Journal of Computational and Applied Mathematics 235 (2010) 1053–1064

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Subspace-restricted singular value decompositions for linear discrete ill-posed problems

Michiel E. Hochstenbach[a], Lothar Reichel[b,*]

[a] Department of Mathematics and Computer Science, Eindhoven University of Technology, PO Box 513, 5600 MB, The Netherlands

[b] Department of Mathematical Sciences, Kent State University, Kent, OH 44242, USA

ARTICLE INFO

Article history:
Received 16 February 2010
Received in revised form 24 May 2010

Dedicated to Adhemar Bultheel on the occasion of his 60th birthday.

MSC:
65R30
65R32
65F10

Keywords:
Ill-posed problem
Inverse problem
Modified SVD
TSVD
SRSVD
Tikhonov regularization

ABSTRACT

The truncated singular value decomposition is a popular solution method for linear discrete ill-posed problems. These problems are numerically underdetermined. Therefore, it can be beneficial to incorporate information about the desired solution into the solution process. This paper describes a modification of the singular value decomposition that permits a specified linear subspace to be contained in the solution subspace for all truncations. Modifications that allow the range to contain a specified subspace, or that allow both the solution subspace and the range to contain specified subspaces also are described.

1. Introduction

The truncated singular value decomposition is commonly used to solve linear discrete ill-posed problems with matrices of small to moderate size. The truncated subspace-restricted singular value decomposition of this paper is a modification, that allows a user to choose subspaces of the domain and range, which can be used in the solution process for all truncations. Our interest in the truncated subspace-restricted singular value decomposition stems from its applicability to the solution of linear discrete ill-posed problems. We first describe this application to motivate our modification of the standard truncated singular value decomposition.

We are concerned with the solution of linear systems of equations

$$A\boldsymbol{x} = \boldsymbol{b}, \quad A \in \mathbb{R}^{m\times n},\ \boldsymbol{x} \in \mathbb{R}^n,\ \boldsymbol{b} \in \mathbb{R}^m, \tag{1.1}$$

with a matrix A of ill-determined rank. Such systems are often referred to as linear discrete ill-posed problems. The singular values of A cluster at the origin and this makes the matrix severely ill-conditioned. In particular, the matrix may be singular. We consider (1.1) a least-squares problem in case the system is inconsistent. The right-hand side $\boldsymbol{b}$ is assumed to be

* Corresponding author.
E-mail address: reichel@math.kent.edu (L. Reichel).
URL: http://www.win.tue.nl/~hochsten (M.E. Hochstenbach).

0377-0427/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.cam.2010.06.016

contaminated by an error $\boldsymbol{e} \in \mathbb{R}^m$, which may stem from discretization or measurement inaccuracies. For notational simplicity, we will assume that $m \geq n$; however, the method of this paper, suitably modified, also can be applied when $m < n$.

Let $\hat{\boldsymbol{b}}$ denote the unknown error-free vector associated with $\boldsymbol{b}$, i.e.,

$$\boldsymbol{b} = \hat{\boldsymbol{b}} + \boldsymbol{e}, \tag{1.2}$$

and assume that the linear system

$$A\boldsymbol{x} = \hat{\boldsymbol{b}} \tag{1.3}$$

is consistent. We would like to determine the solution $\hat{\boldsymbol{x}}$ of (1.3) of minimal Euclidean norm. Since the right-hand side $\hat{\boldsymbol{b}}$ is not available, we seek to determine an approximation of $\hat{\boldsymbol{x}}$ by computing an approximate solution of the available linear system of (1.1). When the linear system (1.1) is of small to moderate size, this is often done with the aid of the Singular Value Decomposition (SVD) of A,

$$A = U\Sigma V^T. \tag{1.4}$$

Here $U = [\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_m] \in \mathbb{R}^{m\times m}$ and $V = [\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_n] \in \mathbb{R}^{n\times n}$ are orthogonal matrices, and the singular values are the diagonal entries of $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_n) \in \mathbb{R}^{m\times n}$. They are ordered according to

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_\ell > \sigma_{\ell+1} = \cdots = \sigma_n = 0, \quad \ell = \mathrm{rank}(A); \tag{1.5}$$

see, e.g., [1] for details on the SVD. Using (1.4), the system (1.1) can be expressed as

$$\Sigma\boldsymbol{y} = U^T\boldsymbol{b}, \qquad \boldsymbol{x} = V\boldsymbol{y}. \tag{1.6}$$

Let $\Sigma_k = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_k, 0, \ldots, 0)$ be obtained by setting the last $n - k$ diagonal entries of Σ to zero. The Truncated SVD (TSVD) method replaces Σ by Σ_k in (1.6) and determines the least-squares solutions $\boldsymbol{y}_k$ of minimal Euclidean norm of the system so obtained. The associated approximate solutions $\boldsymbol{x}_k$ of (1.1) are given by

$$\boldsymbol{x}_k = V\boldsymbol{y}_k = \sum_{j=1}^{k} \frac{\boldsymbol{u}_j^T\boldsymbol{b}}{\sigma_j}\boldsymbol{v}_j, \quad k = 1, 2, \ldots, \ell. \tag{1.7}$$

We note that $\boldsymbol{x}_k \in \mathrm{span}\{\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_k\}$ and define $\boldsymbol{x}_0 = \boldsymbol{0}$. The singular values σ_j and the Fourier coefficients $\boldsymbol{u}_j^T\boldsymbol{b}$ provide valuable insight into the properties of the linear discrete ill-posed problem (1.1); see, e.g., [2] for a discussion on the application of the TSVD to linear discrete ill-posed problems.

Let $\|\cdot\|$ denote the Euclidean vector norm or the associated induced matrix norm, and consider the sequence $\eta_k = \|\boldsymbol{x}_k - \hat{\boldsymbol{x}}\|$, $k = 0, 1, \ldots, \ell$. Generally, the η_k decrease when k increases and k is fairly small. Due to the error $\boldsymbol{e}$ in the right-hand side $\boldsymbol{b}$ and the ill-conditioning of A, the η_k typically increase rapidly with k when k is large. Let $k_* \geq 0$ be the smallest index, such that

$$\|\boldsymbol{x}_{k_*} - \hat{\boldsymbol{x}}\| = \min_{0\leq k\leq \ell} \|\boldsymbol{x}_k - \hat{\boldsymbol{x}}\|. \tag{1.8}$$

The index k_* generally is not explicitly known.

In the computed examples of Section 5, we assume that an estimate δ of the norm of the error $\boldsymbol{e}$ in $\boldsymbol{b}$ is available. The norm of the residual vectors

$$\boldsymbol{r}_k = \boldsymbol{b} - A\boldsymbol{x}_k$$

is a decreasing function of k, with $\boldsymbol{r}_k = P_{\mathcal{N}(A^T)}\boldsymbol{b}$ for $\ell < k \leq n$, where $P_{\mathcal{N}(A^T)}$ denotes the orthogonal projector onto the null space $\mathcal{N}(A^T)$ of A^T. The discrepancy principle suggests that the smallest integer $k \geq 0$, such that

$$\|\boldsymbol{r}_k\| \leq \gamma\delta, \tag{1.9}$$

be used as an approximation of k_*, where $\gamma > 1$ is a user-supplied constant. We denote this integer by k_{discr} and the associated approximation of $\hat{\boldsymbol{x}}$ by $\boldsymbol{x}_{k_{\mathrm{discr}}}$; see, e.g., [3] for further discussion on the discrepancy principle.

