

RESCALING THE GSVD WITH APPLICATION TO ILL-POSED PROBLEMS

L. DYKES*, S. NOSCHESSE†, AND L. REICHEL‡

Abstract. The generalized singular value decomposition (GSVD) of a pair of matrices expresses each matrix as a product of an orthogonal, a diagonal, and a nonsingular matrix. The nonsingular matrix, which we denote by X^T , is the same in both products. Available software for computing the GSVD scales the diagonal matrices and X^T so that the squares of corresponding diagonal entries sum to one. This paper proposes a scaling that seeks to minimize the condition number of X^T . The rescaled GSVD gives rise to new truncated GSVD methods, one of which is well suited for the solution of linear discrete ill-posed problems. Numerical examples show this new truncated GSVD method to be competitive with the standard truncated GSVD method as well as with Tikhonov regularization with regard to the quality of the computed approximate solution.

Key words. generalized singular value decomposition, truncated generalized singular value decomposition, ill-posed problem

1. Introduction. Let $A \in \mathbb{R}^{m \times n}$ and $L \in \mathbb{R}^{p \times n}$, and assume that the null spaces of A and L intersect trivially, i.e., that

$$(1.1) \quad \mathcal{N}(A) \cap \mathcal{N}(L) = \{\mathbf{0}\}.$$

For notational simplicity, we will assume that $m \geq n \geq p$; however, these constraints can be removed.

The generalized singular value decomposition (GSVD) of the matrix pair $\{A, L\}$ are factorizations of the form

$$(1.2) \quad A = U \Sigma X^T, \quad L = V M X^T,$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{p \times p}$ are orthogonal, $X \in \mathbb{R}^{n \times n}$ is nonsingular, and

$$(1.3) \quad \begin{aligned} \Sigma &= \text{diag}[\sigma_1, \sigma_2, \dots, \sigma_p, 1, 1, \dots, 1] \in \mathbb{R}^{m \times n}, \\ M &= [\text{diag}[\mu_1, \mu_2, \dots, \mu_p], \mathbf{0}, \mathbf{0}, \dots, \mathbf{0}] \in \mathbb{R}^{p \times n}. \end{aligned}$$

Thus, Σ is a diagonal, possibly rectangular, matrix and M consists of a leading $p \times p$ diagonal matrix to which $n - p$ columns of zeros are appended. The entries of Σ and M are ordered according to

$$0 \leq \sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_p \leq 1, \quad 0 \leq \mu_p \leq \mu_{p-1} \leq \dots \leq \mu_1 \leq 1$$

and satisfy

$$(1.4) \quad \sigma_j^2 + \mu_j^2 = 1, \quad 1 \leq j \leq p;$$

see Golub and Van Loan [11, Sections 8.7.4 and 6.1.6] as well as Bai and Demmel [2, 3] for properties of the GSVD and its computation. The GSVD finds application, e.g., in data analysis, such as microarray analysis, and in the solution of linear discrete

*Department of Mathematical Sciences, Kent State University, Kent, OH 44242, USA, and University School, Hunting Valley, OH 44022, USA. E-mail: ldykes@math.kent.edu.

†SAPIENZA Università di Roma, P.le A. Moro, 2, I-00185 Roma, Italy. E-mail: noschese@mat.uniroma1.it. Research supported by a grant from SAPIENZA Università di Roma.

‡Department of Mathematical Sciences, Kent State University, Kent, OH 44242, USA. E-mail: reichel@math.kent.edu. Research supported in part by NSF grant DMS-1115385.

ill-posed problems; see, e.g., [1, 4, 12, 13, 18]. The constraint (1.1) is not necessary for the GSVD to exist, but is required for the application of the GSVD to the solution of linear discrete ill-posed problems described below. We will assume that (1.1) holds unless explicitly stated otherwise.

Due to the scaling (1.4), the σ_j are not the singular values of A when L is the identity. The quotients

$$(1.5) \quad \gamma_j = \sigma_j / \mu_j, \quad 1 \leq j \leq p,$$

are referred to as generalized singular values of the matrix pair $\{A, L\}$. They are the singular values of AL^{-1} when L is square and invertible. In particular, up to reordering, the γ_j agree with the singular values of A when $L = I$.

The present paper proposes a rescaling of the matrices in the decompositions (1.2) with the purpose of making the factor X^T better conditioned. Many applications require the multiplication of a computed vector by X^T or the solution of a linear system of equations with the matrix X^T . In both situations, it is beneficial that X^T be well conditioned. We describe the proposed scaling in Section 2. The application of the GSVD to the solution of linear discrete ill-posed problems is discussed in Section 3. The computation of a meaningful approximate solution to these problems generally requires regularization. The truncated GSVD (TGSVD) method based on the standard scaling (1.4) is a simple and popular regularization method. We describe alternative truncated GSVD methods that are based on our rescaled GSVD. These methods are discussed in Section 4. Computed examples that illustrate the performance of the methods considered in this paper can be found in Section 5, and Section 6 contains concluding remarks.

2. A rescaled GSVD. We propose to scale the columns of the matrix X in (1.2) so that they all have Euclidean norm one. This scaling is known to make the condition number of the rescaled matrix near-minimal over all column scalings. Throughout this paper $\|\cdot\|_2$ denotes the Euclidean vector norm or the associated induced matrix norm. The following propositions collect some properties of the rescaled GSVD.

