepts PLB and PUB are somewhat misleading. For example, regarding the PUB, this concept actually refers to an uncountable infinity of probabilistic upper bounds y which are arguments of the cdf of some random variable X. In Theorem 2.2.1, $X = \|(MH)^+\|$. The fact that the right-hand side of the inequality in (2.2.16) serves as a deterministic lower bound to the cdf gives intuitive plausibility to applying the phrase PUB to any y which yields a positive value to the right-hand side of the inequality. E.g., in (2.2.16), $y \geq \frac{2.35\sqrt{\hat{r}}}{\sigma_{\rho(M)}(M)\sigma}$.

Chapter 3

RANDOMIZED ADDITIVE PRECONDITIONING

In this Chapter, unlike Chapter 2, in which we allowed an $m \times n$ matrix M, we presume an $n \times n$ matrix. In Chapter 4, however, we suggest a generalization to the $m \times n$ case. Most of the key theorems here also demand that an $n \times n$ matrix A is singular and well conditioned. That said, one motivation for us begins with the assumption that an $n \times n$ matrix A^* exists, which is both nonsingular and ill conditioned of numerical rank $\rho num(A^*) < \rho(A^*) = n$. The matrix A is then produced by zeroing out the smallest singular values of A^*, a set of cardinality $n - \rho num(A^*) \equiv numnull(A^*)$.

An equivalent way of viewing the transition from A^* to A is by subtracting a perturbation matrix E from A^* to produce $A \equiv A^* - E$, in which the norm of E is $\sigma_{\rho num(A^*)+1}$, presumed small. We will precondition the matrix A in the additive sense to produce a matrix $C \equiv A+UV^T$ which is non-singular with probability = 1 and likely to be well conditioned. The preconditioning summand P will be defined precisely below. We show, in fact, that $\kappa(C)$ and $\kappa(A)$ are of the same probabilistic order, or $\kappa(C) \overset{po}{\approx} \kappa(A)$, also defined below. Moreover, if we define $C^* \equiv C + E$, this implies $\kappa(C^*) \approx \kappa(C)$. And $\kappa(A) \ll \kappa(A^*)$, by assumption. Thus we define $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$. This implies that computations with C^* in lieu of A^* are expected to produce smaller output error for a given input error. In fact, suppose we are given a nonsingular ill conditioned linear system $A^*x = b$. The following sequence : $A^* \overset{1}{\rightarrow} A \overset{2}{\rightarrow} C \overset{3}{\rightarrow} C^* \overset{4}{\rightarrow} C^{*-1} \overset{5}{\rightarrow} A^{*-1}$ should provide a more computationally efficient way of generating the system's solution, since $\kappa(A^{*-1}) = \kappa(A^*) \gg \kappa(A) \overset{po}{\approx} \kappa(C) \approx \kappa(C^*) = \kappa(C^{*-1})$, or what we shorten to $\kappa(C^{*-1}) \overset{po}{\ll} \kappa(A^{*-1})$. In implementation, an algorithm should take A^* directly to C^*, thus replacing steps 1, 2 and 3 with one step; however, the proof that C^* in fact serves as a better conditioned proxy of A^* is best understood via the sequence outlined above, and is what we focus on in Chapter 3.

We do not dwell in this thesis on the savings in computational cost versus standard algorithms, other than to note that preconditioning allows us to limit highly accurate computations to $O(n^2)$ flops versus Gaussian elimination, which requires highly accurate computations of order $O(n^3)$ flops.

The first and third transitions (arrows) are discussed above, while transition 2 is the focus of our first key theorem in Chapter 3. Transition 4 may be accomplished using standard techniques and transition 5 is achieved via the SMW formula, which was recalled in Chapter 1.

We will initially prove three basic theorems in this Chapter. Taken together, they will first show that the smallest positive singular value of A, $\sigma_{\rho(A)}(A)$ and the smallest positive singular value of C, $\sigma_n(C)$ are of the same probabilistic order or $\sigma_{\rho(A)}(A) \overset{po}{\approx} \sigma_n(C)$. With the additional fact that the norm of A and the norm of C are also of the same probabilistic order, we will have the result that $\kappa(A)$ and $\kappa(C)$ are of the same probabilistic order. Then the above argument will show that $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$, which is the intuitive content behind the Theorem in the Preface [and restated in Chapter 1], our Key Result; the Theorem will be proven in Chapter 4. In the next subsection, we define precisely what it means for two random variables, each defined on $\mathbb{R}^p$ for some $p \in \mathbb{N}$ to be of the same probabilistic order.

3.1 Probabilistic Order

Definition 3.1.1 *Suppose Z is a real random variable with domain $D \subseteq \mathbb{R}^p$ for some $p \in \mathbb{N}$ and with associated cdf F_Z. (For example, if Z is the $||M||_2$ of $M \in \mathbb{R}^{r\times r}$, then $p = r^2$). Suppose there exists a real valued function f with domain $\tilde{D} \subseteq \mathbb{R}^q$ for some $q \in \mathbb{N}$, a q - tuple of real numbers. Suppose the first component of f is z, which is an element of the range of Z. Further suppose that each of the remaining $q-1$ components of the domain of f is an independent variable which either represents a parameter associated with the dimension of the underlying measure space of which D is a subset, $\mathbb{R}^p$, (for example r when $p = r^2$), or with a parameter associated with the probability density function of a component of $\mathbb{R}^p$,(for example σ, the standard deviation). In the latter case, we identify the component of $\mathbb{R}^p$ itself as a real-valued random variable with domain $\mathbb{R}$, and range in $(0, 1)$. Suppose $F_Z(z) \geq f(z)$. The dependence of f on the $q-1$ other components is suppressed. We say that z is a Probabilistic Upper Bound (PUB) for Z. Suppose $F_Z(z) \leq f(z)$. The dependence of f on the $q-1$ other components is suppressed. We say that z is a Probabilistic Lower Bound (PLB) for Z.*

In other words, deterministic lower and upper bounds f for the cumulative distribution function of a random variable imply respectively probabilistic upper and lower bounds z for the random variable itself.

Assume the following collection of real non-negative random variables; in particular, assume Y is positive. X, Y, $W_i, i \in 1, 2, \cdots, k_1$, $Z_i, i \in 1, 2, \cdots, k_2$; $k_1, k_2 \in \mathbb{N}$.

