Harmonic and refined extraction methods for the singular value problem, with applications in least squares problems

Michiel E. Hochstenbach*

December 17, 2002

Abstract. For the accurate approximation of the minimal singular triple (singular value and left and right singular vector), we may use two separate search spaces, one for the left, and one for the right singular vector. In Lanczos bidiagonalization, for example, such search spaces are constructed. In [3], the author proposes a Jacobi–Davidson type method for the singular value problem, where solutions to certain correction equations are used to expand the search spaces.

As noted in [3], the standard Galerkin subspace extraction works well for the computation of large singular triples, but may lead to unsatisfactory approximations to small and interior triples. To overcome this problem for the smallest triples, we propose three harmonic and a refined approach. All methods are derived in a number of different ways. Two of these methods can also be applied when we are interested in interior singular triples. Theoretical results as well as numerical experiments indicate that the results of the alternative extraction processes are often better than the standard approach. We show that when Lanczos bidiagonalization is used to approximate the smallest singular triples, the standard, harmonic, and refined extraction methods are all essentially equivalent. This gives more insight in the success of the use of Lanczos bidiagonalization to find the smallest singular triples.

Finally, we present a novel method for the least squares problem, the success of which is based on a good extraction process for the smallest singular triples. The truncated SVD is also discussed in this context.

Key words: SVD, singular value problem, subspace method, subspace extraction, two-sided approach, harmonic extraction, refined extraction, Rayleigh quotient, Lanczos bidiagonalization, Saad's theorem, least squares problem, truncated SVD.

AMS subject classification: 65F15, 65F50, (65F35, 93E24).

1 Introduction

We study subspace methods for the computation of some singular triples (i.e., singular values and their corresponding singular left and right vectors) for large sparse matrices. The methods we consider are *two-sided*, i.e., they work with two search spaces: a search space $\mathcal{U}$ for the left singular vector, and a search space $\mathcal{V}$ for the right singular vector.

In short, such a subspace method works as follows. We start with given search subspaces from which approximations to the singular triples are computed, this is called the *subspace extraction*. Then we expand the subspaces by new directions, the *subspace expansion*. The idea is that as the search subspaces grow, the approximations will converge to a singular triple of the original

*Mathematical Institute, Utrecht University, P.O. Box 80.010, NL-3508 TA Utrecht, The Netherlands. `www.math.uu.nl/people/hochstenbach`, `hochstenbach@math.uu.nl`. Version: December 17, 2002.

problem. In order to keep computation costs low, we usually do not expand the search spaces to the whole space. If the process does not converge in a certain number of iterations, then the method is *restarted* with a few selected approximations as the basis for new search spaces. If one or more singular triples have been detected, and we want to find other triples, we can *deflation* techniques to avoid finding the same triple again.

A well known example of a two-sided subspace method for the singular value problem is the (Golub–Kahan–)Lanczos bidiagonalization ([1], see also [2, p. 495]). A recent method, JDSVD, based on Jacobi–Davidson type expansion techniques, is introduced in [3]. Besides a new subspace expansion process, some new nonstandard extraction techniques are suggested there. In this paper we analyze these, and other alternative extractions in more detail.

Let us first introduce some notations. Let A be a real (large sparse) $m \times n$ matrix with singular value decomposition (SVD) $A = X\Sigma Y^T$ and singular values

$$0 \le \sigma_{\min} = \sigma_p \le \sigma_{p-1} \le \cdots \le \sigma_2 \le \sigma_1 = \sigma_{\max},$$

where $p := \min\{m, n\}$. The assumption that A is real is made for convenience only, the adaptations for complex A are not difficult. Denote the left and right singular vectors by x_j $(1 \le j \le m)$ and y_j $(1 \le j \le n)$, respectively. For later use, we also introduce a second labeling for the singular values:

$$\sigma_{-1} \le \sigma_{-2} \le \cdots \le \sigma_{-p+1} \le \sigma_{-p}, \tag{1.1}$$

so that, if $\sigma_{\min} > 0$,

$$\sigma_{-j}(A) = \sigma_j^{-1}(A^+),$$

where A^+ is the pseudoinverse of A. (Such a labeling is also used for eigenvalues in [7]; one of its benefits is the matrix size independency of the indices of the smallest values. This facilitates the formulation of results for those values.) By $\|\cdot\|$ we denote the Euclidean norm, while $\kappa(A)$ is the condition number of A. For a subspace $\mathcal{U}$, let $P_{\mathcal{U}}$ denote the orthogonal projection onto $\mathcal{U}$; U denotes a "*search matrix*" whose columns form an orthonormal basis for $\mathcal{U}$. We write $\mathcal{N}(A)$ for the nullspace of A, and e_j for the jth canonical vector. For a positive definite matrix B, the B-inner product is defined by

$$(x, y)_B := y^T B x.$$

We denote the situation where $(x, y)_B = 0$ as $x \perp_B y$.

This paper has been organized as follows. Section 2 recalls the standard subspace extraction and some of its properties from [3], and presents a theorem like Saad's theorem on Rayleigh–Ritz approximations, that gives insight in the strength and weakness of the standard extraction. Section 3 explores three variants of harmonic extraction, based on certain Galerkin conditions on the inverse of the matrix. In Section 4, we propose a refined extraction method. Section 5 discusses the Rayleigh quotient for the singular value problem. We discuss the possibilities for interior singular values in Section 6, and those for nonsquare or singular matrices in Section 7. In Section 8, we study the extraction methods for the special case of the Lanczos bidiagonalization, and show that all extraction methods for the smallest singular triples are essentially equivalent. The application of the methods to the least squares problem and the truncated SVD is the subject of Section 9. Numerical experiments are presented in Section 10, and some conclusions are collected in Section 11.

2 Standard extraction

Given a left search space $\mathcal{U}$ and a right search space $\mathcal{V}$, we would like to determine an approximate left singular vector $u \in \mathcal{U}$ and an approximate right singular vector $v \in \mathcal{V}$. Throughout the paper, we assume that both AV and A^TU are of full rank. (When they are not, the search spaces $\mathcal{U}$ or $\mathcal{V}$ contain a singular vector corresponding to the singular value 0.)

Given (possibly different) approximations θ and η to the same singular value, and approximate left and right singular vectors u and v, the *residual* r is defined as

$$r = \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} = \begin{bmatrix} Av - \theta u \\ A^T u - \eta v \end{bmatrix}.$$

These approximate values and vectors are determined by a *double Galerkin condition* as follows. Write $u = Uc$ and $v = Vd$, where it is understood that $c, d \neq 0$. The *standard extraction* process is now derived from one of the two following equivalent conditions:

$$\text{(i)} \left\{ \begin{array}{lll} AVd - \theta Uc & \perp & \mathcal{U}, \\ A^T Uc - \eta Vd & \perp & \mathcal{V}. \end{array} \right.$$

$$\text{(ii)} \left\{ \begin{array}{lll} U^T AVd & = & \theta c, \\ V^T A^T Uc & = & \eta d. \end{array} \right.$$

Here, (i) can be regarded as the standard Galerkin conditions for the singular value problem; the word "standard" reflects the fact that we choose the test spaces (on the right hand side of (i)) equal to the search spaces. Choosing the scaling $\|c\| = \|d\| = 1$, we see from (ii) that c and d are left and right singular vectors of

$$H := U^T AV$$

with singular value $\theta = \eta$; stated differently: $u = Uc$ and $v = Vd$ are left and right singular vectors of A with singular value $\theta = \eta$ with respect to the subspaces $\mathcal{U}$ and $\mathcal{V}$. Note that in this extraction process we have $\theta = \eta$, for some other methods in Section 3 and 4 this will not be the case.