PROPOSITION 2.1. *Let the matrices U, V, Σ, M, X be those of the GSVD (1.2) of the matrix pair $\{A, L\}$ with the entries of Σ and M scaled to satisfy (1.4). Let the diagonal entries of the matrix*

$$(2.1) \quad D = \text{diag}[d_1, d_2, \dots, d_n] \in \mathbb{R}^{n \times n}$$

be positive and such that every column of the matrix

$$\tilde{X} = XD$$

is of unit Euclidean norm. Then the condition number of \tilde{X} , given by $\kappa(\tilde{X}) = \|\tilde{X}\|_2 \|\tilde{X}^{-1}\|_2$, is within a factor n of its smallest value over all choices of diagonal matrices D . Introduce the rescaled diagonal matrices

$$(2.2) \quad \begin{aligned} \tilde{\Sigma} &= \text{diag}[\tilde{\sigma}_1, \tilde{\sigma}_2, \dots, \tilde{\sigma}_p, d_{p+1}^{-1}, d_{p+2}^{-1}, \dots, d_n^{-1}] = \Sigma D^{-1}, \\ \tilde{M} &= [\text{diag}[\tilde{\mu}_1, \tilde{\mu}_2, \dots, \tilde{\mu}_p], \mathbf{0}, \mathbf{0}, \dots, \mathbf{0}] = M D^{-1}. \end{aligned}$$

Then the rescaled GSVD of the matrix pair $\{A, L\}$ is given by

$$(2.3) \quad A = U \tilde{\Sigma} \tilde{X}^T, \quad L = V \tilde{M} \tilde{X}^T.$$

The generalized singular values $\tilde{\gamma}_j = \tilde{\sigma}_j/\tilde{\mu}_j$, $1 \leq j \leq p$, associated with this decomposition are the same as the generalized singular values (1.5) associated with (1.2). When $L = I$, the $\tilde{\sigma}_j$ agree, up to ordering, with the singular values of A . Similarly, when $A = I$, the $\tilde{\mu}_j$ agree with the singular values of L .

Proof. It follows from (1.1) that the matrix X is nonsingular and, therefore, the diagonal matrix (2.1) is invertible. The property of $\kappa(\tilde{X})$ is shown by Demmel [7].

We turn to the situation when $L = I$. It follows from the representation (2.3) of L that

$$I = L^T L = \tilde{X} \tilde{M}^T V^T V \tilde{M} \tilde{X}^T.$$

Thus, $\tilde{X}^T \tilde{X} = \tilde{M}^{-2}$. The columns of \tilde{X} are scaled to be of unit length. Therefore $\tilde{M} = I$ and \tilde{X} is orthogonal. Hence, the representation (2.3) of A is a (standard) singular value decomposition up to the ordering of the singular values.

Let, instead, $A = I$. Then the representation (2.3) of A yields

$$I = A^T A = \tilde{X} \tilde{\Sigma}^T \tilde{\Sigma} \tilde{X}^T,$$

from which we conclude that $\tilde{\Sigma} = I$ and that \tilde{X} is orthogonal. It follows that the representation (2.3) of L is a (standard) singular value decomposition of L . \square

We remark that while, generally, $\kappa(\tilde{X}) < \kappa(X)$, the rescaling may result in that $\kappa(\tilde{\Sigma}) > \kappa(\Sigma)$. However, the condition number of $\tilde{\Sigma}$ typically is of no significance in the application to the solution of discrete ill-posed problems of concern in the present paper. We will comment further on this in Sections 3 and 4.

When condition (1.1) does not hold, the matrix X in (1.2) may have a zero column. We then set the corresponding diagonal entry of (2.1) to one.

The following result shows that the absolute and relative condition numbers for simple finite generalized singular values are not affected by our rescaling of the GSVD.

PROPOSITION 2.2. *Let $\gamma_j = \sigma_j/\mu_j = \tilde{\sigma}_j/\tilde{\mu}_j$ be a simple, finite, nonzero generalized singular value of the matrix pair $\{A, L\}$, and let $\mathbf{y}_j \in \mathbb{R}^n$ be the j th column of $Y = (X^T)^{-1}$ and $\tilde{\mathbf{y}}_j \in \mathbb{R}^n$ the j th column of $\tilde{Y} = (\tilde{X}^T)^{-1}$. Thus, \mathbf{y}_j and $\tilde{\mathbf{y}}_j$ are generalized singular vectors of the decompositions (1.2) and (2.3), respectively. Then the absolute condition number $\kappa_{abs}(\gamma_j)$ can be expressed as*

$$\kappa_{abs}(\gamma_j) = \frac{\|\mathbf{y}_j\|_2}{\mu_j}(1 + \gamma_j) = \frac{\|\tilde{\mathbf{y}}_j\|_2}{\tilde{\mu}_j}(1 + \gamma_j),$$

and the relative condition number $\kappa_{rel}(\gamma_j)$ is given by

$$\kappa_{rel}(\gamma_j) = \|\mathbf{y}_j\|_2 \left(\frac{\|A\|_2}{\sigma_j} + \frac{\|L\|_2}{\mu_j} \right) = \|\tilde{\mathbf{y}}_j\|_2 \left(\frac{\|A\|_2}{\tilde{\sigma}_j} + \frac{\|L\|_2}{\tilde{\mu}_j} \right).$$

Proof. The j th generalized singular vector \mathbf{y}_j associated with the decomposition (1.2) is divided by d_j in the rescaling of the decomposition, since $\tilde{Y} = YD^{-1}$. Analogously, σ_j and μ_j are replaced by $\tilde{\sigma}_j = \sigma_j/d_j$ and $\tilde{\mu}_j = \mu_j/d_j$, respectively. Thus, the formulas derived by Sun [23] for the absolute and relative condition numbers remain the same. \square

3. The solution of linear discrete ill-posed problems. Consider the computation of an approximate solution of the minimization problem

$$(3.1) \quad \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2,$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a matrix with many singular values of different orders of magnitude close to the origin. Minimization problems (3.1) with a matrix of this kind often are referred to as discrete ill-posed problems. They arise, for example, from the discretization of linear ill-posed problems, such as Fredholm integral equations of the first kind with a continuous kernel. The vector $\mathbf{b} \in \mathbb{R}^m$ in discrete ill-posed problems that arise in science and engineering often represents error-contaminated data.