For many linear discrete ill-posed problems (1.1), the approximate solution $\boldsymbol{x}_{k_{\mathrm{discr}}}$ furnished by TSVD and the discrepancy principle is a fairly accurate approximation of $\hat{\boldsymbol{x}}$. However, there are linear discrete ill-posed problems (1.1) for which this is not the case. The latter situation arises when the subspace $\mathrm{span}\{\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_{k_{\mathrm{discr}}}\}$ does not contain an accurate approximation of $\hat{\boldsymbol{x}}$. A choice of $k > k_{\mathrm{discr}}$ is often not feasible, since for these k-values the propagated error, due to the error $\boldsymbol{e}$ in $\boldsymbol{b}$, generally destroys the accuracy in $\boldsymbol{x}_k$.

Various examples which illustrate that other solution methods may determine approximations of $\hat{\boldsymbol{x}}$ of higher accuracy than TSVD can be found in the literature; see, e.g., [4–7]. Further illustrations are provided in Section 5. Here we only note that the cause for poor accuracy generally is not the choice k_{discr} of the truncation index; the difference $\boldsymbol{x}_{k_*} - \hat{\boldsymbol{x}}$ is often not much smaller than $\boldsymbol{x}_{k_{\mathrm{discr}}} - \hat{\boldsymbol{x}}$.

This paper describes modifications of the SVD, such that truncated versions can give more accurate approximations of $\hat{\boldsymbol{x}}$ than the TSVD. A user may choose a subspace $\mathcal{W} \subset \mathbb{R}^n$ that allows the representation of known important features of $\hat{\boldsymbol{x}}$. Let $p = \dim(\mathcal{W})$ and assume that $p < n$. Typically, p is quite small in applications, say, $1 \le p \le 5$. An orthonormal basis of $\mathcal{W}$ makes up the p last columns of the matrix $\tilde{V}$ of the SVD-like decomposition

$$A = \tilde{U}\tilde{S}\tilde{V}^T, \tag{1.10}$$

where $\tilde{U} \in \mathbb{R}^{m\times m}$ and $\tilde{V} \in \mathbb{R}^{n\times n}$ are orthogonal matrices. The leading $n-p$ columns of the matrix $\tilde{S} \in \mathbb{R}^{m\times n}$ form a diagonal matrix. We refer to (1.10) as a subspace-restricted SVD (SRSVD). We also allow some columns of the matrix $\tilde{U}$ in (1.10) to be prescribed. Details of these decompositions are described in Section 2.

It is often meaningful to require that

$$\mathcal{W} \cap \mathcal{N}(A) = \{\mathbf{0}\} \tag{1.11}$$

to avoid that the computed approximate solution of (1.1) contains a significant component in $\mathcal{N}(A)$. The linear discrete ill-posed problems considered in this paper are discretizations of linear compact operator equations. For these kinds of problems, vectors in $\mathcal{N}(A)$ typically represent discretizations of highly oscillatory functions. We are interested in spaces $\mathcal{W}$ that represent slowly varying functions.

Example 1.1. If $\hat{\boldsymbol{x}}$ is known to model a nearly constant function, then it may be beneficial to let

$$\mathcal{W} = \text{range}\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} \tag{1.12}$$

be in the solution subspace. □

Example 1.2. Let $\hat{\boldsymbol{x}}$ be the discretization of a function that can be well approximated by a linear function. Then the SRSVD with

$$\mathcal{W} = \text{range}\begin{bmatrix} 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ 1 & n \end{bmatrix}$$

in the solution subspace may yield a more accurate approximations of $\hat{\boldsymbol{x}}$ than can be determined with the TSVD. □

A variety of modifications and extensions of the SVD, among them the Generalized Singular Value Decomposition (GSVD), have been applied to the solution of linear discrete ill-posed problems with the aim of obtaining decompositions that are more suitable for particular linear discrete ill-posed problems than the SVD; see, e.g., [8,6,7] and the references therein. The GSVD determines factorizations of the matrices in the pair $\{A, L\}$, where A is the matrix in (1.1) and $L \in \mathbb{R}^{(n-p)\times n}$ is a user-chosen regularization operator with $0 \le p < n$. The solution subspace determined by the GSVD contains $\mathcal{N}(L)$; see [1,8].

Example 1.3. The bidiagonal matrix

$$L = \begin{bmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \end{bmatrix} \in \mathbb{R}^{(n-1)\times n}$$

is a commonly used regularization operator. The solution subspace determined by the GSVD contains $\mathcal{N}(L)$, which is given by (1.12). □

The above example illustrates that we can make the solution subspace determined by the GSVD contain a desired subspace $\mathcal{W}$ by choosing a regularization operator L with $\mathcal{N}(L) = \mathcal{W}$. This can be achieved in a more straightforward manner with the SRSVD.

The SRSVD of a matrix is determined by computing the SVD of an orthogonal projection of the matrix. A different way to enforce the solution subspace to contain a user-specified subspace $\mathcal{W}$ by initial orthogonal projection is described in [7]. We comment on the differences between these approaches in Section 2.

This paper is organized as follows. Section 2 describes the SRSVD. Application of the SRSVD to Tikhonov regularization is considered in Section 3 and Section 4 provides theoretical comparisons to other methods. Computed examples can be found in Section 5 and Section 6 contains concluding remarks.

The present paper blends linear algebra and signal processing, areas in which Adhemar Bultheel over the years has made numerous important contributions; see, e.g., [9–12]. It is a pleasure to dedicate this paper to him.

2. Subspace-restricted singular value decompositions

This section introduces several SRSVDs and discusses their application in the Truncated SRSVD (TSRSVD) method to the solution of linear discrete ill-posed problems (1.1). Assume that we would like the solution subspace for any truncation to contain the subspace $\mathcal{W}$ of dimension p. Let the columns of the matrix $W \in \mathbb{R}^{n\times p}$ form an orthonormal basis of $\mathcal{W}$ and let (1.11) hold. Define the orthogonal projectors

$$P_W = WW^T, \qquad P_W^{\perp} = I - P_W,$$

where I denotes the identity. Introduce the singular value decomposition

$$AP_W^{\perp} = \tilde{U}\tilde{\Sigma}\tilde{V}^T, \tag{2.1}$$

where $\tilde{U} \in \mathbb{R}^{m\times m}$ and $\tilde{V} \in \mathbb{R}^{n\times n}$ are orthogonal matrices, and the singular values are the diagonal entries of $\tilde{\Sigma} = \text{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_n) \in \mathbb{R}^{m\times n}$, ordered so that

$$\tilde{\sigma}_1 \geq \tilde{\sigma}_2 \geq \cdots \geq \tilde{\sigma}_{\tilde{\ell}} > \tilde{\sigma}_{\tilde{\ell}+1} = \cdots = \tilde{\sigma}_n = 0$$

for some integer $\tilde{\ell}$, such that

$$\max\{0, \ell - p\} \leq \tilde{\ell} \leq \min\{\ell, n - p\},$$

where $\tilde{\ell} = 0$ when $AP_W^{\perp} = O$. The lower bound is achieved when the space $\mathcal{W}$ is orthogonal to $\mathcal{N}(A)$. We may choose the trailing $n \times p$ submatrix of $\tilde{V}$ to be W, i.e., $\tilde{V}$ is of the form

$$\tilde{V} = [\tilde{V}_1, W], \quad \tilde{V}_1 \in \mathbb{R}^{n\times(n-p)}. \tag{2.2}$$

Theorem 2.1. *Let the matrices $\tilde{U}$, $\tilde{\Sigma}$, and $\tilde{V}$ be determined by (2.1) and (2.2). Then the matrix $\tilde{S} \in \mathbb{R}^{m\times n}$ in the subspace-restricted singular value decomposition*

$$A = \tilde{U}\tilde{S}\tilde{V}^T \tag{2.3}$$

is of the form

$$\tilde{S} = [\tilde{\Sigma}_1,\ \tilde{B}], \tag{2.4}$$

where $\tilde{\Sigma}_1 = \text{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_{n-p})$ is the leading $m \times (n-p)$ submatrix of $\tilde{\Sigma}$ and $\tilde{B} = \tilde{U}^T AW$. Moreover,

$$\sigma_j \geq \tilde{\sigma}_j \geq \sigma_{j+p}, \quad 1 \leq j \leq n, \tag{2.5}$$

where the σ_j, for $1 \leq j \leq n$, are the singular values of A, cf. (1.5), and we define $\sigma_j = 0$ for $j > n$.

