2013

TR-2013001: Randomized Augmentation and Additive Preprocessing

Victor Y. Pan

Guoliang Qian

Ai-Long Zheng

Randomized Augmentations and Additive Preprocessing *

Victor Y. Pan[1,2],[a], Guoliang Qian[2],[b], and Ai-Long Zheng[2],[c]

[1] Department of Mathematics and Computer Science
Lehman College of the City University of New York
Bronx, NY 10468 USA
[2] Ph.D. Programs in Mathematics and Computer Science
The Graduate Center of the City University of New York
New York, NY 10036 USA
[a] victor.pan@lehman.cuny.edu
http://comet.lehman.cuny.edu/vpan/
[b] gqian@gc.cuny.edu
[c] azheng-1999@yahoo.com

Abstract

Random matrices tend to be well conditioned, and so one can expect that appending properly scaled random rows and columns or adding a scaled random matrix of a fixed rank can decrease the condition number of an ill conditioned matrix. We prove probabilistic estimates for this decrease by using Gaussian random matrices as the preprocessors, but our tests showed equally strong impact on the condition numbers in the case where the preprocessors were random sparse and structured matrices, defined by much fewer random parameters. For sample applications of randomized preprocessing to matrix computations, we precondition an ill conditioned matrix, approximate its singular spaces associated with its largest and smallest singular values, approximate this matrix with low-rank matrices, and yield its 2×2 block diagonalization. Combining our present techniques with randomized matrix multiplication (which we study elsewhere) should lead to further progress in matrix computations.

2000 Math. Subject Classification: 15A52, 15A12, 65F22

Key Words: Random matrices, Preconditioning, Numerical rank

1 Introduction

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 matrix computations, by continuing our study in [PGMQ], [PIMR10], [PQ10], [PQ12], [PQZC], [PQZa], and [PY09]. We prove that with probability 1 or near 1 our techniques of scaled Gaussian randomized preprocessing regularize and precondition a large and important class of ill conditioned matrices. By employing randomization

*Some results of this paper have been presented at the ACM-SIGSAM International Symposium on Symbolic and Algebraic Computation (ISSAC '2011), San Jose, CA, 2011, the 3nd International Conference on Matrix Methods in Mathematics and Applications (MMMA 2011) in Moscow, Russia, June 22-25, 2011, the 7th International Congress on Industrial and Applied Mathematics (ICIAM 2011), in Vancouver, British Columbia, Canada, July 18-22, 2011, the SIAM International Conference on Linear Algebra, in Valencia, Spain, June 18-22, 2012, and the Conference on Structured Linear and Multilinear Algebra Problems (SLA2012), in Leuven, Belgium, September 10-14, 2012

we precondition an ill conditioned linear system of equations, approximate the singular spaces of an ill conditioned matrix A associated with its largest and smallest singular values, approximate this matrix by low-rank matrices, and achieve its 2×2 block diagonalization. Our analysis and experiments show substantial progress versus the known algorithms. In our tests (the contribution of the second and third coauthors) our techniques have fully preserved their power when we employed sparse and structured preprocessors, defined by much fewer random parameters. Combining our present techniques with the results of [PGMQ, Section 12.2], [PQZa], and [PQb] on the impacts of randomized matrix multiplication should lead to further advance in matrix computations.

We organize the paper as follows. We recall some definitions and preliminaries in the next section and some estimates for the condition numbers of Gaussian random matrices and randomized matrix products in Section 3. In Sections 4 and 5 we formally estimate how much the scaled randomized preprocessing of two kinds (that is augmentation and additive modification) can decrease the condition number of an ill conditioned matrix, and we apply randomized augmentation to the solution of ill conditioned Toeplitz linear systems of equations. In Sections 6 and 7 we cover some other applications of scaled randomized preprocessing to matrix computations. In Section 8 we present numerical tests, which constitute the contribution of the second and the third authors. We leave Section 9 for conclusions. In Appendix A we estimate the probability that a random matrix has full rank under the uniform probability distribution. In Appendix B we comment on the extension of our probabilistic estimates to the case of complex matrices.

2 Some definitions and basic results

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

Hereafter "flop" stands for "arithmetic operation"; "expect" and "likely" mean "with probability 1 or close to 1" (we do not use the concept of the expected value), and the concepts "large", "small", "near", "closely approximate", "ill conditioned" and "well conditioned" are quantified in the context. For two scalars a and b we write $a \ll b$ and $b \gg a$ if the ratio $|b/a|$ is large. We write $a \approx b$ if $|a - b| \ll |a| + |b|$. Next we recall and extend some customary definitions of matrix computations [GL96], [S98].

2.1 Some basic definitions on matrix computations

$\mathbb{R}^{m \times n}$ is the class of real $m \times n$ matrices $A = (a_{i,j})_{i,j}^{m,n}$.

$(B_1 \mid \ldots \mid B_k) = (B_j)_{j=1}^{k}$ is a $1 \times k$ block matrix with blocks $B_1, \ldots, B_k$. $\mathrm{diag}(B_1, \ldots, B_k) = \mathrm{diag}(B_j)_{j=1}^{k}$ is a $k \times k$ block diagonal matrix with diagonal blocks $B_1, \ldots, B_k$.

$\mathbf{e}_i$ is the ith coordinate vector of dimension n for $i = 1, \ldots, n$. These vectors define the identity matrix $I_n = (\mathbf{e}_1 \mid \ldots \mid \mathbf{e}_n)$ and the reflection matrix $J_n = (\mathbf{e}_n \mid \ldots \mid \mathbf{e}_1)$, both of size $n \times n$. $O_{k,l}$ is the $k \times l$ matrix filled with zeros. $\mathbf{0}_k$ is the vector $O_{k,1}$. We write I, J, O, and $\mathbf{0}$ where the size of a matrix or a vector is not important or is defined by context.

A^T is the transpose of a matrix A. A^H is its Hermitian transpose. A matrix A is symmetric if $A = A^T$ and is symmetric positive definite if $A = B^T B$ for a real nonsingular matrix B.

A real matrix Q is called *orthogonal* if $Q^T Q = I$ or $QQ^T = I$. More generally, over the complex field $\mathbb{C}$ a matrix U is called *unitary* if $U^H U = I$ or $UU^H = I$. Hereafter $Q(A)$ denote a unique orthogonal matrix specified by the following result.

Fact 2.1. *[GL96, Theorem 5.2.2]. QR factorization $A = QR$ of a matrix A having full column rank into the product of an orthogonal matrix $Q = Q(A)$ and an upper triangular matrix $R = R(A)$ is unique provided that the factor R is a square matrix with positive diagonal entries.*

2.2 Range, null space, rank, nullity, and nmbs

$\mathcal{R}(A)$ 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 $\{\mathbf{v} : A\mathbf{v} = \mathbf{0}\}$, $\mathrm{rank}(A) = \dim \mathcal{R}(A)$ its rank, and

$\mathrm{nul}(A) = \dim \mathcal{N}(A) = n - \mathrm{rank}(A)$ its right nullity or just *nullity*, whereas $\mathrm{nul}(A^T) = m - \mathrm{rank}(A)$ is the left nullity of A, equal to $\mathrm{nul}(A)$ if and only if $m = n$. $\mathbf{v}$ is the null vector of A if $A\mathbf{v} = \mathbf{0}$.

Fact 2.2. *The set $\mathbb{M}$ of $m \times n$ matrices of rank ρ is an algebraic variety of dimension $(m+n-\rho)\rho$.*

Proof. Let M be an $m \times n$ matrix of a rank ρ with a nonsingular leading $\rho \times \rho$ block M_{00} and write $M = \begin{pmatrix} M_{00} & M_{01} \\ M_{10} & M_{11} \end{pmatrix}$. Then the $(m-\rho) \times (n-\rho)$ *Schur complement* $M_{11} - M_{10}M_{00}^{-1}M_{01}$ must vanish, which imposes $(m-\rho)(n-\rho)$ algebraic equations on the entries of M. Similar argument can be applied where any $\rho \times \rho$ submatrix of the matrix M (among $\binom{m}{\rho}\binom{n}{\rho}$ such submatrices) is nonsingular. Therefore $\dim \mathbb{M} = mn - (m-\rho)(n-\rho) = (m+n-\rho)\rho$. □

A matrix that has full column rank is a *matrix basis* for its range. A matrix basis B for the null space $\mathcal{N}(A)$ is a *null matrix basis* or a *nmb* for the matrix A, and we write $B = \mathrm{nmb}(A)$. In other words $B = \mathrm{nmb}(A)$ if the matrix B has full column rank and if $\mathcal{R}(B) = \mathcal{N}(A)$, which can be equivalently rewritten as $\mathcal{R}(A) = \mathcal{N}(B)$. $\mathcal{N}(A^T)$ is the left null space of a matrix A, and similarly the map $A \Longrightarrow A^T$ defines left null vectors, left nmbs, and the left nullity of a matrix A.

2.3 Norms, SVD, and singular spaces

We write $||A|| = ||A||_2 = ||A^T||$ and $||\mathbf{v}|| = \sqrt{\mathbf{v}^T\mathbf{v}} = ||\mathbf{v}||_2$ and recall that

$$||AB|| \le ||A||\ ||B|| \text{ for any matrix product } AB, \tag{2.1}$$

$$||\operatorname{diag}(M_j)_j|| = \max_j ||M_j|| \text{ for any set of matrices } M_j. \tag{2.2}$$

A matrix A is *normalized* if $||A|| = 1$. A normalized vector is orthogonal (unitary), and we call it *unit*. We write $A \approx B$ if $||A-B|| \ll ||A|| + ||B||$.

Define an *SVD* or *full SVD* of an $m \times n$ matrix A of a rank ρ as follows,

$$A = S_A \Sigma_A T_A^T. \tag{2.3}$$

Here $S_A S_A^T = S_A^T S_A = I_m$, $T_A T_A^T = T_A^T T_A = I_n$, $\Sigma_A = \operatorname{diag}(\widehat{\Sigma}_A, O_{m-\rho,n-\rho})$, $\widehat{\Sigma}_A = \operatorname{diag}(\sigma_j(A))_{j=1}^{\rho}$, $\sigma_j = \sigma_j(A) = \sigma_j(A^T)$ is the jth largest singular value of a matrix A for $j = 1, \dots, \rho$, $\sigma_1 = ||A||$, $\sigma_\rho > 0$, and we write $\sigma_j = 0$ for $j > \rho$.

Lemma 2.1. *$\sigma_j(SM) = \sigma_j(MT) = \sigma_j(M)$ for all j if S and T are square orthogonal matrices.*

Lemma 2.2. *If A_0 is a submatrix of a matrix A, then $\sigma_j(A) \ge \sigma_j(A_0)$ for all j.*

Proof. [GL96, Corollary 8.6.3] implies the claimed bound where A_0 is any block of columns of the matrix A. Transposition of a matrix and permutations of its rows and columns do not change singular values, and thus we can extend the bounds to all submatrices A_0. □

Theorem 2.1. *We have $|\sigma_j(C) - \sigma_j(C+E)| \le ||E||$ for all $m \times n$ matrices C and E and all j.*

Proof. See [GL96, Corollary 8.6.2] or [S98, Corollary 4.3.2]. □

Hereafter we use the following definitions. For every integer k in the range $1 \le k < \mathrm{rank}(A)$ define the partition $S_A = (S_{k,A} \mid S_{A,m-k})$ and $T_A = (T_{k,A} \mid T_{A,n-k})$ where the submatrices $S_{k,A}$ and $T_{k,A}$ are formed by the first k columns of the matrices S_A and T_A, respectively. Write $\Sigma_{k,A} = \operatorname{diag}(\sigma_j(A))_{j=1}^{k}$, $\mathbb{S}_{k,A} = \mathcal{R}(S_{k,A})$ and $\mathbb{T}_{k,A} = \mathcal{R}(T_{k,A})$. If $\sigma_k > \sigma_{k+1}$, then $\mathbb{S}_{k,A}$ and $\mathbb{T}_{k,A}$ are the left and right *leading singular spaces*, respectively, associated with the k largest singular values of the matrix A, whereas their orthogonal complements $\mathbb{S}_{A,m-k} = \mathcal{R}(S_{A,m-k})$ and $\mathbb{T}_{A,n-k} = \mathcal{R}(T_{A,n-k})$ are the left and right *trailing singular spaces*, respectively, associated with the other singular values of A. The pairs of subscripts $\{k, A\}$ versus $\{A, m-k\}$ and $\{A, n-k\}$ mark the leading versus trailing singular spaces. The left singular spaces of A are the right singular spaces of A^T and vice

versa. All matrix bases for the singular spaces $\mathbb{S}_{k,A}$ and $\mathbb{T}_{k,A}$ are given by matrices $S_{k,A}X$ and $T_{k,A}Y$, respectively, for nonsingular $k \times k$ matrices X and Y. Orthogonal matrices X and Y define orthogonal matrix bases for these spaces. B is an *approximate matrix basis* for a space $\mathbb{S}$ within a relative error norm bound τ if there exists a matrix E such that $B + E$ is a matrix basis for this space $\mathbb{S}$ and if $||E|| \le \tau||B||$.

2.4 Inverses, generalized inverses, and perturbation bounds

$A^+ = T_A \operatorname{diag}(\widehat{\Sigma}_A^{-1}, O_{n-\rho,m-\rho})S_A^T$ is the Moore–Penrose pseudo-inverse of the matrix A of (2.3), and

$$||A^+|| = 1/\sigma_\rho(A) \tag{2.4}$$

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. Theorem 2.1 implies the following bound.

Theorem 2.2. *Suppose two matrices $C, C + E \in \mathbb{C}^{m\times n}$ have full rank. Then $||(C + E)^+ - C^+|| \le ||E||\ ||(C + E)^+\ C^+||$.*

This bound can be improved where the matrices C and $C + E$ are nonsingular.

Theorem 2.3. *Suppose C and $C+E$ are two nonsingular matrices of the same size and $||C^{-1}E|| = \theta < 1$. Then $||I - (C + E)^{-1}C|| \le \frac{\theta}{1-\theta}$ and $|||(C + E)^{-1} - C^{-1}|| \le \frac{\theta}{1-\theta}||C^{-1}||$, in particular $|||(C + E)^{-1} - C^{-1}|| \le 0.5||C^{-1}||$ if $\theta \le 1/3$.*

Proof. See [S98, Corollary 1.4.19] for $P = -C^{-1}E$. □

Theorem 2.4. *[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 $G = I_r - V^TC^{-1}U$ is nonsingular and we have the Sherman–Morrison–Woodbury (hereafter* SMW*) formula*

$$A^{-1} = C^{-1} + C^{-1}UG^{-1}V^TC^{-1}.$$

2.5 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 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], [H02, Chapter 15], [KL94], [KW92], and [S98, Section 5.3] on the estimation of matrix norms and condition numbers.

An $m \times n$ matrix A has *numerical rank* $\operatorname{nrank}(A)$, not exceeding $\operatorname{rank}(A)$, and has the right numerical nullity $\operatorname{nnul}(A) = n - \operatorname{nrank}(A)$ or just *numerical nullity* if the ratios $\sigma_j(A)/||A||$ are small for $j > \operatorname{nrank}(A)$ but not for $j \le \operatorname{nrank}(A)$. The *left numerical nullity* of the matrix A equals the numerical nullity $\operatorname{nnul}(A^T) = m - \operatorname{nrank}(A)$ of the $n \times m$ transpose A^T and coincides with the numerical nullity of A if and only if $m = n$.

Remark 2.1. *One can specify the adjective "small" above as "smaller than a fixed positive tolerance". The choice of the tolerance can be a challenge, e.g., for the matrix* $\operatorname{diag}(1.1^{-j})_{j=0}^{999}$.

If a well conditioned $m \times n$ matrix A has a rank $\rho < l = \min\{m, n\}$, then all its close neighbours have numerical rank ρ and almost all of them have full rank l (see Section 3.2). Conversely, suppose a matrix A has a positive numerical rank $\rho = \operatorname{nrank}(A)$ and *truncate its SVD* by setting to 0 all its singular values, except for the ρ largest ones. Then the resulting matrix $A - E$ is well conditioned and has rank ρ and $||E|| = \sigma_{\rho+1}(A)$, and so $A - E$ is a rank-ρ approximation to the matrix A within the error norm bound $\sigma_{\rho+1}(A)$. At a lower computational cost we can obtain rank-ρ approximations of the matrix A from its rank-revealing factorizations [GE96], [HP92], [P00a], and we can further decrease the computational cost by applying randomized algorithms (cf. Remark 6.4).

Remark 2.2. *How large is the class of $m \times n$ matrices having a numerical rank ρ? We characterize it indirectly, by noting that by virtue of Fact 2.2 the nearby matrices of rank ρ form a variety of dimension $(m+n-\rho)\rho$, which increases as ρ increases.*

2.6 The computation of approximate matrix bases for the leading singular spaces of a matrix and its rank-ρ approximation.