Definition 2.1 *$\mathcal{U}$ and $\mathcal{V}$ are* invariant singular subspaces *if*

$$A\mathcal{V} \subset \mathcal{U} \quad \textit{and} \quad A^T \mathcal{U} \subset \mathcal{V}.$$

In [3], the following argument is given why this H is optimal in the following sense. Given U and V, define the *residual matrices* as

$$R_1(K) := AV - UK \quad \text{and} \quad R_2(L) := A^T U - VL. \tag{2.1}$$

When $\mathcal{U}$ and $\mathcal{V}$ are invariant singular subspaces, then, with $H = U^T AV$, we have $R_1(H) = R_2(H^T) = 0$. When $\mathcal{U}$ and $\mathcal{V}$ are not invariant, we may be interested in the H that minimizes $\|R_1(K)\|$ and $\|R_2(L^T)\|$. The following theorem from [3] states that in both cases the minimizing matrix is (also) given by $H = U^T AV$.

Theorem 2.2 *[3, Theorem 4.1] For given $m \times k$ matrix U and $n \times k$ matrix V, let $H = U^T AV$.*
(a) If the columns of U are orthonormal, then for all $k \times k$ matrices K we have $\|R_1(H)\| \leq \|R_1(K)\|$. Moreover, H is unique with respect to the Frobenius norm: $\|R_1(H)\|_F \leq \|R_1(K)\|_F$ with equality only when $K = H$.
(b) If the columns of V are orthonormal, then $H^T = V^T A^T U$ minimizes the norm of $R_2(L)$, and H^T is unique with respect to the Frobenius norm.

In two-sided subspace methods for the singular value problem, we construct increasing sequences

$$\mathcal{U}_k \subset \mathcal{U}_{k+1} \quad \text{and} \quad \mathcal{V}_k \subset \mathcal{V}_{k+1}.$$

Because of Theorem 2.2, it is a natural idea to take the k singular values $\theta_k^{(k)} \leq \cdots \leq \theta_1^{(k)}$ of $H_k := U_k^T AV_k$ as approximations to the singular values of A. Then the following theorem from [3] states that all $\theta_j^{(k)}$ converge monotonically increasing to the σ_j. (As usual, by "*asymptotic convergence*" in the context of subspace methods, we mean the convergence behavior of these methods in a situation where we have a (very) good approximation to an eigenpair, rather than the situation where the dimension of the subspace goes to infinity.) Moreover, we have elegant error bounds that can be proved using Weyl's monotonicity theorem. We will sometimes omit the index k for convenience.

Theorem 2.3 *[3, Theorem 4.3]*

(a) For fixed $j \leq p$, and $k \geq j$, the $\theta_j^{(k)}$ converge monotonically (up)to σ_j:

$$\theta_j^{(k)} \leq \theta_j^{(k+1)} \leq \sigma_j.$$

(b) For each $j = 1, \ldots, k$, there exist singular values $\sigma_{j'}$ of A which can be put in one-one correspondence with the singular values θ_j of H in such a way that

$$|\sigma_{j'} - \theta_j| \leq \max \{\|R_1(H)\|, \|R_2(H^T)\|\}.$$

Moreover,

$$\sum_{j=1}^{k} (\sigma_{j'} - \theta_j)^2 \leq \|R_1(H)\|_F^2 + \left\|R_2(H^T)\right\|_F^2.$$

Two remarks regarding this theorem are in place. First, as also noticed in [3], the standard extraction leads to monotonic convergence for the largest singular value (part (a)), but it can imply irregular behavior for the smallest singular value.

Second, if the norms of the residual matrices $R_1(H)$ and $R_2(H^T)$ are small, then the largest singular values of H must be good approximations to the largest singular values of A (part (b)). The smallest singular values of H are not necessarily good approximations to the small singular values of A. So H will tend to be a better approximation to the "top" of the singular spectrum of A than to the "bottom". See also Section 9.2.

Now, we present another result that sheds more light on the standard extraction: a theorem that expresses the quality of the approximate singular vectors produced by this extraction in terms of the quality of the search spaces. For the Hermitian eigenvalue problem, such a result is proved by Saad (see [8, p. 136]). This result can be extended to non-Hermitian matrices ([11, p. 286]). We first give a new short proof of Saad's result in terms of orthogonal projections, before we prove a similar result for the singular value problem along the same lines.

Let P_u denote the orthogonal projection onto span(u). Note that because $u \in \mathcal{U}$, the projections satisfy

$$P_{\mathcal{U}} P_u = P_u P_{\mathcal{U}} = P_u.$$

Suppose that we have a search space $\mathcal{U}$ for the Hermitian eigenproblem. The Rayleigh–Ritz process for the eigenvalue problem $Bx = \lambda x$ (see, for instance, [7]) ensures that

$$P_{\mathcal{U}} B P_{\mathcal{U}} P_u = \theta P_u. \tag{2.2}$$

Theorem 2.4 *(Saad, [8, p. 136]) Suppose that B is a Hermitian matrix with eigenpair (λ, x). Let (θ, u) be the Ritz pair (with respect to the search space $\mathcal{U}$), for which θ is the Ritz value closest to λ. Then*

$$\sin(u, x) \leq \sqrt{1 + \frac{\gamma^2}{\delta^2}}\, \sin(\mathcal{U}, x),$$

where

$$\begin{array}{rcl} \gamma & = & \|P_{\mathcal{U}}(B - \lambda I)(I - P_{\mathcal{U}})\|, \\ \delta & = & \min_{\theta_j \neq \theta} |\theta_j - \lambda|, \end{array}$$

where θ_j ranges over all Ritz values not equal to θ.