Further assume that all above random variables W_i are functions of the components of Gaussian random matrices of order $\tilde{m_i} \times \tilde{n_i}$, $i \in \{1, 2, \cdots, k_1\}$.

Similarly all above random variables Z_i are functions of the components of Gaussian random matrices of order $m_i \times n_i$, $i \in \{1, 2, \cdots, k_2\}$.

Suppose $W = \Pi_{i=1}^{k_1} W_i$ and $Z = \Pi_{i=1}^{k_2} Z_i$. Moreover assume that at least one of the following conditions hold:

$W \leq \frac{X}{Y}$ or $\frac{X}{Y} \leq Z$.

Definition 3.1.2 *We define X and Y to be of the same probabilistic order, $X \overset{po}{\approx} Y$ if the above conditions hold, and if in addition there exist either real valued functions e_i, $i \in \{1, 2, \cdots, k_1\}$, in which e_i are respectively upper bounds to the cdfs of each W_i, thus providing $PLBs$ to each W_i or there exist real valued functions f_i, $i \in \{1, 2, \cdots, k_2\}$, and f_i are respectively lower bounds to the cdfs of each Z_i, thus providing $PUBs$ to each Z_i.*

The domains of the f_i are more explicitly addressed in a forthcoming remark.

Definition 3.1.3 *We define for random variables X^*, Y^*, $X^* \ll Y^*$ up to probabilistic order or, more succintly, $X^* \overset{po}{\ll} Y^*$, if there exist X, Y, random variables, such that $X^* \approx X$, $X \overset{po}{\approx} Y$, $Y \ll Y^*$.*

Remark: In the sequel, $X = \kappa(C)$, $X^* = \kappa(C^*)$, $Y = \kappa(A)$, and $Y^* = \kappa(A^*)$. Also in the sequel we address the Zs rather than the Ws since we are primarily interested in bounding from above the ratio $\frac{X}{Y} = \frac{\kappa(C)}{\kappa(A)}$.

Remark: Suppose there exist T_i parameters t_{j_i}, $j_i \in \{1, 2, \cdots, T_i\}$ associated with some (or all) of the p pdfs of the components of $\mathbb{R}^p$, in which p represents the dimension of the domain of Z_i, say $p = m_i \times n_i$, where Z_i is one of the Z random variables cited above. If there is a $PUB = z$ for Z_i, this means that there exists a function f_i such that $F_{Z_i}(z) \geq f_i(z, m_i, n_i, t_{j_i})$, for all $j_i \in \{1, 2, \cdots, T_i\}, i \in \{1, 2, \cdots, k_2\}$. In this example, $q = 1 + 1 + 1 + T_i$; if $T_i = 1$, then, a priori, $q = 1 + 1 + 1 + 1 = 4$; if the dependence on m_i, for example, is not explicit, then $q = 3$.

We show below that $\kappa(C) \overset{po}{\approx} \kappa(A)$, which will then immediately imply $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$. For example, we will provide a PUB for the norm of the inverse of a particular matrix, $||U_r^{-1}||$, by again associating a deterministic lower bound to the cdf of $||U_r^{-1}||$ which will represent one of the factors Z_i, which plays a central role in estimating the ratio of $\sigma_{\rho(A)}(A)$ and $\sigma_n(C)$, which is proportional to $\frac{\kappa(C)}{\kappa(A)}$, the ratio of primary interest. In this case, it will be seen that $p = n \times r$, although the dependence on n is not explicit in the argument of the deterministic lower bound. $m_i = n$; $n_i = r$; $j_i = T_i = 1$ and $t_1 = \sigma$. There will in fact be six Z factors in this example, although only four distinct factors; two of the four distinct factors will be shown to have multiplicity two. Thus $k_2 = 6$ and we will label $||U_r^{-1}||$ as Z_5; thus i above may be taken as $i = 5$. We are here ignoring a technicality; in fact Z_5 will actually be shown to be $\max\{1, ||U_r^{-1}||\}$. Again, the dependence of the lower bound f_5 of the cdf of $||U_r^{-1}||$ on $m_5 = n$ vanishes from the argument of f_5.

It will be seen that $F_{Z_5}(z) \geq f_5(z, n_5, t_{j_5})$; that is, $F_{||U_r^{-1}||}(z) \geq f_5(z, r, \sigma) = 1 - \frac{2.35\sqrt{r}}{\sigma z}$.

Remark: It will again be emphasized that we will focus on the PUB of $\frac{\kappa(C)}{\kappa(A)}$ rather than the PLB. We are interested in controlling the growth of the condition number of C relative to the condition number of A. Other simplifications of the above analysis applied to our forthcoming example: $\frac{\kappa(C)}{\kappa(A)}$ is proportional to $\frac{\sigma_{\rho(A)}(A)}{\sigma_n(C)}$. The proportionality constant is $\frac{||C||}{||A||}$ and we may assume $||A|| = 1$, through scaling. Also $\sigma_{\rho(A)}(A)$ is a constant, since A is given as fixed. When we move from first examining the PUB of $\frac{\sigma_{\rho(A)}(A)}{\sigma_n(C)}$ which requires $k_2 = 4$, to studying the related PUB of $\frac{\kappa(C)}{\kappa(A)}$, we will initially require one extra Z factor; then to avoid issues surrounding the cdf of a product of random variables, will require one additional Z factor for a final $k_2 = 6$.

3.2 Details of Preconditioning

In this section we provide the random variables factors Z_i, $i \in \{1, 2, \cdots, k_2 = 6\}$ whose product provides an upper bound on the ratio $\frac{\kappa(C)}{\kappa(A)}$. We next deduce f_i, $i \in \{1, 2, \cdots, k_2 = 6\}$ which act as deterministic lower bounds to the the cdfs of the respective Z_is, thus implying, according to the prior section, PUBs to the respective Z_is.

We initially deduce that $\sigma_{\rho_A}(A)$ and $\sigma_n(C)$ are of the same probabilistic order. We further conclude that $\kappa(A)$ and $\kappa(C)$ are of the same probabilistic order.