Theorem 2.5. *Cf. [PQb]. Suppose A is an $m \times n$ matrix, $S_A \Sigma_A T_A^T$ is its SVD of (2.3), q is a positive integer, $q \le \min\{m,n\}$, and T and S are matrix bases for the spaces $\mathbb{T}_{q,A}$ and $\mathbb{S}_{q,A}$, respectively. Then*

$$||A - AT(T^TT)^{-1}T^T|| = ||A - S(S^TS)^{-1}S^TA|| = \sigma_{q+1}(A). \tag{2.5}$$

For orthogonal matrices T and S we have $T^TT = S^TS = I_q$ and

$$||A - ATT^T|| = ||A - SS^TA|| = \sigma_{q+1}(A). \tag{2.6}$$

The theorem reduces the computation of a rank-q approximation of a matrix A to computing or approximating a matrix basis for its leading singular space $\mathbb{T}_{q,A}$ or $\mathbb{S}_{q,A}$. The following result from [PQb] supports randomized computation of such approximate matrix bases.

Theorem 2.6. *Suppose a matrix $A \in \mathbb{R}^{m\times n}$ has a numerical rank ρ, $H \in \mathcal{G}_{0,1}^{n\times\rho_+}$ and $G \in \mathcal{G}_{0,1}^{m\times\rho_+}$ for $\rho_+ \ge \rho$. Then the matrices $T = A^TG$ and $S = AH$ have full rank with probability 1 and are likely to have numerical rank ρ and to satisfy the following matrix equations,*

$$S + \Delta = S_{\rho,A}U \text{ and } T + \Delta' = T_{\rho,A}V \tag{2.7}$$

for two matrices Δ and Δ' having norms of order $\sigma_{\rho+1}(A)$ and for two nonsingular matrices U and V having condition numbers of at most order $||A||/(\sigma_\rho(A)\sqrt{\rho})$.

For a matrix basis B of the leading singular space $\mathbb{S}_{\rho,A}$ or $\mathbb{T}_{\rho,A}$ of an $m \times n$ matrix A, a nmb(B) is a matrix basis for the trailing singular space $\mathbb{S}_{A,m-\rho}$ or $\mathbb{T}_{A,n-\rho}$, respectively, and vice versa. Clearly, similar relationships hold for approximate matrix bases.

2.7 Toeplitz matrices

A *Toeplitz* $m \times n$ matrix $T_{m,n} = (t_{i-j})_{i,j=1}^{m,n}$ is defined by its first row and column, that is by the vector $(t_h)_{h=1-n}^{m-1}$ of dimension $m+n-1$. We write $T_n = T_{n,n} = (t_{i-j})_{i,j=1}^{n,n}$ (see (2.8)).

A lower *triangular Toeplitz* $n \times n$ matrix $Z(\mathbf{t}) = (t_{i-j})_{i,j=1}^{n}$ (where $t_k = 0$ for $k < 0$) is defined by its first column $\mathbf{t} = (t_h)_{h=0}^{n-1}$. We write $Z(\mathbf{t})^T = (Z(\mathbf{t}))^T$. $Z = Z(\mathbf{e}_2)$ is the downshift $n \times n$ (see (2.8)). We have $Z\mathbf{v} = (v_i)_{i=0}^{n-1}$ and $Z(\mathbf{v}) = \sum_{i=1}^n v_i Z^{i-1}$ for $\mathbf{v} = (v_i)_{i=1}^n$ and $v_0 = 0$,

$$T_n = \begin{pmatrix} t_0 & t_{-1} & \cdots & t_{1-n} \\ t_1 & t_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & t_{-1} \\ t_{n-1} & \cdots & t_1 & t_0 \end{pmatrix}, \quad Z = \begin{pmatrix} 0 & & \cdots & & 0 \\ 1 & \ddots & & & \\ \vdots & \ddots & \ddots & & \vdots \\ & & \ddots & 0 & \\ 0 & & \cdots & 1 & 0 \end{pmatrix}. \tag{2.8}$$

Theorem 2.7. *[GS72]. Let the matrix T_n of (2.8) be nonsingular and write $\mathbf{p} = T_n^{-1}\mathbf{e}_1$ and $\mathbf{q} = T_n^{-1}\mathbf{e}_n$. If $p_1 = \mathbf{e}_1^T\mathbf{p} \ne 0$, then $p_1 T_n^{-1} = Z(\mathbf{p})Z(J\mathbf{q})^T - Z(Z\mathbf{q})Z(ZJ\mathbf{p})^T$.*

3 Ranks and conditioning of Gaussian random matrices

3.1 Random variables and Gaussian random matrices

Definition 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*

$$\mu - 4\sigma \le y \le \mu + 4\sigma \text{ with a probability near 1.} \tag{3.1}$$

Definition 3.2. *A matrix or a vector is a* Gaussian 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 $\sigma^2 = 1$). By restricting this set to Toeplitz matrices we obtain the set $\mathcal{T}_{\mu,\sigma}^{m\times n}$ of* Gaussian random Toeplitz matrices.

Lemma 3.1. *[SST06, Proposition 2.2]. Suppose $H \in \mathcal{G}_{\mu,\sigma}^{m\times n}$, $SS^T = S^TS = I_m$, $TT^T = T^TT = I_n$. Then $SH \in \mathcal{G}_{\mu,\sigma}^{m\times n}$ and $HT \in \mathcal{G}_{\mu,\sigma}^{m\times n}$.*

3.2 Nondegeneration of Gaussian random matrices

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

Lemma 3.2. *[DL78], [S80], [Z79]. 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 Gaussian random variables range over infinite sets Δ, e.g. the real line or its interval. Then the lemma implies that a nonzero polynomial vanishes with probability 0. Consequently a square Gaussian random general or Toeplitz matrix is nonsingular with probability 1 because its determinant is a polynomials in the entries, and so rectangular Gaussian random general and Toeplitz matrices have generic rank profile with probability 1. Furthermore all entries of such matrix A and of its adjoint $\mathrm{adj}\, A$ are subdeterminants and thus are nonzeros with probability 1. Clearly this property of the adjoint also holds for the inverse $A^{-1} = \frac{\mathrm{adj}\, A}{\det A}$ if the matrix A is nonsingular. Hereafter, wherever this causes no confusion, we assume by default that *Gaussian random general and Toeplitz matrices have full rank, and their inverses (if defined) have nonzero entries.* These properties can be readily extended to the products of the latter matrices by nonsingular matrices, and further to various other functions of Gaussian random general and structured matrices. Moreover similar properties hold with probability near 1 where the random variables are sampled under the uniform probability distribution from a finite set of a large cardinality (see Appendix A).

3.3 Extremal singular values of Gaussian random matrices and of randomized matrix products

Besides having full rank with probability 1, Gaussian random matrices in Definition 3.2 are likely to be well conditioned [D88], [E88], [ES05], [CD05], [B11], and even the sum $M + A$ for $M \in \mathbb{R}^{m\times n}$ and $A \in \mathcal{G}_{\mu,\sigma}^{m\times n}$ is likely to be well conditioned unless the ratio $\sigma/||M||$ is small or large [SST06].

The following theorem states an upper bound proportional to y on the cdf $F_{1/||A^+||}(y)$, that is on the probability that the smallest positive singular value $1/||A^+|| = \sigma_l(A)$ of a Gaussian random matrix A is less than a nonnegative scalar y (cf. (2.4)) and consequently on the probability that the norm $||A^+||$ exceeds a positive scalar x. The stated bound still holds if we replace the matrix A by $A - B$ for any fixed matrix B, and for $B = O_{m,n}$ the bounds can be strengthened by a factor $y^{|m-n|}$ [ES05], [CD05].

Theorem 3.1. *Suppose $A \in \mathcal{G}_{\mu,\sigma}^{m\times n}$, $B \in \mathbb{R}^{m\times n}$, $l = \min\{m,n\}$, $x > 0$, and $y \ge 0$. Then $F_{\sigma_l(A-B)}(y) \le 2.35\sqrt{l}y/\sigma$, that is $Probability\{||(A-B)^+|| \ge 2.35x\sqrt{l}/\sigma\} \le 1/x$.*

Proof. For $m = n$ this is [SST06, Theorem 3.3]. Apply Lemma 2.2 to extend it to any pair $\{m, n\}$. □

The following two theorems supply lower bounds $F_{||A||}(z)$ and $F_{\kappa(A)}(y)$ on the probabilities that $||A|| \le z$ and $\kappa(A) \le y$ for two scalars y and z, respectively, and a Gaussian random matrix A. We do not use the second theorem, but state it for the sake of completeness and only for square $n \times n$ matrices A. The theorems imply that the functions $1 - F_{||A||}(z)$ and $1 - F_{\kappa(A)}(y)$ decay as $z \to \infty$ and $y \to \infty$, respectively, and that the decays are exponential in $-z^2$ and proportional to $\sqrt{\log y}/y$, respectively. For small values $y\sigma$ and a fixed n the lower bound of Theorem 3.3 becomes negative, in which case the theorem becomes trivial. Unlike Theorem 3.1, in both theorems we assume that $\mu = 0$.

Theorem 3.2. *[DS01, Theorem II.7]. Suppose $A \in \mathcal{G}_{0,\sigma}^{m\times n}$, $h = \max\{m, n\}$ and $z \ge 2\sigma\sqrt{h}$. Then $F_{||A||}(z) \ge 1 - \exp(-(z - 2\sigma\sqrt{h})^2/(2\sigma^2))$, and so the norm $||A||$ is likely to have order $\sigma\sqrt{h}$.*

Theorem 3.3. *[SST06, Theorem 3.1]. Suppose $0 < \sigma \le 1$, $y \ge 1$, $A \in \mathcal{G}_{0,\sigma}^{n\times n}$. Then the matrix A has full rank with probability 1 and $F_{\kappa(A)}(y) \ge 1 - (14.1 + 4.7\sqrt{(2\ln y)/n})n/(y\sigma)$.*

Proof. See [SST06, the proof of Lemma 3.2]. □

The following theorem from [PQb] implies that multiplication by standard Gaussian random matrix is unlikely to decrease the smallest positive singular value of a matrix dramatically, even though $UV = O$ for some pairs of rectangular orthogonal matrices U and V.

Theorem 3.4. *Suppose $G' \in \mathcal{G}_{\mu,\sigma}^{r\times m}$, $H' \in \mathcal{G}_{\mu,\sigma}^{n\times r}$, $M \in \mathbb{R}^{m\times n}$, $G = G' + U$, $H = H' + V$ for some matrices U and V, $r(M) = \operatorname{rank}(M)$, $x > 0$ and $y \ge 0$. Then $F_{1/||(GM)^+||}(y) \le F(y, M, \sigma)$ and $F_{1/||(MH)^+||}(y) \le F(y, M, \sigma)$ for $F(y, M, \sigma) = 2.35y\sqrt{\widehat{r}}||M^+||/\sigma$ and $\widehat{r} = \min\{r, r(M)\}$, that is Probability$\{||P^+|| \ge 2.35x\sqrt{\widehat{r}}||M^+||/\sigma\} \le 1/x$ for $P = GM$ and $P = MH$.*

4 Randomized augmentation

4.1 Augmentation and an extension of the SMW formula

The solution of a nonsingular linear system of n equations $A\mathbf{y} = \mathbf{b}$ can be readily recovered from a null vector $\begin{pmatrix} -1/\beta \\ \mathbf{y} \end{pmatrix}$ of the matrix $K = (\beta\mathbf{b} \mid A)$ for a nonzero scalar β. If the matrix A has numerical nullity 1 and if the ratio $||A||/||\beta\mathbf{b}||$ is neither large nor small, then the matrix K is well conditioned for the average vector $\mathbf{b}$ [PQ12, Section 13.1]. The above map $A \Longrightarrow K$ is a special case of more general augmentation

$$K = \begin{pmatrix} W & V^T \\ -U & A \end{pmatrix}, \tag{4.1}$$

which we study next, beginning with the following extension of the SMW formula of Theorem 2.4.

Theorem 4.1. *Suppose equation (4.1) holds, $m = n$ and the matrices A, W and K are nonsingular. Write $S = A + UW^{-1}V^T$. Then the matrix S is nonsingular, S^{-1} is the trailing (southwestern) $n \times n$ block of K^{-1}, and we have*

$$A^{-1} = S^{-1} + S^{-1}UW^{-1}R^{-1}V^TS^{-1} \text{ for } R = I - V^TS^{-1}UW^{-1}. \tag{4.2}$$

Proof. Apply the SMW formula of Theorem 2.4 for C replaced by S, U by UW^{-1}, and G by R. □

4.2 Bounds on the norms, ranks, numerical ranks, and condition numbers

Theorem 4.2. *Suppose m, n, q and r are four positive integers, $l = \min\{m, n\}$, a real normalized $m \times n$ matrix A has a rank ρ, $U \in \mathcal{G}_{0,1}^{m\times q}$, $V \in \mathcal{G}_{0,1}^{n\times r}$, $W \in \mathbb{R}^{q\times r}$, the matrix W has full rank, and K denotes the matrix of (4.1). Write $l_q = \min\{m-\rho, q\}$, $l_r = \min\{n-\rho, r\}$, and $\rho_+ = \min\{\rho+q+r, m+q, n+r\}$. Then (i) $||K|| \le ||A|| + ||W|| + ||U|| + ||V||$ and (ii) $\operatorname{rank}(K) = \rho_+$ with probability 1, that is $\operatorname{rank}(K) = \rho+q+r$ if $l_q = q$ and $l_r = r$; otherwise $\operatorname{rank}(K) = n+q$ if $l_r = r$, and $\operatorname{rank}(K) = m+r$ if $l_q = q$, and so with probability 1 the matrix K is rank deficient if and only if $l_q < q$ and $l_r < r$. (iii) Furthermore unless $l_q < q$ and $l_r < r$ we have $\kappa(K) \le ||K||(N_1+N_2+N_3)$ with probability 1 where*

$$N_1 = \max\{||V_1^{-1}||, ||U_1^{-1}||, 1/\sigma_\rho(A)\},$$

$$N_2 \le \max\{||V_1^{-1}||\ ||V_0||, ||U_0||\ ||U_1^{-1}||\}/\sigma_\rho(A),$$

$$N_3 \le ||V_1^{-1}||\ ||U_1^{-1}||(||W|| + ||V_0||\ ||U_0||/\sigma_\rho(A)),$$

for some matrices $U_0 \in \mathcal{G}_{0,1}^{(l-l_q)\times q}$, $V_0 \in \mathcal{G}_{0,1}^{(l-l_r)\times r}$, $U_1 \in \mathcal{G}_{0,1}^{l_q\times l_q}$, $V_1 \in \mathcal{G}_{0,1}^{l_r\times l_r}$.

Apply Theorems 3.1 and 3.2 to bound the norms $||U_0||$, $||V_0||$, $||U_1^{-1}||$, and $||V_1^{-1}||$ and then deduce from part (iii) that the ratio $\kappa(K)/||K||$ is likely to be of at most order $1/\sigma_\rho(A)$).

Proof. Part (i) is verified immediately. Now suppose $A = S_A\Sigma_A T_A^T$ and $W = S_W\Sigma_W T_W^T$ are the SVDs and write $\bar{K} = \operatorname{diag}(J_q S_W^T, S_A^T)K \operatorname{diag}(T_W J_r, T_A)$ and $\bar{W} = J_q\Sigma_W J_r$. Then

$$\bar{K} = \begin{pmatrix} \bar{W} & \bar{V}^T \\ -\bar{U} & \Sigma_A \end{pmatrix}.$$

Here $\sigma_j(K) = \sigma_j(\bar{K})$ for all j by virtue of Lemma 2.1, whereas $\bar{U} \in \mathcal{G}_{0,1}^{m\times q}$ and $\bar{V} \in \mathcal{G}_{0,1}^{n\times r}$ by virtue of Lemma 3.1 because the matrices J_q, J_r, S_W, T_W, S_A, and T_A are square and orthogonal.