Proof: We start with

$$x = P_u x + (P_{\mathcal{U}} - P_u)x + (I - P_{\mathcal{U}})x.$$

Apply $P_{\mathcal{U}}(B - \lambda I)$ on both sides, and use (2.2) to get

$$0 = (\theta - \lambda)P_u x + P_{\mathcal{U}}(B - \lambda I)(P_{\mathcal{U}} - P_u)x + P_{\mathcal{U}}(B - \lambda I)(I - P_{\mathcal{U}})x, \tag{2.3}$$

so

$$(P_{\mathcal{U}} - P_u)(B - \lambda I)(P_{\mathcal{U}} - P_u)x + P_u\left((\theta - \lambda) + (B - \lambda I)(P_{\mathcal{U}} - P_u)\right)x = -P_{\mathcal{U}}(B - \lambda I)(I - P_{\mathcal{U}})x.$$

Taking the square of the norms and using Pythagoras' theorem leads to

$$\delta^2 \, \|(P_{\mathcal{U}} - P_u)x\|^2 \le \|(P_{\mathcal{U}} - P_u)(B - \lambda I)(P_{\mathcal{U}} - P_u)x\|^2 \le \gamma^2 \, \|(I - P_{\mathcal{U}})x\|^2.$$

Since $\|(I - P_{\mathcal{U}})x\| = \sin(\mathcal{U}, x)$ and $\|(I - P_u)x\| = \sin(u, x)$, the result now follows from

$$\|(I - P_u)x\|^2 = \|(I - P_{\mathcal{U}})x\|^2 + \|(P_{\mathcal{U}} - P_u)x\|^2.$$

□

For the singular value problem, the standard extraction gives

$$P_{\mathcal{U}} A P_{\mathcal{V}} P_v = \theta P_u \quad \text{and} \quad P_{\mathcal{V}} A^T P_{\mathcal{U}} P_u = \theta P_v.$$

We are now in a position to prove a similar result for the standard extraction for the singular value problem.

Theorem 2.5 *(cf. Theorem 2.4) Let (σ, x, y) be a singular triple of A, and (θ, u, v) be the approximate triple (derived with the standard extraction with respect to the search spaces $\mathcal{U}$ and $\mathcal{V}$), for which θ is the value closest to σ. Then*

$$\max\{\sin(u, x), \sin(v, y)\} \le \sqrt{1 + 2\frac{\widetilde{\gamma}^2}{\widetilde{\delta}^2}} \, \max\{\sin(\mathcal{U}, x), \sin(\mathcal{V}, y)\},$$

where

$$\begin{array}{rcl} \widetilde{\gamma} & = & \max\{\|P_{\mathcal{U}} A (I - P_{\mathcal{V}})\|, \|(I - P_{\mathcal{U}}) A P_{\mathcal{V}})\|\}, \\ \widetilde{\delta} & = & \left\{ \begin{array}{ll} \min_{\theta_j \ne \theta} |\theta_j - \sigma| & \text{when } H \text{ is square,} \\ \min(\min_{\theta_j \ne \theta} |\theta_j - \sigma|, \sigma) & \text{when } H \text{ is nonsquare,} \end{array} \right. \end{array}$$

where θ_j ranges over all approximate singular values of H not equal to θ.

Proof: The proof follows the same line as the proof of Theorem 2.4, where we take for B the augmented matrix

$$\begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}, \tag{2.4}$$

and make the following other substitutions in the proof of Theorem 2.4: replace

$$x \text{ by } \begin{bmatrix} x \\ y \end{bmatrix}, \qquad P_{\mathcal{U}} \text{ by } \begin{bmatrix} P_{\mathcal{U}} & 0 \\ 0 & P_{\mathcal{V}} \end{bmatrix}, \quad \text{and} \quad P_u \text{ by } \begin{bmatrix} P_u & 0 \\ 0 & P_v \end{bmatrix}.$$

One may check that we get (cf. (2.3))

$$\begin{aligned} 0 = (\theta - \sigma) \begin{bmatrix} P_u x \\ P_v y \end{bmatrix} &+ \begin{bmatrix} P_{\mathcal{U}} & 0 \\ 0 & P_{\mathcal{V}} \end{bmatrix} \begin{bmatrix} -\sigma I & A \\ A^T & -\sigma I \end{bmatrix} \begin{bmatrix} P_{\mathcal{U}} - P_u & 0 \\ 0 & P_{\mathcal{V}} - P_v \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \\ &+ \begin{bmatrix} P_{\mathcal{U}} & 0 \\ 0 & P_{\mathcal{V}} \end{bmatrix} \begin{bmatrix} -\sigma I & A \\ A^T & -\sigma I \end{bmatrix} \begin{bmatrix} I - P_{\mathcal{U}} & 0 \\ 0 & I - P_{\mathcal{V}} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \end{aligned}$$

Since the last term on the right-hand side can be written as

$$\begin{bmatrix} 0 & P_{\mathcal{U}} A (I - P_{\mathcal{V}}) \\ (I - P_{\mathcal{U}}) A^T P_{\mathcal{V}} & 0 \end{bmatrix} \begin{bmatrix} (I - P_{\mathcal{U}})x \\ (I - P_{\mathcal{V}})y \end{bmatrix},$$

the norm of this term is bounded by

$$\sqrt{2}\, \widetilde{\gamma} \max\{\|(I - P_{\mathcal{U}})x\|, \|(I - P_{\mathcal{V}})y\|\}.$$

Furthermore, the smallest singular value of

$$\begin{bmatrix} P_{\mathcal{U}} - P_u & 0 \\ 0 & P_{\mathcal{V}} - P_v \end{bmatrix} \begin{bmatrix} -\sigma I & A \\ A^T & -\sigma I \end{bmatrix} \begin{bmatrix} P_{\mathcal{U}} - P_u & 0 \\ 0 & P_{\mathcal{V}} - P_v \end{bmatrix}$$

is $\widetilde{\delta}$. The result now follows in the same way as in the proof of Theorem 2.4. □

The theorem states that when $\angle(\mathcal{U}, x) \to 0$ and $\angle(\mathcal{V}, y) \to 0$, the standard extraction gives good Ritz vectors, unless the singular values of H are poorly separated ($\widetilde{\delta} \approx 0$). In fact, $\widetilde{\delta}$ becomes 0 in case of double singular values of H, see also Example 7.1. The phenomenon of poorly separated Ritz values is also encountered in the Rayleigh–Ritz method for eigenvalue problem; it is often not very serious: we may just continue with the subspace method (by expanding the search spaces), at the next step the singular values of H may be well separated.

A more serious problem of the standard extraction is that the theorem does not predict *which* singular triple is the best: it is a problem of *selection*. In [3], the following example is given. Suppose that $u = \sum_{j=1}^{m} \gamma_j x_j$ and $v = \sum_{j=1}^{n} \delta_j y_j$ are approximate singular vectors of unit length; then $\theta = u^T A v = \sum_{j=1}^{p} \gamma_j \delta_j \sigma_j$. (We may assume θ is nonnegative; otherwise, take $-u$ instead of u.) Now suppose that $\theta \approx \sigma_1$, in the sense that $\sigma_2 < \theta < \sigma_1$, and that $\sigma_1 - \theta$ is (much) smaller than $\theta - \sigma_2$. Then we conclude that $\gamma_1 \approx 1$ and $\delta_1 \approx 1$, so u and v are good approximations to x_1 and y_1. But when $\theta \approx \sigma_p$, u and v are not necessarily good approximations to x_p and y_p. For example, u could have a large component of x_{p-1} and a small component of x_1, and v could have a large component of y_{p-2} and a small component of y_1. In conclusion, when we search for the largest singular value, it is asymptotically safe to select the largest singular triple of H, but for the smallest triple it is not safe to select the smallest approximate triple. See also Example 7.1.

Failure to select the best approximate vectors is especially dangerous when we use *restarts*. At the moment of restart, selection of bad approximate singular vectors may spoil the whole process.

