BIT
2001, Vol. xx, No. x, pp. 1–10
0006-3835/98/3804-0101 $12.00

TRUNCATED VSV SOLUTIONS TO SYMMETRIC RANK-DEFICIENT PROBLEMS*

Ricardo D. FIERRO[1] and Per Christian HANSEN[2]

[1] *Department of Mathematics, California State University San Marcos CA 92096, USA. email: fierro@csusm.edu*

[2] *Informatics and Mathematical Modelling, Technical University of Denmark Building 321, DK-2800 Lyngby, Denmark. email: pch@imm.dtu.dk*

Abstract.

Symmetric VSV decompositions are new rank-revealing decompositions that exploit and preserve symmetry. Truncated VSV solutions are stabilized solutions computed by neglecting blocks in the VSV decomposition with small norm. We compare the truncated VSV solutions with truncated SVD solutions and give perturbation bounds for the VSV solutions. Numerical examples illustrate our results.

AMS subject classification: 65F25.

Key words: Rank-revealing decompositions, accuracy, perturbation bounds.

1 Introduction

Symmetric rank-revealing decompositions are among the most recent members of the growing family of numerical tools for rank reduction in signal processing and other applications. They were first considered by Luk and Qiao [5] and, independently, by Baker and DeGroat [1]. A recent survey of algorithms based on an initial triangular factorization is given in [4].

One motivation for developing specialized algorithms for symmetric rank-deficient matrices is the reduction in computational effort, compared to general rank-revealing algorithms. The second motivation is that a symmetric decomposition allows one to compute symmetric matrix approximations with specified rank, which may be important in connection with, e.g., correlation matrices in signal processing and Hessian matrices in optimization.

Along with the definition of any new decomposition and algorithm goes a careful study of its accuracy, numerical stability, efficient implementation, and related issues. In this paper we focus on the accuracy of rank-revealing VSV decompositions: we compare truncated VSV solutions with solutions based on the SVD and we present perturbation bounds.

*Received ???. Revised ???. Communicated by ???.

Let A be an $n \times n$ symmetric and possibly indefinite matrix whose eigenvalue decomposition (EVD) is given by

$$A = W\,\Lambda\,W^T = (\,W_k\,,\,W_0\,)\begin{pmatrix}\Lambda_k & 0\\ 0 & \Lambda_0\end{pmatrix}(\,W_k\,,\,W_0\,)^T \tag{1.1}$$

with $\Lambda = \mathrm{diag}(\lambda_1,\ldots,\lambda_n)$. The singular values of A are $\sigma_i = |\lambda_i|$, $i = 1,\ldots,n$, and we can always choose the right singular vectors to be equal to the columns of W. The VSV decomposition of A is a symmetric decomposition of the form

$$A = V\,S\,V^T = (\,V_k\,,\,V_0\,)\begin{pmatrix}S_k & J^T\\ J & K\end{pmatrix}(\,V_k\,,\,V_0\,)^T \tag{1.2}$$

where V is orthogonal, V_k has k columns, S_k is a $k \times k$ nonsingular matrix, and k is the numerical rank of A – see [6] for definitions of "numerical rank" as well as practical issues. The VSV decomposition is rank-revealing if $\sigma_k(S_k)$, the smallest singular value of S_k, is of the order σ_k and if the 2-norm of the block $(\,J\,,\,K\,)$ is of the order σ_{k+1}. Then the columns of V_k and V_0 span approximations to the signal and noise subspaces, defined in terms of the SVD as the ranges of W_k and W_0, respectively.

When using VSV decompositions in practice, it is inconvenient to work directly with the matrix S_k, and instead we assume that we compute a factorization of S_k of the form

$$S_k = T_k\,\Omega_k\,T_k^T, \tag{1.3}$$

where T_k is a triangular matrix, and Ω_k is a signature matrix with elements ± 1 on the diagonal and zeroes elsewhere ($\Omega_k = I_k$ if A is semidefinite); see [4].

Given the EVD and the VSV decomposition, we can compute stabilized solutions to the system $A\,x = b$ by truncating the decompositions at k. Hence, the truncated EVD solution – which for symmetric matrices is identical to the truncated SVD solution – is given by

$$x_k = W_k\,\Lambda_k^{-1}\,W_k^T b \tag{1.4}$$

where W_k and Λ_k are from (1.1). Similarly, the truncated VSV solution is given by

$$x_{Sk} = V_k\,S_k^{-1} V_k^T b \tag{1.5}$$

where V_k and S_k are from (1.2). The latter solution is easily computed via the factorization (1.3) of S_k. When A is ill conditioned, then both solutions are guaranteed to have smaller norm and variance than $A^{-1}b$, at the cost of introducing some bias into the solution – these are classical issues in the regularization of ill-conditioned problems, and we refer to [3] for a detailed discussion.