Write $\Sigma_\rho = \operatorname{diag}(\sigma_j(A))_{j=1}^{\rho}$, $\Sigma_A = \operatorname{diag}(\Sigma_\rho, O_{m-\rho,n-\rho})$, $\bar{U} = \begin{pmatrix}\bar{U}_0\\ \bar{U}_1\end{pmatrix}$, and $\bar{V} = \begin{pmatrix}\bar{V}_0\\ \bar{V}_1\end{pmatrix}$ where $\bar{U}_0 \in \mathcal{G}_{0,1}^{\rho\times q}$, $\bar{V}_0 \in \mathcal{G}_{0,1}^{\rho\times r}$, $\bar{U}_1 \in \mathcal{G}_{0,1}^{(m-\rho)\times q}$, and $\bar{V}_1 \in \mathcal{G}_{0,1}^{(n-\rho)\times r}$, and with probability 1 obtain

$$\bar{K} = \begin{pmatrix} \bar{W} & \bar{V}_0^T & \bar{V}_1^T \\ \bar{U}_0 & \Sigma_\rho & O_{\rho,n-\rho} \\ \bar{U}_1 & O_{m-\rho,\rho} & O_{m-\rho,n-\rho} \end{pmatrix}. \tag{4.3}$$

Therefore $\operatorname{rank}(K) = \operatorname{rank}(\bar{K}) \ge \operatorname{rank}(\Sigma_\rho) + \operatorname{rank}(\bar{U}_1) + \operatorname{rank}(\bar{V}_1)$, and so with probability 1 we have $\operatorname{rank}(K) \ge \rho + l_q + l_r = \min\{\rho+q+r, m+r, n+q, m+n-\rho\}$. Clearly $\operatorname{rank}(K) = \operatorname{rank}(\bar{K}) \le \min\{m+r, n+q\}$. By counting the numbers of linearly independent rows and columns of the matrix $\bar{K}$, we obtain that $\operatorname{rank}(\bar{K}) \le \rho+q+r$ with probability 1. Consequently $\operatorname{rank}(K) = \operatorname{rank}(\bar{K}) = \rho_+$ (proving part (ii) of the theorem) unless we have both bounds $q > l_q = m-\rho$ and $r > l_r = n-\rho$.

In the latter case we delete the first $\rho+q-m$ columns and the first $\rho+r-n$ rows of the matrix $\bar{K}$ and write $\bar{A}$ to denote the resulting trailing $(m+n-\rho)\times(m+n-\rho)$ submatrix $\bar{A}$ of the matrix $\bar{K}$. Now we apply part (ii) of Theorem 4.2 replacing the matrix A by Σ_A and replacing the matrix K by the matrix $\bar{A}$. Accordingly we replace the integers m by $m' = m+n-\rho-r$, n by $n' = m+n-\rho-q$, q by $q' = m-\rho$ and r by $r' = n-\rho$ and keep ρ unchanged. Now we have $q' = m'-\rho$ and $r' = n-\rho$, apply part (ii), and obtain that the submatrix $\bar{A}$ is nonsingular with probability 1.

To complete the proof of part (ii), it remains to deduce that with probability 1 the trailing $s\times s$ submatrix K' of the matrix $\bar{K}$ is nonsingular for $s = \min\{m+r, n+q\}$ provided that its submatrix $\bar{A}$ is nonsingular. We only need to prove that the Schur complement $S = W' - V'^T\bar{A}^{-1}U'$ of this matrix in the matrix K' is nonsingular with probability 1 where $U', V' \in \mathcal{G}_{0,1}^{(m+n-\rho)\times h}$, W' is the trailing $h\times h$ block of the matrix $\bar{W} = J_q\Sigma_W J_r$, and $h = \min\{m+r, n+q\} - (m+n-\rho)$. Clearly

the matrix S is nonsingular (that is has nonvanishing determinant) where the entries of the matrices U' and V' are indeterminates. By applying Lemma 3.2 we deduce that the matrix S is nonsingular with probability 1 for Gaussian random matrices U' and V'.

Toward proving part (iii), suppose $\mathrm{rank}(K) = n+q$, $l_r = n-\rho$, and $l_q = q$ and examine the leading $(n+q)\times(n+q)$ submatrix $\tilde{K}$ of the matrix $\bar{K}$. Truncate the expression (4.3) to obtain

$$\tilde{K} = \begin{pmatrix} W' & V_0^T & V_1^T \\ U_0 & \Sigma_\rho & O_{\rho,n-\rho} \\ U_1 & O_{n-\rho,\rho} & O_{n-\rho,n-\rho} \end{pmatrix} \tag{4.4}$$

where $U_0 = \bar{U}_0$, $U_1 \in \mathcal{G}_{0,1}^{q\times q}$, $V_0 \in \mathcal{G}_{0,1}^{l_r\times\rho}$, and $V_1 \in \mathcal{G}_{0,1}^{l_r\times l_r}$, $l_r = n-\rho$, U_1, V_0, and V_1 are submatrices of the matrices $\bar{U}_1$, $\bar{V}_0$, and $\bar{V}_1$, respectively. The matrix Σ_ρ is nonsingular, and with probability 1 so are the matrices U_1 and V_1 as well. Therefore $\mathrm{rank}(\tilde{K}) = n+q = \mathrm{rank}(K)$ with probability 1, and so, by virtue of Lemma 2.2, we have $\sigma_j(K) = \sigma_j(\bar{K}) \geq \sigma_j(\tilde{K})$ for all j, and in particular for $j = n+q$. Consequently $||\tilde{K}||^{-1} = \kappa(\tilde{K})/||\tilde{K}|| = 1/\sigma_{n+q}(\tilde{K}) \geq 1/\sigma_{n+q}(K) = \kappa(K)/||K||$.

Suppose the matrix $\tilde{K}$ is indeed nonsingular and readily deduce from (4.4) that

$$\tilde{K}^{-1} = \begin{pmatrix} O_{l_q,l_r} & O_{l_q,\rho} & -U_1^{-1} \\ O_{\rho,l_r} & \Sigma_\rho^{-1} & -\Sigma_\rho^{-1}U_0U_1^{-1} \\ V_1^{-T} & -V_1^{-T}V_0^T\Sigma_\rho^{-1} & V_1^{-T}(\bar{W} + V_0^T\Sigma_\rho^{-1}U_0)U_1^{-1} \end{pmatrix}.$$

Deduce from this equation and from (2.2) that $\kappa(K)/||K|| \leq N_1+N_2+N_3$ for N_1, N_2, and N_3 of part (iii) of the theorem because $||W'|| = ||W||$, $||\Sigma_\rho^{-1}|| = 1/\sigma_\rho(A)$, $||V_0^T|| = ||V_0||$, and $||V_1^{-T}|| = ||V_1^{-1}||$. Therefore part (iii) of the theorem follows provided $\mathrm{rank}(K) = n+q$ and $l_r = n-\rho$.

Similarly we deduce part (iii) in the two remaining cases, that is where $\mathrm{rank}(K) = m+r$ and $l_q = m-\rho$ and where $\mathrm{rank}(K) = \rho+q+r$, $l_r = n-\rho$, and $l_q = m-\rho$. □

Corollary 4.1. *Keep the assumptions of Theorem 4.2, except that now suppose that K is a square matrix, whereas the matrix A has numerical rank ρ, rather than rank ρ. Then one can extend accordingly all the estimates of Theorem 4.2 to the norm and the numerical rank of the matrix K and up to a factor 1.5 to its condition number.*

Proof. We immediately extend part (i) of Theorem 4.2. To extend parts (ii) and (iii) truncate the SVD of the matrix A by setting to 0 all its singular values except for the ρ largest ones. Obtain the matrices $A-E \approx A$ of the rank ρ and $\widehat{K} = \begin{pmatrix} W & V^T \\ -U & A-E \end{pmatrix}$ such that $||\widehat{K}-K|| \leq ||E|| = \sigma_{\rho+1}(A)$. The value $\sigma_{\rho+1}(A)$ is small because the matrix A has numerical rank ρ. Therefore the estimates of part (ii) of Theorem 4.2 for the rank of the matrix K are extended into the estimates for the numerical rank of the perturbed matrix.

Let us finally extend part (iii). It is likely that $||E\widehat{K}^{-1}|| \leq 1/3$ by virtue of Theorem 2.3, because the norm $||E||$ is small, whereas the random value $||\widehat{K}^{-1}||$ is not likely to be large by virtue of Theorem 4.2. It follows that probabilistically $||K^{-1}|| \leq 1.5||\widehat{K}^{-1}||$. Consequently, up to a factor 1.5, we extend the bound of part (iii) of Theorem 4.2 on the ratio $\kappa(K)/||K||$. □

We can relax the assumption of the corollary that K is a square matrix by applying the weaker estimates of Theorem 2.2 instead of the stronger ones of Theorem 2.3, but we can avoid this deterioration by combining Corollary 4.1 and Lemma 2.2 to obtain the following result.

Corollary 4.2. *Keep the assumptions of Theorem 4.2, except that now suppose that the matrix A has numerical rank ρ, rather than rank ρ. Then one can extend accordingly all the estimates of Theorem 4.2 to the norm and the numerical rank of the matrix K and up to a factor 1.5 to its condition number.*

Proof. Part (i) of Theorem 4.2 is extended immediately. To extend parts (ii) and (iii) apply Corollary 4.1 to the square submatrices of the input matrix K having the largest size; for extending part (iii) also apply Lemma 2.2. □

4.3 Condition estimates for full rank matrices

Next we show an alternative derivation of the probabilistic upper bound of order $1/\sigma_{m-q}(A)$ on the condition number $\kappa(K)$ provided that randomized augmentation produces a matrix K of full rank. We begin with the important special case where $K = (-U \mid A)$.

Theorem 4.3. *Keep the assumptions of Theorem 4.2, but restrict them to the case where $r = 0$, $K = (-U \mid A)$, and $m = \mathrm{rank}(A) \le n + q$. Then $1/\sigma_m(K) \le \sigma_{m-q}(A)/((1 + ||U_1||)||U_0^{-1}||)$ where $U_0 \in \mathcal{G}_{0,1}^{(m-q)\times q}$ and $U_1 \in \mathcal{G}_{0,1}^{q\times q}$.*

Proof. At first reduce the original task to the case of an $m \times m$ matrix $\bar{K}$, which is nonsingular with probability 1 and for which we have $\sigma_m(K) = \sigma_m(\bar{K})$; then estimate the value $\sigma_m(\bar{K})$ as the reciprocal $1/||\bar{K}^{-1}||$. Namely assume the SVD $A = S_A \Sigma_A T_A^T$ of (2.3) and write $K' = S_A^T K \,\mathrm{diag}(I_q, T_A) = (U' \mid \Sigma_A)$. Note that S_A, T_A and $\mathrm{diag}(T_A, I_q)$ are square orthogonal matrices and infer that $\sigma_{l'}(K) = \sigma_m(K')$ by virtue of Lemma 2.1, whereas $U' = S_A^T U \in \mathcal{G}_{0,1}^{m\times q}$ by virtue of Lemma 3.1. The $m \times (n+q)$ matrix K' has the $m \times m$ leading submatrix $\bar{K} = \begin{pmatrix} U_0 & \Sigma_{m-q} \\ U_1 & O_{q,m-q} \end{pmatrix}$ where $U_0 \in \mathcal{G}_{0,1}^{(m-q)\times q}$, $U_1 \in \mathcal{G}_{0,1}^{q\times q}$, $\mathrm{rank}(\bar{K}) = \mathrm{rank}(K) = m$, $\Sigma_{m-q} = \mathrm{diag}(\sigma_j(A))_{j=1}^{m-q}$, and so $\mathrm{rank}(\Sigma_{m-q}) = m - q$ and $\sigma_{l'}(K) = \sigma_m(K) \ge \sigma_m(\bar{K}) = 1/||\bar{K}^{-1}||$. We have

$$\bar{K}^{-1} = \begin{pmatrix} O_{q,m-q} & U_1^{-1} \\ \Sigma_{m-q}^{-1} & -\Sigma_{m-q}^{-1} U_0 U_1^{-1} \end{pmatrix} = \mathrm{diag}(I_q, \Sigma_{m-q}^{-1}) \begin{pmatrix} O_{q,n} & I_q \\ I_{m-q} & -U_0 \end{pmatrix} \mathrm{diag}(I_{m-q}, U_1^{-1}).$$

Therefore $||\bar{K}^{-1}|| \le ||\Sigma_{m-q}^{-1}||(1 + ||U_0||)||U_1^{-1}|| = ||(1 + ||U_0||)||U_1^{-1}||/\sigma_{m-q}(A)$. □

Theorems 3.1 and 3.2 together bound the norms $||U_1^{-1}||$ and $||U_0||$, and consequently the value $1/\sigma_{l'}(K) = ||\bar{K}^{-1}||$ is likely to have at most order $1/\sigma_{m-q}(A)$.

Theorem 4.4. *Under the assumptions of Theorem 4.2 suppose that the matrix K has full rank with probability 1, that is $l' = \min\{m + r, n + q\} \le \rho + q + r$. Write $\rho' = \min\{m - q, n - r\}$. Then*

$$\kappa(K) \le ||K||\, ||B_1^{-1}||(1 + ||B_0||)||U_1^{-1}||(||W|| + ||U_0||)/\sigma_{\rho'}(A)$$

where $B_0 \in \mathcal{G}_{0,1}^{r\times m}$, $B_1 \in \mathcal{G}_{0,1}^{r\times r}$, $U_0 \in \mathcal{G}_{0,1}^{(m-q)\times q}$, $U_1 \in \mathcal{G}_{0,1}^{q\times q}$ if $\rho' = m - q$, whereas $B_0 \in \mathcal{G}_{0,1}^{q\times n}$, $B_1 \in \mathcal{G}_{0,1}^{q\times q}$, $U_0 \in \mathcal{G}_{0,1}^{(n-r)\times r}$, $U_1 \in \mathcal{G}_{0,1}^{r\times r}$ if $\rho' = n - r$.

Proof. We can assume that $l' = m + r = \mathrm{rank}(K)$ and $\rho' = m - q$ (with probability 1). Otherwise we can prove the theorem for the transposed matrix A^T replacing A.

At first reduce the task to the case of a matrix $\widehat{K} \in \mathbb{R}^{l'\times l'}$ such that $\sigma_j(K) \ge \sigma_j(\widehat{K})$ for all j. Then estimate the value $\sigma_{l'}(\widehat{K})$ as the reciprocal $1/||\widehat{K}^{-1}||$. Namely write $K = \begin{pmatrix} B \\ F \end{pmatrix}$ where $B = (W \mid V^T)$, $F = (-U \mid A)$ and, by virtue of Theorem 4.3, $||F^+|| \le ||(1 + ||U_0||)||U_1^{-1}||/\sigma_{m-q}(A)$. Let $F = S_F \Sigma_F T_F^T$ be SVD and write $K'' = \mathrm{diag}(I_r, S_F^T) K T_F = \begin{pmatrix} B_0 & B_1 \\ \widehat{\Sigma}_F & O_{m,n+q-m} \end{pmatrix}$ where $B_0 \in \mathcal{G}_{0,1}^{r\times m}$, $B_1 \in \mathcal{G}_{0,1}^{r\times(n+q-m)}$, $\widehat{\Sigma}_F = \mathrm{diag}(\sigma_j(F))_{j=1}^m$, $\mathrm{rank}(K'') = \mathrm{rank}(K)$, and so with probability 1 the matrix $\widehat{\Sigma}_F$ is nonsingular and $||\widehat{\Sigma}_F^{-1}|| = ||F^+||$. We have $\sigma_j(K'') = \sigma_j(K)$ for all j because the matrices $\mathrm{diag}(I_s, S_F^T)$ and T_K are square and orthogonal.

Delete the last $n + q - l'$ columns of the matrix K'' and obtain the $l' \times l'$ submatrix $\widehat{K} = \begin{pmatrix} B_0 & \bar{B}_1 \\ \widehat{\Sigma}_F & O_{m,n+q-m} \end{pmatrix}$. We have $\sigma_{l'}(K) = \sigma_{l'}(K'') \ge \sigma_{l'}(\widehat{K})$. The Gaussian random $r \times r$ matrix $\bar{B}_1$ is nonsingular with probability 1. We assume that it is nonsingular, and then so is the matrix $\widehat{K}$ as well, and consequently $\sigma_{l'}(K) = \sigma_{l'}(\widehat{K}) = 1/||\widehat{K}^{-1}||$. Observe that

$$\widehat{K}^{-1} = \begin{pmatrix} O_{m,n+q-m} & \widehat{\Sigma}_F^{-1} \\ \bar{B}_1^{-1} & -\bar{B}_1^{-1} B_0 \widehat{\Sigma}_F^{-1} \end{pmatrix} = \mathrm{diag}(I_q, \bar{B}_1^{-1}) \begin{pmatrix} O_{q,n} & I_q \\ I_{m-q} & -B_0 \end{pmatrix} \mathrm{diag}(I_{m-q}, \Sigma_F^{-1}).$$

Therefore $||\widehat{K}^{-1}|| \le ||\bar{B}_1^{-1}||(1+||B_0||)||\widehat{\Sigma}_F^{-1}||$. To complete the proof substitute the upper bound of Theorem 4.3 on $||\widehat{\Sigma}_F^{-1}|| = ||F^+||$. □

Similarly to Corollary 4.2 in the previous subsection we deduce the following extension of Theorem 4.4.

Corollary 4.3. *Keep the assumptions of Theorem 4.4, except that now suppose that the matrix A has numerical rank ρ, rather than rank ρ. Then one can extend accordingly all the estimates of Theorem 4.4 to the norm and the numerical rank of the matrix K and up to a factor 1.5 to its condition number.*

4.4 A randomized Toeplitz solver

Let us apply Theorem 2.7 to support randomized augmentation for solving a nonsingular Toeplitz linear system $T\mathbf{y} = \mathbf{b}$ of n equations provided the matrix T has numerical nullity 1.