Let $\mathbf{e} \in \mathbb{R}^m$ denote the (unknown) error in \mathbf{b} , and let $\check{\mathbf{b}} \in \mathbb{R}^m$ be the (unknown) error-free vector associated with \mathbf{b} , i.e.,

$$(3.2) \quad \mathbf{b} = \check{\mathbf{b}} + \mathbf{e}.$$

Let $\check{\mathbf{x}}$ denote the solution of minimal Euclidean norm of the least-squares problem obtained by replacing \mathbf{b} by $\check{\mathbf{b}}$ in (3.1). We would like to determine an approximation of $\check{\mathbf{x}}$ by computing a suitable approximate solution of the available least-squares problem (3.1). However, the presence of singular values of \mathbf{A} close to the origin makes the solution of (3.1) sensitive to the error \mathbf{e} in \mathbf{b} . In particular, the solution of (3.1) of minimal Euclidean norm typically is not a useful approximation of $\check{\mathbf{x}}$.

A common approach to remedy this difficulty is to replace the least-squares problem (3.1) by a nearby problem that is less sensitive to the error \mathbf{e} in \mathbf{b} . One of the most popular approaches is known as Tikhonov regularization, which replaces (3.1) by a penalized least-squares problem of the form

$$(3.3) \quad \min_{\mathbf{x} \in \mathbb{R}^n} \{\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{L}\mathbf{x}\|_2^2\}.$$

The matrix $\mathbf{L} \in \mathbb{R}^{p \times n}$ is referred to as the regularization matrix and the scalar $\lambda \geq 0$ as the regularization parameter. For thorough discussions on ill-posed and discrete ill-posed problems, as well as on Tikhonov regularization and other regularization methods, we refer to Engl et al. [10] and Hansen [13]. When (1.1) holds, the minimization problem (3.3) has a unique solution $\mathbf{x}_{L,\lambda}$ for any $\lambda > 0$. We required above that $p \leq n$. This inequality can be assumed to hold in the present application, because when $p > n$, we may replace \mathbf{L} in (3.3) by the upper triangular $n \times n$ matrix of its QR factorization.

The regularization parameter $\lambda > 0$ can be determined in many ways, including by the discrepancy principle, the L-curve criterion, generalized cross validation, and extrapolation; see, e.g., [5, 6, 10, 13, 16, 17, 21] for discussions on these and other methods. The discrepancy principle requires that a bound ε for the error in (3.2) be known and that $\check{\mathbf{b}}$ is in the range of \mathbf{A} . We will assume these requirements to hold, and will apply the discrepancy principle to determine $\lambda > 0$ in the computed examples of Section 5. Thus,

$$(3.4) \quad \|\mathbf{e}\|_2 \leq \varepsilon.$$

The discrepancy principle prescribes that the parameter $\lambda > 0$ be chosen so that

$$(3.5) \quad \|\mathbf{A}\mathbf{x}_{L,\lambda} - \mathbf{b}\|_2 = \eta\varepsilon$$

is satisfied, where $\eta \geq 1$ is a user-supplied parameter independent of ε . Typically, η is chosen close to unity when ε is known to be a fairly sharp upper bound for $\|e\|_2$.

Commonly used regularization matrices include the identity and finite difference matrices, such as

$$(3.6) \quad L' = \frac{1}{2} \begin{bmatrix} 1 & -1 & & & 0 \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ 0 & & & 1 & -1 \end{bmatrix} \in \mathbb{R}^{(n-1) \times n}$$

and

$$L'' = \frac{1}{4} \begin{bmatrix} -1 & 2 & -1 & & 0 \\ & -1 & 2 & -1 & \\ & & \ddots & \ddots & \ddots \\ 0 & & & -1 & 2 & -1 \end{bmatrix} \in \mathbb{R}^{(n-2) \times n}.$$

Also more complicated regularization matrices are discussed in the literature. For instance, for n even, the regularization matrix

$$(3.7) \quad \hat{L} = L''(I - \mathbf{w}\mathbf{w}^T),$$

where

$$\mathbf{w} = c [\cos(t_{n/2}), \cos(t_{n/2-1}), \dots, \cos(t_1), \cos(t_1), \dots, \cos(t_{n/2-1}), \cos(t_{n/2})]^T$$

with $t_j = \frac{2j-1}{n}\pi$, $j = \pm 1, \pm 2, \dots, \pm n/2$, and the scaling factor c chosen so that \mathbf{w} is of unit length, is considered in [9, Example 3.1]. Many other examples of regularization matrices can be found in [8, 9, 19, 22].