The purpose of this paper is to compare the truncated SVD and VSV solutions, and to derive perturbation bounds for the truncated VSV solutions, as defined above. Our bounds give insight into the factors that influence the accuracy of the VSV solutions. The comparison of the truncated solutions is presented in §2, while the perturbation bounds are given in §3. We give some illustrative numerical examples in §4, and then summarize our contributions in §5.

2 Comparison of Truncated SVD and VSV Solutions

Before presenting our main result about $\|x_k - x_{Sk}\|_2$ we need to bound the distance between the SVD-based signal subspace

$$\mathcal{W}_k = \mathrm{range}(W_k) \tag{2.1}$$

and the approximation

$$\mathcal{V}_k = \mathrm{range}(V_k) \tag{2.2}$$

computed by the VSV decomposition. The distance between these two subspaces is measured by the sine to the largest principal angle θ_k associated with the subspace pair, and it is given by

$$\sin\theta_k = \|W_k\, W_k^T - V_k\, V_k^T\|_2 = \|V_k^T W_0\|_2 = \|V_0^T W_k\|_2. \tag{2.3}$$

See, e.g., §2.6.3 in [2] for more details about distances between subspaces.

THEOREM 2.1. *The distance $\sin\theta_k$ between the two subspaces $\mathcal{W}_k$ and $\mathcal{V}_k$ in (2.1) and (2.2) is bounded by*

$$\sin\theta_k \le \frac{\|J\|_2}{\sigma_k - \|K\|_2}. \tag{2.4}$$

Proof. For a proof, see Thm. 3.1 in [4].

This theorem says that the smaller the norm of the off-diagonal block J in S the closer the two subspaces, provided there is a gap between σ_k and $\|K\|_2$. This is an expected result for rank-revealing decompositions.

THEOREM 2.2. *If x_k and x_{Sk} denote the truncated SVD and VSV solutions defined in (1.4) and (1.5), then*

$$\frac{\|x_k - x_{Sk}\|_2}{\|x_k\|_2} \le \sin\theta_k \left(1 + \Psi_k\, \frac{\|J\|_2}{\|A\|_2} + \Psi_k\, \frac{\|r_k\|_2}{\|A\, x_k\|_2}\right) \tag{2.5}$$

where $\Psi_k = \|A\|_2/\sigma_k(S_k)$, $r_k = b - A\, x_k$ and $\sigma_k(S_k)$ is the smallest singular value of S_k.

Proof. To simplify notation we introduce the matrices $A_k = W_k\, \Lambda_k\, W_k^T$ and $A_{Sk} = V_k\, S_k\, V_k^T$, with pseudoinverses $A_k^\dagger = W_k\, \Lambda_k^{-1} W_k^T$ and $A_{Sk}^\dagger = V_k\, S_k^{-1} V_k^T$; and then (since $A_k^\dagger r_k = 0$):

$$\begin{aligned} x_k - x_{Sk} &= \big(A_k^\dagger - A_{Sk}^\dagger\big)\, b \\ &= \big(A_k^\dagger - A_{Sk}^\dagger\big)(r_k + A\, x_k) \\ &= -A_{Sk}^\dagger r_k + A_k^\dagger A\, x_k - A_{Sk}^\dagger A\, x_k. \end{aligned}$$

Clearly $A_k^\dagger A\, x_k = A_k^\dagger A_k\, x_k$, and using (1.2) the last term takes the form

$$A_{Sk}^\dagger A\, x_k = A_{Sk}^\dagger (A_{Sk} + V_k\, J^T V_0^T)\, x_k.$$

Thus

$$x_k - x_{Sk} = -A_{Sk}^{\dagger} r_{Sk} + \big(A_k^{\dagger} A_k - A_{Sk}^{\dagger} A_{Sk}\big) x_k - A_{Sk}^{\dagger} V_k \, J^T V_0^T x_k.$$

Using $r_k = W_0 \, W_0^T r_k$ and (2.3), the norm of the first term in this expression can be bounded as

$$\begin{aligned} \|A_{Sk}^{\dagger} r_k\|_2 &= \|V_k \, S_k^{-1} V_k^T W_0 \, W_0^T r_k\|_2 \\ &\leq \|S_k^{-1}\|_2 \, \|V_k^T W_k\|_2 \, \|r_k\|_2 \\ &= \frac{\sin\theta_k}{\sigma_k(S_k)} \, \|r_k\|_2. \end{aligned}$$