Proof. The decomposition (2.3) with $\tilde{S}$ given by (2.4) follows from (2.1) and the fact that $\tilde{\sigma}_j = 0$ for $j > n - p$. To show the inequalities (2.5), we observe that the matrices A and $A\tilde{V} = [A\tilde{V}_1, AW]$ have the same singular values. Moreover, the matrices $AP_W^{\perp}$ and $AP_W^{\perp}\tilde{V} = [A\tilde{V}_1, O]$ have the same singular values. The latter matrix is obtained by replacing the submatrix $AW \in \mathbb{R}^{m\times p}$ of $A\tilde{V}$ by the zero matrix. The inequalities (2.5) now follow from inequalities for singular values of a submatrix; see, e.g., [13, Corollary 3.1.3] for a proof. □

Truncation of the decomposition (2.3) can be used to determine approximate solutions of (1.1) similarly as with the TSVD. The TSRSVD method so obtained proceeds as follows. We may assume that the subspace $\mathcal{W}$ is chosen so that the restriction of A to $\mathcal{W}$ is well conditioned, i.e., that the matrix AW has a small to moderate condition number

$$\kappa(AW) = \frac{\max\limits_{\|\boldsymbol{y}\|=1} \|AW\boldsymbol{y}\|}{\min\limits_{\|\boldsymbol{y}\|=1} \|AW\boldsymbol{y}\|}; \tag{2.6}$$

if $\kappa(AW)$ is large, then we choose a different space $\mathcal{W}$. Since $\kappa(\tilde{B}) = \kappa(AW)$, it follows that the columns of $\tilde{B}$ are not nearly linearly dependent. Moreover, we would like the spaces $A\mathcal{W}$ and $A\mathcal{W}^{\perp}$ to be fairly well separated. These requirements typically are satisfied when the matrix A is the discretization of a compact operator and the space $\mathcal{W}$ represents slowly varying functions. When these conditions on $\mathcal{W}$, $A\mathcal{W}$, and $A\mathcal{W}^{\perp}$ are satisfied, the least-squares problem (2.7) below can be solved rapidly with the aid of Givens rotation in a straightforward way. Of course, for small matrices A, it may be attractive to solve (2.7) by computing the SVD of the matrix $\tilde{S}^{(k)}$. This approach is more expensive, but no conditions on the spaces $\mathcal{W}$, $A\mathcal{W}$, and $A\mathcal{W}^{\perp}$ have to be imposed.

Similarly to the representation (1.6) of (1.1) based on (1.4), we have the representation

$$\tilde{S}\boldsymbol{y} = \tilde{U}^T\boldsymbol{b}, \qquad \boldsymbol{x} = \tilde{V}\boldsymbol{y},$$

determined by (2.3). Introduce the truncated versions of the matrix $\tilde{S}$ defined by (2.4),

$$\tilde{S}^{(k)} = [\tilde{\Sigma}_1^{(k)},\ \tilde{B}], \quad k = 1, 2, \ldots, n-p,$$

where $\tilde{\Sigma}_1^{(k)} = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_k, 0, \ldots, 0)$ is obtained by setting the last $n-p-k$ diagonal entries of $\tilde{\Sigma}_1$ to zero. Let $\tilde{\boldsymbol{y}}_k$ denote the minimal-norm least-squares solution of

$$\min_{\tilde{\boldsymbol{y}} \in \mathbb{R}^n} \|\tilde{S}^{(k)}\tilde{\boldsymbol{y}} - \tilde{U}^T\boldsymbol{b}\|. \tag{2.7}$$

Then the associated TSRSVD solutions of (1.1) are

$$\tilde{\boldsymbol{x}}_k = \tilde{V}\tilde{\boldsymbol{y}}_k, \quad k = 1, 2, \ldots, n-p. \tag{2.8}$$

The discrepancy principle prescribes that we choose k to be the smallest nonnegative integer, such that the associated residual vector

$$\tilde{\boldsymbol{r}}_k = \boldsymbol{b} - A\tilde{\boldsymbol{x}}_k$$

satisfies (1.9). The computations of the $\tilde{\boldsymbol{y}}_k$ and $\|\tilde{\boldsymbol{r}}_k\|$ can be carried out efficiently for decreasing values of k by applying Givens rotations.

Instead of prescribing columns of $\tilde{V}$ in the decomposition (2.3), we may also specify columns of $\tilde{U}$. The corresponding decomposition can be derived by replacing A by A^T in (2.3). We outline the decomposition. Let the matrix $\hat{W} \in \mathbb{R}^{m\times\hat{p}}$ have orthonormal columns and introduce the orthogonal projectors

$$P_{\hat{W}} = \hat{W}\hat{W}^T, \qquad P_{\hat{W}}^{\perp} = I - P_{\hat{W}}.$$

Let $\mathcal{R}(\hat{W})$ denote the range of $\hat{W}$ and assume that

$$q = \dim(\mathcal{R}(\hat{W}) \cap \mathcal{R}(A)). \tag{2.9}$$

Consider the singular value decomposition

$$P_{\hat{W}}^{\perp}A = \tilde{U}\tilde{\Sigma}\tilde{V}^T, \tag{2.10}$$

where the matrices $\tilde{U} \in \mathbb{R}^{m\times m}$ and $\tilde{V} \in \mathbb{R}^{n\times n}$ are orthogonal, and the singular values are the diagonal entries of $\tilde{\Sigma} = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_n) \in \mathbb{R}^{m\times n}$ with

$$\tilde{\sigma}_1 \geq \tilde{\sigma}_2 \geq \cdots \geq \tilde{\sigma}_{\tilde{\ell}} > \tilde{\sigma}_{\tilde{\ell}+1} = \cdots = \tilde{\sigma}_n = 0, \quad \tilde{\ell} = \ell - q.$$

The value of $\tilde{\ell}$ follows from (1.5) and (2.9). Introduce the matrix

$$\hat{U} = [\tilde{U}_1,\ \hat{W},\ \tilde{U}_2], \tag{2.11}$$

where $\tilde{U}_1$ is the leading $m \times (n-q)$ submatrix of $\tilde{U}$ and $\tilde{U}_2 \in \mathbb{R}^{m\times(m-n+q-\hat{p})}$ has orthonormal columns that are orthogonal to the columns of the matrices U_1 and $\hat{W}$. In particular, $\mathcal{R}(\tilde{U}_2) \perp \mathcal{R}(A)$. The following result is analogous to Theorem 2.1 and can be shown in a similar manner.