To compute the vector $\mathbf{y} = T^{-1}\mathbf{b}$, we first embed the matrix T into a Toeplitz $(n+1)\times(n+1)$ matrix $K = \begin{pmatrix} w & \mathbf{v}^T \\ \mathbf{f} & T \end{pmatrix}$. We write $w = \mathbf{e}_1^T T \mathbf{e}_1$ and fill the vectors $\mathbf{f} = (f_i)_{i=1}^n$ and $\mathbf{v} = (v_i)_{i=1}^n$ with appropriate entries of the matrix T except for the two coordinates f_n and v_n, which we choose at random and then scale to have the ratios $\frac{|f_n|}{||K||}$ and $\frac{|v_n|}{||K||}$ neither large nor small.

Part (b) of Theorem 2.7 expresses the inverse T^{-1} via the vectors $\mathbf{v} = K^{-1}\mathbf{e}_1$ and $\mathbf{w} = K^{-1}\mathbf{e}_{n+1}$.

In view of Section 3.2 and Appendix A, this policy is likely to produce a nonsingular matrix K whose inverse is likely to have a nonzero entry $\mathbf{e}_1^T K^{-1}\mathbf{e}_1$. In good accordance with these formal results our tests have always produced nonsingular and well conditioned matrices K such that $\mathbf{e}_1^T K^{-1}\mathbf{e}_1 \neq 0$.

To summarize, we reduce the solution of a nonsingular ill conditioned Toeplitz linear system $T\mathbf{y} = \mathbf{b}$ to computing highly accurate solutions of two linear systems $K\mathbf{x} = \mathbf{e}_1$ and $K\mathbf{z} = \mathbf{e}_{n+1}$, both expected to be well conditioned. High accuracy shall counter the magnification of the input and rounding errors, expected to occur in the case of ill conditioned input.

In the important special case where a Toeplitz matrix T is real symmetric, we choose real scalars w and $f_n = v_n$ to yield a real symmetric matrix $K = \begin{pmatrix} w & \mathbf{v}^T \\ \mathbf{v} & T \end{pmatrix}$. In this case $J_{n+1}K^{-1}J_{n+1} = K^{-1}$, and so $K^{-1}\mathbf{e}_{n+1} = J_{n+1}K^{-1}\mathbf{e}_1$ because $J_{n+1}\mathbf{e}_{n+1} = \mathbf{e}_1$. Thus we only need to solve a single linear system with the matrix K. For the transition back to the solution of the original problem, we can employ expression (4.2) or Theorem 2.7. Hereafter we refer to the resulting algorithm for the linear system $T\mathbf{y} = \mathbf{b}$ as **Algorithm 6.1**. In Section 8.4 we test this algorithm for solving an ill conditioned real symmetric Toeplitz linear system.

One can readily extend the approach of this section to the case of Toeplitz-like, Hankel and Hankel-like inputs and to augmenting the input matrix with r rows and r columns for $r > 1$.

5 Randomized additive preconditioning

In this section we study randomized additive preprocessing $A \to C = A + UMV^T = A + UW^{-1}V^T$ and its impact on the condition number of a matrix A provided that U and V are Gaussian random matrices and $M = W^{-1}$ is a nonsingular matrix.

5.1 Link between additive preprocessing and augmentation

The following simple theorem links additive preprocessing $A \Longrightarrow C = A + UW^{-1}V^T$ and the augmentation $A \Longrightarrow K$ for the matrix K of (4.1).

Theorem 5.1. *Suppose $A \in \mathbb{R}^{m\times n}$, $W \in \mathbb{R}^{r\times r}$, $U \in \mathbb{R}^{m\times r}$, $V \in \mathbb{R}^{n\times r}$, the matrix W is nonsingular, $l = \min\{m,n\}$, a matrix K in $\mathbb{R}^{(m+r)\times(n+r)}$ is defined by (4.1), and $C = A + UW^{-1}V^T$. Then*

$$K = \widehat{U}\,\mathrm{diag}(C, I_r)\widehat{V}\,\mathrm{diag}(W, I_n) \tag{5.1}$$

for $\widehat{U} = \begin{pmatrix} O_{r,m} & I_r \\ I_m & -UW^{-1} \end{pmatrix}$, $\widehat{V} = \begin{pmatrix} O_{n,r} & I_n \\ I_r & V^T \end{pmatrix}$, $\widehat{U}^{-1} = \begin{pmatrix} UW^{-1} & I_n \\ I_r & O_{r,n} \end{pmatrix}$, $\widehat{V}^{-1} = \begin{pmatrix} -V^T & I_r \\ I_n & O_{n,r} \end{pmatrix}$. *Furthermore the matrix* C *has full rank if and only if the matrix* K *has full rank, and both matrices are rank deficient if* $r + \mathrm{rank}(A) < l$. *If* $m = n$ *and if the matrices* C *and* K *are nonsingular, then* $C^{-1} = (I_n \mid O_{n,r})\widehat{V}\,\mathrm{diag}(W, I_n)K^{-1}\widehat{U}(I_n \mid O_{n,r})^T$ *and* $K^{-1} = \mathrm{diag}(W^{-1}, I_n)\widehat{V}^{-1}\,\mathrm{diag}(C^{-1}, I_r)\widehat{U}^{-1}$.

Remark 5.1. *Assume a symmetric and positive definite matrix* K. *Then the Interlacing Property of its eigenvalues [GL96, Theorem 8.6.3] implies that the map* $A \Longrightarrow K = \begin{pmatrix} W & V^T \\ V & A \end{pmatrix}$ *cannot decrease the condition number* $\kappa(A)$, *whereas the opposit is true for the scaled randomized symmetric additive preprocessing* $A \Longrightarrow C = A + VV^T$ *(cf. (5.6) for* $U = V$ *and [W07]). This limitation of the link between augmentation and additive preprocessing, however, is minor as we will see next.*

5.2 Estimating the norms, ranks, numerical ranks, and condition numbers via the link to augmentation

The following norm bounds for additive preprocessing are obvious.

Theorem 5.2. *Given five matrices* A, C, U, M, *and* V *such that* $C = A + UMV$, *we have* $||A|| - ||U||\,||M\,||V|| \le ||C|| \le ||A|| + ||U||\,||M||\,||V||$.

Theorem 5.1 enables us to extend the results and techniques from the study of randomized augmentation to randomized additive preprocessing and vice versa. In particular we combine Theorems 4.2 and 4.4 for $r = q$ with Theorem 5.1 to obtain the following corollary.

Corollary 5.1. *Assume four matrices* $A \in \mathbb{R}^{m\times n}$, $W \in \mathbb{G}_{0,1}^{r\times r}$, $U \in \mathbb{G}_{0,1}^{m\times r}$, *and* $V \in \mathbb{R}^{n\times r}$, *where the matrix* W *is nonsingular, and* A *is a real normalized* $m \times n$ *matrix having a rank* ρ. *Write* $C = A + UW^{-1}V^T$. *Then with probability* 1 *(i) the matrix* C *has rank* $\rho' = \min\{m, n, \rho + r\}$. *Furthermore suppose* $\rho' = m = n$ *with probability* 1. *Then (ii)* $\kappa(C) \le ||C||\,||B_1^{-1}||\,||B_1^{-1}||(1 + ||B_0||)||U_1^{-1}||(||W|| + ||U_0||)/\sigma_{\rho'}(A)$ *where* $B_0 \in \mathcal{G}_{0,1}^{r\times m}$, $B_1 \in \mathcal{G}_{0,1}^{r\times r}$, $U_0 \in \mathcal{G}_{0,1}^{(m-r)\times r}$, *and* $U_1 \in \mathcal{G}_{0,1}^{r\times q}$.

Suppose the ratios $||W^{-1}||/||A||$ and $||W||/||A||$ are neither large nor small, apply Theorems 3.1 and 3.2 to estimate the norms of the matrices B_0, U_0, B_1^{-1}, and U_1^{-1}, and deduce from the corollary that the ratio $\kappa(C)/||C||$ is likely to have at most order $1/\sigma_{\rho'}(A)$.

Similarly to Corollaries 4.2 and 4.1 in the previous section we deduce the following extension of Corollary 5.1.

Corollary 5.2. *Keep the assumptions of Corollary 5.1, except that now suppose that the matrix* A *has a numerical rank* ρ, *rather than a rank* ρ. *Then one can extend accordingly all the estimates of Corollary 5.1 to the norm and the numerical rank of the matrix* K *and up to a factor 1.5 to its condition number.*

5.3 An alternative way to estimating the condition numbers

In this subsection we describe an alternative way to estimating the impact of randomized additive preprocessing on the condition number $\kappa(K)$. For simplicity we let $C = A + UV^T$. The estimated ratio $\kappa(C)/||C||$ is still likely to have at most order $1/\sigma_{\rho'}(A)$ provided that $||A|| = 1$ and U and V are standard Gaussian random matrices.

Theorem 5.3. *Suppose* $A, C, S, T \in \mathbb{R}^{n\times n}$ *and* $U, V \in \mathbb{R}^{n\times r}$ *for two positive integers* r *and* n, $r \le n$, $A = S\Sigma T^T$ *is full SVD of the matrix* A *(cf. (2.3)),* S *and* T *are square orthogonal matrices,* $\Sigma = \mathrm{diag}(\sigma_j)_{j=1}^n$, *the matrix* $C = A + UV^T$ *is nonsingular, and so* $\rho = \mathrm{rank}(A) = n - r$ *and* $\sigma_\rho > 0$. *Write*

$$S^T U = \begin{pmatrix} \bar{U} \\ U_r \end{pmatrix},\ T^T V = \begin{pmatrix} \bar{V} \\ V_r \end{pmatrix},\ R_U = \begin{pmatrix} I_\rho & \bar{U} \\ O_{r,\rho} & U_r \end{pmatrix},\ R_V = \begin{pmatrix} I_\rho & \bar{V} \\ O_{r,\rho} & V_r \end{pmatrix}, \tag{5.2}$$

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

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

$$C = S R_U D R_V^T T^T,\ D = \Sigma + \operatorname{diag}(O_{\rho,\rho}, I_r) = \operatorname{diag}(d_j)_{j=1}^n \tag{5.3}$$

where $d_j = \sigma_j$ *for* $j = 1, \dots, \rho$, $d_j = 1$ *for* $j = \rho+1, \dots, n$.

Furthermore suppose that the matrix A *has been normalized so that* $||A|| = 1$ *and that the* $r \times r$ *matrices* U_r *and* V_r *are nonsingular, which holds with probability* 1 *where* U *and* V *are Gaussain random matrices (cf. Section 3.2). Write*

$$p = ||R_U^{-1}||\ ||R_V^{-1}|| \text{ and } f_r = \max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}. \tag{5.4}$$

Then

(b) *the matrix* C *is nonsingular,*

(c) $1 \le \sigma_\rho(A)/\sigma_n(C) \le p$,

(d) $p \le (1 + ||U||)(1 + ||V||) f_r$.

Proof. Parts (a) and (b) are readily verified.

(c) Combine the equations $S^{-1} = S^T$, $T^{-1} = T^T$ and (5.3) and obtain $C^{-1} = T R_V^{-T} D^{-1} R_U^{-1} S^T$. It follows that $||C^{-1}|| = ||R_V^{-T} D^{-1} R_U^{-1}||$. Apply bound (2.1), substitute $||S|| = ||S^T|| = ||T|| = ||T^T|| = 1$ and obtain $||C^{-1}|| \le ||R_V^{-T}||\ ||D^{-1}||\ ||R_U^{-1}||$. Substitute the equations (5.4), $||D^{-1}|| = 1/\sigma_\rho(A)$ (implied by the equations $||A|| = 1$ and (5.3)), and $||C^{-1}|| = 1/\sigma_n(C)$ and obtain that $\sigma_\rho(A)/\sigma_n(C) \le p$. Next deduce from (5.2) and (5.3) that

$$R_V^{-T} = \begin{pmatrix} I_\rho & O_{\rho,r} \\ -V_r^{-T} \bar{V}^T & V_r^{-T} \end{pmatrix}, D^{-1} = \Sigma^{-1} + \operatorname{diag}(O_{\rho,\rho}, I_r),\ R_U^{-1} = \begin{pmatrix} I_\rho & -\bar{U} U_r^{-1} \\ O_{r,\rho} & U_r^{-1} \end{pmatrix}, \tag{5.5}$$

substitute these expressions into the matrix product $R_V^{-T} D^{-1} R_U^{-1}$, and obtain that $R_V^{-T} D^{-1} R_U^{-1} = \begin{pmatrix} \Sigma^{-1} & X \\ Y & Z \end{pmatrix}$ and consequently $1/\sigma_n(C) = ||C^{-1}|| = ||R_V^{-T} D^{-1} R_U^{-1}|| \ge ||\Sigma^{-1}|| = 1/\sigma_\rho(A)$. This completes the proof of parts (c).

(d) 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. □

Corollary 5.3. *Under the assumptions of Theorem 5.3 we have the following bound,*

$$\frac{\kappa(C)}{\kappa(A)} \le (1 + ||U||\ ||V||)(1 + ||U||)(1 + ||V||)\ \max\{1, ||U_r^{-1}||\}\ \max\{1, ||V_r^{-1}||\}, \tag{5.6}$$

which turn into the bound $\frac{\kappa(C)}{\kappa(A)} \le (1 + ||U||^2)(1 + ||U||)^2 \max\{1, ||U_r^{-1}||^2\}$ *if* $U = V$.

Proof. We have $\frac{\kappa(C)}{\kappa(A)} \le \frac{||C||}{||A||} \frac{\sigma_\rho(A)}{\sigma_n(C)}$, $||A|| = 1$, $||C|| \le ||A|| + ||U||\ ||V||$, and so parts (c) and (d) together imply the corollary. □

Suppose U and V are standard Gaussian random matrices. Then so are the matrices U_r and V_r as well by virtue of Lemma 3.1, because S and T are square orthogonal matrices. In this case we can estimate the norms $||U||$, $||V||$, $||U_r^{-1}||$, and $||V_r^{-1}||$ by applying Theorems 3.1 and 3.2 to prove that with probability near 1 the condition numbes of the matrices A and C have the same order.

Corollary 5.4. *The bound of Corollary 5.3 increases by at most a factor 1.5 where the matrix* A *has a numerical rank* ρ *rather than a rank* ρ.

Proof. Reapply the techniques of the proofs of Corollaries 4.2 and 4.3. □

Corollary 5.5. *The estimates of Corollaries 5.3 and 5.4 can be extended to the case where* A *and* C *are* $m \times n$ *rectangular matrices,* $||A|| = 1$, $U \in G_{0,1}^{m \times r}$, $V \in G_{0,1}^{n \times r}$, *and* $\operatorname{rank}(A) \ge l - r$ *for* $l = \min\{m, n\}$.

Proof. Reapply the techniques of the proof of Corollary 4.2. □

5.4 Can we weaken randomness?

Would the results of this section and of the two previous ones still hold if we weaken randomness of the matrices U and V by allowing them to be sparse and structured, to share some or all their entries, or generally to be defined by a smaller number of independent parameters, possibly under other probability distributions rather than Gaussian? We have some progress with our analytical study in this direction (see [PQa]), but empirically all the presented randomized techniques remain as efficient where randomization is very weak in the above sense (see Sections 4.4, 8.1, and 8.4 and Tables 8.1, 8.4, and 8.8). Similar phenomenon has been observed in the study of randomized matrix multiplication [T11], [PQZa], [PQb].

6 The computation of nmbs and the bases for the trailing singular spaces

At first we apply randomized additive preprocessing and augmentation to compute a nmb of a rank deficient matrix A having a rank ρ, and then extend the techniques to computing an approximate matrix basis B for the trailing singular space $\mathbb{T}_{A,\rho}$ of an ill conditioned matrix A having numerical rank ρ. This provides an alternative to the efficient methods of Section 2.6 for computing approximate matrix bases for leading and trailing singular spaces of the latter matrix. Clearly we can compute left nmbs and approximate matrix bases for the left trailing singular spaces by applying the same algorithms to the matrix A^T, and we can obtain approximate matrix bases for the leading singular spaces as nmbs of such bases for its trailing singular spaces. Theorem 2.5 enables us to extend these computations to obtain a rank-ρ approximation of a matrix A.