The norm of the second term is bounded by $\sin\theta_k \, \|x_k\|_2$ because

$$A_k^{\dagger} A_k - A_{Sk}^{\dagger} A_{Sk} = W_k \, W_k^T - V_k \, V_k^T.$$

Finally, using $x_k = W_k \, W_k^T x_k$ and (2.3), the norm of the third term is bounded as

$$\begin{aligned} \|A_{Sk}^{\dagger} V_k \, J^T V_0^T W_k \, W_k^T x_k\|_2 &\leq \|S_k^{-1}\|_2 \, \|J\|_2 \, \|V_0^T W_k\|_2 \, \|x_k\|_2 \\ &= \frac{\sin\theta_k \, \|J\|_2 \, \|x_k\|_2}{\sigma_k(S_k)}. \end{aligned}$$

Eq. (2.5) follows from division by $\|x_k\|_2$ and using $\|A \, x_k\|_2 \leq \|A\|_2 \, \|x_k\|_2$. □

The above theorem shows that the difference between the two solutions depends on several factors: $\|J\|_2$, the bound for $\sin\theta_k$, as well as on the TSVD residual norm $\|r_k\|_2$. We conclude that $\|J\|_2$ plays an important role in determining the accuracy of the truncated VSV solution, especially when the norm of the residual (i.e., $\|r_k\|_2$) is small. Therefore, a VSV algorithm is most useful when it takes into account the user's upper bound on $\|J\|_2 / \|A\|_2$.

3 The Perturbation Bound

We now turn to the perturbation bound associated with the truncated VSV solution. Throughout this section, we use the notation

$$\widetilde{A} = A + E, \qquad \widetilde{b} = b + e \tag{3.1}$$

for the perturbed matrix and right-hand side, and all quantities associated with the VSV decomposition of $\widetilde{A}$ are equipped with a tilde in an obvious manner. We first bound the perturbation of the approximate signal subspace.

THEOREM 3.1. *Let $\widetilde{\mathcal{V}}_k = \mathrm{range}(\widetilde{V}_k)$ denote the perturbed signal subspace, and let $\sin\phi_k$ denote the distance between the subspaces $\mathcal{V}_k$ and $\widetilde{\mathcal{V}}_k$. If $\|E\|_2 < \sigma_k(S_k) - \|K\|_2$ then*

$$\sin\phi_k \leq \frac{\|J\|_2 + \|\widetilde{J}\|_2 + \|E\|_2}{\sigma_k(\widetilde{S}_k) - \|\widetilde{K}\|_2}. \tag{3.2}$$

Proof. We use that $\sin\phi_k = \|\widetilde{V}_0^T V_k\|_2$ and therefore we consider the matrix $\widetilde{V}_0^T V_k$. It follows from (1.2) that

$$A\,V_k = V_k\,S_k + V_0\,J \quad\Longleftrightarrow\quad A\,V_k\,S_k^{-1} = V_k + V_0\,J\,S_k^{-1}$$

and hence $\widetilde{V}_0^T V_k = \widetilde{V}_0^T\,(A\,V_k - V_0\,J)\,S_k^{-1}$. Inserting (3.1) we obtain

$$\begin{aligned}\widetilde{V}_0^T V_k &= \widetilde{V}_0^T\left(\widetilde{A}\,V_k - E\,V_k - V_0\,J\right)S_k^{-1}\\ &= \left(\widetilde{J}\,\widetilde{V}_k^T V_k + \widetilde{K}\,\widetilde{V}_0^T V_k - \widetilde{V}_0^T E\,V_k - \widetilde{V}_0^T V_0\,J\right)S_k^{-1}\end{aligned}$$

and thus

$$\sin\phi_k = \|\widetilde{V}_0^T V_k\|_2 \le \left(\|\widetilde{J}\|_2 + \|\widetilde{K}\|_2\,\sin\phi_k + \|E\|_2 + \|J\|_2\right)\|S_k^{-1}\|_2.$$

Collecting terms we then arrive at the bound in (3.2). ☐

Note that the bound on $\sin\phi_k$ depends on both the norm of the perturbation E and the norm of the off-diagonal block J. The condition $\|E\|_2 < \sigma_k(S_k) - \|K\|_2$ guarantees that $\widetilde{A}$ is an acute matrix perturbation of A in the sense the principal angles between $\mathcal{V}_k$ and $\widetilde{\mathcal{V}}_k$ are less than $\pi/2$; see [7, p. 139]. This mild restriction on $\|E\|_2$ is necessary to make meaningful conclusions about the truncated VSV solutions to $Ax = b$ and $\widetilde{A}\,x = \widetilde{b}$. The next theorem bounds the perturbation of the truncated VSV solution.