In the following theorem, we may assume that the well conditioned singular matrix A has been obtained from the ill conditioned nonsingular matrix A^* via the methodology described in Chapter 1.

3.2.1 Proof of Theorem

Theorem 3.2.1 *Suppose $A, C, S, T \in \mathbb{R}^{n\times n}$ and $U, V \in \mathcal{G}_{\mu,\sigma}^{n\times r}$ for two positive integers r and n, $r = n - \rho(A)$, $0 < \rho(A) < n$. Define $\rho(A) \equiv \rho$. Define $\sigma_j(A) \equiv \sigma_j$. $A = S\Sigma T^T$ is the SVD of the matrix A. So S and T are orthogonal matrices, $\Sigma = \mathrm{diag}(\sigma_j)_{j=1}^n$, $\sigma_j = 0$ for all $j \in \rho+1, \cdots, n$. Let $C \equiv A + UV^T$. Implicitly define the matrices $\bar{U}, \bar{V}, U_r, V_r$ by writing:*

$$S^T U = \begin{pmatrix} \bar{U} \\ U_r \end{pmatrix},\ T^T V = \begin{pmatrix} \bar{V} \\ V_r \end{pmatrix},\ R_U = \begin{pmatrix} I_\rho & \bar{U} \\ O_{r,\rho} & U_r \end{pmatrix},\ R_V = \begin{pmatrix} I_\rho & \bar{V} \\ O_{r,\rho} & V_r \end{pmatrix}, \tag{3.2.2}$$

where U_r and V_r are $r \times r$ matrices. [$\bar{U}$ and $\bar{V}$, both of order $\rho \times r$, are in some sense residual matrices which play only an implicit role here]. Define $D \equiv \Sigma + \mathrm{diag}(O_{\rho,\rho}, I_r) = \mathrm{diag}(d_j)_{j=1}^n$ where $d_j = \sigma_j$ for $j = 1, \ldots, \rho$, $d_j = 1$ for $j = \rho+1, \ldots, n$. Then:

(a) *$R_U \Sigma R_V^T = \Sigma$ and $R_U\,\mathrm{diag}(O_{\rho,\rho}, I_r) R_V^T = S^T U V^T T$, and so*

$$C = S R_U D R_V^T T^T \tag{3.2.3}$$

Assume that $||A|| = 1$ *and* U_r *and* V_r *are nonsingular, which holds with probability* 1 *from Chapter 2. Then*

(b) *the matrix* C *is nonsingular with probability* 1,

(c) $||C|| \le 1 + ||U||\ ||V||$,

(d) $\sigma_\rho(A)/\sigma_n(C) \le ||R_U^{-1}||\ ||R_V^{-1}||$,

(e) $||R_U^{-1}||\ ||R_V^{-1}|| \le (1+||U||)(1+||V||)\max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}$,

and thus conclude

(f) $\sigma_\rho(A)/\sigma_n(C) \le (1 + ||U||)(1 + ||V||)\max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}$.

Moreover,

(g) $\frac{\kappa(C)}{\kappa(A)} \le (1 + ||U||)^2(1 + ||V||)^2\ \max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}$.

Proof. Part (a) is readily verified; part (b) is well established.

(c) Combine the relationships $||C|| \le ||A|| + ||U||\ ||V^T||$, $||A|| = 1$ and $||V^T|| = ||V||$.

(d) $D = \mathrm{diag}(\hat{\Sigma}_A, I_r)$; so D is nonsingular. Utilize $S^{-1} = S^T$, $T^{-1} = T$ and obtain $C^{-1} = TR_V^{-T}D^{-1}R_U^{-1}S^T$. Thus $||C^{-1}|| \le ||R_V^{-T}||\ ||D^{-1}||\ ||R_U^{-1}||$ because S and$T$ are square orthogonal matrices. So $\frac{||C^{-1}||}{||D^{-1}||} \le ||R_U^{-1}||\ ||R_V^{-1}||$. But $D^{-1} = \mathrm{diag}(\hat{\Sigma}_A^{-1}, I_r)$ and recall that$||A|| = 1$. So $||D^{-1}|| = \frac{1}{\sigma_{\rho_{(A)}}(A)} = ||A^+||$. And $||C^{-1}|| = 1/\sigma_n(C)$, and thus (d) is proven.

(e) Observe that $R_U^{-1} = \begin{pmatrix} I_\rho & -\bar{U} \\ O & I_r \end{pmatrix}\begin{pmatrix} I_\rho & O \\ O & U_r^{-1} \end{pmatrix}$, $R_V^{-1} = \begin{pmatrix} I_\rho & -\bar{V} \\ O & I_r \end{pmatrix}\begin{pmatrix} I_\rho & O \\ O & V_r^{-1} \end{pmatrix}$, $||\bar{U}|| \le ||U||$ and $||\bar{V}|| \le ||V||$. Then combine these relationships.

(f) Trivially combine the bounds of parts (d) and (e).

(g) We have $\frac{\kappa(C)}{\kappa(A)} \equiv \frac{||C||}{||A||}\frac{\sigma_\rho(A)}{\sigma_n(C)}$, and so parts (c) and (f) bound the ratio $\frac{\kappa(C)}{\kappa(A)}$ in terms of the norms $||U||$, $||V||$, $||U_r^{-1}||$ and $||V_r^{-1}||$ as follows:

$\frac{\kappa(C)}{\kappa(A)} \le (1+||U||\,||V||)(1+||U||)(1+||V||)\ \max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}$.

Noting that $(1 + ||U||\ ||V||) \le (1 + ||U||)(1 + ||V||)$, we have: $\frac{\kappa(C)}{\kappa(A)} \le (1 + ||U||)^2(1 + ||V||)^2\ \max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}$. ∎

In particular, if $U = V$, $\frac{\kappa(C)}{\kappa(A)} \le (1 + ||U||^2)(1 + ||U||)^2 \max\{1, ||U_r^{-1}||^2\} \le (1+||U||)^4\ \max\{1, ||U_r^{-1}||\}^2$. So, in general, define $Z_1 = Z_2 = 1+||U||$, $Z_3 = Z_4 = 1 + ||V||$, $Z_5 = ||U_r^{-1}||$, and $Z_6 = ||V_r^{-1}||$. Z_5 and Z_6 are the interesting cases in which $max\{1, ||U_r^{-1}||\} = ||U_r^{-1}||$ and $max\{1, ||V_r^{-1}||\} = ||V_r^{-1}||$.