3 Harmonic extractions

As seen in the previous section, the standard extraction process is satisfactory in the quest for the largest singular values, but the approximations to the smallest singular values often display a irregular convergence (see also [3]).

> Here and in the next section we assume that A is nonsingular (which implies that A is square);

we will treat the general case in Section 7. Based on the observation that the smallest singular values of A are the largest ones of A^{-1}, it was suggested in [3] to consider modified Galerkin conditions on A^{-T} and A^{-1}. Here we work out this idea in detail.

The following are equivalent:

$$\text{(i)} \quad \left\{ \begin{array}{lll} A^{-T} V \widetilde{d} - \widetilde{\eta}^{-1} U \widetilde{c} & \perp & \widetilde{\mathcal{U}}, \\ A^{-1} U \widetilde{c} - \widetilde{\theta}^{-1} V \widetilde{d} & \perp & \widetilde{\mathcal{V}}. \end{array} \right.$$

$$\text{(ii)} \quad \left\{ \begin{array}{lll} \widetilde{\eta}\, \widetilde{U}^T A^{-T} V \widetilde{d} & = & \widetilde{U}^T U \widetilde{c}, \\ \widetilde{\theta}\, \widetilde{V}^T A^{-1} U \widetilde{c} & = & \widetilde{V}^T V \widetilde{d}. \end{array} \right.$$

The idea is now to choose the test spaces $\widetilde{\mathcal{U}}$ and $\widetilde{\mathcal{V}}$ in such a way, that we do not have to work with the inverse of the (large sparse) matrices A and A^T. Some terminology: in line with nomenclature for the Rayleigh–Ritz procedure, $\widetilde{u} := U\widetilde{c}$ and $\widetilde{v} := V\widetilde{d}$ are called *left and right harmonic singular vectors*, $\widetilde{\theta}$ and $\widetilde{\eta}$ *harmonic singular values*, and $(\widetilde{\theta}, \widetilde{\eta}, \widetilde{u}, \widetilde{v})$ a *harmonic singular tuple*. We remark that, similar to the harmonic Ritz values for the eigenvalue problem, the harmonic Ritz values may be ∞, see Example 7.1, and Section 5 for a way to overcome this difficulty.

Now we can make the following four choices for $(\widetilde{\mathcal{U}}, \widetilde{\mathcal{V}})$:

- $(A\mathcal{V}, A^T\mathcal{U})$ gives the standard extraction of Section 2.
- $(AA^T\mathcal{U}, A^T\mathcal{U})$ leads to the *$\mathcal{U}$-harmonic extraction*, see Section 3.1.
- $(A\mathcal{V}, A^TA\mathcal{V})$ is the *$\mathcal{V}$-harmonic extraction* to be discussed in Section 3.1.
- $(AA^T\mathcal{U}, A^TA\mathcal{V})$ gives the *double harmonic extraction*, examined in Section 3.2.

3.1 $\mathcal{U}$- and $\mathcal{V}$-harmonic extraction

In this subsection we only treat the $\mathcal{V}$-harmonic extraction, the $\mathcal{U}$-harmonic extraction is derived by interchanging the roles of $\mathcal{U}$ and $\mathcal{V}$, and those of A and A^T. The following are equivalent:

$$\text{(i)} \quad \begin{cases} A^{-T}V\widetilde{d} - \widetilde{\eta}^{-1}U\widetilde{c} & \perp \quad A\mathcal{V}, \\ A^{-1}U\widetilde{c} - \widetilde{\theta}^{-1}V\widetilde{d} & \perp \quad A^TA\mathcal{V}. \end{cases}$$

$$\text{(ii)} \quad \begin{cases} AV\widetilde{d} - \widetilde{\theta}U\widetilde{c} & \perp \quad A\mathcal{V}, \\ A^TU\widetilde{c} - \widetilde{\eta}V\widetilde{d} & \perp \quad \mathcal{V}. \end{cases}$$

$$\text{(iii)} \quad \begin{cases} V^TA^TAV\widetilde{d} & = \quad \widetilde{\theta}V^TA^TU\widetilde{c}, \\ V^TA^TU\widetilde{c} & = \quad \widetilde{\eta}\widetilde{d}. \end{cases}$$

Here, (i) expresses that the $\mathcal{V}$-harmonic method arises from Galerkin conditions on A^{-1} and A^{-T} with respect to modified test spaces. Item (ii) gives a derivation in terms of Galerkin conditions on A and A^T, but with different test spaces, compared with the standard extraction of Section 2. From (iii), we see that $V^TA^TAV\widetilde{d} = (\widetilde{\theta}\widetilde{\eta})\widetilde{d}$, that is, $(\widetilde{\theta}\widetilde{\eta}, \widetilde{v})$ is a Ritz pair of A^TA with respect to the search space $\mathcal{V}$. This is just the well-known Raleigh–Ritz approach on A^TA. The "secret", however, is in the formation of the vector $\widetilde{c}$, and hence the approximate left singular vector $\widetilde{u}$. We have $\widetilde{c} = H^{-T}\widetilde{d}$, up to scaling (see Section 7 for the case that H is nonsquare or singular). Since H^{-T} can be considered as a projected A^{-T}, this suggests that the vector $\widetilde{u} = UH^{-T}\widetilde{d}$ may be a much better approximation to the left singular vector than the "usual" approximation $A\widetilde{v} = AV\widetilde{d}$. The following lemma substantiates this.

Lemma 3.1 *Suppose that v is an approximation to the "smallest" right singular vector $y_{\min}$. Then, denoting $\varepsilon := \tan(v, y_{\min})$, we have*

$$\kappa(A)^{-1}\varepsilon \le \tan(A^{-T}v, x_{\min}) \le \varepsilon \le \tan(Av, x_{\min}) \le \kappa(A)\varepsilon,$$

where the inequalities are sharp.

Proof: Write $v = y_{\min} + \varepsilon e$, where $e \perp y_{\min}$ and $\|e\| = 1$. Note that $\tan(v, y_{\min}) = \varepsilon$. Then for $Av = \sigma_{\min}x_{\min} + \varepsilon Ae$, we have $\tan(Av, x_{\min}) = \varepsilon\sigma_{\min}^{-1}\|Ae\|$, from which the last two inequalities follow (sharp if $e = y_{p-1}$ and $\sigma_{\min} = \sigma_{p-1}$, respectively $e = y_{\max}$). Since $A^{-T}v = \sigma_{\min}^{-1}x_{\min} + \varepsilon A^{-T}e$, we have $\tan(A^{-T}v, x_{\min}) = \varepsilon\sigma_{\min}\|A^{-T}e\|$, from which we get the first two inequalities (sharp in the same circumstances as above). $\square$

Concluding from the lemma, it would be ideal, given an approximate right vector $\widetilde{v}$, to take $A^{-T}\widetilde{v}$ as approximate left vector. In practice, the action with A^{-T} is often too expensive, but H^{-T} is a, much cheaper, projected approximation to A^{-T}. Numerical experiments (see Section 10) confirm that $\widetilde{u} = UH^{-T}\widetilde{d}$ may be a much more accurate than $AV\widetilde{d}$.