Theorem 2.2. *Assume that (2.9) holds and let the matrices $\tilde{U}$, $\tilde{\Sigma}$, and $\tilde{V}$ be determined by (2.10) and (2.11). Then the matrix $\tilde{S} \in \mathbb{R}^{m\times n}$ in the subspace-restricted singular value decomposition*

$$A = \hat{U}\tilde{S}\tilde{V}^T \tag{2.12}$$

is of the form

$$\tilde{S} = \begin{bmatrix} \tilde{\Sigma}_1 \\ \tilde{B} \\ O \end{bmatrix}, \tag{2.13}$$

where $\tilde{\Sigma}_1 = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_{n-q})$ is the leading $(n-q) \times n$ submatrix of $\tilde{\Sigma}$, $\tilde{B} = \hat{W}^T A\tilde{V}$, and O denotes the $(m-\hat{p}+q) \times n$ zero matrix. Moreover,

$$\sigma_j \geq \tilde{\sigma}_j \geq \sigma_{j+\hat{p}}, \quad 1 \leq j \leq n,$$

where we let $\sigma_j = 0$ for $j > n$.

The TSRSVD method for the solution of (1.1) based on the decomposition (2.12) is analogous to the TSRSVD method based on (2.3). Specifically, the approximate solution $\tilde{\boldsymbol{x}}_k$ is given by (2.8) with $\tilde{V}$ defined by (2.10), and $\tilde{\boldsymbol{y}}_k$ is the minimal-norm least-squares solution of (2.7) with $\tilde{U}$ replaced by $\hat{U}$, defined by (2.11). The matrix $\tilde{S}^{(k)}$ in (2.7) is given by

$$\tilde{S}^{(k)} = \begin{bmatrix} \tilde{\Sigma}_1^{(k)} \\ \tilde{B} \\ O \end{bmatrix},$$

where $\tilde{\Sigma}_1^{(k)} = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_k, 0, \ldots, 0) \in \mathbb{R}^{(n-\hat{p})\times n}$ is obtained by setting the last $n - \hat{p} - k$ diagonal entries of the matrix $\tilde{\Sigma}_1$ in (2.13) to zero.

The application of the decomposition (2.12) with $\hat{W}$ chosen to be an approximation of a normalized denoised version of $\boldsymbol{b}$ may be of interest when the first few coefficients $\boldsymbol{u}_j^T\boldsymbol{b}$ in the sum (1.7) obtained from the SVD (1.4) are much smaller than $\|\boldsymbol{b}\|$.

A decomposition which combines the properties of the factorizations (2.3) and (2.12) also can be derived. Let the matrices W and $\hat{W}$, as well as the projectors $P_W^\perp$ and $P_{\hat{W}}^\perp$, be as above. Introduce the singular value decomposition

$$P_{\hat{W}}^\perp A P_W^\perp = \tilde{U}\tilde{\Sigma}\tilde{V}^T, \tag{2.14}$$

where the matrices $\tilde{U} \in \mathbb{R}^{m\times m}$ and $\tilde{V} \in \mathbb{R}^{n\times n}$ are orthogonal. The singular values are the diagonal entries of $\tilde{\Sigma} = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_n) \in \mathbb{R}^{m\times n}$ with

$$\tilde{\sigma}_1 \geq \tilde{\sigma}_2 \geq \cdots \geq \tilde{\sigma}_{\tilde{\ell}} > \tilde{\sigma}_{\tilde{\ell}+1} = \cdots = \tilde{\sigma}_n = 0.$$

The matrix $\tilde{V}$ may be assumed to be of the form (2.2). Define

$$\hat{U} = [\tilde{U}_1, \hat{W}, \tilde{U}_2], \tag{2.15}$$

where $\tilde{U}_1$ is the leading $m \times \tilde{\ell}$ submatrix of $\tilde{U}$ and $\tilde{U}_2 \in \mathbb{R}^{m\times(m-\tilde{\ell}-\hat{p})}$ has orthonormal columns that are orthogonal to the columns of the matrices $\tilde{U}_1$ and $\hat{W}$. Similarly as for the decomposition (2.12), we have $\tilde{U}_2^T A = O$.

Theorem 2.3. *Let the matrices $\tilde{U}$, $\tilde{\Sigma}$, $\tilde{V}$, and $\hat{U}$ be determined by (2.14) and (2.15), and assume that the columns of $\tilde{V}$ are ordered according to (2.2). Then the matrix $\tilde{S} \in \mathbb{R}^{m\times n}$ in the subspace-restricted singular value decomposition*

$$A = \hat{U}\tilde{S}\tilde{V}^T \tag{2.16}$$

is of the form

$$\tilde{S} = \begin{bmatrix} \tilde{\Sigma}_1 & \tilde{B}_{12} \\ \tilde{B}_{21} & \tilde{B}_{22} \\ O & O \end{bmatrix}, \tag{2.17}$$

where $\tilde{\Sigma}_1$ is the leading $\tilde{\ell} \times (n-p)$ submatrix of $\tilde{\Sigma}$, $\tilde{B}_{12} \in \mathbb{R}^{\tilde{\ell}\times p}$, $\tilde{B}_{21} \in \mathbb{R}^{\hat{p}\times(n-p)}$, and $\tilde{B}_{22} \in \mathbb{R}^{\hat{p}\times p}$. The last $m - \tilde{\ell} - \hat{p}$ rows of $\tilde{S}$ vanish. Moreover,

$$\sigma_j \geq \tilde{\sigma}_j \geq \sigma_{j+p+\hat{p}}, \quad 1 \leq j \leq n, \tag{2.18}$$

where $\sigma_j = 0$ for $j > n$.

Proof. The structure of (2.17) is a consequence of the relation $\tilde{S} = \hat{U}^T A \tilde{V}$ as well as of the structure (2.15) of $\hat{U}$ and (2.2) of $\tilde{V}$. The inequalities (2.18) follow by observing that the matrix $\tilde{S}$ is a modification of $\tilde{\Sigma}$ of at most rank $p + \hat{p}$. □

A TSRSVD method based on the decomposition (2.16) is obtained by setting the smallest diagonal entries of the matrix $\tilde{\Sigma}_1$ in (2.17) to zero.

The computational effort required to compute the subspace-restricted singular value decompositions (2.3), (2.12) and (2.16) is dominated by the computation of the singular value decompositions (2.1), (2.10) and (2.14), respectively, because typically the matrices W and $\hat{W}$ have few columns, only. Therefore, the computation of the projections of the matrix A in (2.3), (2.12) and (2.16) is inexpensive when compared to the computation of the SVD.

We conclude this section with some comments on a different technique described in [7] to enforce the solution subspace to contain a user-specified subspace $\mathcal{W} = \mathcal{R}(W)$. Let $P_{\mathcal{R}(AW)}^\perp$ be the orthogonal projector onto the complement of $\mathcal{R}(AW)$. The approach in [7] is based on first solving

$$P_{\mathcal{R}(AW)}^\perp A P_W^\perp \boldsymbol{x} = P_{\mathcal{R}(AW)}^\perp \boldsymbol{b}$$

with the aid of the SVD and then updating the computed solution to include solution components in $\mathcal{R}(W)$. We find the approach of the present paper attractive because of its versatility. We may within the same framework impose that the solution subspace and/or the range contain chosen subspaces for all truncations. The relation of SRSVD to GSVD is explored in Section 4. The numerical examples of Section 5 show the accuracy in the computed approximations of $\hat{\boldsymbol{x}}$ determined by TSRSVD to compare well with the accuracy of approximants determined by TGSVD.