Let us estimate the norms $||U||$, $||V||$, $||U_r^{-1}||$ and $||V_r^{-1}||$.

Theorem 3.2.4 *Suppose $W \in \mathcal{G}_{0,\sigma}^{n\times r}$. E.g., let $W = U$ or V from Theorem 3.2.1. Let $y \ge 2\sigma\sqrt{n}$. Then $F_{||W||}(y) \ge 1 - \exp(-(y - 2\sigma\sqrt{n})^2/(2\sigma^2))$.*

Proof. The theorem follows from Theorem 1.3.7 of Chapter 1 applied for $A = W$ (in which case $h = n$). Namely, let m and n in the cited Theorem be replaced respectively with n and r here. ∎

We have thus provided lower bounds to the cdfs of $||U||$ and $||V||$ and thus to the respective cdfs of $1 + ||U||$ and $1 + ||V||$, i.e., to the cdfs of $Z_i, i \in 1, 2, \cdots, 4$. That is, in Theorem 1.3.7, simply replace y with $y-1$ and n with r in $1-\exp(-(y-2\sigma\sqrt{n})^2/(2\sigma^2))$ to obtain $1-\exp(-(y-1-2\sigma\sqrt{r})^2/(2\sigma^2))$. The latter expression provides the required f_i, $i \in \{1, 2, \cdots, 4\}$, which act as lower bounds to the respective Z_i.

Theorem 3.2.5 *Suppose that A, U, V, U_r and V_r denote the five matrices of Theorem 3.2.1 so $U, V \in \mathcal{G}_{\mu,\sigma}^{n\times r}$. Then* $\max\{F_{1/||U_r^{-1}||}(y), F_{1/||V_r^{-1}||}(y)\} \leq 2.35\ y\sqrt{r}/\sigma$ *for $y \geq 0$. Or, to conform more precisely with Theorem 1.3.6: $F_{||U_r^{-1}||}(y) \geq 1 - 2.35\sqrt{r}/\sigma y >$ for $y > 2.35\sqrt{r}/\sigma$.*

Proof. Lemma 2.1.4 implies that $S^T U, T^T V \in \mathcal{G}_{\mu,\sigma}^{n\times r}$ because S and T are orthogonal matrices. Therefore $U_r, V_r \in \mathcal{G}_{\mu,\sigma}^{r\times r}$. Apply Theorem 1.3.6 for $M = U_r$ and $M = V_r$ where in both cases $m = n = r$. ∎

Remark: Here is an alternative, slightly lengthier, but perhaps more instructive proof of Theorem 3.2.5 which connects Theorem 2.2.1, focused on multiplicative preprocessing, with the additive preprocessing which is here of paramount interest: Let M in Theorem 2.2.1 equal the bottom r rows of S^T here, so m of Theorem 2.2.1 turns into r here. Let H in Theorem 2.2.1 equal U here, so n of Theorem 2.2.1 remains n here also. Then $MH = U_r$. But $\rho(M) = r$, since S^T is of full rank n, so its bottom r rows form a matrix of full rank r. Morever, $\rho(H) = r$, since U is of full rank r. So $\hat{r}$ in Chapter 2 equals r here. Furthermore $\sigma_{\rho(M)}(M) = 1$. So, $F_{||U_r^{-1}||}(y) \geq 1 - 2.35\sqrt{r}/\sigma y$.

Chapter 4

CONCLUSION

In this chapter we slightly extend the analysis in Chapter 3, first by discussing the consequences of the Chapter 3 theorems with respect to the basic Theorem, our key result, stated in the Preface, in Chapter 1, and repeated below. This discussion entails some philosophical intuition, in addition to mathematics. Next we mention an extension of Theorem 3.2.1 to the case of rectangular matrices. Finally we briefly discuss an extension of Theorem 3.2.1 to the case of complex matrices.

4.1 Proof of Key Result

The following is a [final] restatement and proof of our Key Result.

Theorem 4.1.1 *Suppose A^* is a nonsingular ill conditioned matrix with real components. Suppose $C^* \equiv A^* + UV^T$. Suppose further that $\|A^*\|_2 \approx ||UV^T||_2$, where $U, V \in \mathcal{G}_{0,\sigma}^{n\times r}$, Gaussian random matrices with mean zero and standard deviation σ of order $n \times r$. Let r be the numerical nullity of A^*, denoted numnull(A^*). Then with Probability $= 1$, C^* is nonsingular and $\kappa(C^*) \ll \kappa(A^*)$ up to probabilistic order, or $\kappa(C^*) \overset{po}{\ll} \kappa(A^*)$.*

Proof.

Definitions 3.1.1, 3.1.2, 3.1.3, the three Remarks following these Definitions, and Theorems 3.2.1, 3.2.4, 3.2.5 prove this Key Result. At first we recall Theorem 3.2.1, which provides the requisite product which acts as an upper bound to $\frac{\kappa(C)}{\kappa(A)}$.

Now, again, consider the result of Theorem 3.2.4. By replacing y with $y - 1$ in the right-hand side of the inequality, we obtain a lower bound to the cdf of the random variable $1 + ||W||$. So, $F_{1+||W||}(y) \geq 1 - \exp(-(y - 1 - 2\sigma\sqrt{n})^2/(2\sigma^2))$, with associated new bounds on y. Upon replacing W with U, then V, from Chapter 3, we then have two of the four factors whose product acts as an upper bound to $\frac{\sigma_{\rho(A)}(A)}{\sigma_n(C)}$, or these same two, not counting multiplicities, of the six factors whose product acts as an upper bound to $\frac{\kappa(A)}{\kappa(C)}$, as further detailed below.

Counting multiplicities we have four of the six factors whose product acts as an upper bound to $\frac{\kappa(A)}{\kappa(C)}$. Consistent with the notation in Chapter 3, we may label these factors $1 + ||U||$ as $Z_1 \equiv Z_2$ and $1 + ||V||$ as $Z_3 \equiv Z_4$. The lower bounds to the cdfs of Z_1 thru Z_4, respectively, are what we have defined as f_1 thru f_4, and each of these is precisely $1 - \exp(-(y-1-2\sigma\sqrt{n})^2/(2\sigma^2))$.