Since the $(\widetilde{\theta}\widetilde{\eta})_j$ are Ritz values of A^TA with respect to the subspace $\mathcal{V}$, we can invoke well-known results. We label the values in two different ways (cf. (1.1)):

$$(\widetilde{\theta}\widetilde{\eta})_k \le \cdots \le (\widetilde{\theta}\widetilde{\eta})_1 \quad \text{and} \quad (\widetilde{\theta}\widetilde{\eta})_{-1} \le \cdots \le (\widetilde{\theta}\widetilde{\eta})_{-k}.$$

Theorem 3.2 *Let $R_V = A^TAV - V(V^TA^TAV) = (I - VV^T)A^TAV$. Then:*

(a) for fixed $j \le p$, and $k \ge j$, the $((\theta\eta)_j^{(k)})^2$ converge monotonically (up)to σ_j^2:

$$((\theta\eta)_j^{(k)})^2 \le ((\theta\eta)_j^{(k+1)})^2 \le \sigma_j^2;$$

the $((\theta\eta)_{-j}^{(k)})^2$ converge monotonically (down)to σ_{-j}^2:

$$\sigma_{-j}^2 \le ((\theta\eta)_{-j}^{(k+1)})^2 \le ((\theta\eta)_{-j}^{(k)})^2;$$

(b) for each $j = 1, \ldots, k$, there exist singular values $\sigma_{j'}$ of A which can be put in one-one correspondence with the $(\theta\eta)_j$ in such a way that

$$|(\widetilde{\theta}\widetilde{\eta})_j - \sigma_{j'}^2| \le \|R_V\| \quad \textit{and} \quad \sum_{j=1}^{k}((\widetilde{\theta}\widetilde{\eta})_j - \sigma_{j'}^2)^2 \le \|R_V\|_F^2;$$

(c) let $(\theta\eta, v)$ be the Ritz pair of A^TA where $\theta\eta$ is the Ritz value closest to σ^2. Then

$$\sin(v, y) \le \sqrt{1 + \frac{\gamma_V^2}{\delta_V^2}}\,\sin(\mathcal{V}, y),$$

where

$$\begin{array}{rcl} \gamma_V & = & \|P_{\mathcal{V}}(A^TA - \sigma^2 I)(I - P_{\mathcal{V}})\|, \\ \delta_V & = & \min_{(\theta\eta)_j \ne \theta\eta} |(\theta\eta)_j - \sigma^2|. \end{array}$$

Proof: For (a) and (b), apply [7, Theorems 11.5.1 and 11.5.2] to A^TA and AA^T. Part (c) is a corollary to Theorem 2.4, when we take $B = A^TA$. □

3.2 Double harmonic extraction

We now give a number of possible derivations for the double harmonic extraction. The following are equivalent:

$$\text{(i)} \quad \left\{ \begin{array}{lcl} A^{-T}V\widetilde{d} - \widetilde{\eta}^{-1}U\widetilde{c} & \perp & AA^T\mathcal{U}, \\ A^{-1}U\widetilde{c} - \widetilde{\theta}^{-1}V\widetilde{d} & \perp & A^TA\mathcal{V}. \end{array} \right.$$

$$\text{(ii)} \quad \left\{ \begin{array}{lcl} A^{-T}V\widetilde{d} - \widetilde{\eta}^{-1}U\widetilde{c} & \perp_{AA^T} & \mathcal{U}, \\ A^{-1}U\widetilde{c} - \widetilde{\theta}^{-1}V\widetilde{d} & \perp_{A^TA} & \mathcal{V}. \end{array} \right.$$

$$\text{(iii)} \quad \left\{ \begin{array}{lcl} A^{-T}\widetilde{V}\widetilde{c} - \widetilde{\theta}^{-1}\widetilde{U}\widetilde{d} & \perp & \widetilde{\mathcal{U}} = A\mathcal{V}, \\ A^{-1}\widetilde{U}\widetilde{d} - \widetilde{\eta}^{-1}\widetilde{V}\widetilde{c} & \perp & \widetilde{\mathcal{V}} = A^T\mathcal{U}. \end{array} \right.$$

$$\text{(iv)} \quad \left\{ \begin{array}{lcl} AV\widetilde{d} - \widetilde{\theta}U\widetilde{c} & \perp & A\mathcal{V}, \\ A^TU\widetilde{c} - \widetilde{\eta}V\widetilde{d} & \perp & A^T\mathcal{U}. \end{array} \right.$$

$$\text{(v)} \quad \left\{ \begin{array}{lcl} \widetilde{\eta}U^TAV\widetilde{d} & = & U^TAA^TU\widetilde{c}, \\ \widetilde{\theta}V^TA^TU\widetilde{c} & = & V^TA^TAV\widetilde{d}. \end{array} \right.$$

Here, (i) states that the double harmonic method arises from Galerkin conditions on A^{-1} and A^{-T} with respect to modified test spaces. Item (ii) formulates the result with respect to the standard search and test spaces, but with respect to a different inner product; note that both $(\cdot,\cdot)_{A^TA}$ and $(\cdot,\cdot)_{AA^T}$ are inner products because of the assumption that A is nonsingular. Item (iii) derives the approach from the situation where we take modified search and test spaces; as in the standard extraction, the test spaces are equal to the search spaces. This notation in item is different from the other items (exchange of $\widetilde{c}$ and $\widetilde{d}$, and of $\widetilde{\theta}$ and $\widetilde{\eta}$). The reason for this is the following. The vectors $\widetilde{u} = \widetilde{U}\widetilde{d} = AV\widetilde{d}$ and $\widetilde{v} = \widetilde{V}\widetilde{d} = A^TU\widetilde{c}$ will tend to be deficient in the direction of the smallest singular vectors. So it is better to do a "free" step of inverse iteration and take $\widetilde{u} = A^{-T}(A^TU\widetilde{c}) = U\widetilde{c}$ and $\widetilde{v} = A^{-1}(AV\widetilde{d}) = V\widetilde{d}$ as approximate singular vectors. Similar remarks for the harmonic vectors in the eigenvalue problem can be found in [5]. Item (iv) gives a derivation in terms of Galerkin conditions on A and A^T, but with respect to different test spaces. Finally, (v) can be interpreted as: $(\widetilde{\theta}^{-1} = \widetilde{\eta}^{-1}, \widetilde{u}, \widetilde{v})$ is a singular tuple of U^TAV, where A^TU and AV are orthogonal, as the following analysis shows.