Theorem 6.1. *[PQ10, Theorem 3.1 and Corollary 3.1]. Suppose a matrix $A \in \mathbb{R}^{m\times n}$ has a rank ρ, $U \in \mathbb{R}^{m\times r}$, $V \in \mathbb{R}^{n\times r}$, and the matrix $C = A + UV^T$ has full rank n. Write $B = C^{(I)}U$. Then $r \geq n - \rho$, $\mathcal{R}(B) \supseteq \mathcal{N}(A)$; moreover if $r = n - \rho$, then $C^{(I)}U = \mathrm{nmb}(A)$. Furthermore $\mathcal{R}(BX) = \mathcal{N}(A)$ if $\mathcal{R}(X) = \mathcal{N}(AB)$. (Note that $AB = U(I_r V^T C^{-1} U)$ for $m = n$.)*

Theorem 6.2. *Suppose that $A \in \mathbb{R}^{m\times n}$, $U \in \mathbb{R}^{m\times q}$, $V \in \mathbb{R}^{n\times s}$, $W \in \mathbb{R}^{s\times q}$, $K = \begin{pmatrix} W & V^T \\ -U & A \end{pmatrix}$, $\mathrm{rank}(W) = q \geq \mathrm{nul}(A)$, $\mathrm{rank}(K) = n + q$, $m \geq n$. Write $Y = (O_{n,q} \mid I_n)K^{(I)} \begin{pmatrix} O_{s,q} \\ U \end{pmatrix}$. Then*

(a) $\mathcal{N}(A) \subseteq \mathcal{R}(Y)$ and if $\mathrm{rank}(U) = \mathrm{nul}(A)$, then $\mathcal{N}(A) = \mathcal{R}(Y)$,
(b) $\mathcal{R}(YZ) = \mathcal{N}(A)$ if $\mathcal{R}(Z) = \mathcal{N}(AY)$, whereas
(c) $\mathcal{R}(Z) = \mathcal{N}(AY)$ if $\mathcal{R}(YZ) = \mathcal{N}(A)$ and if $\mathrm{rank}(Y) = q$.

Proof. See [PQ12, Theorems 11.2 and 11.3]. □

Remark 6.1. *Both theorems define* aggregation processes *(cf. [MP80]). For $r > n - \rho$, Theorem 6.1 reduces the computation of a $\mathrm{nmb}(A)$ to the same nmb task, but for the input BX of a smaller size $n \times (r - n + \rho)$. Furthermore, suppose that the matrices U and Y have full rank q. Then part (a) of Theorem 6.2 implies that Y is a $\mathrm{nmb}(A)$ if $q = \mathrm{nul}(A)$, but otherwise parts (b) and (c) reduce the original task of computing a nmb(A) to the case of the input AY of a smaller size $m \times (q - \mathrm{nul}(A))$.*

Theorem 6.3. *Assume that $U \in \mathbb{R}^{m\times r_+}$, $V \in \mathbb{R}^{n\times r_+}$, $m \geq n$, a real $m \times n$ matrix A has a numerical rank $\rho = n - r$, and the matrix $C = A + UV^T$ has full rank and is well conditioned. Then $\rho \geq n - r_+$ and there is a scalar c independent of A, U, V, m, n and ρ such that $||C^+UX - T_{A,r}|| \leq c\sigma_{\rho+1}(A)||U||$ where $X = I_r$ if $r = r_+$ and otherwise $X \in \mathbb{R}^{r_+\times r}$, $X = \mathrm{nmb}(AC^+U + \Delta)$, $||\Delta|| \leq c\sigma_{\rho+1}(A)||U||$.*

Proof. The theorem turns into Theorem 6.1 if $\rho = \mathrm{nrank}(A) = \mathrm{rank}(A)$. If $\rho = \mathrm{nrank}(A) < \mathrm{rank}(A)$, set to zero all but the ρ largest singular values in the SVD of the matrix A. Then $\rho = \mathrm{nrank}(A - E) = \mathrm{rank}(A - E)$ and the theorem holds for the resulting matrix $A - E$ and the matrix $C - E = A - E + UV^T$. Therefore $T_{A-E,r} = (C - E)^+UX$ where for X we choose an orthogonal $\mathrm{nmb}((A - E)(C - E)^+U)$, of size $r_+ \times r$. Clearly $||T_{A-E,r} - T_{A,r}Q|| = O(\sigma_{\rho+1}(A))$ for some $r \times r$ orthogonal matrix Q, and it remains to estimate the norm $||(C - E)^+UX - C^+UX||$. We have

$||((C-E)^+ - C^+)UX|| \le ||(C-E)^+ - C^+||\ ||U||$. The norm $||E|| = \sigma_{\rho+1}(A)$ is small because the matrix A has numerical rank ρ, whereas the norm $||C-E)^+||$ is not large because the matrix C has full rank and is well conditioned. Therefore the value $\tau = ||C-E)^+|| - ||C^+||$ has at most order $\sigma_{\rho+1}(A)$ by virtue of Theorem 2.2. □

Corollary 6.1. *Suppose a normalized real $m \times n$ matrix A has a numerical rank $\rho = n - r$, $U \in \mathcal{G}_{0,1}^{m\times r_+}$, $V \in \mathcal{G}_{0,1}^{n\times r_+}$, $m \ge n$, and $C = A + UV^T$. Then (i) the matrix C is singular or ill conditioned if $r_+ < r$ but otherwise (ii) has full rank with probability 1, and (iii) we can expect that the matrix C^+UX is an approximate matrix basis for the singular space $\mathbb{T}_{A,r}$ within an error norm of at most order $\sigma_{\rho+1}(A)$ where $X = I_r$ if $r = r_+$ and otherwise X is an orthogonal $\mathrm{nmb}(AC^+U+\Delta)$ of the size $r_+ \times r$ and $||\Delta|| \le c\sigma_{\rho+1}(A)$.*

Proof. Part (i) is immediately verified. Furthermore by virtue of Corollaries 5.2 and 5.5 the matrix C has full rank with probability 1 and is expected to be well conditioned, whereas the norm $||U||$ is expected to be not large by virtue of Theorem 3.2. Therefore Corollary 6.1 follows from Theorem 6.3. □

Likewise by employing Theorems 4.4 and 6.2 instead of Theorem 6.1 and Corollaries 5.2 and 5.5, we obtain the following result.

Corollary 6.2. *Suppose that a normalized real $m \times n$ matrix A has a numerical nullity $r = \mathrm{nnul}(A)$, $U \in \mathcal{G}_{0,1}^{m\times q}$, $V \in \mathcal{G}_{0,1}^{n\times s}$, $W \in \mathcal{G}_{0,1}^{s\times q}$, $K = \begin{pmatrix} W & V^T \\ -U & A \end{pmatrix}$, $\mathrm{rank}(W) = q$, $\mathrm{rank}(K) = n + q$, $Y = (O_{n,q} \mid I_n)K^+ \begin{pmatrix} O_{s,q} \\ U \end{pmatrix}$, and $m \ge n$. Then (i) the matrix K is rank deficient or ill conditioned where $q < r$ but otherwise has full rank with probability 1 and is expected to be well conditioned. Furthermore we can expect that a matrix basis for the singular space $\mathbb{T}_{A,q}$ is approximated within an error norm of at most order $\sigma_{n-q+1}(A)$ by (ii) the matrix Y if $r = q$ or (iii) the matrix YZ if $q > r$ where $Z \in \mathbb{R}^{q\times r}$, $Z = \mathrm{nmb}(AY + \Delta)$, and $||\Delta|| \le c\sigma_{n-q+1}(A)$.*

Corollaries 6.1 and 6.2 (for $s = q$) imply correctness of the following two Prototype Algorithms, which output the numerical nullity r of a matrix A and an approximate matrix basis of its trailing singular space $\mathbb{T}_{A,r}$.

Proto-Algorithm 6.1. An approximate basis for a trailing singular space by using randomized additive preprocessing.

INPUT: *A matrix $A \in \mathbb{R}^{m\times n}$ for $m \ge n$ with $||A|| \approx 1$, an upper bound r_+ on its unknown numerical nullity $r = \mathrm{nnul}(A)$, and two tolerances τ and τ' of order $\sigma_{n-r+1}(A)$. (The tolerances are defined by the requested output accuracy. In a variation of the algorithm one can reapply it with a decreased tolerance τ' instead of outputing FAILURE at Stage 4.)*

OUTPUT: *FAILURE (with a low probability) or the numerical nullity r and, within an error norm in $O(\sigma_{n-r+1}(A))$, an approximate matrix basis B of the trailing singular space $\mathbb{T}_{A,r}$.*

INITIALIZATION: *Generate two matrices $U \in \mathcal{G}_{0,1}^{m\times r_+}$ and $V \in \mathcal{G}_{0,1}^{n\times r_+}$ for σ of order $||A||$.*

COMPUTATIONS:

1. *Compute the matrix $C = A + UV^T$.*
2. *Stop and output FAILURE if this matrix is rank deficient or ill conditioned. Otherwise compute the matrices $Y = C^+U$ and AY.*
3. *Output $r = r_+$ and $B = Y$ and stop if $||AY|| \le \tau||A||\ ||Y||$.*
4. *Otherwise apply an algorithm (e.g. employing SVD, rank revealing factorization [GE96], [HP92], [P00a], a technique from [PQ10] or [PQ12], or one of Proto-Algorithms 6.1 and 6.2) that for the matrix AY and a fixed tolerance τ' computes an integer r and an orthogonal approximate matrix basis X (of size $r_+ \times r$) for the space $\mathbb{T}_{AY,r}$.*

5. *Compute the matrix $B = YX$. Output r and B and stop if $||AB|| \le \tau ||A||\ ||B||$. Otherwise output FAILURE and stop.*

Proto-Algorithm 6.2. An approximate basis for a trailing singular space by using randomized augmentation.

INPUT, OUTPUT *and Stages 3, 4 and 5 of* COMPUTATIONS *are as in Proto-Algorithm 6.1.*

INITIALIZATION*:* *Generate three matrices $U \in \mathcal{G}_{0,1}^{m\times r_+}$, $V \in \mathcal{G}_{0,1}^{n\times r_+}$, and $W \in \mathcal{G}_{0,1}^{r_+\times r_+}$ for σ of order $||A||$.*

COMPUTATIONS*:*

1. *Stop and output FAILURE if the matrix $K = \begin{pmatrix} W & V^T \\ -U & A \end{pmatrix}$ is rank deficient or ill conditioned.*
2. *Otherwise compute the matrices $Y = (O_{n,r_+} \mid I_n)K^+ \begin{pmatrix} O_{r_+,\ r_+} \\ U \end{pmatrix}$ and AY.*

Remark 6.2. *Approximation of the leading and trailing singular spaces as well as the computation of numerical rank and numerical nullity are facilitated as the gaps increase between the singular values of the input matrix A. This motivates using the power transforms $A \Longrightarrow B_h = (AA^T)^h A$ for positive integers h because $\sigma_j(B_h) = (\sigma_j(A))^{2h+1}$ for all j.*

Remark 6.3. *In the case where $m = n$ the computations are simplified and stabilized, and furthermore we can apply Theorem 6.1 or 6.2 to both A and A^T to define both left and right nmbs. We can reduce to this case the computation for a rectangular matrix A in various ways, e.g., by observing that (a) $\mathcal{N}(A) = \mathcal{N}(A^TA)$, (b) $\mathcal{N}(A) = \mathcal{N}(B^TA)$ if $A, B \in \mathbb{R}^{m\times n}$ and the matrix B has full rank $m \le n$, and (c) $(A \mid O_{n,m-n})\mathbf{u} = \mathbf{0}_m$ if and only if $A\widehat{\mathbf{u}} = \mathbf{0}_m$ provided $m \ge n$ and $\widehat{\mathbf{u}} = (I_n \mid O_{n,m-n})\mathbf{u}$, whereas $(A^T \mid O_{n,m-n})\mathbf{v} = \mathbf{0}_n$ if and only if $\widehat{\mathbf{v}} = \mathbf{0}_n^T$ provided $m < n$ and $\widehat{\mathbf{v}} = (I_m \mid O_{n-m,m})\mathbf{v}$. Furthermore given an $m \times n$ matrix A for $m > n$, we can represent it as the block vector $A = (B_1^T \mid B_2^T \mid \ldots \mid B_h^T)^T$ where B_i are $k_i \times n$ blocks for $i = 1, \ldots, h$, $\sum_{i=1}^h k_i = m$, and observe that $\mathcal{N}(A) = \cap_{i=1}^h \mathcal{N}(B_i)$, and we can compute the intersection of null spaces by applying [GL96, Theorem 12.4.1]. One can extend these comments to the tasks of the approximation of the singular spaces of ill conditioned matrices.*

Remark 6.4. *Let B be an approximate matrix basis of the trailing singular space $\mathbb{T}_{A,r}$ of an $m \times n$ matrix A that has a numerical nullity r. Then $B' = \mathrm{nmb}(B)$ is an approximate matrix basis of the leading singular space $\mathbb{T}_{n-r,A}$. Furthermore write $Q = Q(B')$ and deduce from Theorem 2.5 that the matrix $A - AQQ^T$ of the rank $n - r$ approximates the matrix A. Alternatively we can compute such bases B' and rank-$(n - r)$ approximations to the matrix A by applying randomized matrix multiplication [HMT11], [PQb]. The latter technique is particularly effective where the numerical rank $n - r$ of the matrix A is small.*

Remark 6.5. *If we are given the numerical nullity $r = \mathrm{nnul}(A)$ or the numerical rank $n - r = \mathrm{nrank}(A)$ of an input matrix A, we can substantially simplify Proto-Algorithms 6.1 and 6.2 by replacing their Stages 4 and 5 by the following stage:*

4. *Otherwise output FAILURE.*

Customary algorithms determine numerical rank by computing SVD or rank revealing factorization of an input matrix; in Appendix C we use randomization to avoid orthogonalization and pivoting.

7 2×2 block diagonalization of an ill conditioned matrix

Next we apply the results of the previous sections to compute 2×2 block diagonalization of a nonsingular ill conditioned matrix where we are given its numerical rank. In Remarks 7.2 and 7.3 we comment on some variations of this approach.

Having an ill conditioned input we must perform part of the computations with high or infinite precision to ensure obtaining a meaningful output. Accordingly we can partition the algorithms of this section into *symbolic* and *numerical* stages. At the symbolic stage we would perform computations with infinite precision, but this part of the computations would cover only a small fraction of all arithmetic operations involved (cf. Remark 7.1).

Proto-Algorithm 7.1. Randomized block diagonalization with orthogonalization.

INPUT: *A matrix $A \in \mathbb{R}^{n \times n}$ whose norm $||A||$ is neither large nor small, its numerical rank q satisfying $0 < q = n - r < n$, and a Subroutine* LIN·SOLVE *that either solves a linear system of equations if it is nonsingular and well condtioned or outputs FAILURE otherwise.*

OUTPUT: *FAILURE (with a low probability) or four orthogonal matrices K_0 and L_0 in $\mathbb{R}^{n \times q}$ and K_1 and L_1 in $\mathbb{R}^{n \times r}$ such that with a probability near 1 the $q \times q$ block submatrix $W_{00} = K_0^T A L_0$ of the matrix $W = (K_0 \mid K_1)^T A (L_0 \mid L_1) = \begin{pmatrix} W_{00} & W_{01} \\ W_{10} & W_{11} \end{pmatrix}$ is nonsingular, well conditioned, and strongly dominant such that $\sigma_q(W_{00}) \gg \max\{||W_{01}||,\ ||W_{10}||,\ ||W_{11}||\}$.*

COMPUTATIONS *(see Remark 7.1):*

1. *Generate two matrices $U, V \in \mathbb{G}_{0,1}^{n \times r}$.*
2. *Compute the matrix $C = A + UV^T$, expected to be nonsingular and well conditioned.*
3. *Apply the Subroutine* LIN·SOLVE *to compute the matrices $C^{-T}V$ and $C^{-1}U$. Stop and output FAILURE if so does the subroutine.*
4. *Compute and output two orthogonal matrices $K_1 = Q(C^{-1}U)$ and $L_1 = Q(C^{-T}V)$.*
5. *Compute and output two orthogonal nmbs $K_0 = \mathrm{nmb}(K_1)$ and $L_0 = \mathrm{nmb}(L_1)$.*

The algorithm can only fail with a low probability by virtue of Theorems 6.3 and Corollaries 6.1, 5.2 and 5.5. We prove correctness of the algorithm by applying the following theorem.

Theorem 7.1. *For a matrix $A \in \mathbb{R}^{m \times n}$ and $0 < q < l = \min\{n, m\}$, write $r = n - q$ and $\bar{r} = m - q$. Let $K_0 \in \mathbb{R}^{m \times q}$, $L_0 \in \mathbb{R}^{n \times q}$, $K_1 \in \mathbb{R}^{m \times \bar{r}}$, $L_1 \in \mathbb{R}^{n \times r}$, and $Q_K, Q_L \in \mathbb{R}^{r \times r}$ be six orthogonal matrices such that $K_1 = S_{A,\bar{r}} Q_K$, $L_1 = T_{A,r} Q_L$, $K_1^T K_0 = O_{\bar{r},q}$ and $L_1^T L_0 = O_{r,q}$. Then $||K_1^T A|| \le \sigma_{q+1}(A)$, $||AL_1|| \le \sigma_{q+1}(A)$, $||K_0 A L_0|| = \sigma_1(A)$, and $\kappa(K_0 A L_0) = \sigma_1(A)/\sigma_q(A)$.*

Proof. Suppose $A = S_A \Sigma_A T_A^T$ is SVD of (2.3). Then $AL_1 = S_A \Sigma_A T_A^T T_{A,r} Q_L = S_A \Sigma_A \begin{pmatrix} O_{n,q} \\ Q_L \end{pmatrix} = S_A \,\mathrm{diag}(O_{m-r,n-r}, (\sigma_j(A))_{j=q+1}^{n} Q_L)$, and so $||AL_1|| \le \sigma_{q+1}(A)$ because S_A and Q_L are orthogonal matrices. Similarly obtain that $||K_1^T A|| \le \sigma_{q+1}(A)$.