Similarly, if $||U_r^{-1}|| > 1$ and $||V_r^{-1}|| > 1$, which is the interesting case, then the right-hand side of the inequality in Theorem 3.2.5, $1 - 2.35\sqrt{r}/\sigma y$ provides a lower bound to the cdfs of both $||U_r^{-1}||$ and $||V_r^{-1}||$. Then we have the remaining two of the four factors whose product acts as an upper bound to $\frac{\sigma_\rho(A)}{\sigma_n(C)}$, namely $||U_r^{-1}||$ and $||V_r^{-1}||$, or these same two of the remaining six factors whose product acts as an upper bound to $\frac{\kappa(A)}{\kappa(C)}$. Consistent with the notation in Chapter 3, we may label these factors $||U_r^{-1}||$ as Z_5 and $||V_r^{-1}||$ as Z_6. Similarly, we may label $f_5 \equiv f_6$ as $1 - 2.35\sqrt{r}/\sigma y$.

Finally Definitions 3.1.1, 3.1.2, 3.1.3, and the associated Remarks imply the Key Result. ■

Remark: Note the intuitive content underlying the right-hand side of the inequalities in Theorems 3.2.4 and 3.2.5; in particular, consider Theorem 3.2.5. $1 - 2.35\sqrt{r}/\sigma y$ is that right-hand side, which has been labeled as $f_5 = f_6$, which we here define simply as f. f is a function of (y, σ, r); the following are partial derivatives. Note that $f'(r) < 0, f'(\sigma) > 0, f'(y) > 0$. $f'(r) = -2.35/2\sigma y\sqrt{r}$, $f'(\sigma) = 2.35\sqrt{r}/\sigma^2 y$, $f'(y) = 2.35\sqrt{r}/\sigma y^2$.

The negative derivative in r is perhaps counterintuitive, since it suggests that fewer random parameters provides a stronger (larger) lower bound to the cdf of $||U_r^{-1}||$ or $||V_r^{-1}||$. Although, in as much as r here ultimately represents the numerical nullity of A^*, this suggests that the better behaved is A^*, that is, the smaller its numerical nullity r, the stronger (larger) is the lower bound to the cdf of $||U_r^{-1}||$ or $||V_r^{-1}||$. That said, it is of course also true that Theorem 3.2.1 has no memory that r represents the numerical nullity of A^* and r may also simply be viewed as the nullity of A. The positive derivative in σ seems natural, suggesting more randomness in each component of U or V gives a stronger (larger) lower bound to the cdf of $||U_r^{-1}||$ or $||V_r^{-1}||$. Finally, the positive derivative in y says there is no free lunch. If one decreases the argument y of the cdf of $||U_r^{-1}||$ or $||V_r^{-1}||$, this gives a weaker (smaller) lower bound to the cdf of $||U_r^{-1}||$ or $||V_r^{-1}||$.

Remark: A random variable that acts as an upper bound to $||C||$ with a cdf with a well defined lower bound, would then generate a PUB for $||C||$. This situation is slightly more problematic. [C is not Gaussian, in fact even UV^T is not Gaussian, thus Theorems cited in [**CD05**], [**E88**]and [**ES05**] are not applicable for directly providing a deterministic lower bound to the cdf of $\frac{\kappa(C)}{\kappa(A)}$]. However, from Theorem 3.2.1 (c), it is true that $||C|| \leq 1 + ||UV|| \leq 1 + ||U||\ ||V|| \leq [[1 + ||U||] \times [1 + ||V||]]$. So, the right hand side of the inequality serves as an upper bound to $||C||$, and each factor has already been shown to have a cdf bounded below by a known function.

The above comment serves to display the multiplicity 2 associated with both Z_1 and Z_3.

4.2 Sufficiency of Our Approach

4.2.1 Why Choose this Route?

A priori, there are at least several other possible ways to approach the problem of ascertaining the growth of $\kappa(C)$ relative to $\kappa(A)$ which, in principle, would be preferable to the approach taken in this monograph. Unfortunately, each of these alternatives entails either exceedingly difficult or even impossible solutions. We list these alternatives in decreasing order of preference.

- The best possible solution would be to know explicitly the cumulative distribution function of $\frac{\kappa(C)}{\kappa(A)}$. However, if $n < 4$, the problem is difficult, to say the least, since this entails computing the cdf of the ratio of $\sigma_1(C)$ to $\sigma_n(C)$, each singular value being the square root of an eigenvalue of CC^T. If $n \geq 5$, this is generally impossible, due to the insolvability of the quintic.

- The next best alternative would be to at least bound from below the cdf of $\frac{\kappa(C)}{\kappa(A)}$ by a known function f. Unfortunately, the key results in [**CD05**], [**E88**]and [**ES05**] are not applicable here since neither UV^T nor, more important, C, are Gaussian. If $\mu = 0$ and $\sigma = 1$, then UV^T is χ square with r degrees of freedom; if $\sigma \neq 1$, then instead we achieve a type of scaled χ square. Moreover, the addition of A to UV^T shifts $\mu = 0$ to a collection of $\mu_{i,j}$. In any event, no theorem is known to this author which would readily compute an appropriate f.

- Another possibility would be to bound from above $\frac{\kappa(C)}{\kappa(A)}$ by a random variable Z, which we have done in Chapter 3, where $Z = \Pi_{i=1}^{k_2} Z_i$, and the cdf of Z is known. In our situation this would entail computing the cdf of a product of random variables in which each random variable factor Z_i itself has a difficulty to ascertain pdf and cdf. This paper presents the next best thing:

- We bound from above $\frac{\kappa(C)}{\kappa(A)}$ by $Z = \Pi_{i=1}^{k_2} Z_i$, and produce deterministic lower bounds f_i to the cdf of each Z_i.

- Simulations performed in [**PQZ11**] and [**PQZ14**] suggest that, on average,
$\frac{\kappa(C^*)}{\kappa(A^*)}$ is small if $\kappa(A^*)$ is large.