We introduce the QR decompositions

$$AV = Q_UG_U \quad \text{and} \quad A^TU = Q_V^TG_V.$$

(We choose for the letter 'G', because the letter 'R' is already "overloaded", and because $G_U^TG_U$ and $G_V^TG_V$ are Choleski decompositions of V^TA^TAV and U^TAA^TU, respectively. We choose for these subscripts since AV "lives in the U-space", i.e., A maps right to left singular vectors. Similarly, A^T maps left to right singular vectors.) Note that G_U and G_V are nonsingular because of the assumption that AV and A^TU are of full rank. Then

$$Q_U = AVG_U^{-1} \quad \text{and} \quad Q_V = A^TUG_V^{-1}$$

are orthogonal and span $A\mathcal{V}$ and $A^T\mathcal{U}$, respectively. Then characterization (v) can be written as

$$\begin{cases} \widetilde{\eta}G_V^{-T}U^TAVG_U^{-1}(G_U\widetilde{d}) &= (G_V\widetilde{c}), \\ \widetilde{\theta}G_U^{-T}V^TA^TUG_V^{-1}(G_V\widetilde{c}) &= (G_U\widetilde{d}). \end{cases}$$

When we normalize $\widetilde{c}$ and $\widetilde{d}$ such that $\|G_V\widetilde{c}\| = \|G_U\widetilde{d}\| = 1$, we see that $G_V\widetilde{c}$ and $G_U\widetilde{d}$ are left and right singular vectors of

$$\widetilde{H}^{-1} := Q_V^TA^{-1}Q_U = G_V^{-T}U^TAVG_U^{-1} = G_V^{-T}HG_U^{-1}$$

corresponding to singular value $\widetilde{\theta}^{-1} = \widetilde{\eta}^{-1}$. Analogously to (2.1), we define the *residual matrices*

$$\widetilde{R}_1(K) := A^{-1}Q_U - Q_VK, \qquad \widetilde{R}_2(L) := A^{-T}Q_V - Q_UL.$$

Then

$$\begin{aligned} \widetilde{R}_1(\widetilde{H}^{-1}) &= (I - Q_VQ_V^T)A^{-1}Q_U &&= (I - Q_VQ_V^T)VG_U^{-1}, \\ \widetilde{R}_2(\widetilde{H}^{-T}) &= (I - Q_UQ_U^T)A^{-T}Q_V &&= (I - Q_UQ_U^T)UG_V^{-1}. \end{aligned}$$

Note that a multiplication by A^{-1} or A^{-T} is not necessary to compute the residual matrices (which, in practice, will not be done anyway). The following theorem can be proved applying Theorem 2.2 to A^{-1} instead of A. Informally, it states that $\widetilde{H}^{-1}$ can be considered as the best approximation to A^{-1} over $A\mathcal{V}$ and $A^T\mathcal{U}$.

Theorem 3.3 *For given $m \times k$ matrix Q_U and $n \times k$ matrix Q_V with orthogonal columns, let $\widetilde{H}^{-1} = Q_V^TA^{-1}Q_U$.*
(a) For all $k \times k$ matrices K we have $\|\widetilde{R}_1(\widetilde{H}^{-1})\| \leq \|\widetilde{R}_1(K)\|$. Moreover, $\widetilde{H}^{-1}$ is unique with respect to the Frobenius norm: $\|\widetilde{R}_1(\widetilde{H}^{-1})\|_F \leq \|\widetilde{R}_1(K)\|_F$ with equality only when $K = \widetilde{H}^{-1}$.
(b) $\widetilde{H}^{-T}$ minimizes the norm of $\widetilde{R}_2(L)$, and $\widetilde{H}^{-T}$ is unique with respect to the Frobenius norm.

Since $\widetilde{H}^{-1}$ approximates A^{-1}, we may take the singular values $\widetilde{\theta}_{-1}^{(k)} \leq \cdots \leq \widetilde{\theta}_{-k}^{(k)}$ (notation: cf. (1.1)) of

$$\widetilde{H}_k = G_UH_k^{-1}G_V^T \tag{3.1}$$

as approximations to the $\sigma_j(A^{-1}) = \sigma_{-j}(A)$. By applying Theorem 2.3 to A^{-1}, we get the following result, which states that all $\widetilde{\theta}_{-j}^{(k)}$ converge monotonically decreasing to the σ_{-j}.

Theorem 3.4 *(cf. Theorem 2.3).*

(a) For fixed $j \le p$, and $k \ge j$, the $\widetilde{\theta}_{-j}^{(k)}$ converge monotonically (down)to σ_{-j}:

$$\sigma_{-j} \le \widetilde{\theta}_{-j}^{(k+1)} \le \widetilde{\theta}_{-j}^{(k)}.$$

(b) For each $j = 1, \dots, k$, there exist singular values $\sigma_{-j'}$ of A which can be put in one-one correspondence with the singular values $\widetilde{\theta}_{-j}$ of $\widetilde{H}$ in such a way that

$$|\sigma_{-j'}^{-1} - \widetilde{\theta}_{-j}^{-1}| \le \max\,\{\|\widetilde{R}_1(\widetilde{H}^{-1})\|, \|\widetilde{R}_2(\widetilde{H}^{-T})\|\}.$$

Moreover,

$$\sum_{j=1}^{k} (\sigma_{-j'}^{-1} - \widetilde{\theta}_{-j}^{-1})^2 \le \|\widetilde{R}_1(\widetilde{H}^{-1})\|_F^2 + \|\widetilde{R}_2(\widetilde{H}^{-T})\|_F^2.$$

The previous two theorems indicate that the double harmonic approach indeed has favorable properties in the quest for the smallest singular triples. For instance, if the norms of the residual matrices $\widetilde{R}_1(\widetilde{H}^{-1})$ and $\widetilde{R}_2(\widetilde{H}^{-T})$ are small, then, since

$$|\sigma_{-j'}^{-1} - \widetilde{\theta}_{-j}^{-1}| = \frac{|\sigma_{-j'} - \widetilde{\theta}_{-j}|}{\sigma_{-j'}\widetilde{\theta}_{-j}},$$

the smallest singular values of $\widetilde{H}$ must be good approximations to the small singular values of A. So while H tends to approximate A well with respect to the largest singular values, $\widetilde{H}$ tends to approximate A well with respect to the smallest singular values. See also Section 9.2.

4 Refined extraction

The key idea in this section is that the "minimal" left and right singular vector of A minimize $\|A^Tx\|$ and $\|Ay\|$, respectively. When we have search spaces $\mathcal{U}$ and $\mathcal{V}$, it is a natural idea to extract those vectors that minimize the norm of the matrix and its transpose over the search spaces. This approach amounts to computing two SVDs of tall skinny matrices AV and A^TU. It is somewhat similar to the refined extraction in the Ritz method in the eigenvalue problem (see, for instance [11, p. 289]), in the sense that we look for the minimal singular vector in a search space; therefore we choose the name *refined extraction.*

With $\widehat{u} = U\widehat{c}$ and $\widehat{v} = V\widehat{d}$, the following are equivalent (remind that we assume in this section that A is nonsingular):

$$\text{(i)} \quad \begin{cases} \widehat{\theta} = \min_{\substack{u \in \mathcal{U} \\ \|u\| = 1}} \|A^T u\| \\ \widehat{\eta} = \min_{\substack{v \in \mathcal{V} \\ \|v\| = 1}} \|Av\| \end{cases} \quad \text{with } \widehat{u}, \widehat{v}\text{, respectively, as minimizing argument,}$$

$$\text{(ii)} \quad \begin{cases} \left(A^TA\right)^{-1}\widehat{V}\widehat{c} - \widehat{\theta}^{-2}\widehat{V}\widehat{c} \perp \widehat{\mathcal{V}} := A^T\mathcal{U} \\ \left(AA^T\right)^{-1}\widehat{U}\widehat{d} - \widehat{\eta}^{-2}\widehat{U}\widehat{d} \perp \widehat{\mathcal{U}} := A\mathcal{V} \end{cases} \quad \text{where } \widehat{\theta}^{-2} \text{ and } \widehat{\eta}^{-2} \text{ are maximal,}$$

$$\text{(iii)} \quad \begin{cases} U^TAA^TU\widehat{c} = \widehat{\theta}^2\widehat{c} \\ V^TA^TAV\widehat{d} = \widehat{\eta}^2\widehat{d} \end{cases} \quad \text{where } \widehat{\theta}^2 \text{ and } \widehat{\eta}^2 \text{ are minimal.}$$