THEOREM 3.2. *Let $\widetilde{x}_{S,k} = \widetilde{V}_k\,\widetilde{S}_k^{-1}\widetilde{V}_k^T\widetilde{b}$ denote the perturbed truncated VSV solution. If $\|E\|_2 < \sigma_k(S_k) - \|K\|_2$ then*

$$\frac{\|x_{Sk} - \widetilde{x}_{Sk}\|_2}{\|x_{Sk}\|_2} \le \Psi_k\left(\frac{\|e\|_2}{\|A\,x_{Sk}\|_2} + \frac{\|E\|_2}{\|A\|_2} + \frac{\|\widetilde{J}\|_2}{\|A\|_2} + \sin\phi_k\,\frac{\|r_{Sk}\|_2}{\|A\,x_{Sk}\|_2}\right) + \sin\phi_k \tag{3.3}$$

where $r_{Sk} = b - A\,x_{Sk}$ and $\Psi_k = \sigma_1/\sigma_k(S_k)$.

Proof. Using (3.1) and the definition (1.5) of x_{Sk} we get

$$\begin{aligned} x_{,k} - \widetilde{x}_{Sk} &= A_{Sk}^\dagger b - \widetilde{A}_{Sk}^\dagger\widetilde{b}\\ &= \left(A_{Sk}^\dagger - \widetilde{A}_{Sk}^\dagger\right)b - \widetilde{A}_{Sk}^\dagger e\\ &= \left(A_{Sk}^\dagger - \widetilde{A}_{Sk}^\dagger\right)r_k + \left(A_{Sk}^\dagger - \widetilde{A}_{Sk}^\dagger\right)A\,x_{Sk} - \widetilde{A}_{Sk}^\dagger e\\ &= -\widetilde{A}_{Sk}^\dagger r_{Sk} + \left(A_{Sk}^\dagger - \widetilde{A}_{Sk}^\dagger\right)A\,x_{Sk} - \widetilde{A}_{S,k}^\dagger e.\end{aligned}$$

Since $\widetilde{A}_{Sk}^\dagger = \widetilde{V}_k\,\widetilde{S}_k^{-1}\widetilde{V}_k^T$ and $r_{Sk} = V_0\,V_0^T r_{Sk}$ we can bound the norm of the first term as follows (using that $A\,x_{Sk} = A_{Sk}\,x_{Sk}$):

$$\|\widetilde{A}_{Sk}^\dagger r_{Sk}\|_2 = \|\widetilde{S}_k^{-1}\widetilde{V}_k^T V_0\,V_0^T r_{Sk}\|_2 \le \frac{\sin\phi_k\,\|r_{Sk}\|_2}{\sigma_k(\widetilde{S}_k)}.$$

The norm of the third term is simply bounded by $\|e\|_2/\sigma_k(\widetilde{S}_k)$. The second term can be rewritten as

$$\begin{aligned}
&(A_{Sk}^\dagger - \widetilde{A}_{Sk}^\dagger)A\,x_{Sk} \\
&\quad= A_{Sk}^\dagger A_{Sk}\,x_{Sk} - \widetilde{A}_{Sk}^\dagger A\,x_{Sk} \\
&\quad= A_{Sk}^\dagger A_{Sk}\,x_{Sk} - \widetilde{A}_{Sk}^\dagger \widetilde{A}_{Sk}\,x_{Sk} + \widetilde{A}_{Sk}^\dagger \widetilde{A}_{Sk}\,x_{Sk} - \widetilde{A}_{Sk}^\dagger A\,x_{Sk} \\
&\quad= (A_{Sk}^\dagger A_{Sk} - \widetilde{A}_{Sk}^\dagger \widetilde{A}_{Sk})x_{Sk} + \widetilde{A}_{Sk}^\dagger \widetilde{A}_{Sk}\,x_{Sk} - \widetilde{A}_{Sk}^\dagger \widetilde{A}\,x_{Sk} + \widetilde{A}_{Sk}^\dagger E\,x_{Sk} \\
&\quad= (V_k\,V_k^T - \widetilde{V}_k\,\widetilde{V}_k^T)x_{Sk} + \widetilde{A}_{Sk}^\dagger(\widetilde{A}_{Sk} - \widetilde{A})x_{Sk} + \widetilde{A}_{Sk}^\dagger E\,x_{Sk} \\
&\quad= (V_k\,V_k^T - \widetilde{V}_k\,\widetilde{V}_k^T)x_{Sk} - \widetilde{V}_k\,\widetilde{S}_k^{-1}\widetilde{J}\,\widetilde{V}_0^T x_{Sk} + \widetilde{A}_{Sk}^\dagger E\,x_{Sk}
\end{aligned}$$

and the norm of this quantity is bounded by

$$\sin\phi_k\,\|x_{Sk}\|_2 + \sigma_k(\widetilde{S}_k)^{-1}\|\widetilde{J}\|_2\,\|x_{Sk}\|_2 + \sigma_k(\widetilde{S}_k)^{-1}\|E\|_2\,\|x_{Sk}\|_2 .$$