4.3 Generalization to the Case of Rectangular Matrices

Theorem 4.3.1 *Suppose the hypotheses of Theorem 3.2.1 hold except that* $A, C \in \mathbb{R}^{m\times n}$, $S \in \mathbb{R}^{m\times m}$, $T \in \mathbb{R}^{n\times n}$, $U \in \mathcal{G}_{\mu,\sigma}^{m\times r}$, *and* $V \in \mathcal{G}_{\mu,\sigma}^{n\times r}$. *Let* $l = \min\{m, n\}$. *Then the conclusion of part* (f) *of Theorem 3.2.1 still holds.*

Proof. Write: $C \equiv A+UV^T$. Now pre - multiply both sides of the above equation by $(I|0)$ where $I \in \mathbb{R}^{l\times l}$ and $0 \in \mathbb{R}^{l\times(m-l)}$, and post - multiply by the transpose of $(I|0)$ where $I \in \mathbb{R}^{l\times l}$ and $0 \in \mathbb{R}^{l\times(n-l)}$. This generates $\tilde{C} \equiv \tilde{A} + \widetilde{(UV^T)}$ in which $\tilde{C}$, $\tilde{A}$, and $\widetilde{(UV^T)}$ are all $\in \mathbb{R}^{l\times l}$. In particular, the SVD of $\tilde{A}$ yields an associated diagonal $\Sigma \in \mathbb{R}^{l\times l}$, whose leading $\rho \times \rho$ northeastern block consists of the singular values $\tilde{A}$ on the main diagonal, and $0s$ elsewhere in Σ. Apply Theorem 3.2.1 to $\tilde{C}$ and $\tilde{A}$ and note that if $\sigma_\rho(\tilde{A})/\sigma_n(\tilde{C}) \geq \sigma_\rho(A)/\sigma_n(C)$, the result follows. ■.

4.4 Conditioning of Random Complex Matrices

We have assumed dealing with real random matrices and vectors throughout the paper, but most of our study can be readily extended to the computations in the field $\mathbb{C}$ of complex numbers if we replace the transposes A^T by the Hermitian transposes A^H. Below we elaborate upon the respective extension of our probabilistic bounds on the norms and singular values.

Definition 4.4.1 *The set $\mathcal{G}^{m\times n}_{\mathbb{C},\mu,\sigma}$ of complex Gaussian random $m\times n$ matrices with a mean μ and a variance σ is the set $\{A+B\sqrt{-1}\}$ for $(A \mid B) \in \mathcal{G}^{m\times 2n}_{\mu,\sigma}$.*

Lemma 4.4.2 *Suppose y is a positive number; $\mathbf{w} \in \mathbb{R}^{n\times 1}$ is any fixed real unit vector, $||\mathbf{w}|| = 1$, $A \in \mathcal{G}^{n\times n}_{\mu,\sigma}$ and therefore is nonsingular with probability 1. Then* Probability$\{||A^{-1}\mathbf{w}|| > 1/y\} \leq \sqrt{\frac{2}{\pi}}\frac{y}{\sigma}$ for $j = 1, \ldots, n$.

Lemma 4.4.3 *The bound $F_{|||\mathbf{t}^H\mathbf{b}|||(y)} \leq y$ holds provided $\mathbf{t} = \mathbf{q} + \mathbf{r}\sqrt{-1}$ is a fixed complex unit vector and $\mathbf{b} = \mathbf{f} + \mathbf{g}\sqrt{-1} \in \mathcal{G}^{n\times 1}_{\mathbb{C},\mu,\sigma}$ is a complex vector such that $\mathbf{f}$, $\mathbf{g}$, $\mathbf{q}$ and $\mathbf{r}$ are real vectors, $||\mathbf{t}|| = 1$, and the vectors $\mathbf{f}$ and $\mathbf{g}$ are in $\mathcal{G}^{n\times 1}_{\mu,\sigma}$.*

Proof. We have $\mathbf{t}^H\mathbf{b} = \mathbf{q}^T\mathbf{f} + \mathbf{r}^T\mathbf{g} + (\mathbf{q}^T\mathbf{g} - \mathbf{r}^T\mathbf{f})\sqrt{-1}$, and so $|\mathbf{t}^H\mathbf{b}|^2 = |\mathbf{q}^T\mathbf{f} + \mathbf{r}^T\mathbf{g}|^2 + |\mathbf{q}^T\mathbf{g} - \mathbf{r}^T\mathbf{f}|^2$. Hence $|\mathbf{t}^H\mathbf{b}| \geq |\mathbf{q}^T\mathbf{f} + \mathbf{r}^T\mathbf{g}| = |\mathbf{u}^T\mathbf{v}|$ where $\mathbf{u}^T = (\mathbf{q}^T \mid \mathbf{r}^T)$ and $\mathbf{v} = (\mathbf{f}^T \mid \mathbf{g}^T)^T$. Note that $\mathbf{v} \in \mathcal{G}^{2n\times 1}_{\mu,\sigma}$ and $||\mathbf{u}|| = ||\mathbf{t}|| = 1$ and consider real vectors $\mathbf{u}$ and $\mathbf{v}$ replacing $\mathbf{b}$ and $\mathbf{t}$. ∎

Corollary 4.4.4 *Suppose y is a positive number and suppose a matrix $A \in \mathcal{G}^{n\times n}_{\mathbb{C},\mu,\sigma}$ and therefore is nonsingular with probability 1. Then*

$$\text{Probability}\{||A^{-1}\mathbf{e}_j|| > 1/y\} \leq \sqrt{\frac{2}{\pi}}\frac{y}{\sigma} \text{ for } j = 1, \ldots, n.$$

Proof. In the case of real matrices A the corollary is supported by the argument in the proof of [**SST06**, Lemma 3.2]. Now we employ Lemma 4.4.3 instead of this estimate, otherwise keep the same argument as in [**SST06**], and arrive at Corollary 4.4.4. ∎

Corollary 4.4.5 *Under the assumptions of Corollary 4.4.4 we have* $||A^{-1}|| \leq \sum_{j=1}^{n} X_j$ *where* X_j *are nonnegative random variables such that*

$$\text{Probability}\{X_j > 1/y\} \leq \sqrt{\frac{2}{\pi}}\frac{y}{\sigma} \text{ for } j = 1, \ldots, n.$$