Next deduce from the assumptions about L_0 and L_1 that $L_0 = T_{q,A} Q_0'$ for an orthogonal matrix $Q_0' \in \mathbb{R}^{q \times q}$ and similarly that $K_0 = S_{q,A} Q_0$ for an orthogonal matrix $Q_0 \in \mathbb{R}^{q \times q}$. Therefore

$$K_0 A L_0 = Q_0' S_{q,A}^T S_A \Sigma_A T_A^T T_{q,A} Q_0 = (Q_0' \mid O_{m,\bar{r}}) \Sigma_A (Q_0 \mid O_{r,n})^T = Q_0' \,\mathrm{diag}(\sigma_j(A))_{j=1}^{q} Q_0,$$

and so $||K_0 A L_0|| = \sigma_1(A)$, $\kappa(K_0 A L_0) = \sigma_1(A)/\sigma_q(A)$. □

In Proto-Algorithm 7.1 we expect to have $\mathcal{R}(L_1) \approx \mathbb{T}_{A,r}$ by virtue of Theorem 6.3 and similarly to have $\mathcal{R}(K_1) \approx \mathbb{S}_{A,r}$. Theorem 7.1 implies that the norms $||K_1^T A||$ and $||AL_1||$ have an upper bound close to $\sigma_{q+1}(A)$, whereas $\kappa(K_0^T A L_0) \approx \sigma_1(A)/\sigma_q(A)$. Now correctness of the algorithm follows because the matrix A has numerical rank q.

From this correctness proof we can see that the approach would work based on any algorithm that computes approximate bases for the trailing singular spaces $\mathbb{T}_{A,r}$ and $\mathbb{S}_{A,r}$ or similarly for the leading singular spaces $\mathbb{T}_{n-r,A}$ and $\mathbb{S}_{m-r,A}$. Instead of Proto-Algorithm 7.1 we can apply randomized matrix multiplication or scaled randomized augmentation. Let us cover some of such variations.

We can proceed with nonorthogonal matrices K_0, K_1, L_0, L_1, Q_K, and Q_l, preserving matrix structure at the expense of weakening numerical stability a little. Then we can still expect that the norms $||W_{01}||$, $||W_{10}||$, and $||W_{11}||$ have at most order $\sigma_{q+1}(A)$, the norm $||W_{00}||$ has order $\sigma_1(A)$, and the condition number $\kappa(W_{00})$ has order $\sigma_1(A)/\sigma_q(A)$. Moreover we can choose random matrices $K_0 \in \mathcal{G}_{0,1}^{q\times n}$ and $L_0 \in \mathcal{G}_{0,1}^{n\times q}$. By virtue of Theorems 3.1 and 3.2 combined they are likely to be well conditioned, and so is the matrixK_0AL_0 by virtue of Theorem 3.4. This implies that we can extend our probabilistic estimates for the values $||W_{i,j}||$ for $i, j = 1, 2$ and $\kappa(W_{00})$. Here is the resulting simplified algorithm. Our tests in Section 8.2 show its efficiency.

Proto-Algorithm 7.2. Block diagonalization with randomized additive preprocessing and randomized multiplication.

INPUT, OUTPUT *and Stages 1 and 2 of* COMPUTATIONS *are the same as in Proto-Algorithm 7.1 except that the output matrices K_0, L_0, K_1 and L_1 are no longer assumed to be orthogonal.*

COMPUTATIONS:

3. Generate and output two random matrices $K_0, L_0 \in \mathcal{G}_{0,1}^{n\times q}$.

4. Apply the Subroutine LIN·SOLVE *to compute and to output the matrices $K_1 = C^{-T}V$ and $L_1 = C^{-1}U$. Output FAILURE and stop if so does the subroutine.*

We can further simplify the computations as follows.

Proto-Algorithm 7.3. Simplified randomized block diagonalization.

INPUT: *A matrix $A \in \mathbb{R}^{n\times n}$ whose norm $||A||$ is neither large nor small, and its numerical rank q satisfying $0 < q = n - r < n$.*

OUTPUT: *Four matrices K_0 and L_0 in $\mathbb{R}^{n\times q}$ and K_1 and L_1 in $\mathbb{R}^{n\times r}$ such that with a probability near 1 the $q \times q$ block submatrix $W_{00} = K_0^T AL_0$ of the matrix $W = (K_0 \mid K_1)^T A(L_0 \mid L_1) = \begin{pmatrix} W_{00} & W_{01} \\ W_{10} & W_{11} \end{pmatrix}$ is nonsingular, well conditioned, and strongly dominant such that $\sigma_q(W_{00}) \gg \max\{||W_{01}||,\ ||W_{10}||,\ ||W_{11}||\}$.*

COMPUTATIONS:

1. *Generate two random matrices $K_0^T, L_0 \in \mathcal{G}_{0,1}^{n\times q}$ and output the matrices K_0 and L_0.*
2. *Compute and output the matrix $K_1 = \mathrm{nmb}(AK_0)$ and the transpose L_1 of the matrix $\mathrm{nmb}(A^T L_0)$.*

To prove correctness of the algorithm, note that the matrices AK_0 and $A^T L_0$ are approximate matrix bases of the leading singular spaces $\mathbb{S}_{q,A}$ and $\mathbb{T}_{q,A}$ (see Theorem 2.6). Consequently the transposes of their nmbs are approximate matrix bases of the trailing singular spaces $\mathbb{S}_{A,r}$ and $\mathbb{T}_{A,r}$ for $r = n - q$, and correctness of the algorithm is readily verified.

Proto-Algorithms 7.1–7.3 do not produce block diagonalization but prepare its computation. Having strong domination of the block W_{00}, we can readily compute the block factorizations

$$W = \begin{pmatrix} I & O \\ W_{10}W_{00}^{-1} & I \end{pmatrix} \begin{pmatrix} W_{00} & O \\ O & G \end{pmatrix} \begin{pmatrix} I & W_{00}^{-1}W_{01} \\ O & I \end{pmatrix} \tag{7.1}$$

for $G = W_{11} - W_{10}W_{00}^{-1}W_{01}$ and

$$W^{-1} = \begin{pmatrix} I & -W_{00}^{-1}W_{01} \\ O & I \end{pmatrix} \begin{pmatrix} W_{00}^{-1} & O \\ O & G^{-1} \end{pmatrix} \begin{pmatrix} I & O \\ -W_{10}W_{00}^{-1} & I \end{pmatrix}. \tag{7.2}$$

Combined with these factorizations Proto-Algorithms 7.1–7.3 reduce the inversion of the matrix A to the inversion of the matrices W_{00} and G of smaller sizes, where both matrices W_{00} and G are expected to be nonsingular and better conditioned than the matrix A (cf. [PGMQ, Section 9]).

Remark 7.1. *We expect to arrive at the matrices W_{01}, W_{10} and W_{11} having small norms. To counter the expected cancellation of the leading digits of the $2rn - r^2$ entries of these matrices, we should compute the matrices C, K_1 and L_1, their products by the blocks of the matrix A, and the Schur complement G with a high precision p_+ (possibly by employing symbolic computations with infinite precision). These computations involve $O(n^2r)$ flops, that is a r/n fraction of order n^3 flops in high precision p_+ required by Gaussian elimination. See further study in [PGMQ, Section 9]. Having implemented this part of the computations with higher precision, we have outperformed the standard algorithms (see Section 8.3 and Tables 8.5 and 8.6).*

Remark 7.2. *We can simplify Proto-Algorithms 7.1–7.3 by computing only one of the two factors, $(K_0 \mid K_1)^T$ or $(L_0 \mid l_1)$. Then we would still readily compute at first block triangulation of its product with the matrix A, and then the solution of the linear system $A\mathbf{y} = \mathbf{b}$.*

Remark 7.3. *One can combine scaled randomized additive preprocessing with the SMW formula of Theorem 2.4 to compute an approximate inverse A_- of an input matrix A^{-1} to employ this inverse as a preconditioner for the linear system $A\mathbf{y} = \mathbf{b}$. Then again we would require extended or infinite precision at the SMW stage of the transition from C^{-1} to A^{-1}, but this only applies to a fraction r/n of all arithmetic operations involved. One can refine the preconditioner A_- by reapplying the same techniques or by applying Newton's iteration initialized with the current approximate inverse (cf. [P01, Chapter 6], [P10], and the bibliography therein). One can similarly combine scaled randomuzed augmentation with extension (4.2) of the SMW formula. It is not clear whether these alternatives have any advantage versus randomized computation of block factorizations (7.1) and (7.2).*

8 Numerical Experiments

Our theorems formally support our randomized additive preprocessing and augmentation where we use randomization with Gaussian random matrices, but our tests have provided consistent empirical support where the same techniques employed Gaussian random Toeplitz and other structured matrices defined by much fewer random parameters compared to the number of the entries.

Our numerical experiments with random general, Hankel, Toeplitz and circulant matrices have been performed in the Graduate Center of the City University of New York on a Dell server with a dual core 1.86 GHz Xeon processor and 2G memory running Windows Server 2003 R2. The test Fortran code has been compiled with the GNU gfortran compiler within the Cygwin environment. Random numbers have been generated with the random_number intrinsic Fortran function, assuming the uniform probability distribution over the range $\{x : -1 \le x < 1\}$. The tests have been designed by the first author and performed by his coauthors.

8.1 Preconditioning tests

Table 8.1 covers our tests for the preconditioning power of additive preprocessing in [PIMR10] and augmentaion. We 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 $\frac{A}{||A||} + \beta I$ where the recipes for defining the matrices A and scalars β are 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 $i - j < n - 1$. The entry $a_{n,1}$ is selected to ensure that the last row is linearly expressed through the other rows.

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 $|i-j| < n-1$, whereas the entry $a_{1,n} = a_{n,1}$ is a root of the quadratic equation $\det A = 0$. We repeatedly generated the matrices A until we arrived at the quadratic equation having real roots.

We 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 we defined 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 8.1 displays the average values of the condition numbers $\kappa(C)$ of the matrices $C = A + UU^T$ over 100,000 tests for the inputs in the above classes, $r = 1, 2, 4, 8$ and $n = 100$. Here $V = 2I - U$, U and W are normalized matrices of the sizes $n \times r$ and $r \times r$, respectively, $U = \bar{U}/||\bar{U}||$, $W = \bar{W}/||\bar{W}||$, $\bar{U}^T = (\pm I_r \mid O_{r,r} \mid \pm I_r \mid O_{r,r} \mid \ldots \mid O_{r,r} \mid \pm I_r \mid O_{r,s})$, we have chosen the integer s to obtain $n \times r$ matrices $\bar{U}$, have chosen the signs for the matrices $\pm I_r$ at random, and have chosen circulant matrices $\bar{W}$ filled with ± 1 and each defined by its first column where again the signs $\pm$ have been chosen at random. In our further tests the condition numbers of the matrices $C = A + 10^p UV^T$ for $p = -10, -5, 5, 10$ were steadily growing within a factor $10^{|p|}$ as the value $|p|$ was growing. This showed the importance of proper scaling of the additive preprocessor UV^T.

8.2 Approximation of the tails and heads of SVDs, computation of numerical ranks, and low-rank approximation of a matrix

At some specified stages of our tests of this subsection and Section 8.3 we performed additions, subtractions and multiplications with infinite precision (hereafter referred to as *error-free ring operations*). At the other stages we performed computations with double precision, and we rounded to double precision all random values. We performed two refinement iterations for the computed solution of every linear system of equations and matrix inverse.

Table 8.2 shows the data from our tests on the approximation of trailing singular spaces of an $n \times n$ matrix A having numerical nullity $r = n - q$ and on the approximation of this matrix with a matrix of rank $q = n - r$. For $n = 64, 128, 256$ and $q = 1, 8, 32$ we generated $n \times n$ random orthogonal matrices S and T and diagonal matrices $\Sigma = \mathrm{diag}(\sigma_j)_{j=1}^{n}$ such that $\sigma_j = 1/j$, $j = 1, \ldots, q$, $\sigma_j = 10^{-10}$, $j = q+1, \ldots, n$ (cf. [H02, Section 28.3]). Then we applied error-free ring operations to compute the input matrices $A = S_A \Sigma_A T_A^T$, for which $||A|| = 1$ and $\kappa(A) = 10^{10}$. Furthermore we generated pairs of random $n \times r$ matrices U and V for $r = 1, 8, 32$, scaled them to have $||UV^T|| \approx 1$, and computed the matrices $C = A + UV^T$ (by applying error-free ring operations), $B_{A,r} = C^{-1}U$ (by using two refinement iterations), $T_{A,r}$, $B_{A,r}Y_{A,r}$ as a least-squares approximation to $T_{A,r}$, $Q = Q(B_{A,r})$, and AQQ^T (by applying error-free ring operations). Table 8.2 summarizes the data on the values $\kappa(C)$ and the residual norms $\mathrm{rn}_1 = ||B_{A,r}Y_{A,r} - T_{A,r}||$ and $\mathrm{rn}_2 = ||AQQ^T||$ observed in 100 runs of our tests for every pair of n and r.

We performed similar tests on the approximation of leading singular spaces of the same $n \times n$ matrices A having numerical rank q and numerical nullity $r = n - q$ and on the approximation of this matrix with a matrix of rank q. We have generated random $n \times q$ matrices U (for $q = 1, 8, 32$) and random Toeplitz $n \times q$ matrices $\bar{U}$ (for $q = 8, 32$) and then successively computed the matrices $B_{q,A} = A^T U$ and $B_{q,A} = A^T \bar{U}$ (to obtain approximate matrix bases for the leading singular space $\mathbb{T}_{q,A}$),

$T_{q,A}$, $B_{q,A}Y_{q,A}$ as a least-squares approximation to $T_{q,A}$, $Q_{q,A} = Q(B_{q,A})$, and $A - AQ_{q,A}(Q_{q,A})^T$ (by applying error-free ring operations). Tables 8.3 and 8.4 summarize the data on the condition numbers $\kappa(C_-)$ and the residual norms $\mathrm{rn}^{(1)} = ||B_{q,A}Y_{q,A} - T_{q,A}||$ and $\mathrm{rn}^{(2)} = ||A - AQ_{q,A}(Q_{q,A})^T||$ obtained in 100 runs of our tests for every pair of n and q. In both cases where we have chosen $B_{q,A} = A^T U$ and $B_{q,A} = A^T \bar{U}$ the computed residual norms were equally small and were about as small as in Table 8.2.

8.3 Randomized 2×2 block factorization, preconditioning and the solution of linear systems of equations

For $n = 32, 64$, $r = 1, 2, 4$ and for every pair $\{n, r\}$ we generated 100 instances of random vectors $\mathbf{b}$ of dimension n and of matrices A, U, and V as follows. We generated random real orthogonal $n \times n$ matrices S and T, and random $n \times r$ matrices U and V, which we scaled to have $||U|| = ||V|| = 1$. Then we defined the matrices $\Sigma = \mathrm{diag}(\sigma_j)_{j=1}^n$, with $\sigma_{n-j} = 10^{-17}$ for $j = 0, 1, \dots, r-1$, and $\sigma_{n-j} = 1/(n-j)$ for $j = r, \dots, n-1$, and then applied error-free ring operations to compute the matrices $A = S\Sigma T^T$. Note that $||A|| = 1$ and $\kappa(A) = ||A^{-1}|| = 10^{17}$.

We applied Proto-Algorithm 7.1 to every such input for $n = 32, 64$ to obtain block diagonalization of the matrix A. Then we solved the linear systems $A\mathbf{y} = \mathbf{b}$. We first generated random $n \times (n-r)$ matrices K_0 and L_0 and then computed the matrices $C = A + UV^T$ (by applying error-free ring operations), $K_1 = C^{-T}V$ and $L_1 = C^{-1}U$ (by using two refinement iterations), and $W = (K_0 \mid K_1)^T A(L_0 \mid L_1) = \begin{pmatrix} W_{00} & W_{01} \\ W_{10} & W_{11} \end{pmatrix}$ (by applying error-free ring operations). In all our tests the matrices C were nonsingular and well conditioned, and the leading principal $(n-r)\times(n-r)$ blocks $W_{00} = K_0^T A L_0$ were well conditioned and strongly dominated the three other blocks W_{01}, W_{10}, and W_{11} in the 2×2 block matrices W, as we expected to see by virtue of our analysis in Section 7. Then we computed the vector $\widehat{\mathbf{b}} = (K_0 \mid K_1)^T \mathbf{b}$ (by applying error-free ring operations) and solved the linear system $W\mathbf{x} = \widehat{\mathbf{b}}$ (by using two refinement iterations). Finally we computed the vector $\mathbf{y} = (L_0 \mid L_1)\mathbf{x}$ (by applying error-free ring operations). Table 8.5 shows the average (mean) values of the relative residual norms $||A\mathbf{y} - \mathbf{b}||/||\mathbf{b}||$ of the output vectors $\mathbf{y}$ (these values range about 10^{-10}) as well as the minimums, maximums, and standard deviations in these tests. For the same ill conditioned inputs the Subroutine MLDIVIDE(A,B) for Gaussian elimination from MATLAB has produced corrupted outputs, as can be seen from Table 8.6.