Combining all the bounds, we get

$$\frac{\|x_{Sk} - \widetilde{x}_{Sk}\|_2}{\|x_{Sk}\|_2} \le \frac{1}{\sigma_k(\widetilde{S}_k)}\left(\sin\phi_k\,\frac{\|r_{Sk}\|_2}{\|x_{Sk}\|_2} + \|\widetilde{J}\|_2 + \|E\|_2 + \frac{\|e\|_2}{\|x_{Sk}\|_2}\right) + \sin\phi_k .$$

Eq. (3.3) follows from division by $\|x_k\|_2$ and using $\|A\,x_k\|_2 \le \|A\|_2\,\|x_k\|_2$. $\square$

The upper bound for the relative perturbation of x_{Sk} depends on $\sin\phi_k$, the norm of the perturbations E and e, as well as the norm of the block J. The factor $\Psi_k = \sigma_1/\sigma_k(S_k)$ is approximately equal to the "effective" condition number σ_1/σ_k associated with the truncated SVD solution.

We emphasize that upper bounds of the form (3.2) and (3.3) are often pessimistic. One reason is their generality; for example, in (3.2) we must allow for the possibility that $\|J\|_2 = 0$ and $\|\widetilde{J}\|_2 \neq 0$. Another reason is the repeated use of the triangle inequality to gain insight into identifying the influencing parameters – which is the purpose, after all, of perturbation theory.

4 Numerical Results

In this section we illustrate the above theory with some numerical examples. All matrices are symmetric indefinite, and they are generated as

$$A = W\,\Lambda\,W^T, \qquad \Lambda = D\,\Sigma,$$

where W is a random orthogonal matrix, $D = \mathrm{diag}(\pm1,\ldots,\pm1)$ with the signs chosen randomly, and the singular values in Σ are chosen such that $\sigma_1 = 1$ and $\sigma_n = 10^{-10}$. Moreover, the singular values are geometrically distributed between σ_1 and σ_k, and also between σ_{k+1} and σ_n. In all the experiments we used $n = 20$ and numerical rank $k = 15$. The VSV algorithm for symmetric indefinite matrices computes a middle matrix S of the form $S = L\,\Omega\,L^T$ where L is lower triangular, and the matrix T_k in (1.3) is therefore the leading $k \times k$ lower triangular submatrix of L.

Table 4.1: Average results from 50 experiments with symmetric indefinite 20×20 matrices with numerical rank $k=15$; N is the maximum number of iterations used in the condition estimator in the VSV decomposition algorithm.

$\frac{\sigma_k}{\sigma_{k+1}}$	N	$\|J\|_2$	$\sin\theta_k$	bound (2.4)	$\frac{\|x_k-x_{Sk}\|_2}{\|x_k\|_2}$	bound (2.5)
10	2	$4.4\cdot10^{-5}$	$2.9\cdot10^{-2}$	$4.9\cdot10^{-2}$	$5.4\cdot10^{-4}$	$3.5\cdot10^{-2}$
	4	$3.2\cdot10^{-7}$	$3.0\cdot10^{-4}$	$3.5\cdot10^{-4}$	$6.9\cdot10^{-6}$	$3.0\cdot10^{-4}$
	6	$4.3\cdot10^{-9}$	$4.0\cdot10^{-6}$	$4.7\cdot10^{-6}$	$9.7\cdot10^{-8}$	$4.0\cdot10^{-6}$
2	4	$1.3\cdot10^{-4}$	$1.2\cdot10^{-1}$	$2.2\cdot10^{-1}$	$3.8\cdot10^{-3}$	$1.7\cdot10^{-1}$
	6	$3.0\cdot10^{-5}$	$2.9\cdot10^{-2}$	$5.6\cdot10^{-2}$	$9.1\cdot10^{-4}$	$3.3\cdot10^{-2}$
	10	$4.0\cdot10^{-6}$	$4.2\cdot10^{-3}$	$7.9\cdot10^{-3}$	$9.9\cdot10^{-5}$	$4.3\cdot10^{-3}$

In our first numerical experiments we illustrate the accuracy of the VSV signal subspace $\mathcal{V}_k$ and the truncated VSV solution x_{Sk}. We used two values of σ_{k+1} corresponding to a small gap, $\sigma_k/\sigma_{k+1}=2$, and a larger gap, $\sigma_k/\sigma_{k+1}=10$, respectively. A rank-revealing algorithm cannot be expected to reveal relative gaps much smaller than about 2. The condition estimator in the VSV algorithm is allowed to use at most N inverse iterations, and by varying N we can study the influence of the norm $\|J\|_2$ on the results.