3. Tikhonov regularization

One of the most popular approaches to regularization is due to Tikhonov. The simplest form of Tikhonov regularization replaces the linear discrete ill-posed problem (1.1) by the least-squares problem

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} \left\{ \|A\boldsymbol{x}-\boldsymbol{b}\|^2 + \lambda\,\|\boldsymbol{x}\|^2 \right\}, \tag{3.1}$$

where $\lambda > 0$ is a regularization parameter. The value of λ determines how sensitive the solution $\boldsymbol{x}_\lambda$ of (3.1) is to the error in $\boldsymbol{b}$ and how well $\boldsymbol{x}_\lambda$ approximates $\hat{\boldsymbol{x}}$; see, e.g., [3,14]. The normal equations associated with (3.1) are given by

$$(A^TA + \lambda I)\,\boldsymbol{x} = A^T\boldsymbol{b}. \tag{3.2}$$

The solution $\boldsymbol{x}_\lambda$ can be easily computed by substituting the SVD (1.4) into (3.1) or (3.2).

Two approaches to use the decomposition (2.3) in Tikhonov regularization suggest themselves. Substituting (2.3) into (3.2) yields

$$(\tilde{S}^T\tilde{S} + \lambda I)\,\boldsymbol{y} = \tilde{S}^T\tilde{U}^T\boldsymbol{b}, \quad \boldsymbol{y} = \tilde{V}^T\boldsymbol{x}. \tag{3.3}$$

In actual computations, the solution $\boldsymbol{y}_\lambda$ should be determined by solving a least-squares problem for which (3.3) are the normal equations.

Alternatively, since the vectors in the chosen subspace $\mathcal{W}$ are assumed to represent important features of the solution and $\mathcal{W}$ is chosen so that A is well conditioned on this subspace, we may modify (3.3) so that only the solution component in $\mathbb{R}^n \setminus \mathcal{W}$ is regularized. This yields the equation

$$\left(\tilde{S}^T\tilde{S} + \lambda \begin{bmatrix} I_{n-p} & O \\ O & O \end{bmatrix}\right)\boldsymbol{y} = \tilde{S}^T\tilde{U}^T\boldsymbol{b}, \quad \boldsymbol{y} = \tilde{V}^T\boldsymbol{x}, \tag{3.4}$$

where I_{n-p} denotes the identity matrix of order $n-p$. Note that the leading $(n-p)\times(n-p)$ principal submatrix of $\tilde{S}^T\tilde{S}$ is $\tilde{\Sigma}_1^T\tilde{\Sigma}_1$; cf. (2.4). Similarly as for (3.3), the solution $\boldsymbol{y}_\lambda$ of (3.4) should be computed by solving a least-squares problem for which (3.4) are the normal equations.

4. Relations to singular value decompositions

In the experiments of Section 5, we will observe that the computed TSRSVD solution with respect to the right space $\mathcal{R}(W)$ is often close to the computed TGSVD solution with respect to the matrix pair $\{A, I - WW^T\}$, although the former often seems to approximate $\hat{\boldsymbol{x}}$ at least slightly better. If $\mathcal{R}(W)$ is spanned by right singular vectors of A, then the computed solutions are in fact exactly the same. We will now prove some results that imply this fact. In the following, π denotes a permutation of the numbers $1, 2, \ldots, m$ or $1, 2, \ldots, n$, as appropriate.

Proposition 4.1. *Let the prescribed right vectors $W = [\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_p]$ in the SRSVD (see Theorem 2.1) be p right singular vectors of A. Then the SRSVD $A = \tilde{U}\tilde{S}\tilde{V}^T$ can be chosen to be a permuted singular value decomposition: $\tilde{S} = \mathrm{diag}(\tilde{\sigma}_1, \tilde{\sigma}_2, \ldots, \tilde{\sigma}_n)$, where there exists a permutation π such that $\sigma_{\pi(j)} = \tilde{\sigma}_j$, $\boldsymbol{v}_{\pi(j)} = \tilde{\boldsymbol{v}}_j$, $\boldsymbol{u}_{\pi(j)} = \tilde{\boldsymbol{u}}_j$, for $j = 1, 2, \ldots, n$.*

Proof. Since W is made up of singular vectors, the singular vectors of $A(I - WW^T)$ are the same as of A (or can be selected to be the same in the case of multiple singular values or in the case $p > 1$). The singular vectors that define W are (or can be chosen to be, in case A has zero singular values) the last columns of the singular vector matrix. □

The meaning of Proposition 4.1 is that, for this special choice of W, the SRSVD differs from the SVD in that for the former a user may preselect certain singular vectors that are required to be involved in the solution of (1.1).

If we prescribe left singular vectors, i.e., if we let the columns of the matrix $\hat{W}$ be left singular vectors of A, then we have the following analogous result. It can be shown similarly as Proposition 4.1.

Proposition 4.2. *Suppose that the prescribed left vectors $\hat{W} = [\hat{\boldsymbol{w}}_1, \hat{\boldsymbol{w}}_2, \ldots, \hat{\boldsymbol{w}}_p]$ in the SRSVD (see Theorem 2.2) are p left singular vectors of A. Then the SRSVD $A = \tilde{U}\tilde{S}\tilde{V}^T$ can be chosen to be a permuted singular value decomposition: $\tilde{S} = \mathrm{diag}(\tilde{\sigma}_1, \ldots, \tilde{\sigma}_n)$, where there exists a permutation π such that $\sigma_{\pi(i)} = \tilde{\sigma}_i$, $\boldsymbol{v}_{\pi(i)} = \tilde{\boldsymbol{v}}_i$, $\boldsymbol{u}_{\pi(i)} = \tilde{\boldsymbol{u}}_i$, for $i = 1, 2, \ldots, m$.*

Finally, we mention that we also can formulate an analogous result for the choice of both right singular vectors $W = [\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_p]$ and left singular vectors $\hat{W} = [\hat{\boldsymbol{w}}_1, \hat{\boldsymbol{w}}_2, \ldots, \hat{\boldsymbol{w}}_q]$. The details are notationally quite involved. We therefore prefer to omit them.

Instead, we now proceed towards the final result of this section: if the columns of W are right singular vectors of A, then the TGSVD solutions with respect to the matrix pair $\{A,\ I - WW^T\}$ of (1.1) coincides with the TSRSVD solution with respect to the subspace $\mathcal{R}(W)$. The following lemma can be shown by direct verification.

Table 5.1
Overview of the relative errors in approximate solutions determined by TSVD, TGSVD, and TSRSVD for Experiments 5.1 and 5.2.