We have also performed similar tests by applying Proto-Algorithm 7.3 for $n = 32, 64$ and $n \times n$ matrices A and vectors $\mathbf{b}$, which we generated as before, but for $q = n - r = 1, 2, 4$. Then we generated random $n \times q$ matrices U and V, successively computed the matrices $A^T V$ and AU (by applying error-free ring operations), $K_0 = Q(A^T V)$, $L_0 = Q(AU)$, $K_1 = Q(\mathrm{nmb}(K_0^T))$ and $L_1 = Q(\mathrm{nmb}(L_0^T))$, and continued as in the tests for Table 8.5. Table 8.7 displays the data on the residual norms computed in these tests. They have order 10^{-25} on the average. Then again for the same ill conditioned inputs the Subroutine MLDIVIDE(A,B) from MATLAB produced corrupted outputs, as in Table 8.6.

8.4 Solution of a real symmetric Toeplitz linear system of equations with randomized augmentation

We have solved 100 real symmetric linear systems of equations $T\mathbf{y} = \mathbf{b}$ for each n where we used vectors $\mathbf{b}$ with random coordinates from the range $[-1, 1)$ and Toeplitz matrices $T = S + 10^{-9} I_n$ for a singular symmetric Toeplitz $n \times n$ matrices S having rank $n-1$ and nullity 1 and generated according to the recipe in [PQ10, Section 10.1b]. Table 8.8 shows the average CPU time of the solutions by our Algorithm 6.1 and, for comparison, based on the QR factorization and SVD, which we computed by applying the LAPACK procedures DGEQRF and DGESVD, respectively. To solve the auxiliary Toeplitz linear system $K\mathbf{x} = \mathbf{e}_1$ in Algorithm 6.1, we first employed the Toeplitz linear solver of [KV99], [V99], [VBHK01], and [VK98] and then applied iterative refinement with double precision.

The abbreviations "Alg. 6.1", "QR", and "SVD" indicate the respective algorithms. The last two columns of the table display the ratios of these data in the first and the two other columns. We measured the CPU time with the mclock function by counting cycles. One can convert them into seconds by dividing their number by a constant CLOCKS_PER_SEC, which is 1000 on our platform. We marked the table entries by a "-" where the tests were running too long and have not been completed. We obtained the solutions $\mathbf{y}$ with the relative residual norms of about 10^{-15} in all three algorithms, which showed that Algorithm 6.1 employing iterative refinement was as reliable as the QR and SVD based solutions but ran much faster. We refer the reader to [PQZC, Table 3] on similar test results for the solution of ill conditioned homogeneous Toeplitz linear systems.

Table 8.1: Preconditioning tests

Type	r	Cond (C)
1n	1	3.21E+2
1n	2	4.52E+3
1n	4	2.09E+5
1n	8	6.40E+2
1s	1	5.86E+2
1s	2	1.06E+4
1s	4	1.72E+3
1s	8	5.60E+3
2n	1	8.05E+1
2n	2	6.82E+3
2n	4	2.78E+4
2n	8	3.59E+3
2s	1	1.19E+3
2s	2	1.96E+3
2s	4	1.09E+4
2s	8	9.71E+3
3n	1	2.02E+4
3n	2	1.53E+3
3n	4	6.06E+2
3n	8	5.67E+2
3s	1	2.39E+4
3s	4	1.69E+3
3s	8	6.74E+3
4n	1	4.93E+2
4n	2	4.48E+2
4n	4	2.65E+2
4n	8	1.64E+2
4s	1	1.45E+3
4s	2	5.11E+2
4s	4	7.21E+2
4s	8	2.99E+2

9 Conclusions

It is well known that random matrices tend to be well conditioned, and one can expect that scaled randomized augmentation of an ill conditioned matrix, that is appending to it some properly scaled random rows and columns, is likely to decrease its condition number. We prove quantitative es-

Table 8.2: Tails of the SVDs and lower-rank approximations with additive preprocessing (cf. [PQ10])

r	$\kappa(C)$ or rrn_i	n	**min**	**max**	**mean**	**std**
1	$\kappa(C)$	64	$2.38 \times 10^{+02}$	$1.10 \times 10^{+05}$	$6.25 \times 10^{+03}$	$1.68 \times 10^{+04}$
1	$\kappa(C)$	128	$8.61 \times 10^{+02}$	$7.48 \times 10^{+06}$	$1.32 \times 10^{+05}$	$7.98 \times 10^{+05}$
1	$\kappa(C)$	256	$9.70 \times 10^{+02}$	$3.21 \times 10^{+07}$	$3.58 \times 10^{+05}$	$3.21 \times 10^{+06}$
1	rn_1	64	4.01×10^{-10}	1.50×10^{-07}	5.30×10^{-09}	1.59×10^{-08}
1	rn_1	128	7.71×10^{-10}	5.73×10^{-07}	1.58×10^{-08}	6.18×10^{-08}
1	rn_1	256	7.57×10^{-10}	3.2×10^{-07}	1.69×10^{-08}	5.02×10^{-08}
1	rn_2	64	4.01×10^{-10}	1.50×10^{-07}	5.30×10^{-09}	1.59×10^{-08}
1	rn_2	128	7.71×10^{-10}	5.73×10^{-07}	1.58×10^{-08}	6.18×10^{-08}
1	rn_2	256	7.57×10^{-10}	3.22×10^{-07}	1.69×10^{-08}	5.02×10^{-08}
8	$\kappa(C)$	64	$1.26 \times 10^{+03}$	$1.61 \times 10^{+07}$	$2.68 \times 10^{+05}$	$1.71 \times 10^{+06}$
8	$\kappa(C)$	128	$2.92 \times 10^{+03}$	$3.42 \times 10^{+06}$	$1.58 \times 10^{+05}$	$4.12 \times 10^{+05}$
8	$\kappa(C)$	256	$1.39 \times 10^{+04}$	$8.75 \times 10^{+07}$	$1.12 \times 10^{+06}$	$8.74 \times 10^{+06}$
8	rn_1	64	3.39×10^{-10}	2.27×10^{-06}	2.74×10^{-08}	2.27×10^{-07}
8	rn_1	128	4.53×10^{-10}	1.91×10^{-07}	1.03×10^{-08}	2.79×10^{-08}
8	rn_1	256	8.74×10^{-10}	1.73×10^{-07}	7.86×10^{-09}	1.90×10^{-08}
8	rn_2	64	1.54×10^{-09}	7.59×10^{-06}	8.87×10^{-08}	7.58×10^{-07}
8	rn_2	128	1.82×10^{-09}	7.27×10^{-07}	2.95×10^{-08}	8.57×10^{-08}
8	rn_2	256	2.62×10^{-09}	3.89×10^{-07}	2.27×10^{-08}	5.01×10^{-08}
32	$\kappa(C)$	64	$1.77 \times 10^{+03}$	$9.68 \times 10^{+06}$	$1.58 \times 10^{+05}$	$9.70 \times 10^{+05}$
32	$\kappa(C)$	128	$1.65 \times 10^{+04}$	$6.12 \times 10^{+07}$	$1.02 \times 10^{+06}$	$6.19 \times 10^{+06}$
32	$\kappa(C)$	256	$3.57 \times 10^{+04}$	$2.98 \times 10^{+08}$	$4.12 \times 10^{+06}$	$2.98 \times 10^{+07}$
32	rn_1	64	2.73×10^{-10}	3.29×10^{-08}	2.95×10^{-09}	4.93×10^{-09}
32	rn_1	128	3.94×10^{-10}	1.29×10^{-07}	7.18×10^{-09}	1.64×10^{-08}
32	rn_1	256	6.80×10^{-10}	4.00×10^{-07}	1.16×10^{-08}	4.27×10^{-08}
32	rn_2	64	2.10×10^{-09}	1.49×10^{-07}	1.55×10^{-08}	2.18×10^{-08}
32	rn_2	128	2.79×10^{-09}	3.80×10^{-07}	3.81×10^{-08}	6.57×10^{-08}
32	rn_2	256	5.35×10^{-09}	1.05×10^{-06}	5.70×10^{-08}	1.35×10^{-07}

timates for such an impact as well as for the impact of scaled randomized additive preprocessing, that is for adding a properly scaled random matrix of a small rank. We showed some applications of these results to matrix computations. Our numerical tests are in good accordance with the formally proved estimates, but the tests also showed the same impact of the randomized preprocessing on the condition number where we used sparse and structured preprocessors, which had much fewer random parameters. Formal explanation of this empirical observation remains a research challenge (see [T11] on some recent important progress). Another interesting research direction is the study of the impact of combining our current techniques of randomized augmentation and additive preprocessing with randomized matrix multiplication of [PGMQ, Section 12.2], [PQZa], and [PQb].

Table 8.3: Heads of SVDs and low-rank approximation with random multipliers

q	rrn_i	n	**min**	**max**	**mean**	**std**
1	$\mathrm{rn}^{(1)}$	64	2.35×10^{-10}	1.32×10^{-07}	3.58×10^{-09}	1.37×10^{-08}
1	$\mathrm{rn}^{(1)}$	128	4.41×10^{-10}	3.28×10^{-08}	3.55×10^{-09}	5.71×10^{-09}
1	$\mathrm{rn}^{(1)}$	256	6.98×10^{-10}	5.57×10^{-08}	5.47×10^{-09}	8.63×10^{-09}
1	$\mathrm{rn}^{(2)}$	64	8.28×10^{-10}	1.32×10^{-07}	3.86×10^{-09}	1.36×10^{-08}
1	$\mathrm{rn}^{(2)}$	128	1.21×10^{-09}	3.28×10^{-08}	3.91×10^{-09}	5.57×10^{-09}
1	$\mathrm{rn}^{(2)}$	256	1.74×10^{-09}	5.58×10^{-08}	5.96×10^{-09}	8.47×10^{-09}
8	$\mathrm{rn}^{(1)}$	128	2.56×10^{-09}	1.16×10^{-06}	4.30×10^{-08}	1.45×10^{-07}
8	$\mathrm{rn}^{(1)}$	256	4.45×10^{-09}	3.32×10^{-07}	3.40×10^{-08}	5.11×10^{-08}
8	$\mathrm{rn}^{(2)}$	64	1.46×10^{-09}	9.56×10^{-08}	5.77×10^{-09}	1.06×10^{-08}
8	$\mathrm{rn}^{(2)}$	128	1.64×10^{-09}	4.32×10^{-07}	1.86×10^{-08}	5.97×10^{-08}
8	$\mathrm{rn}^{(2)}$	256	2.50×10^{-09}	1.56×10^{-07}	1.59×10^{-08}	2.47×10^{-08}
32	$\mathrm{rn}^{(1)}$	64	6.80×10^{-09}	2.83×10^{-06}	1.01×10^{-07}	3.73×10^{-07}
32	$\mathrm{rn}^{(1)}$	128	1.25×10^{-08}	6.77×10^{-06}	1.28×10^{-07}	6.76×10^{-07}
32	$\mathrm{rn}^{(1)}$	256	1.85×10^{-08}	1.12×10^{-06}	1.02×10^{-07}	1.54×10^{-07}
32	$\mathrm{rn}^{(2)}$	64	1.84×10^{-09}	6.50×10^{-07}	2.30×10^{-08}	8.28×10^{-08}
32	$\mathrm{rn}^{(2)}$	128	3.11×10^{-09}	1.45×10^{-06}	2.87×10^{-08}	1.45×10^{-07}
32	$\mathrm{rn}^{(2)}$	256	4.39×10^{-09}	2.16×10^{-07}	2.37×10^{-08}	3.34×10^{-08}

Table 8.4: Heads of SVDs and low-rank approximations with random Toeplitz multipliers

q	$\mathrm{rrn}^{(i)}$	n	**min**	**max**	**mean**	**std**
8	$\mathrm{rrn}^{(1)}$	64	2.22×10^{-09}	7.89×10^{-06}	1.43×10^{-07}	9.17×10^{-07}
8	$\mathrm{rrn}^{(1)}$	128	3.79×10^{-09}	4.39×10^{-05}	4.87×10^{-07}	4.39×10^{-06}
8	$\mathrm{rrn}^{(1)}$	256	5.33×10^{-09}	3.06×10^{-06}	6.65×10^{-08}	3.12×10^{-07}
8	$\mathrm{rrn}^{(2)}$	64	1.13×10^{-09}	3.66×10^{-06}	6.37×10^{-08}	4.11×10^{-07}
8	$\mathrm{rrn}^{(2)}$	128	1.81×10^{-09}	1.67×10^{-05}	1.90×10^{-07}	1.67×10^{-06}
8	$\mathrm{rrn}^{(2)}$	256	2.96×10^{-09}	1.25×10^{-06}	2.92×10^{-08}	1.28×10^{-07}
32	$\mathrm{rrn}^{(1)}$	64	6.22×10^{-09}	5.00×10^{-07}	4.06×10^{-08}	6.04×10^{-08}
32	$\mathrm{rrn}^{(1)}$	128	2.73×10^{-08}	4.88×10^{-06}	2.57×10^{-07}	8.16×10^{-07}
32	$\mathrm{rrn}^{(1)}$	256	1.78×10^{-08}	1.25×10^{-06}	1.18×10^{-07}	2.03×10^{-07}
32	$\mathrm{rrn}^{(2)}$	64	1.64×10^{-09}	1.26×10^{-07}	9.66×10^{-09}	1.48×10^{-08}
32	$\mathrm{rrn}^{(2)}$	128	5.71×10^{-09}	9.90×10^{-07}	5.50×10^{-08}	1.68×10^{-07}
32	$\mathrm{rrn}^{(2)}$	256	4.02×10^{-09}	2.85×10^{-07}	2.74×10^{-08}	4.48×10^{-08}

Table 8.5: Relative residual norms for ill conditioned linear systems, with Proto-Algorithm 7.1

n	**r**	**min**	**max**	**mean**	**std**
32	1	1.49×10^{-13}	1.36×10^{-9}	4.25×10^{-11}	1.56×10^{-10}
32	2	3.70×10^{-13}	2.13×10^{-8}	3.83×10^{-10}	2.35×10^{-9}
32	4	9.33×10^{-13}	1.08×10^{-8}	3.37×10^{-10}	1.26×10^{-9}
64	1	1.11×10^{-12}	6.87×10^{-9}	2.03×10^{-10}	7.49×10^{-10}
64	2	1.53×10^{-12}	1.21×10^{-8}	5.86×10^{-10}	1.77×10^{-9}
64	4	2.21×10^{-12}	1.27×10^{-7}	1.69×10^{-9}	1.28×10^{-8}

Table 8.6: Relative residual norms for ill conditioned linear systems, with MLDIVIDE(A,B)

n	**r**	**min**	**max**	**mean**	**std**
32	1	6.34×10^{-3}	7.44×10^{1}	1.74×10^{0}	7.53×10^{0}
32	2	2.03×10^{-2}	1.32×10^{1}	9.19×10^{-1}	1.62×10^{0}
32	4	4.57×10^{-2}	1.36×10^{1}	1.14×10^{0}	1.93×10^{0}
64	1	3.82×10^{-3}	9.93×10^{0}	1.03×10^{0}	1.66×10^{0}
64	2	1.96×10^{-2}	1.27×10^{2}	3.09×10^{0}	1.40×10^{1}
64	4	7.13×10^{-3}	6.63×10^{0}	8.23×10^{-1}	1.20×10^{0}

Table 8.7: Relative residual norms for ill conditioned linear system, with Proto-Algorithm 7.3

n	q	**min**	**max**	**mean**	**std**
32	1	7.08×10^{-30}	4.00×10^{-23}	4.52×10^{-25}	4.01×10^{-24}
32	2	7.49×10^{-30}	2.29×10^{-21}	2.77×10^{-23}	2.33×10^{-22}
32	4	1.46×10^{-28}	1.63×10^{-07}	4.83×10^{-25}	2.73×10^{-24}
64	1	1.13×10^{-29}	1.01×10^{-24}	2.31×10^{-26}	1.11×10^{-25}
64	2	6.60×10^{-29}	6.90×10^{-24}	1.45×10^{-25}	7.73×10^{-25}
64	4	2.60×10^{-28}	1.41×10^{-21}	1.61×10^{-23}	1.42×10^{-22}

Table 8.8: The CPU time (in cycles) for solving an ill conditioned real symmetric Toeplitz linear system

n	**Alg. 6.1**	**QR**	**SVD**	**QR/Alg. 6.1**	**SVD/Alg. 6.1**
512	56.3	148.4	4134.8	2.6	73.5
1024	120.6	1533.5	70293.1	12.7	582.7
2048	265.0	11728.1	–	44.3	–
4096	589.4	–	–	–	–
8192	1304.8	–	–	–	–

Appendix

A Uniform random sampling and nonsingularity of random matrices

Uniform random sampling of elements from a finite set Δ is their selection from this set at random, independently of each other and under the uniform probability distribution on the set Δ.