Table 4.1 shows average results of 50 experiments for each combination of σ_k/σ_{k+1} and N. In these experiments the relative residual norm $\|r_k\|_2/\|A\,x_k\|_2$ is of the order 10^{-6}. Evidently, the more inverse iterations the better the condition estimation and the smaller the norm $\|J\|_2$ of the off-diagonal block in S. Clearly, as $\|J\|_2$ decreases the quality of the VSV signal subspace as well as the truncated VSV solution increases, in the sense that $\mathcal{V}_k$ approaches $\mathcal{W}_k$ and x_{Sk} approaches x_k. We note that in this example the bound (2.4) for $\sin\theta_k$ is always very accurate, while the bound (2.5) for the relative error in the solution is approximately a factor 100 too large (which is due to the use of the repeated use of the triangle inequality in deriving the bound).

Table 4.2 shows average results of 50 experiments for various combinations of σ_k and $\|r_k\|_2/\|A\,x_k\|_2$. The gap in the singular values was $\sigma_k/\sigma_{k+1}=10$, and at most $N=2$ inverse iterations are used in the condition estimation. The table shows that as $\sigma_{k+1}=\sigma_k/10$ decreases, the norm $\|J\|_2$ of the off-diagonal block decreases, and consequently the subspace angle θ_k also decreases. The table also shows that the accuracy of the solution improves as the residual norm decreases, until the accuracy settles at a level determined by the subspace angle θ_k (and thus by the norm $\|J\|_2$).

To explain the influence of σ_k on the accuracy of the solution, note that $\sigma_k(S_k)$, the smallest singular value of S_k, is approximately equal to σ_k. Since the third term in (2.5) is proportional to $\sigma_k(S_k)^{-1}$, the influence of this term increases

Table 4.2: Average results from 50 experiments with symmetric indefinite 20×20 matrices with numerical rank $k = 15$, and different combinations of σ_k and relative residual $\|r_k\|_2/\|A\,x_k\|_2$. At most $N = 2$ iterations are used in the condition estimator.

σ_k	$\frac{\|r_k\|_2}{\|A\,x_k\|_2}$	$\|J\|_2$	$\sin\theta_k$	$\frac{\|x_k - x_{Sk}\|_2}{\|x_k\|_2}$	bound (2.5)
10^{-3}	10^{-2}	$3.3 \cdot 10^{-5}$	$2.6 \cdot 10^{-2}$	$3.7 \cdot 10^{-2}$	$3.0 \cdot 10^{-1}$
	10^{-4}	$3.2 \cdot 10^{-5}$	$2.4 \cdot 10^{-2}$	$6.8 \cdot 10^{-3}$	$2.9 \cdot 10^{-2}$
	10^{-6}	$2.1 \cdot 10^{-5}$	$1.6 \cdot 10^{-2}$	$4.4 \cdot 10^{-3}$	$1.8 \cdot 10^{-2}$
	0	$3.3 \cdot 10^{-5}$	$2.5 \cdot 10^{-2}$	$3.8 \cdot 10^{-3}$	$2.8 \cdot 10^{-2}$
10^{-6}	10^{-5}	$1.9 \cdot 10^{-8}$	$1.6 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$	$1.7 \cdot 10^{-1}$
	10^{-7}	$2.5 \cdot 10^{-8}$	$2.3 \cdot 10^{-2}$	$5.6 \cdot 10^{-2}$	$2.6 \cdot 10^{-2}$
	10^{-9}	$2.2 \cdot 10^{-8}$	$2.0 \cdot 10^{-2}$	$5.0 \cdot 10^{-3}$	$2.2 \cdot 10^{-2}$
	0	$2.2 \cdot 10^{-8}$	$1.9 \cdot 10^{-2}$	$5.1 \cdot 10^{-3}$	$2.0 \cdot 10^{-2}$

when σ_k decreases. Although the second term in (2.5) is also proportional to $\sigma_k(S_k)^{-1}$, the influence from this term remains essentially unchanged because the factor $\|J\|_2$ decreases proportionally with σ_k.