Substituting the rescaled GSVD (2.3) into (3.3) and letting $\mathbf{z} = \tilde{X}^T \mathbf{x}$ yields the penalized least-squares problem with diagonal matrices

$$\min_{\mathbf{z} \in \mathbb{R}^n} \{ \|\tilde{\Sigma} \mathbf{z} - U^T \mathbf{b}\|_2^2 + \lambda \|\tilde{M} \mathbf{z}\|_2^2 \},$$

whose solution can be determined by solving

$$(3.8) \quad \min_{\mathbf{z} \in \mathbb{R}^n} \left\| \begin{bmatrix} \tilde{\Sigma} \\ \lambda^{1/2} \tilde{M} \end{bmatrix} \mathbf{z} - \begin{bmatrix} U^T \mathbf{b} \\ \mathbf{0} \end{bmatrix} \right\|_2.$$

Assume for the moment that $m = n = p$. Then

$$\kappa \left(\begin{bmatrix} \tilde{\Sigma} \\ \lambda^{1/2} \tilde{M} \end{bmatrix} \right) = \frac{\max_{1 \leq j \leq p} (\tilde{\sigma}_j^2 + \lambda \tilde{\mu}_j^2)^{1/2}}{\min_{1 \leq j \leq p} (\tilde{\sigma}_j^2 + \lambda \tilde{\mu}_j^2)^{1/2}}.$$

Let the maximum and minimum of $j \rightarrow (\tilde{\sigma}_j^2 + \lambda \tilde{\mu}_j^2)^{1/2}$ be achieved for the indices j_{\max} and j_{\min} , respectively, and assume that $\tilde{\sigma}_{j_{\max}} \geq \tilde{\sigma}_{j_{\min}}$ and $\tilde{\mu}_{j_{\min}} > 0$. These inequalities typically hold. Then

$$\frac{(\tilde{\sigma}_{j_{\max}}^2 + \lambda \tilde{\mu}_{j_{\max}}^2)^{1/2}}{(\tilde{\sigma}_{j_{\min}}^2 + \lambda \tilde{\mu}_{j_{\min}}^2)^{1/2}} \leq \frac{\tilde{\sigma}_{j_{\max}}}{\tilde{\sigma}_{j_{\min}}}$$

and it follows that

$$\kappa \left(\begin{bmatrix} \tilde{\Sigma} \\ \lambda^{1/2} \tilde{M} \end{bmatrix} \right) \leq \kappa(\tilde{\Sigma}).$$

We conclude that due to regularization, the matrix in (3.8) generally has a smaller condition number than $\tilde{\Sigma}$. In particular, the scaling in (2.2) typically does not result in a least-squares problem with a very ill-conditioned matrix. We refer to [24, 25] for discussions on the conditioning of least-squares problems.

Let $\mathbf{z}_{L,\lambda}$ denote the solution of (3.8). Then the solution of (3.3) can be determined from

$$\mathbf{x}_{L,\lambda} = (\tilde{X}^T)^{-1} \mathbf{z}_{L,\lambda},$$

i.e., by solving a linear system of equations with the matrix \tilde{X}^T and right-hand side $\mathbf{z}_{L,\lambda}$. The condition number of the task of computing $\mathbf{x}_{L,\lambda}$ from $\mathbf{z}_{L,\lambda}$ is $\kappa(\tilde{X})$. It is beneficial if this condition number is small. If, instead, the standard GSVD (1.2) is substituted into (3.3), then the solution $\mathbf{x}_{L,\lambda}$ of (3.3) is determined in an analogous fashion by solving a linear system of equations with the matrix X^T . The use of the rescaled GSVD is advantageous since, generally, $\kappa(\tilde{X}) < \kappa(X)$.

4. TGSVD-type methods. Truncated GSVD (TGSVD) is a regularization method in which a truncation index is a discrete regularization parameter. Insightful discussions on this method are provided by Hansen [12, 13]. Let $U = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m]$ and X be the matrices in (1.2), and define $Y = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n] = X^{-T}$. Substituting the decomposition (1.2) of A into (3.1) yields

$$\min_{\mathbf{z} \in \mathbb{R}^n} \|\Sigma \mathbf{z} - U^T \mathbf{b}\|_2, \quad \mathbf{z} = [z_1, z_2, \dots, z_n]^T = X^T \mathbf{x}.$$

The TGSVD method restricts the above minimization to vectors \mathbf{z} whose first $p - k$ entries z_1, z_2, \dots, z_{p-k} are required to vanish for a user-chosen truncation index k . Then the computed solution only depends on the k largest σ_j . The solution of the restricted minimization problem is given by

$$(4.1) \quad \mathbf{z}_{L,k} = \left[0, 0, \dots, 0, \frac{\mathbf{u}_{p-k+1}^T \mathbf{b}}{\sigma_{p-k+1}}, \frac{\mathbf{u}_{p-k+2}^T \mathbf{b}}{\sigma_{p-k+2}}, \dots, \frac{\mathbf{u}_p^T \mathbf{b}}{\sigma_p}, \mathbf{u}_{p+1}^T \mathbf{b}, \mathbf{u}_{p+2}^T \mathbf{b}, \dots, \mathbf{u}_n^T \mathbf{b} \right]^T,$$

which yields the approximate solution

$$\mathbf{x}_{L,k} = Y \mathbf{z}_{L,k} = \sum_{j=p-k+1}^p \frac{\mathbf{u}_j^T \mathbf{b}}{\sigma_j} \mathbf{y}_j + \sum_{j=p+1}^n (\mathbf{u}_j^T \mathbf{b}) \mathbf{y}_j$$

of (3.1). The truncation index k is a regularization parameter and can be determined by the same methods commonly used to determine the regularization parameter λ in Tikhonov regularization (3.3). For instance, the discrepancy principle prescribes that the truncation index k be chosen as small as possible such that

$$(4.2) \quad \|A \mathbf{x}_{L,k} - \mathbf{b}\|_2 \leq \eta \varepsilon.$$

This inequality is analogous to (3.5) for Tikhonov regularization and requires a bound (3.4) to be known.