W	deriv2		baart		
	1	t^2	$[1,\ t]$	t^2	$\sin(t)$
TSVD	$2.94 \cdot 10^{-1}$		$1.66 \cdot 10^{-1}$		
TGSVD	$1.38 \cdot 10^{-1}$	$1.64 \cdot 10^{-1}$	$1.74 \cdot 10^{-1}$	$2.73 \cdot 10^{-1}$	$3.43 \cdot 10^{-3}$
TSRSVD	$1.36 \cdot 10^{-1}$	$1.34 \cdot 10^{-1}$	$1.41 \cdot 10^{-1}$	$6.56 \cdot 10^{-2}$	$3.43 \cdot 10^{-3}$

Lemma 4.3. *Let A have the SVD $A = U\Sigma V^T$, and let $W = [\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_p]$ be a subset of right singular vectors of A with corresponding partial SVD $AW = U_2\Sigma_2$. Let Σ_1 be the diagonal matrix containing the remaining singular values in decreasing order with corresponding left and right singular vectors contained in U_1 and V_1 (cf. Section 2). Then the pair $\{A,\ I - WW^T\}$ has the GSVD,*

$$A = [U_1,\ U_2]\ \mathrm{diag}([\Sigma_1(\Sigma_1^2 + I)^{-1/2},\ I_p])\ [V_1(\Sigma_1^2 + I)^{1/2},\ W\Sigma_2]^T,$$
$$I - WW^T = [V_1,\ W]\ \mathrm{diag}([(\Sigma_1^2 + I)^{-1/2},\ 0_p])\ [V_1(\Sigma_1^2 + I)^{1/2},\ W\Sigma_2]^T.$$

Theorem 4.4. *Let $W = [\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_p]$ be p right singular vectors of A corresponding to nonzero singular values. Then the TGSVD solutions of (1.1) corresponding to the matrix pair $\{A,\ I - WW^T\}$ are equal to the TSRSVD solutions of (1.1) with respect to W.*

Proof. We use the notations and result of the preceding lemma. Since

$$[V_1(\Sigma_1^2 + I)^{1/2},\ W\Sigma_2]^{-T} = [V_1(\Sigma_1^2 + I)^{-1/2},\ W\Sigma_2^{-1}],$$

we have that the TGSVD solution of the pair $\{A,\ I - WW^T\}$ to the problem (1.1) is given by (see, e.g., [14])

$$V_k(\Sigma_k^2 + I)^{-1/2}(\Sigma_k(\Sigma_k^2 + I)^{-1/2})^{-1}U_k^T\boldsymbol{b} + W\Sigma_2^{-1}U_2^T\boldsymbol{b}, \tag{4.1}$$

where Σ_k, U_k, and V_k are truncated versions of Σ_1, U_1, and V_1, respectively.

On the other hand, the TSRSVD solution with special vectors W satisfies

$$\boldsymbol{x}_k = [V_k\ W]\,\boldsymbol{y}_k, \quad \boldsymbol{y}_k = \underset{\boldsymbol{y}}{\mathrm{argmin}}\ \|A\,[V_k\ W]\boldsymbol{y} - \boldsymbol{b}\|.$$

Since $\mathcal{R}(A[V_k,\ W]) = \mathcal{R}([U_k,\ U_2])$ and

$$[U_k,\ U_2]^T A[V_k,\ W] = \begin{bmatrix} \Sigma_k & 0 \\ 0 & \Sigma_2 \end{bmatrix},$$

we have

$$V_k\Sigma_k^{-1}U_k^T\boldsymbol{b} + W\Sigma_2^{-1}U_2^T\boldsymbol{b},$$

from which, in view of (4.1), the result now follows. □

5. Numerical experiments

This section presents a few computed examples which illustrate the performance of the SRSVD. The right-hand sides in the examples below are contaminated by an error $\boldsymbol{e}$ of relative norm ε, i.e.,

$$\|\boldsymbol{e}\|/\|\hat{\boldsymbol{b}}\| = \varepsilon. \tag{5.1}$$

We take $\varepsilon = 0.01$, which means 1% noise. The entries of $\boldsymbol{e}$ are normally distributed pseudorandom numbers with zero mean, generated by the MATLAB function `randn`. They are scaled so that (5.1) holds. The constant γ in the discrepancy principle (1.9) is set to 1.1 and we let $\delta = \varepsilon\,\|\hat{\boldsymbol{b}}\|$ in (1.9).

Let W contain user-selected orthonormal columns. The columns of W play a special role both in the SRSVD as presented in this paper, and in the GSVD of the matrix pair $\{A,\ I - WW^T\}$. Therefore, in Experiments 5.1 and 5.2 we compare the TSVD of A, the TGSVD of the pair $\{A,\ I - WW^T\}$, and the SRSVD of A with specified right vectors given by the columns of W. We present results for two well-known examples, deriv2 and baart from [15], of size $m = n = 500$. Table 5.1 provides an overview of the relative errors $\|\boldsymbol{x}_{\mathrm{discr}} - \hat{\boldsymbol{x}}\|/\|\hat{\boldsymbol{x}}\|$ for some choices of W. We discuss the details below.

Experiment 5.1. Consider the Fredholm integral equation of the first kind,

$$\int_0^1 k(s,t)x(t)\mathrm{d}t = e^s + (1-e)s - 1, \quad 0 \le s \le 1, \tag{5.2}$$

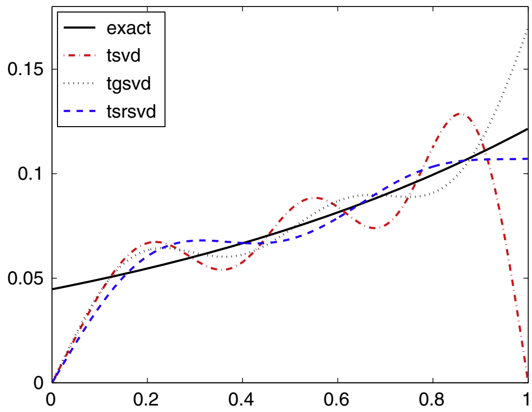

Fig. 5.1. Experiment 5.1: Exact solution $\hat{\boldsymbol{x}}$ (continuous curve), approximate solutions using the discrepancy principle as determined by TSVD (dash-dotted curve), TGSVD (dotted curve), and TSRSVD (dashed curve).

where

$$k(s,t) = \begin{cases} s\,(t-1), & s < t, \\ t\,(s-1), & s \geq t. \end{cases}$$

We discretize the integral equation by a Galerkin method with orthonormal box functions as test and trial functions using the MATLAB program deriv2 from [15]. This program yields a symmetric indefinite matrix $A \in \mathbb{R}^{500\times500}$ and a scaled discrete approximation $\hat{\boldsymbol{x}} \in \mathbb{R}^{500}$ of the solution $x(t) = \exp(t)$ of (5.2). The condition number $\kappa(A)$, defined analogously to (2.6), is $3.0 \cdot 10^5$. Fig. 5.1 shows $\hat{\boldsymbol{x}}$ (continuous curve). The error-free right-hand side vector is given by $\hat{\boldsymbol{b}} = A\hat{\boldsymbol{x}}$, and the right-hand side vector $\boldsymbol{b}$ in (1.1) is determined by (1.2) with $\varepsilon = 1 \cdot 10^{-2}$ in (5.1).

We first consider approximants $\boldsymbol{x}_k$ of $\hat{\boldsymbol{x}}$ computed by TSVD. The discrepancy principle (1.9) yields $k_{\text{discr}} = 6$. The dash-dotted curve of Fig. 5.1 displays $\boldsymbol{x}_6$. The relative error in $\boldsymbol{x}_6$ is seen to be quite large; we have $\|\boldsymbol{x}_6 - \hat{\boldsymbol{x}}\|/\|\hat{\boldsymbol{x}}\| = 2.94 \cdot 10^{-1}$. For comparison, we determine $k_* = 13$ from Eq. (1.8) and obtain $\|\boldsymbol{x}_{13} - \hat{\boldsymbol{x}}\|/\|\hat{\boldsymbol{x}}\| = 2.22 \cdot 10^{-1}$. Thus, the error in $\boldsymbol{x}_{13}$ is not much smaller than the error in $\boldsymbol{x}_9$. The low accuracy obtained by TSVD combined with the discrepancy principle therefore does not depend on a failure of the latter, but instead depends on that linear combinations of the first few columns of the matrix V in (1.4) are not well suited to approximate $\hat{\boldsymbol{x}}$.