Here, (i) expresses that $\widehat{u}$ and $\widehat{u}$ minimize the matrix norm over $\mathcal{U}$ and $\mathcal{V}$. Item (ii) gives a derivation in terms of a Galerkin condition on $\left(A^TA\right)^{-1}$ and $\left(AA^T\right)^{-1}$. Item (iii) states that $\widehat{u}$ and $\widehat{v}$ are the "smallest" Ritz vectors of AA^T and A^TA with respect to $\mathcal{U}$ and $\mathcal{V}$, respectively.

We call $\widehat{\theta}$ and $\widehat{\eta}$ *refined singular values*, $\widehat{u}$ and $\widehat{v}$ *refined singular vectors*, and $(\widehat{\theta}, \widehat{\eta}, \widehat{u}, \widehat{v})$ a *refined singular tuple.*

A difference between the refined approach and the harmonic approaches is that the first leads to only one approximate triple instead of k ones; however, this can easily be modified by computing more than just one smallest singular value and corresponding vectors of AV and A^TU, respectively. We label these values in two different ways (cf. (1.1)):

$$\widehat{\theta} = \widehat{\theta}_{-1} \leq \cdots \leq \widehat{\theta}_{-k} \quad \text{and} \quad \widehat{\theta} = \widehat{\theta}_{k} \leq \cdots \leq \widehat{\theta}_{1},$$

and the $\widehat{\eta}$'s similarly. These approximations have the desirable property of monotonic convergence.

Theorem 4.1 *Define $R_U = (I - UU^T)AA^TU$. Then:*

(a) for fixed $j \leq p$, and $k \geq j$, both $\widehat{\theta}_{-j}^{(k)}$ and $\widehat{\eta}_{-j}^{(k)}$ converge monotonically (down)to σ_{-j}:

$$\sigma_{-j} \leq \widehat{\theta}_{-j}^{(k+1)} \leq \widehat{\theta}_{-j}^{(k)} \quad \textit{and} \quad \sigma_{-j} \leq \widehat{\eta}_{-j}^{(k+1)} \leq \widehat{\eta}_{-j}^{(k)};$$

both $\widehat{\theta}_{j}^{(k)}$ and $\widehat{\eta}_{j}^{(k)}$ converge monotonically (up)to σ_j:

$$\widehat{\theta}_{j}^{(k)} \leq \widehat{\theta}_{j}^{(k+1)} \leq \sigma_j \quad \textit{and} \quad \widehat{\eta}_{j}^{(k)} \leq \widehat{\eta}_{j}^{(k+1)} \leq \sigma_j;$$

(b) for each $j = 1, \ldots, k$, there exist singular values $\sigma_{j'}$ and $\sigma_{j''}$ of A which can be put in one-one correspondence with the $\widehat{\theta}_j$ and $\widehat{\eta}_j$ in such a way that

$$|\widehat{\theta}_j^2 - \sigma_{j'}^2| \leq \|R_U\| \quad \textit{and} \quad \sum_{j=1}^{k} (\widehat{\theta}_j^2 - \sigma_{j'}^2)^2 \leq \|R_U\|_F^2,$$

$$|\widehat{\eta}_j^2 - \sigma_{j''}^2| \leq \|R_V\| \quad \textit{and} \quad \sum_{j=1}^{k} (\widehat{\eta}_j^2 - \sigma_{j''}^2)^2 \leq \|R_V\|_F^2;$$

(c) $\sin(\widehat{u}, x_{\min}) \leq \sqrt{1 + \frac{\gamma_U^2}{\delta_U^2}}\, \sin(\mathcal{U}, x_{\min})$ *and* $\sin(\widehat{v}, y_{\min}) \leq \sqrt{1 + \frac{\gamma_V^2}{\delta_V^2}}\, \sin(\mathcal{V}, y_{\min})$,
where

$$\begin{aligned} \gamma_U &= \|P_{\mathcal{U}}(AA^T - \sigma_{\min}^2 I)(I - P_{\mathcal{U}})\|, & \gamma_V &= \|P_{\mathcal{V}}(A^TA - \sigma_{\min}^2 I)(I - P_{\mathcal{V}})\|, \\ \delta_U &= |\widehat{\theta}_{-2}^2 - \sigma_{\min}^2|, & \delta_V &= |\widehat{\eta}_{-2}^2 - \sigma_{\min}^2|. \end{aligned}$$

Proof: Follows from characterizations (i) or (iii), using the same techniques as in Theorem 3.2. □

Advantages of this refined approach are good asymptotic vector extraction ($\sigma_{\min}^2$ is an exterior eigenvalue of A^TA), and the fact that we have upper bounds for $\sigma_{\min}$. On the other hand, the left and right singular vector are approximated completely independently. It may thus happen that $u \approx x$ and $v \approx y$ are approximate vectors to singular vectors corresponding to different singular values. In this case the Rayleigh quotient of the vectors (see Section 5) is meaningless as approximate singular value.

Note that we can also formulate a result similar to part (c) of Theorem 4.1 for the largest singular vectors. This also shows that the refined approach is equally useful for the largest singular triples as for the smallest ones: just take the maximum instead of the minimum in characterization (i) and (iii).

The next theorem gives some idea how fast the refined approach converges to the minimal singular value. In particular, it shows that, since $\angle(\mathcal{U}, x) \to 0$ and $\angle(\mathcal{V}, y) \to 0$ as the search spaces expand, convergence is guaranteed.

Theorem 4.2 *(cf. [11, p. 290])*

$$\|A\widehat{v}\| \leq \frac{\sigma_{\min} + \sin(\mathcal{V}, y_{\min})\sigma_{\max}}{\sqrt{1 - \sin^2(\mathcal{V}, y_{\min})}}, \quad \textit{and} \quad \|A^T\widehat{u}\| \leq \frac{\sigma_{\min} + \sin(\mathcal{U}, x_{\min})\sigma_{\max}}{\sqrt{1 - \sin^2(\mathcal{U}, x_{\min})}}.$$

Proof: We only prove the first statement, the proof of the second one being similar. Decompose $y_{\min} = c_V y_V + s_V f_V$, where $y_V := VV^* y_{\min}/\|VV^* y_{\min}\|$ is the orthogonal projection of $y_{\min}$ onto $\mathcal{V}$, $c_V = \cos(\mathcal{V}, y_{\min})$, and $s_V = \sin(\mathcal{V}, y_{\min})$. Since $Ay_V = (\sigma_{\min} x_{\min} - s_V A f_V)/c_V$, we have by definition of a refined singular vector

$$\|A\widehat{v}\| \le \|Ay_V\| \le (\sigma + s_V\|A\|)/c_V.$$

□

The next theorem gives a justification of the new methods: they retrieve singular triples that are exactly present in $\mathcal{U}$ and $\mathcal{V}$.