We turn to a modification of the TGSVD method described above. Introduce the diagonal, possibly rectangular, matrix

$$\Sigma_k = \text{diag}[0, \dots, 0, \sigma_{p-k+1}, \dots, \sigma_p, 1, 1, \dots, 1] \in \mathbb{R}^{m \times n}.$$

The vector $\mathbf{z}_{L,k}$ defined by (4.1) is the solution of minimal Euclidean norm of the minimization problem

$$\min_{\mathbf{z} \in \mathbb{R}^n} \|\Sigma_k \mathbf{z} - U^T \mathbf{b}\|_2.$$

The condition number of Σ_k is given by $\kappa(\Sigma_k) = \sigma_{p-k+1}^{-1}$. Let $\mathbb{M}_{m,n,p}$ be the set of block diagonal matrices with an arbitrary leading principal $p \times p$ block and a trailing principal $(m-p) \times (n-p)$ block that is made up of the first $n-p$ columns of the identity matrix of order $m-p$. The following proposition characterizes a closest matrix to Σ in the set $\mathbb{M}_{m,n,p}$ with the same condition number as Σ_k .

PROPOSITION 4.1. *Let the matrix $\widehat{\Sigma}_{\widehat{k}}$ be given by*

$$(4.3) \quad \widehat{\Sigma}_{\widehat{k}} = \text{diag}[0, \dots, 0, \widehat{\sigma}_{p-\widehat{k}+1}, \dots, \widehat{\sigma}_p, 1, 1, \dots, 1] \in \mathbb{R}^{m \times n},$$

where

$$\begin{aligned} \widehat{\sigma}_j &= 0, & 1 \leq j \leq p - \widehat{k}, \\ \widehat{\sigma}_j &= \sigma_{p-k+1}, & p - \widehat{k} + 1 \leq j \leq p - k, \\ \widehat{\sigma}_j &= \sigma_j, & p - k < j \leq p, \end{aligned}$$

and the index \widehat{k} is determined by the inequalities $\sigma_{p-\widehat{k}+1} \geq \sigma_{p-k+1}/2$ and $\sigma_{p-\widehat{k}} < \sigma_{p-k+1}/2$. Then $\widehat{\Sigma}_{\widehat{k}}$ is a closest matrix to Σ in the set $\mathbb{M}_{m,n,p}$ in the spectral norm with the same condition number as Σ_k .

Proof. The last $m-p$ rows and the last $n-p$ columns of the matrices in the set $\mathbb{M}_{m,n,p}$ agree with the last $m-p$ rows and the last $n-p$ columns, respectively, of the matrix Σ in (1.3). These rows and columns are related to the number of rows of A and the dimension of the null space of L , and should be kept fixed.

Let the leading principal $p \times p$ submatrix M_p of $M \in \mathbb{M}_{m,n,p}$ have the singular value decomposition $M_p = U_p \Sigma_p V_p^T$, where $U_p, V_p \in \mathbb{R}^{p \times p}$ are orthogonal and the nontrivial entries of the diagonal matrix Σ_p are sorted in increasing order. Denote the leading $p \times p$ principal submatrix of Σ , defined by (1.3), by Σ_0 . Then it follows from [15, Theorem 3.4.5] and the Fan dominance theorem [15, Corollary 3.5.9] that

$$\|\Sigma_0 - M_p\|_2 \geq \|\Sigma_0 - \Sigma_p\|_2.$$

Thus, the closest matrix to Σ_0 in the spectral norm is diagonal. It follows that the closest matrix to Σ is of the form (4.3) for certain entries $\widehat{\sigma}_j$ and index \widehat{k} . If $\sigma_{p-k+1} = 0$, then $\widehat{k} = p$ and $\widehat{\Sigma}_{\widehat{k}} = \Sigma$. When $\sigma_{p-k+1} > 0$, a closest matrix $\widehat{\Sigma}_{\widehat{k}}$ to Σ with smallest diagonal entry σ_{p-k+1} is obtained by modifying each diagonal entry of Σ that is smaller than σ_{p-k+1} as little as possible to be either σ_{p-k+1} or zero. Hence, the diagonal entries that are larger than or equal to $\sigma_{p-k+1}/2$ are set to σ_{p-k+1} and the remaining ones are set to zero. This determines a matrix $\widehat{\Sigma}_{\widehat{k}}$ such that $\kappa(\Sigma_k) = \kappa(\widehat{\Sigma}_{\widehat{k}})$ and

$$\begin{aligned} \|\Sigma - \widehat{\Sigma}_{\widehat{k}}\|_2 &= 0, & \text{if } \sigma_{p-k+1} = 0, \\ \|\Sigma - \widehat{\Sigma}_{\widehat{k}}\|_2 &= \sigma_{p-k}, & \text{if } \sigma_{p-k} < \frac{\sigma_{p-k+1}}{2}, \\ \|\Sigma - \widehat{\Sigma}_{\widehat{k}}\|_2 &\leq \frac{\sigma_{p-k+1}}{2}, & \text{if } 0 < \frac{\sigma_{p-k+1}}{2} \leq \sigma_{p-k}. \end{aligned}$$

Comparing with $\|\Sigma - \Sigma_k\|_2 = \sigma_{p-k}$ shows the proposition. \square

Let $\hat{\mathbf{z}}_{L,\hat{k}}$ denote the solution of minimal Euclidean norm of the minimization problem

$$\min_{\mathbf{z} \in \mathbb{R}^n} \|\hat{\Sigma}_{\hat{k}} \mathbf{z} - U^T \mathbf{b}\|_2.$$

Since $\kappa(\Sigma_k) = \kappa(\hat{\Sigma}_{\hat{k}})$ and $\hat{\Sigma}_{\hat{k}}$ is at least as close to Σ as Σ_k , the approximate solution $\hat{\mathbf{x}}_{L,\hat{k}} = Y \hat{\mathbf{z}}_{L,\hat{k}}$ can be expected to be of the same or higher quality as $\mathbf{x}_{L,k}$. This is confirmed by computed examples reported in Section 5. We refer to this method for computing $\hat{\mathbf{x}}_{L,\hat{k}}$ as the truncated modified GSVD (TMGSVD) method. An analogous modification of the truncated (standard) singular value decomposition is proposed in [20].