We turn to the TGSVD and TSRSVD methods. If we let W be the "constant unit vector" $\frac{1}{10\sqrt{5}}[1, 1, \ldots, 1]^T \in \mathbb{R}^{500}$, then both the TGSVD applied to the pair $\{A, I - WW^T\}$, and the TSRSVD give more accurate approximate solutions with relative errors of $1.38 \cdot 10^{-1}$ and $1.36 \cdot 10^{-1}$, respectively; both using 5 vectors.

These approximate solutions are quite similar, and suggest that the desired solution $\hat{\boldsymbol{x}}$ might be fairly well approximated by a parabola. If we apply TGSVD and TSRSVD with

$$\mathcal{W} = \text{range}\begin{bmatrix} 1 \\ 4 \\ \vdots \\ n^2 \end{bmatrix}, \tag{5.3}$$

we obtain the dotted (TGSVD) and dashed (TSRSVD) curves of Fig. 5.1. The relative errors are $1.64 \cdot 10^{-1}$ and $1.34 \cdot 10^{-1}$, respectively; see also Table 5.1. It is clear that the TSRSVD solution is of better quality. To satisfy the discrepancy principle, the TGSVD solution uses one extra vector compared to the TSRSVD: 5 (TGSVD) versus 4 (TSRSVD).

If we increase the number of columns of W, e.g., if we let the columns of W be an orthonormal basis for the subspace

$$\mathcal{W} = \text{range}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ \vdots & \vdots & \vdots \\ 1 & n & n^2 \end{bmatrix}, \tag{5.4}$$

then both the TGSVD and TSRSVD methods give very similar excellent approximate solutions with relative error $5.06 \cdot 10^{-3}$. We remark that TGSVD applied to the pair $\{A, L\}$, with L the 4-diagonal regularization operator

$$L = \begin{bmatrix} -1 & 3 & -3 & 1 & & & \\ & -1 & 3 & -3 & 1 & & \\ & & \ddots & \ddots & \ddots & \ddots & \\ & & & -1 & 3 & -3 & 1 \end{bmatrix} \in \mathbb{R}^{(n-3)\times n},$$

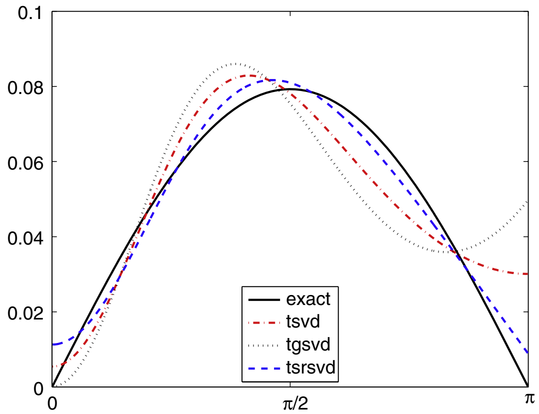

Fig. 5.2. Experiment 5.2: Exact solution $\hat{\boldsymbol{x}}$ (continuous curve), approximate solutions using the discrepancy principle as determined by TSVD (dash-dotted curve), TGSVD (dotted curve), and TSRSVD (dashed curve).

which is a scaled approximation of a third derivative operator with null space (5.4), yields an approximate solution of similar quality. □

Experiment 5.2. We discretize the integral equation

$$\int_0^{\pi} \exp(s\cos(t))x(t)\,\mathrm{d}t = 2\,\frac{\sinh(s)}{s}, \quad 0 \le s \le \frac{\pi}{2},$$

discussed by Baart [16] by a Galerkin method with piecewise constant test and trial functions using the MATLAB code baart from [15]. This yields the nonsymmetric matrix $A \in \mathbb{R}^{500\times 500}$ of ill-determined rank. The code also furnishes the "exact" solution $\hat{\boldsymbol{x}}$, which represents a scaled sine function. We determine the error-free right-hand side $\hat{\boldsymbol{b}}$ of (1.3) and the contaminated right-hand side $\boldsymbol{b}$ of (1.1) similarly as in Experiment 5.1.

We first consider approximants $\boldsymbol{x}_k$ of $\hat{\boldsymbol{x}}$ computed by TSVD. The discrepancy principle (1.9) now yields $k_{\mathrm{discr}} = 3$. The dash-dotted curve of Fig. 5.2 displays $\boldsymbol{x}_3$. The relative error in $\boldsymbol{x}_3$ is $\|\boldsymbol{x}_3 - \hat{\boldsymbol{x}}\|/\|\hat{\boldsymbol{x}}\| = 1.66 \cdot 10^{-1}$. The value of k_{discr} is optimal, i.e., (1.8) yields $k_* = 3$.

Now we turn to the TGSVD and TSRSVD methods for various W. Let W be the "constant unit vector" $\frac{1}{10\sqrt{5}}[1, 1, \ldots, 1]^T \in \mathbb{R}^{500}$. Then TGSVD yields the approximate solution $\boldsymbol{x}_4$ with relative error $1.70 \cdot 10^{-1}$ and TSRSVD gives the approximate solution $\boldsymbol{x}_3$ with relative error $1.58 \cdot 10^{-1}$. Thus, TSRSVD determines the best approximation of $\hat{\boldsymbol{x}}$ and TGSVD the worst.

For W the "constant and linear vectors", i.e., the first two columns of (5.4), orthonormalized, TGSVD determines an approximate solution with relative error $1.74 \cdot 10^{-1}$ and TSRSVD gives an approximate solution with relative error $1.41 \cdot 10^{-1}$; see Table 5.1. When we take W as in (5.3), the difference is even larger: the relative errors of TGSVD and TSRSVD are $2.73 \cdot 10^{-1}$ and $6.56 \cdot 10^{-2}$, respectively. These computed solutions are displayed in Fig. 5.2.

Finally, when we supply a normalization of $\hat{\boldsymbol{x}}$ as W, i.e., a normalization of the vector generated by the MATLAB command `sin((0:n-1)*pi/n)'`, both TGSVD and TSRSVD determine excellent approximations of $\hat{\boldsymbol{x}}$ with relative errors $3.43 \cdot 10^{-3}$. □

Experiment 5.3. Our last experiment is concerned with the restoration of an image, which has been contaminated by Gaussian blur and noise. Fig. 5.3 shows the original image bultheel represented by an array of 512×448 pixels. This image is too large for direct solution methods. We therefore consider the subpicture eye of 71×71 pixels displayed in Fig. 5.4(a). This image is assumed not to be available. The available blur- and noise-contaminated image is shown in Fig. 5.4(b). The pixel values, ordered column-wise, determine the right-hand side $\boldsymbol{b} \in \mathbb{R}^{5041}$. The blurring operator is represented by the symmetric block Toeplitz matrix with Toeplitz blocks,

$$A = (2\pi\sigma^2)^{-1} T \otimes T,$$

where T is a 71×71 symmetric banded Toeplitz matrix, whose first row is given by `[exp(-((0:band-1).^2)/(2*sigma^2)); zeros(1,n-band)]`, and $\otimes$ denotes the Kronecker product. The parameter band is the half-bandwidth of the matrix T and the parameter σ controls the effective width of the underlying Gaussian point spread function

$$h(x, y) = \frac{1}{2\pi\sigma^2}\,\exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right),$$

which models blurring. We chose band $= 16$ and $\sigma = 1.5$. The matrix A so obtained is numerically singular. For further details on image restoration, see, e.g., [17].