Proof. Recall that for any $n \times n$ matrix B we have $||B|| = ||B\mathbf{w}||$ for some unit vector $\mathbf{w} = \sum_{j=1}^{n} w_j \mathbf{e}_j$. We have $|w_j| \leq ||\mathbf{w}|| = 1$ for all j. Substitute $B = A^{-1}$ and obtain $||A^{-1}|| = ||A^{-1}\mathbf{w}|| = ||\sum_{j=1}^{n} A^{-1} w_j \mathbf{e}_j|| \leq \sum_{j=1}^{n} |w_j| \; ||A^{-1}\mathbf{e}_j||$, and so $||A^{-1}|| \leq \sum_{j=1}^{n} X_j$ where $X_j = ||A^{-1}\mathbf{e}_j||$ for all j. It remains to combine this bound with Corollary 4.4.4. ∎

Appendix A

PRECONDITIONING TESTS

Table 1 covers the tests for the preconditioning power of randomized additive preprocessing from [PIMR10]. The tests show great power of additive preprocessing, even though we limited randomization to choosing the signs + and − for the nonzero entries of some very sparse and highly structured matrices U, V, and W. Namely, additive preprocessing consistently decreased the condition numbers of the input matrices from about 10^{16} to the values in the range from 10^2 to $5 * 10^5$. The authors of [PIMR10] have tested the input matrices of the following classes.

1n. *Nonsymmetric matrices of type I with numerical nullity r.* $A = S\Sigma_r T^T$ are $n \times n$ matrices where S and T are $n \times n$ random orthogonal matrices, that is, the factors Q in the QR factorizations of random real

matrices; $\Sigma_r = \mathrm{diag}(\sigma_j)_{j=1}^n$ is the diagonal matrix such that $\sigma_{j+1} \le \sigma_j$ for $j = 1, \dots, n-1$, $\sigma_1 = 1$, the values $\sigma_2, \dots, \sigma_{n-r-1}$ are randomly sampled in the semi-open interval $[0.1, 1)$, $\sigma_{n-r} = 0.1$, $\sigma_j = 10^{-16}$ for $j = n-r+1, \dots, n$, and therefore $\kappa(A) = 10^{16}$ [**H02**, Section 28.3].

1s. *Symmetric matrices of type I with numerical nullity r.* The same as in part 1n, but for $S = T$.

The matrices of the six other classes have been constructed in the form of $\frac{A}{||A||} + \beta I$, with the recipes for defining the matrices A and scalars β specified below.

2n. *Nonsymmetric matrices of type II with numerical nullity r.* $A = (W \mid WZ)$ where W and Z are random orthogonal matrices of sizes $n \times (n-r)$ and $(n-r) \times r$, respectively.

2s. *Symmetric matrices of type II with numerical nullity r.* $A = WW^T$ where W are random orthogonal matrices of size $n \times (n-r)$.

3n. *Nonsymmetric Toeplitz-like matrices with numerical nullity r.* $A = c(T \mid TS)$ for random Toeplitz matrices T of size $n \times (n-r)$ and S of size $(n-r) \times r$ and for a positive scalar c such that $||A|| \approx 1$.

3s. *Symmetric Toeplitz-like matrices with numerical nullity r.* $A = cTT^T$ for random Toeplitz matrices T of size $n \times (n-r)$ and a positive scalar c such that $||A|| \approx 1$.

4n. *Nonsymmetric Toeplitz matrices with numerical nullity* 1. $A = (a_{i,j})_{i,j=1}^n$ is a Toeplitz $n \times n$ matrix. Its entries $a_{i,j} = a_{i-j}$ are random

for $i - j < n - 1$, and so the matrix $A_{n-1} = (a_{i,j})_{i,j=1}^{n-1}$ is nonsingular (with probability 1) and was indeed nonsingular in all our tests. The entry $a_{n,1}$ is selected to annihilate or nearly annihilate $\det A$, that is, to fulfill

$$\det A = 0 \text{ or } \det A \approx 0, \tag{A.0.1}$$

in which case the matrix A is singular or ill conditioned.

4s. *Symmetric Toeplitz matrices with numerical nullity* 1. $A = (a_{i,j})_{i,j=1}^{n}$ is a Toeplitz $n \times n$ matrix. Its entries $a_{i,j} = a_{i-j}$ are random for $|i-j| < n-1$, while the entry $a_{1,n} = a_{n,1}$ was selected to satisfy equation (A.0.1), which is the quadratic equation in this entry. Occasionally it had no real roots, but then one can repeatedly generate the matrix A.

Set $\beta = 10^{-16}$ for symmetric matrices A in the classes 2s, 3s, and 4s, so that $\kappa(A) = 10^{16} + 1$ in these cases. For nonsymmetric matrices A define the scalar β by an iterative process such that $||A|| \approx 1$ and $10^{-18}||A|| \le \kappa(A) \le 10^{-16}||A||$ [**PIMR10**, Section 8.2].

Table 1 displays the average values of the condition numbers $\kappa(C)$ of the matrices $C = A + UV^T$ over 1000 tests for the inputs in the above classes, $r = 1, 2, 4, 8$ and $n = 128$.

Adopt choosing random matrices U,V, and W, both of which produce similar results:

- U, V, and W are Gaussian matrices,
- $U = \bar{U}/||\bar{U}||$, $\bar{U}^T = (\pm I_r \mid O_{r,r} \mid \pm I_r \mid O_{r,r} \mid \ldots \mid O_{r,r} \mid \pm I_r \mid O_{r,s})$, s is such that $\bar{U} \in \mathbb{R}^{n \times r}$,

  $V = \bar{V}/||\bar{V}||$, $\bar{V}^T = (2I_r \mid O_{r,r} \mid 2I_r \mid O_{r,r} \mid \ldots \mid O_{r,r} \mid 2I_r \mid O_{r,s}) - U^T$,

  $W = \bar{W}/||\bar{W}|| \in \mathbb{R}^{r \times r}$, $\bar{W}$ are circulant matrices, each defined by its first column, filled with ± 1, and here as well as in the expression for $\bar{U}$, all signs $\pm$ turn into $+$ and $-$ with the same probability 0.5, independently of each other.