We have introduced the TMGSVD method because it is a simple and natural modification of the TGSVD method. However, we do not advocate its use for the solution of discrete ill-posed problems, because the computed results obtained with this method depend in an intricate way on the scaling of the matrices A and L , and the vector \mathbf{b} in (3.3). For instance, when A and \mathbf{b} are replaced by $10A$ and $10\mathbf{b}$, but L is kept fixed, generally all entries of the diagonal matrices Σ in (1.3) change, but not by a common factor. Therefore, the modification $\hat{\Sigma}_{\hat{k}}$ of Σ of Proposition 4.1, and the computed approximation $\hat{\mathbf{x}}_{L,\hat{k}}$ of $\tilde{\mathbf{x}}$, depend on the scaling of A , L , and \mathbf{b} . This difficulty can be circumvented by modifying the rescaled GSVD instead of the standard GSVD. We now describe this approach.

Consider truncated regularization methods based on the rescaled GSVD. Define the diagonal, possibly rectangular, matrix

$$(4.4) \quad \tilde{\Sigma}' = \text{diag}[\tilde{\sigma}_{j_1}, \tilde{\sigma}_{j_2}, \dots, \tilde{\sigma}_{j_p}, d_{p+1}^{-1}, d_{p+2}^{-1}, \dots, d_n^{-1}] \in \mathbb{R}^{m \times n},$$

where the $i \rightarrow j_i$ defines a renumbering such that

$$\tilde{\sigma}_{j_1} \leq \tilde{\sigma}_{j_2} \leq \dots \leq \tilde{\sigma}_{j_p}.$$

The last $n - p$ diagonal entries of (4.4), which are associated with the null space of L , are not renumbered. The same renumbering also is applied to the diagonal entries of the matrix \tilde{M} to give the new diagonal matrix \tilde{M}' , and we renumber the columns of the matrices \tilde{X} , V , and U in (2.3) accordingly so that the renumbering does not affect the rescaled GSVD Tikhonov solution $\mathbf{x}_{L,\lambda}$. We refer to the matrices with renumbered columns as \tilde{X}' , V' , and U' .

The renumbering may affect the approximate solution determined by the truncated rescaled GSVD (TRGSVD) method. Introduce for some $1 \leq k \leq p$ the matrix

$$(4.5) \quad \tilde{\Sigma}'_k = \text{diag}[0, \dots, 0, \tilde{\sigma}_{j_{p-k+1}}, \tilde{\sigma}_{j_{p-k+2}}, \dots, \tilde{\sigma}_{j_p}, d_{p+1}^{-1}, d_{p+2}^{-1}, \dots, d_n^{-1}] \in \mathbb{R}^{m \times n}$$

and consider the minimization problem

$$(4.6) \quad \min_{\mathbf{z} \in \mathbb{R}^n} \|\tilde{\Sigma}'_k \mathbf{z} - (U')^T \mathbf{b}\|_2.$$

Let $\tilde{\mathbf{z}}_{L,k}$ denote the solution of minimal Euclidean norm. Then $\tilde{\mathbf{x}}_{L,k} = (\tilde{X}')^{-T} \tilde{\mathbf{z}}_{L,k}$ is the TRGSVD approximate solution of (3.1).

The TRGSVD method can be modified similarly as TGSVD was modified to give the TMGSVD method. The modified method so obtained, which we refer to

as TMRGSVD, will be seen to perform the best among the truncated GSVD-type methods considered. Its derivation is completely analogous to the derivation of the TMGSVD method from TGSVD. We therefore omit the details.

We conclude this section with a comment on the conditioning of the matrix in (4.6). Regularization implies that the smallest $p - k$ of the p first diagonal entries of the matrix (4.4) are ignored in the solution process. This generally makes the matrix (4.5) better conditioned than the matrix (4.4). Therefore, the possible increase of the condition number of the matrix (4.4) due to rescaling generally is not a concern.

5. Computed examples. The calculations of this section were carried out in MATLAB with machine epsilon about $2.2 \cdot 10^{-16}$. We use examples from the MATLAB package Regularization Tools [13]. All examples but one are obtained by discretizing Fredholm integral equations of the first kind

$$(5.1) \quad \int_a^b h(s, t)x(t) dt = g(s), \quad c \leq s \leq d,$$

with a smooth kernel h . The discretizations are carried out by Galerkin or Nyström methods and yield linear discrete ill-posed problems (3.1). MATLAB functions in [13] determine discretizations $A \in \mathbb{R}^{n \times n}$ of the integral operators and scaled discrete approximations $\tilde{x} \in \mathbb{R}^n$ of the solution x of (5.1). In all examples, $n = 500$. We generate the “unknown” error-free vector $\tilde{b} = A\tilde{x}$ and add a “noise vector” $e \in \mathbb{R}^n$ with normally distributed random entries with zero mean to obtain the vector b in (3.1); cf. (3.2). The noise vector e is scaled to correspond to a specified noise level $\|e\|_2 / \|\tilde{b}\|_2$, which is assumed to be known.