Fig. 5.3. The 512×448 pixel picture bultheel.

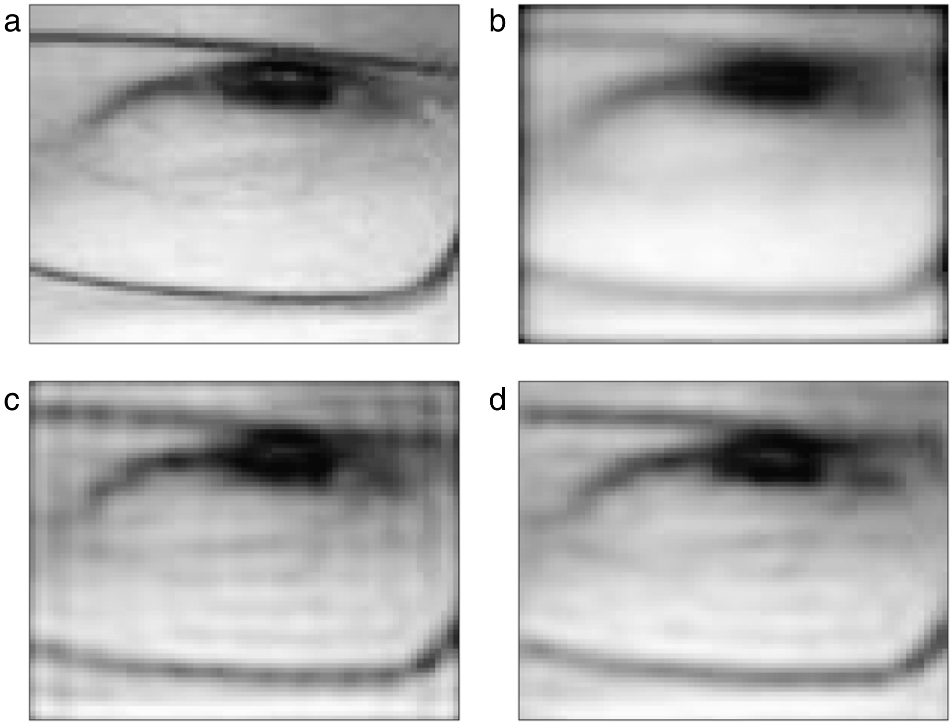

Fig. 5.4. Experiment 5.3: Subpicture eye (71×71) of bultheel: (a) original; (b) blurred and noised; (c) TSVD restoration; (d) TSRSVD restoration with "$W = [1,\ t,\ t^2]$" (see (5.4)).

Fig. 5.4(c) shows the restoration $\boldsymbol{x}_{392}$ obtained with $k = 392$ right singular vectors using the TSVD method. It has relative error $\|\boldsymbol{x}_{392} - \hat{\boldsymbol{x}}\|/\|\hat{\boldsymbol{x}}\| = 6.81 \cdot 10^{-2}$. The restoration $\boldsymbol{x}_{358}$ obtained with $k = 358$ right singular vectors using the TSRSVD method with $\mathcal{W}$ given by (5.4) is displayed in Fig. 5.4(d). The relative error $5.23 \cdot 10^{-2}$ of $\boldsymbol{x}_{358}$ is clearly smaller than that of the TSVD approximation and the restoration looks superior. In particular, Fig. 5.4(d) displays less "ringing" than Fig. 5.4(c). This example illustrates that it is possible to achieve an improved restoration by including vectors that model "polynomial behavior" and are not taylored to the problem at hand. $\square$

6. Conclusions

This paper describes a new SVD-type decomposition, the subspace-restricted SVD (SRSVD), which allows a user to prescribe some of the columns of the U and V matrices of this decomposition. Computed examples illustrate the truncated version of the SRSVD to determine more accurate approximate solutions of linear discrete ill-posed problems than the TSVD.

When W contains an orthonormal basis of user-selected vectors, there are some similarities between the TGSVD of $\{A,\ I - WW^T\}$ and the TSRSVD of A with respect to W. In some cases the quality of the computed approximate solutions of (1.1) is about the same, and in certain special cases the computed solutions are (mathematically) identical.

In most examples we carried out, TSRSVD gave at least as accurate approximate solutions as TGSVD, and sometimes approximate solutions of clearly higher accuracy. In particular, TSRSVD seems to perform better when W has few columns. Moreover, the fact that TSRSVD only requires the standard SVD and not the GSVD may be seen as an advantage; cf. also the remarks in [14, p. 51].

Acknowledgements

LR would like to thank MH for an enjoyable visit to TU/e, during which work on this paper was begun. The authors thank the referee for very useful comments.

References

[1] G.H. Golub, C.F. Van Loan, Matrix Computations, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.
[2] P.C. Hansen, Truncated SVD solutions to discrete ill-posed problems with ill-determined numerical rank, SIAM J. Sci. Stat. Comput. 11 (1990) 503–518.
[3] H.W. Engl, M. Hanke, A. Neubauer, Regularization of Inverse Problems, Kluwer, Dordrecht, 1996.
[4] D. Calvetti, B. Lewis, L. Reichel, GMRES, *L*-curves and discrete ill-posed problems, BIT 42 (2002) 44–65.
[5] L. Eldén, Partial least squares vs. Lanczos bidiagonalization I: analysis of a projection method for multiple regression, Comput. Statist. Data Anal. 46 (2004) 11–31.
[6] P.C. Hansen, T. Sekii, H. Shibahashi, The modified truncated SVD method for regularization in general form, SIAM J. Sci. Stat. Comput. 13 (1992) 1142–1150.
[7] S. Morigi, L. Reichel, F. Sgallari, A truncated projected SVD method for linear discrete ill-posed problems, Numer. Algorithms 43 (2006) 197–213.
[8] P.C. Hansen, Regularization, GSVD and truncated GSVD, BIT 29 (1989) 491–504.
[9] A. Bultheel, Laurent Series and their Padé Approximations, Birkhäuser, Basel, 1987.
[10] A. Bultheel, M. Van Barel, Linear Algebra, in: Rational Approximation and Orthogonal Polynomials, Elsevier, Amsterdam, 1997.
[11] A. Bultheel, M. Van Barel, P. Van gucht, Orthogonal bases in discrete least squares rational approximation, J. Comput. Appl. Math. 164–165 (2004) 175–194.
[12] A. Bultheel, M. Van Barel, Y. Rolain, R. Pintelon, Numerically robust transfer function modelling from noisy frequency domain data, IEEE Trans. Automat. Control 50 (2005) 1835–1839.
[13] R.A. Horn, C.R. Johnson, Topics in Matrix Analysis, Cambridge University Press, Cambridge, England, 1991.
[14] P.C. Hansen, Rank Deficient and Discrete Ill-Posed Problems, SIAM, Philadelphia, 1998.
[15] P.C. Hansen, Regularization tools version 4.0 for MATLAB 7.3, Numer. Algorithms 46 (2007) 189–194.
[16] M.L. Baart, The use of auto-correlation for pseudo-rank determination in noisy ill-conditioned least-squares problems, IMA J. Numer. Anal. 2 (1982) 241–247.
[17] P.C. Hansen, J.G. Nagy, D.P. O'Leary, Deblurring Images: Matrices, Spectra, and Filtering, SIAM, Philadelphia, 2006.