The next numerical examples illustrate the accuracy of perturbed VSV quantities $\widetilde{\mathcal{V}}_k$ and $\widetilde{x}_{Sk}$. In these experiments we used $\sigma_k = 10^{-3}$, $\sigma_{k+1} = 10^{-4}$, and a maximum of $N = 2$ and 6 inverse iterations in the condition estimator.

Table 4.3 shows the influence of the norm $\|E\|_2$ of the matrix perturbation on the subspace angle ϕ_k between $\mathcal{V}_k$ and $\widetilde{\mathcal{V}}_k$, together with the upper bound (3.2). The average size of the norms $\|J\|_2$ and $\|\widetilde{J}\|_2$ of the off-diagonal blocks are also given. Clearly, $\sin\phi_k$ decreases with $\|E\|_2$ for both $N = 2$ and 6. The results for $N = 2$ illustrate that the norms $\|J\|_2$ and $\|\widetilde{J}\|_2$ must enter in a conservative way in the worst-case bound (3.2), as discussed in §3.

Our last experiments, shown in Table 4.4, illustrate the perturbation bound in Eq. (3.3) for two different values of σ_k and N, and a range of perturbation norms δ chosen such that $\|e\|_2 = \|E\|_2 = \delta$. The limiting factor is determined by the largest of the terms in the parenthesis in (3.3), and we chose to keep the relative residual norm small in order to illustrate the influence of σ_k and δ.

When $\sigma_k = 10^{-3}$ (top part of the table) and $N = 2$ then the error decreases with δ until this quantity is somewhat below $\|J\|_2 \approx 10^{-5}$ (see Table 4.3), and then the latter quantity "takes over" and effectively acts as a perturbation. When $N = 6$ then $\|J\|_2$ is of the order machine precision, and the accuracy is determined solely by the size of δ. This is true for both values of the relative residual.

The overall behavior of the error norm and the bound is practically the same

Table 4.3: Average results from 50 experiments with symmetric indefinite 20×20 matrices with numerical rank $k = 15$, $\sigma_k = 10^{-3}$, and $\sigma_{k+1} = 10^{-4}$.

	—— $N = 2$ ——		—— $N = 6$ ——	
	$\|J\|_2 \approx \|\widetilde{J}\|_2 \approx 10^{-5}$		$\|J\|_2 \approx \|\widetilde{J}\|_2 \approx 10^{-9}$	
$\|E\|_2$	$\sin\phi_k$	bound (3.2)	$\sin\phi_k$	bound (3.2)
10^{-4}	$4.9 \cdot 10^{-2}$	$1.8 \cdot 10^{-1}$	$2.8 \cdot 10^{-2}$	$1.1 \cdot 10^{-1}$
10^{-5}	$2.6 \cdot 10^{-2}$	$7.0 \cdot 10^{-2}$	$2.8 \cdot 10^{-3}$	$1.1 \cdot 10^{-2}$
10^{-6}	$1.6 \cdot 10^{-2}$	$6.5 \cdot 10^{-2}$	$2.7 \cdot 10^{-4}$	$1.1 \cdot 10^{-3}$
10^{-8}	$7.5 \cdot 10^{-5}$	$7.5 \cdot 10^{-2}$	$2.8 \cdot 10^{-6}$	$1.6 \cdot 10^{-5}$
10^{-10}	$1.9 \cdot 10^{-7}$	$6.8 \cdot 10^{-2}$	$3.6 \cdot 10^{-8}$	$6.8 \cdot 10^{-6}$
10^{-12}	$3.7 \cdot 10^{-8}$	$4.7 \cdot 10^{-2}$	$3.4 \cdot 10^{-8}$	$7.9 \cdot 10^{-6}$

for $\sigma_k = 10^{-6}$ (bottom part of the table), except that the perturbation norm δ must now be smaller in order to ensure an acute matrix perturbation and thus a meaningful perturbed solution.

5 Conclusion

Rank-revealing VSV decompositions–which may be important in connection with, e.g., correlation matrices in signal processing and Hessian matrices in optimization–were developed recently as alternatives to the SVD and the UTV decompositions. The VSV decomposition preserves symmetry in both the computed decomposition as well as in the reduced-rank matrix approximations derived from it, and symmetry is also maintained in the underlying algorithm. This paper gives bounds for the accuracy of truncated VSV solutions computed by neglecting blocks with small norm in the VSV decomposition. We give bounds for the difference between the truncated VSV solution and the related truncated SVD solution, and bounds for the perturbation in the truncated VSV solution when the matrix and/or the right-hand side is perturbed. Numerical results are given that illustrate the theoretical bounds. Our bounds and numerical examples give insight into the factors that influence the accuracy of the VSV solutions.