We determine approximations of the desired solution \tilde{x} by the TGSVD, TRGSVD, and TMRGSVD methods described in Section 4, as well as by Tikhonov regularization considered in Section 3. The performance of several regularization matrices is illustrated. The regularization parameters are determined by the discrepancy principle with $\varepsilon = \|e\|$ and $\eta = 1.01$ in (3.5) and (4.2). To gain insight into the behavior of the solution methods, we report in every example the average of the relative errors in the computed solution over 1000 runs with different noise vectors e for each specified noise level.

method	average k	average relative error
Tikhonov	--	0.1318
TGSVD	1.000	0.1029
TRGSVD	1.000	0.1026
TMRGSVD	1.000	0.1026

TABLE 5.1

Example 5.1: Average relative errors in the computed approximate solutions for the deriv2 test problem and average truncation indices for the truncated GSVD-type methods for $L = L'$ and noise level 10^{-1} .

Example 5.1. We first consider the problem deriv2 from [13]. The kernel, solution, and right-hand side of (5.1) are given by

$$\begin{aligned} h(s, t) &= \begin{cases} s(t-1), & s < t, \\ t(s-1), & s \geq t, \end{cases} \\ x(t) &= t, \\ g(s) &= \frac{s^3 - s}{6}, \end{aligned}$$

and $a = c = 0$, $b = d = 1$. Thus, the kernel h is the Green's function for the second derivative. Table 5.1 shows the average relative errors in the computed solutions determined by Tikhonov regularization and truncated GSVD-type methods. The average truncation indices for the latter methods also are displayed. The regularization matrix $L = L'$ defined by (3.6) is used and the regularization parameters are determined by the discrepancy principle. The noise level is 10^{-1} . Note that all truncated GSVD-type methods yield better approximations of the desired solution $\tilde{\mathbf{x}}$ than Tikhonov regularization. We have $\kappa(X) = 1 \cdot 10^2$ and $\kappa(\tilde{X}) = 1 \cdot 10^1$. \square

method	average k	average relative error
Tikhonov	--	0.1028
TGSVD	15.66	0.1120
TRGSVD	15.66	0.1120
TMRGSVD	19.67	0.0942

TABLE 5.2

Example 5.2: Average relative errors in the computed approximate solutions for the heat test problem and average truncation indices for the truncated GSVD-type methods for $L = L'$ and noise level 10^{-2} .

Example 5.2. The test problem **heat** from [13] is a discretization of a Volterra integral equation of the first kind on the interval $[0, 1]$ with a convolution kernel. Table 5.2 shows the average relative errors in the computed solutions determined by Tikhonov regularization and truncated GSVD-type methods for the noise level 10^{-2} . Also the average truncation index for the truncated GSVD-type methods is displayed. The regularization matrix is $L = L'$ defined by (3.6). The smallest average error is furnished by TMRGSVD. We have $\kappa(X) = 9 \cdot 10^1$ and $\kappa(\tilde{X}) = 2 \cdot 10^1$.

To further illustrate the performance of the truncated GSVD-type methods, we display in Table 5.3 the relative errors and truncation indices for one run. The TMRGSVD method yields an approximate solution with the largest k -value and provides the most accurate approximation of $\tilde{\mathbf{x}}$. Figure 5.1 depicts the error in the approximate solutions determined by the truncated GSVD-type methods for several k -values. The figure shows that the discrepancy principle when used with the TMRGSVD method determines a k -value that gives an approximate solution with the smallest error. This is not the case for the TGSVD and TRGSVD methods. However, even the best k -value for the latter methods ($k = 21$) yields approximate solutions with larger errors than the solution obtained with the TMRGSVD method.

We conclude this example with an illustration of the performance of the TMGSVD method. As mentioned in Section 4, the computed approximate solutions determined with this method depend on the scaling of A , L , and \mathbf{b} . When we solve the problem of this example 1000 times with different noise realizations, all corresponding to the noise level 10^{-2} , we obtain the average relative error 0.1011 and average truncation index $k = 18.66$. This average relative error is smaller than the average relative errors for all methods of Table 5.2 except for TMRGSVD. If we replace the matrix A and vector \mathbf{b} by $10A$ and $10\mathbf{b}$, respectively, and keep the regularization matrix L fixed, then the TMGSVD method yields the average relative error 0.0996 and average truncation index $k = 19.05$ over 1000 runs. Thus, the scaling reduces the average relative error in the computed approximations of $\tilde{\mathbf{x}}$ determined by TMGSVD. Since it is not obvious how to scale the matrices A and L , and vector \mathbf{b} , to obtain the minimal average relative error in the computed approximate solutions, and because TMRGSVD delivers a smaller average relative error, we propose that the latter method be used instead of




FIG. 5.1. Example 5.2. Errors in the computed approximate solutions given by TGSVD (solid curve), TRGSVD (dash-dotted curve), and TMRGSVD (dashed curve). The curves for the TGSVD and TRGSVD methods are indistinguishable.

TMRGSVD. \square

method	k	relative error
Tikhonov	--	0.0735
TGSVD	18	0.1005
TRGSVD	18	0.1005
TMRGSVD	23	0.0572

TABLE 5.3

Example 5.2: Relative errors in the computed approximate solutions for the **heat** test problem and truncation indices for the truncated GSVD-type methods for $L = L'$ and noise level 10^{-2} .

method	average k	average relative error
Tikhonov	--	0.1073
TGSVD	2.03	0.0938
TRGSVD	2.03	0.0938
TMRGSVD	3.03	0.0628

TABLE 5.4

Example 5.3: Average relative errors in the computed approximate solutions for the **gravity** test problem and average truncation indices for the truncated GSVD-type methods for $L = \hat{L}$ and noise level 10^{-1} .