Table 1: Preconditioning Tests

Type	r	cond(C)-Gaussian
$1n$	1	$1.38 \times 10^{+04}$
$1n$	2	$9.07 \times 10^{+03}$
$1n$	4	$6.91 \times 10^{+04}$
$1n$	8	$2.03 \times 10^{+04}$
$1s$	1	$4.48 \times 10^{+03}$
$1s$	2	$2.32 \times 10^{+04}$
$1s$	4	$2.38 \times 10^{+04}$
$1s$	8	$7.49 \times 10^{+04}$
$2n$	1	$6.75 \times 10^{+03}$
$2n$	2	$1.78 \times 10^{+04}$
$2n$	4	$3.91 \times 10^{+04}$
$2n$	8	$4.57 \times 10^{+04}$
$2s$	1	$1.35 \times 10^{+04}$
$2s$	2	$1.07 \times 10^{+04}$
$2s$	4	$2.01 \times 10^{+04}$
$2s$	8	$2.99 \times 10^{+04}$
$3n$	1	$4.62 \times 10^{+04}$
$3n$	2	$2.68 \times 10^{+06}$
$3n$	4	$4.29 \times 10^{+04}$
$3n$	8	$1.22 \times 10^{+05}$
$3s$	1	$5.34 \times 10^{+05}$
$3s$	2	$2.88 \times 10^{+06}$
$3s$	4	$1.44 \times 10^{+06}$
$3s$	8	$9.63 \times 10^{+05}$
$4n$	1	$4.26 \times 10^{+03}$
$4n$	2	$6.51 \times 10^{+03}$
$4n$	4	$4.22 \times 10^{+03}$
$4n$	8	$4.39 \times 10^{+03}$
$4s$	1	$4.06 \times 10^{+05}$
$4s$	2	$1.34 \times 10^{+06}$
$4s$	4	$1.30 \times 10^{+05}$
$4s$	8	$2.85 \times 10^{+04}$

Bibliography

[B11] C. Beltrán, Estimates on the Condition Number of Random, Rank-deficient Matrices, *IMA Journal of Numerical Analysis*, **31, 1**, 25–39, 2011.

[BP94] D. Bini, V. Y. Pan, *Polynomial and Matrix Computations,* volume 1: Fundamental Algorithms, Birkhäuser, Boston, 1994.

[CD05] Z. Chen, J. J. Dongarra, Condition Numbers of Gaussian Random Matrices, *SIAM. J. on Matrix Analysis and Applications*, **27**, 603–620, 2005.

[D83] J. D. Dixon, Estimating Extremal Eigenvalues and Condition Numbers of Matrices, *SIAM J. on Numerical Analysis*, **20, 4**, 812–814, 1983.

[D88] J. Demmel, The Probability That a Numerical Analysis Problem Is Difficult, *Math. of Computation*, **50**, 449–480, 1988.

[DS01] K. R. Davidson, S. J. Szarek, Local Operator Theory, Random Matrices, and Banach Spaces, in *Handbook on the Geometry of Banach Spaces* (W. B. Johnson and J. Lindenstrauss editors), pages 317–368, North Holland, Amsterdam, 2001.

[E88] A. Edelman, Eigenvalues and Condition Numbers of Random Matrices, *SIAM J. on Matrix Analysis and Applications*, **9**, **4**, 543–560, 1988.

[ES05] A. Edelman, B. D. Sutton, Tails of Condition Number Distributions, *SIAM J. on Matrix Analysis and Applications*, **27**, **2**, 547–560, 2005.

[GL96] G. H. Golub, C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, Maryland, 1996 (third addition).

[H02] N. J. Higham,*Accuracy and Stability in Numerical Analysis*, SIAM, Philadelphia, 2002 (second edition).

[HMT11] N. Halko, P. G. Martinsson, J. A. Tropp, Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions, *SIAM Review*, **53, 2**, 217–288, 2011.

[P01] V. Y. Pan, *Structured Matrices and Polynomials: Unified Superfast Algorithms*, Birkhäuser/Springer, Boston/New York, 2001.

[PGMQ] V. Y. Pan, D. Grady, B. Murphy, G. Qian, R. E. Rosholt, A. Ruslanov, Schur Aggregation for Linear Systems and Determinants, *Theoretical Computer Science*, *Special Issue on Symbolic–Numerical Algorithms* (D. A. Bini, V. Y. Pan, and J. Verschelde editors), **409**, **2**, 255–268, 2008.

[PIMR10] V. Y. Pan, D. Ivolgin, B. Murphy, R. E. Rosholt, Y. Tang, X. Yan, Additive Preconditioning for Matrix Computations, *Linear Algebra and Its Applications*, **432**, 1070–1089, 2010.

[PQ10] V. Y. Pan, G. Qian, Randomized Preprocessing of Homogeneous Linear Systems of Equations, *Linear Algebra and Its Applications*, **432**, 3272–3318, 2010.

[PQ12] V. Y. Pan, G. Qian, Solving Linear Systems of Equations with Randomization, Augmentation, and Aggregation, *Linear Algebra and Its Applications*, **437**, 1851–1876, 2012.

[PQYa] V. Y. Pan, G. Qian, X. Yan, Random Multipliers Numerically Stabilize Gaussian and Block Elimination: Proofs and an Extension to Low-rank Approximation, http://arxiv.org/abs/arxiv1406.5802 math.NA (28 pages), 2014.

[PQZ11] V. Y. Pan, G. Qian,A.-L. Zheng, *Randomized Matrix Computations*, (to appear).

[PQZ13] V. Y. Pan, G. Qian, A. Zheng, Randomized Preprocessing versus Pivoting, *Linear Algebra and Its Applications*, **438, 4**, 1883–1899, 2013.

[PQZ14] V. Y. Pan, G. Qian, A. Zheng, Randomized Augmentations and Additive Preprocessing. arXiv:1412.5864v1 [math.NA] 18 Dec 2014.

[S98] G. W. Stewart, *Matrix Algorithms, Vol I: Basic Decompositions*, SIAM, 1998.

[SST06] A. Sankar, D. Spielman, S.-H. Teng, Smoothed Analysis of the Condition Numbers and Growth Factors of Matrices, *SIAM J. on Matrix Analysis*, **28**, **2**, 446–476, 2006.