Theorem A.1. *Under the assumptions of Lemma 3.2 let the values of the variables of the polynomial be randomly and uniformly sampled from a finite set Δ. Then the polynomial vanishes with a probability at most $\frac{d}{|\Delta|}$.*

Corollary A.1. *Let the entries of a general or Toeplitz $m \times n$ matrix have been randomly and uniformly sampled from a finite set Δ of cardinality $|\Delta|$ (in any fixed ring). Let $l = \min\{m, n\}$. Then (a) any $k \times k$ submatrix M for $k \le l$ is nonsingular with a probability at least $1 - \frac{k}{|\Delta|}$ and (b) is strongly nonsingular with a probability at least $1 - \sum_{i=1}^{k} \frac{i}{|\Delta|} = 1 - \frac{(k+1)k}{2|\Delta|}$. Furthermore (c) if the submatrix M is indeed nonsingular, then any entry of its inverse is nonzero with a probability at least $1 - \frac{k-1}{|\Delta|}$.*

Proof. The claimed properties of nonsingularity and nonvanishing hold for generic matrices. The singularity of a $k \times k$ matrix means that its determinant vanishes, but the determinant is a polynomial of total degree k in the entries. Therefore Theorem A.1 implies parts (a) and consequently (b). Part (c) follows because a fixed entry of the inverse vanishes if and only if the respective entry of the adjoint vanishes, but up to the sign the latter entry is the determinant of a $(k-1) \times (k-1)$ submatrix of the input matrix M, and so it is a polynomial of degree $k-1$ in its entries. □

B Extremal singular values 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. All the results of Section 3.2 apply to this case equally well. Below we elaborate upon the respective extension of our probabilistic bounds on the norms and singular values.

Definition B.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 $(A \mid B) \in \mathcal{G}^{m \times 2n}_{\mu,\sigma}$ (cf. Definition 3.2).*

We can immediately extend Theorem 3.2 to the latter matrices. Let us extend Theorem 3.1. Its original proof in [SST06] employes the following two lemmas.

Lemma B.1. *[SST06, Lemma A.2]. For a nonnegative scalar y, a unit vector $\mathbf{t} \in \mathbb{R}^{n \times 1}$, and a vector $\mathbf{b} \in \mathcal{G}^{n \times 1}_{\mu,\sigma}$, we have $F_{|\mathbf{t}^T \mathbf{b}|}(y) \le \sqrt{\frac{2}{\pi}} \frac{y}{\sigma}$.*

Lemma B.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*

$$\text{Probability}\{||A^{-1}\mathbf{w}|| > 1/y\} \le \sqrt{\frac{2}{\pi}} \frac{y}{\sigma} \text{ for } j = 1, \dots, n.$$

The following lemma and corollary extend these lemmas to the complex case.

Lemma B.3. *The bound of Lemma B.1 also 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 $\mathbf{v} \in \mathcal{G}_{\mu,\sigma}^{1\times 2n}$ and $||\mathbf{u}|| = ||\mathbf{t}|| = 1$ and apply Lemma B.1 to real vectors $\mathbf{u}$ and $\mathbf{v}$ replacing $\mathbf{b}$ and $\mathbf{t}$. □

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

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

Proof. In the case of real matrices A the corollary is supported by the argument in the proof of [SST06, Lemma 3.2], which employs the well known estimate that we state as our Lemma B.1. Now we employ Lemma B.3 instead of this estimate, otherwise keep the same argument as in [SST06], and arrive at Corollary B.1. □

Corollary B.2. *Under the assumptions of Corollary B.1 we have $||A^{-1}|| \leq \sum_{j=1}^n X_j$ where X_j are nonnegative random variables such that*

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

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 B.1. □

The corollary implies that Probability$\{||A^{-1}|| > 1/y\}$ converges to 0 proportionally to y as $y \to 0$, which can be viewed as an extension of Theorem 3.1 to the case of complex inputs. The resulting estimates are a little weaker than in Section 3.3, being overly pessimistic. Actually random complex matrices are a little better conditioned than random real matrices (see [E88], [ES05], [CD05]).

C Computation of numerical ranks

Unlike the customary algorithms we involve neither pivoting nor orthogonalization in the following noncostly algorithms that compute the numerical rank of a matrix. Suppose $m \geq n$ (else we shift to $\tilde{A}^T$), let $[\rho_-, \rho_+] = [0, n]$ unless you know a more narrow range for the numerical rank of A, and successively test the candidate integers in the range $[\rho_-, \rho_+]$ until you find the numerical rank ρ.

Proto-Algorithm C.1. Numerical rank via additive preprocessing *(see Remarks C.1–C.3).*

INPUT: *Two integers ρ_- and ρ_+ and a matrix $\tilde{A} \in \mathbb{R}^{m\times n}$ having unknown numerical rank $\rho = \text{rank}(\tilde{A})$ in the range $[\rho_-, \rho_+]$ such that $0 \leq \rho_- < \rho_+ \leq n \leq m$, a rule for the selection of a candidate integer ρ in a range $[\rho_-, \rho_+]$, and a Subroutine COND that determines whether a given matrix has full rank and is well conditioned or not.*

OUTPUT: *An integer ρ expected to equal numerical rank of the matrix $\tilde{A}$. (The expectation can actually fail, but with a low probability, see Remark C.1.)*

INITIALIZATION: *Compute the integer $r_+ = n - \rho_-$ and a scalar σ of order $||A||$, generate two matrices $U_+ \in \mathcal{G}_{0,\sigma}^{m\times r_+}$ and $V_+ \in \mathcal{G}_{0,\sigma}^{n\times r_+}$, and write $U_s = U_+(I_s \mid O_{s,m-s})^T$ and $V_s = V_+(I_s \mid O_{n-s,s})^T$ for $s = r_-, r_- + 1, \dots, r_+$.*

COMPUTATIONS:

1. *Stop and output $\rho = \rho_+$ if $\rho_- = \rho_+$. Otherwise fix an integer ρ in the range $[\rho_-, \rho_+]$.*
2. *Compute an integer $r = n - s$ and the matrix $C = A + U_rV_r^T$ and apply to it the Subroutine COND.*

3. *If this matrix has full rank and is well conditioned, write* $\rho_+ = \rho$ *and go to Stage 1. Otherwise write* $\rho_- = \rho$ *and go to Stage 1.*

Proto-Algorithm C.2. Numerical rank via randomized augmentation.

INPUT, OUTPUT *and Stages 1 and 3 of* COMPUTATIONS *are the same as in Proto-Algorithm C.1.*

INITIALIZATION*:* *Compute the integer* $r_+ = n - \rho_-$ *and a scalar* σ *of order* $||A||$*, generate three matrices* $U_+ \in \mathcal{G}_{0,\sigma}^{m\times r_+}$*,* $V_+ \in \mathcal{G}_{0,\sigma}^{n\times r_+}$*, and* $W_+ \in \mathcal{G}_{0,\sigma}^{r_+\times r_+}$*, and write* $i = 1$*,* $A_0 = A$*,* $U_r = U_+(I_r \mid O_{r,m-r})^T$*,* $V_r = V_+(I_r \mid O_{n-r,r})^T$*, and* $W_r = (I_r \mid O_{r,r_+-r})W_+(I_r \mid O_{r,r_+-r})^T$ *for* $r = r_-, r_- + 1, \dots, r_+$*.*

COMPUTATIONS*:*

2. *Compute the integer* $r = n - s$*. Compute the matrix* $K = \begin{pmatrix} W_r & V_r^T \\ -U_r & A \end{pmatrix}$ *and apply to it the Subroutine COND.*

For the sake of completeness we also reproduce the following alternative algorithm from [PQb].

Proto-Algorithm C.3. Numerical rank via randomized multiplication *(see Remarks C.1–C.4).*

INPUT *is the same as in Proto-Algorithms C.1 and C.2.*

OUTPUT: *an integer* ρ *expected to equal numerical rank of the matrix* $\tilde{A}$ *and a matrix* B *expected to approximate (within an error norm in* $O(\sigma_{\rho+1}(\tilde{A}))$*) a matrix basis of the singular space* $\mathbb{T}_{\rho,\tilde{A}}$*. (Both expectations can actually fail, but with a low probability, see Remark C.1.)*

INITIALIZATION: *Generate matrix* $G \in \mathcal{G}_{0,1}^{m\times\rho_+}$ *and write* $B = \tilde{A}$*,* $G_\rho = G(I_\rho \mid O_{\rho,m-\rho})^T$ *for* $\rho = \rho_-, \rho_- + 1, \dots, \rho_+$ *(The* $m \times \rho$ *matrix* G_ρ *is formed by the first* ρ *columns of the matrix* G*.)*

COMPUTATIONS*:*

1. *Stop and output* $\rho = \rho_+$ *and the matrix* B *if* $\rho_- = \rho_+$*. Otherwise fix an integer* ρ *in the range* $[\rho_-, \rho_+]$*.*
2. *Compute the matrix* $B' = B^T G_\rho$ *and apply to it the Subroutine COND.*
3. *If this matrix has full rank and is well conditioned, write* $\rho_+ = \rho$ *and* $B = B'$ *and go to Stage 1. Otherwise write* $\rho_- = \rho$ *and go to Stage 1.*

Remark C.1. *The Proto-Algorithms can output a wrong value of the numerical rank, although by virtue of Theorems 2.5 and 2.6 combined this occurs with a low probability. One can decrease this probability by alternating application of the three Proto-Algorithms to the same inputs and choosing distinct random parameters. Furthermore one can extend Proto-Algorithm C.3 to computing rank-*ρ *approximation to the input matrix* A *and compute the error norm. Its large value would imply that* $\rho < \mathrm{nrank}(A)$*.*

Remark C.2. *Unlike the customary condition estimators, we can avoid orthogonalization and pivoting in the Subroutine COND by approximating both extremal eigenvalues of the matrix* $S = \tilde{A}^T\tilde{A}$ *or* $S = \tilde{A}\tilde{A}^T$*. We can apply either the Power Method (cf. [B74], [D83]), first to the matrix* S *to compute a close upper bound* σ_+^2 *on the value* $(\sigma_1(A))^2$ *and then to the matrix* $\sigma_+^2 I - S$ *to approximate the value* $(\sigma_n(A))^2$*, or the Lanczos algorithm, which approximates both values* $(\sigma_1(A))^2$ *and* $(\sigma_n(A))^2$ *and converges much faster (cf. [KW92], [GL96, Sections 9.1.4 and 9.1.5]).*

Remark C.3. *The binary search* $\rho = \lceil(\rho_- + \rho_+)/2\rceil$ *is an attractive policy for choosing the candidate values* ρ *in Proto-Algorithms C.1 and C.2 , but for Proto-Algorithm C.3 one may prefer to move toward* ρ_-*, the left end of the range more rapidly, to decrease the size of the matrix* B'*.*

Remark C.4. *One can simplify Stage 2 of Proto-Algorithm C.3 by applying the Subroutine COND to the matrix* $G'_\rho = F_\rho G_\rho$ *of a smaller size (rather than to* G_ρ*) where* $F_\rho \in \mathcal{G}_{0,1}^{\rho \times m}$*. By virtue of Theorem 3.4 the matrices* G_ρ *and* G'_ρ *are likely to have condition numbers of the same order.*

Acknowledgements: Our work has been supported by NSF Grant CCF–1116736 and PSC CUNY Awards 4512–0042 and 65792–0043. Mr. Jesse Wolf has provided many helpful comments.

References

[B74] D.W. Boyd, The Power method for l^p Norms, *Linear Algebra and Its Applications*, **9**, 95–101, 1974.

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

[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.

[DL78] R. A. Demillo, R. J. Lipton, A Probabilistic Remark on Algebraic Program Testing, *Information Processing Letters*, **7**, **4**, 193–195, 1978.

[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.

[EG99] Y. Eidelman, I. Gohberg, On a New Class of Structured Matrices, *Integral Equations and Operator Theory*, **34**, 293–324, Birkhäuser, Basel, 1999.

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

[GE96] M. Gu, S. C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM Journal on Scientific Computing*, **17**, 848–869, 1996.

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

[GS72] I. Gohberg, A. Sementsul, On the Inversion of Finite Toeplitz Matrices and Their Continuous Analogs, *Matematicheskiie Issledovaniia* (in Russian), **7**, **2**, 187–224, 1972.

[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.

[HP92] Y. P. Hong, C.–T. Pan, The rank revealing QR decomposition and SVD, *Math. of Computation*, **58**, 213–232, 1992.

[KL94] C. S. Kenney, A. J. Laub, Small-Sample Statistical Condition Estimates for General Matrix Functions, *SIAM J. on Scientific and Statistical Computing*, **15**, 36–61, 1994.

[KV99] P. Kravanja, M. Van Barel, Algorithms for Solving Rational Interpolation Problems Related to Fast and Superfast Solvers for Toeplitz Systems, *SPIE*, 359–370, 1999.

[KW92] J. Kuczynski, H. Wozniakowski, Estimating the Largest Eigenvalue by the Power and Lanczos Algorithms with a Random Start, *SIAM J. on Matrix Analysis and Applications*, **13**, 1094–1122, 1992.

[MP80] W. L. Miranker, V. Y. Pan, Methods of Aggregations, *Linear Algebra and Its Applications*, **29**, 231–257, 1980.

[P00a] C.–T. Pan, On the Existence and Computation of Rank-revealing LU Factorization, *Linear Algebra and Its Applications*, **316**, 199–222, 2000.

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

[P10] V. Y. Pan, Newton's Iteration for Matrix Inversion, Advances and Extensions, pp. 364–381, in *Matrix Methods: Theory, Algorithms and Applications* (dedicated to the Memory of Gene Golub, edited by V. Olshevsky and E. Tyrtyshnikov), World Scientific Publishing, New Jersey, ISBN-13 978-981-283-601-4, ISBN-10-981-283-601-2 (2010).

[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**, 2851–1876, 2012.

[PQa] V. Y. Pan, G. Qian, Condition Numbers of Random Toeplitz and Circulant Matrices, Tech. Report TR 2012013, *PhD Program in Comp. Sci.*, *Graduate Center, CUNY*, 2012. Available at

http://www.cs.gc.cuny.edu/tr/techreport.php?id=442

[PQb] V. Y. Pan, G. Qian, More on the Power of Randomized Matrix Multiplication, Tech. Report TR 2012014, *PhD Program in Comp. Sci.*, *Graduate Center, CUNY*, 2012. Available at

http://www.cs.gc.cuny.edu/tr/techreport.php?id=443

[PQZa] V. Y. Pan, G. Qian, A. Zheng, Randomized Preprocessing versus Pivoting, *Linear Algebra and Its Applications*, in print. http://dx.doi.org/10.1016/j.laa.2011.02.052

[PQZC] V. Y. Pan, G. Qian, A. Zheng, Z. Chen, Matrix Computations and Polynomial Root-finding with Preprocessing, *Linear Algebra and Its Applications*, **434**, 854–879, 2011.

[PY09] V. Y. Pan, X. Yan, Additive Preconditioning, Eigenspaces, and the Inverse Iteration, *Linear Algebra and Its Applications*, **430**, 186–203, 2009.

[S80] J. T. Schwartz, Fast Probabilistic Algorithms for Verification of Polynomial Identities, *Journal of ACM*, **27**, **4**, 701–717, 1980.

[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.

[T11] J. A. Tropp, Improved analysis of the subsampled randomized Hadamard transform, *Adv. Adapt. Data Anal.*, **3, 1–2**, Special Issue, "Sparse Representation of Data and Images," 115-126, 2011.

[V99] M. Van Barel, A Superfast Toeplitz Solver, 1999.
Available at http://www.cs.kuleuven.be/˜ marc/software/index.html

[VBHK01] M. Van Barel, G. Heinig, P. Kravanja, A Stabilized Superfast Solver for Nonsymmetric Toeplitz Systems, *SIAM J. on Matrix Analysis and Applications*, **23**, **2**, 494–510, 2001.

[VK98] M. Van Barel, P. Kravanja, A Stabilized Superfast Solver for Indefinite Hankel Systems, *Linear Algebra and its Applications*, **284**, **1–3**, 335–355, 1998.

[W07] X. Wang, Affect of Small Rank Modification on the Condition Number of a Matrix, *Computer and Math. (with Applications)*, **54**, 819–825, 2007.

[Z79] R. E. Zippel, Probabilistic Algorithms for Sparse Polynomials, *Proceedings of EUROSAM'79, Lecture Notes in Computer Science*, **72**, 216–226, Springer, Berlin, 1979.