REFERENCES

1. E. S. Baker and R. D. DeGroat, *A correlation-based subspace tracking algorithm*, IEEE Trans. Signal Proc., 46 (1998), pp. 3112–3116.
2. G. H. Golub and C. F. Van Loan, *Matrix Computations*, Third Edition, Johns Hopkins University Press, Baltimore, MD, 1996.

Table 4.4: Average results from 50 experiments with symmetric indefinite 20×20 matrices with numerical rank $k = 15$ and $\sigma_k/\sigma_{k+1} = 10$. The perturbation norms are $\|e\|_2 = \|E\|_2 = \delta$, and N denotes the maximum number of condition estimation inverse iterations. The results above the double line are for $\sigma_k = 10^{-3}$, and below the double line for $\sigma_k = 10^{-6}$.

		$N = 2$		$N = 6$	
$\frac{\|r_{Sk}\|_2}{\|Ax_{Sk}\|_2}$	δ	$\frac{\|x_{Sk}-\tilde{x}_{Sk}\|_2}{\|x_{Sk}\|_2}$	bound (3.3)	$\frac{\|x_{Sk}-\tilde{x}_{Sk}\|_2}{\|x_{Sk}\|_2}$	bound (3.3)
10^{-8}	10^{-3}	$1.1 \cdot 10^{-1}$	> 1	$1.9 \cdot 10^{-1}$	> 1
	10^{-4}	$2.4 \cdot 10^{-2}$	$2.8 \cdot 10^{-1}$	$1.8 \cdot 10^{-2}$	$2.1 \cdot 10^{-1}$
	10^{-6}	$2.5 \cdot 10^{-3}$	$3.7 \cdot 10^{-2}$	$1.9 \cdot 10^{-4}$	$2.1 \cdot 10^{-3}$
	10^{-8}	$2.1 \cdot 10^{-3}$	$4.1 \cdot 10^{-2}$	$2.2 \cdot 10^{-6}$	$2.5 \cdot 10^{-5}$
10^{-4}	10^{-3}	$1.5 \cdot 10^{-1}$	> 1	$1.4 \cdot 10^{-1}$	> 1
	10^{-4}	$2.1 \cdot 10^{-2}$	$2.7 \cdot 10^{-1}$	$1.7 \cdot 10^{-2}$	$2.1 \cdot 10^{-1}$
	10^{-6}	$3.1 \cdot 10^{-3}$	$4.4 \cdot 10^{-2}$	$2.0 \cdot 10^{-4}$	$2.1 \cdot 10^{-3}$
	10^{-8}	$3.6 \cdot 10^{-3}$	$4.5 \cdot 10^{-2}$	$1.8 \cdot 10^{-6}$	$2.3 \cdot 10^{-5}$
10^{-4}	10^{-7}	> 1	> 1	$2.4 \cdot 10^{-1}$	> 1
	10^{-8}	> 1	> 1	$3.0 \cdot 10^{-2}$	$2.8 \cdot 10^{-1}$
	10^{-9}	$9.5 \cdot 10^{-2}$	> 1	$2.3 \cdot 10^{-3}$	$2.6 \cdot 10^{-2}$
	10^{-10}	$2.7 \cdot 10^{-2}$	$4.3 \cdot 10^{-1}$	$2.7 \cdot 10^{-4}$	$2.9 \cdot 10^{-3}$
	10^{-11}	$2.3 \cdot 10^{-3}$	$1.1 \cdot 10^{-1}$	$2.4 \cdot 10^{-5}$	$2.7 \cdot 10^{-4}$
	10^{-12}	$3.2 \cdot 10^{-4}$	$3.2 \cdot 10^{-2}$	$3.2 \cdot 10^{-6}$	$3.7 \cdot 10^{-5}$

3. P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM, Philadelphia, 1998.

4. P. C. Hansen and P. Y. Yalamov, *Computing symmetric rank-revealing VSV decompositions via triangular factorization*, SIAM J. Matrix Anal. Appl., to appear.

5. F. T. Luk and S. Qiao, *A symmetric rank-revealing Toeplitz matrix decomposition*, J. VLSI Signal Proc., 14 (1996), pp. 19–28.

6. G. W. Stewart, *Rank degeneracy*, SIAM J. Sci. Stat. Comput., 5 (1984), pp. 403–413.

7. G. W. Stewart and J.-G. Sun, *Matrix Perturbation Theory*, Academic Press, San Diego, 1990.