Example 5.3. The test problem **gravity** from [13] is a discretization of a Fredholm integral equation of the first kind with kernel $h(s, t) = \delta(\delta^2 + (s - t)^2)^{-3/2}$ and $a = c = 0$, $b = d = 1$ in (5.1). The right-hand side function g in (5.1) is chosen so that the integral equation has the solution $x(t) = \sin(\pi t) + \sin(2\pi t)/2$. Discretization is carried out by a Nyström method based on the midpoint quadrature rule. The result-

method	average k	average relative error
Tikhonov	—	0.0315
TGSVD	3.53	0.0305
TRGSVD	3.53	0.0305
TMRGSVD	4.53	0.0227

TABLE 5.5

Example 5.3: Average relative errors in the computed approximate solutions for the gravity test problem and average truncation indices for the truncated GSVD-type methods for $L = \hat{L}$ and noise level 10^{-2} .

ing matrix A is symmetric and Toeplitz. The parameter δ is the depth at which the magnetic deposit is located in the model. We use the default value $\delta = 0.25$. Tables 5.4 and 5.5 show the average relative errors in the computed approximate solutions as well as the average truncation indices for the truncated GSVD-type methods for the noise levels 10^{-1} and 10^{-2} . The regularization matrix (3.7) is used. The TMRGSVD method is seen to yield the smallest average relative errors in the computed approximate solutions for both noise levels. We have $\kappa(X) = 1 \cdot 10^2$ and $\kappa(\tilde{X}) = 3 \cdot 10^1$. \square

6. Conclusion. This paper proposes a rescaling of the GSVD that is designed to reduce the condition number of the square nonsingular matrix X in (1.3). The rescaling and the solution of a matrix nearness problem give rise to new truncated GSVD-type regularization methods for the solution of discrete ill-posed problems. Computed examples show one of these methods, TMRGSVD, to generally determine approximate solutions of higher quality than the standard TGSVD method and Tikhonov regularization.

Acknowledgment. We would like to thank the referee for comments that lead to improvements of the presentation.

REFERENCES

- [1] O. Alter, P. O. Brown, and D. Botstein, *Generalized singular value decomposition for comparative analysis of genome-scale expression datasets of two different organisms*, PNAS USA, 100 (2003), pp. 3351–3356.
- [2] Z. Bai, *The CSD, GSVD, their applications and computation*, IMA preprint 958, Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, MN, 1992.
- [3] Z. Bai and J. W. Demmel, *Computing the generalized singular value decomposition*, SIAM J. Sci. Comput., 14 (1993), pp. 1464–1486.
- [4] Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, 1996.
- [5] C. Brezinski, G. Rodriguez, and S. Seatzu, *Error estimates for linear systems with applications to regularization*, Numer. Algorithms, 49 (2008), pp. 85–104.
- [6] C. Brezinski, G. Rodriguez, and S. Seatzu, *Error estimates for the regularization of least squares problems*, Numer. Algorithms, 51 (2009), pp. 61–76.
- [7] J. Demmel, *The condition number of equivalence transformations that block diagonalize matrix pencils*, SIAM J. Numer. Anal., 20 (1983), pp. 599–610.
- [8] M. Donatelli and L. Reichel, *Square smoothing regularization matrices with accurate boundary conditions*, J. Comput. Appl. Math., in press.
- [9] L. Dykes and L. Reichel, *On the reduction of Tikhonov minimization problems and the construction of regularization matrices*, Numer. Algorithms, 60 (2012), pp. 683–696.
- [10] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Kluwer, Dordrecht, 1996.
- [11] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, 2013.
- [12] P. C. Hansen, *Regularization, GSVD and truncated GSVD*, BIT, 29 (1989), pp. 491–504.

- [13] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM, Philadelphia, 1998.
- [14] P. C. Hansen, *Regularization tools version 4.0 for Matlab 7.3*, Numer. Algorithms, 46 (2007), pp. 189–194.
- [15] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.
- [16] S. Kindermann, *Convergence analysis of minimization-based noise level-free parameter choice rules for linear ill-posed problems*, Electron. Trans. Numer. Anal., 38 (2011), pp. 233–257.
- [17] S. Kindermann, *Discretization independent convergence rates for noise level-free parameter choice rules for the regularization of ill-conditioned problems*, Electron. Trans. Numer. Anal., 40 (2013), pp. 58–81.
- [18] C. H. Lee, B. O. Alpert, P. Sankaranarayanan, and O. Alter, *GSVD comparison of patient-matched normal and tumor aCGH profiles reveals global copy-number alterations predicting glioblastoma multiforme survival*, PLoS One, 7 (1) (2012), article e30098.
- [19] S. Noschese and L. Reichel, *Inverse problems for regularization matrices*, Numer. Algorithms, 60 (2012), pp. 531–544.
- [20] S. Noschese and L. Reichel, *A modified TSVD method for discrete ill-posed problems*, Numer. Linear Algebra Appl., in press.
- [21] L. Reichel and G. Rodriguez, *Old and new parameter choice rules for discrete ill-posed problems*, Numer. Algorithms, 63 (2013), pp. 65–87.
- [22] L. Reichel and Q. Ye, *Simple square smoothing regularization operators*, Electron. Trans. Numer. Anal., 33 (2009), pp. 63–83.
- [23] J. G. Sun, *Condition number and backward error for the generalized singular value decomposition*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 323–341.
- [24] L. N. Trefethen and D. Bau, III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [25] A. van der Sluis, *Stability of the solution of linear least squares problems*, Numer. Math., 23 (1975), pp. 241–254.