Theorem 4.3 *Let (σ, x, y) be a singular triple of A with $x = Uc$ and $y = Vd$. Then (σ, σ, Uc, Vd) is both a harmonic and refined singular tuple.*

Proof: This can be verified by calculating the left and right hand sides of (iii) of Section 3.1, (v) of Section 3.2, and (iii) of Section 4. □

The standard extraction method in principle also finds singular triples that are exactly present in the search spaces. However, it may have difficulties selecting them, see Example 7.1.

In the following table we summarize the properties of convergence of the different extraction methods (Theorems 2.3, 3.2, 3.4, 4.1). A "+" stands for monotonic convergence. A "–" means that we do not have monotonic convergence; as a result the convergence can in practice be irregular and very slow.

TABLE 1: Properties of monotonic convergence of the extraction methods to $\sigma_{\max}$ and $\sigma_{\min}$.

extraction	$\sigma_{\max}$	$\sigma_{\min}$
standard	+	–
$\mathcal{U}$-, $\mathcal{V}$-harmonic	+	+
double-harmonic	–	+
refined	+	+

The table suggests that to find the smallest singular triples, one can use all methods except the standard extraction, while for the largest singular values one can use all methods except the double-harmonic extraction. This appears to be a good rule of thumb indeed, see also Sections 10 and 11.

5 Rayleigh quotient for the singular value problem

In the harmonic Ritz approach for the eigenproblem, the harmonic Ritz value may be a bad approximation to the eigenvalue (it can even be ∞), while the harmonic Ritz vector may be of good quality; see, e.g., [10]. Therefore, it is advisable to discard the harmonic Ritz value and, instead, approximate the eigenvalue by the Rayleigh quotient of the harmonic Ritz vector.

In our case, we also encounter the situation that while the harmonic vectors may be good, the harmonic value can be bad, see also Example 7.1. Therefore, we propose to take the Rayleigh quotient (in the sense of the singular value problem) of the left and right approximate singular vector as an approximate singular vector.

Definition 5.1 *For $u, v \neq 0$, we define the* Rayleigh quotient for the singular value problem *of u and v by*

$$\rho(u, v) := \frac{u^T A v}{\|u\| \, \|v\|}.$$

This Rayleigh quotient has the following attractive property.

Proposition 5.2 *Let $u = x + \delta f$ and $v = y + \varepsilon g$ be approximate singular vectors corresponding to the singular value σ, where $f \perp x$ and $g \perp y$ are vectors of unit length. Then $|\rho(u, v) - \sigma| = \mathcal{O}((\delta + \varepsilon)^2)$.*

Proof: This follows directly from $\rho(u, v) = (\sigma + \delta\varepsilon f^T Ag)/\sqrt{(1 + \delta^2)(1 + \varepsilon^2)}$. □

In other words, if u and v are first order approximations to the singular vectors, their Rayleigh quotient is a second order approximation to the singular value.

One may check that we have the expressions as in Table 2 for the Rayleigh quotient of approximate vectors in the standard and harmonic approaches.

TABLE 2: Rayleigh quotients for the standard and harmonic extraction methods.

approach	Rayleigh quotient of approximate vectors
standard	$\theta = \eta$
$\mathcal{V}$-harmonic	$\widetilde{\eta}$
double-harmonic	$\frac{\|AVd\|^2}{\widetilde{\theta}} = \frac{\|A^TUc\|^2}{\widetilde{\eta}} = \frac{\|AVd\|\,\|A^TUc\|}{\sqrt{\widetilde{\theta}\widetilde{\eta}}}$

So for the standard and harmonic methods we can obtain the Rayleigh quotient of the approximate singular vectors at little (double-harmonic approach) or no (standard, $\mathcal{U}$-, and $\mathcal{V}$-harmonic approaches) additional cost. Of course, we can also take the Rayleigh quotient in the refined extraction process. For complex matrices, we should scale u and v in such a way that their Rayleigh quotient is real and nonnegative.

6 Interior singular values

The extraction processes introduced in Section 3 and 4 are tailored for the smallest (and, with exception of the double-harmonic approach, also useful for the largest) singular triples. Now we study the situation where we are interested in interior singular values (and corresponding vectors) near a target $\tau \geq 0$. We present two methods that can be seen as generalizations of the double-harmonic approach (Section 3.2) and the refined approach (Section 4). We will see that for $\tau = 0$, the methods in this section deduce to those of Sections 3 and 4.

For a Hermitian matrix B and a search space $\mathcal{W}$, it is well known that instead of standard Raleigh–Ritz, better results may be expected from the harmonic Ritz approach (see, e.g., [11, p. 292])

$$W^T(B - \tau I)^2 Wc = (\widetilde{\theta} - \tau) W^T (B - \tau I) Wc,$$

where we are interested in the $\widetilde{\theta}$ closest to some (interior) target τ. When we take the augmented matrix (2.4) for B, and "split up" $\mathcal{W}$ into a left and right space $\mathcal{U}$ and $\mathcal{V}$—i.e., we take $W = \begin{bmatrix} U & 0 \\ 0 & V \end{bmatrix}$—then we get the *harmonic extraction for target* $\tau \geq 0$ for the singular value problem: find the eigenpair(s) of the generalized symmetric eigenvalue problem

$$\begin{bmatrix} U^TAA^TU + \tau^2 I_m & -2\tau H \\ -2\tau H^T & V^TA^TAV + \tau^2 I_n \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix} = (\widetilde{\theta} - \tau) \begin{bmatrix} -\tau I & H \\ H^T & -\tau I \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix},$$

for which $\widetilde{\theta}$ is closest to τ. Here the matrix on the left hand side is positive semidefinite. Restated, the problem is to find the smallest eigenpairs of

$$\begin{bmatrix} U^TAA^TU & -\tau H \\ -\tau H^T & V^TA^TAV \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix} = \widetilde{\theta} \begin{bmatrix} -\tau I & H \\ H^T & -\tau I \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix}. \tag{6.1}$$

One may check that for $\tau \to \infty$ we get the standard Galerkin approach of Section 2, while $\tau = 0$ gives the double-harmonic approach of Section 3.2.

A *refined method* for the approximation of interior singular triples is the minimization over $\widehat{c}, \widehat{d} \in \mathbb{R}^k$, $\|\widehat{c}\| = \|\widehat{d}\| = 1$ of

$$\left\| \begin{bmatrix} -\tau I_m & A \\ A^T & -\tau I_n \end{bmatrix} \begin{bmatrix} U\widehat{c} \\ V\widehat{d} \end{bmatrix} \right\| = \left\| \begin{bmatrix} -\tau U & AV \\ A^T U & -\tau V \end{bmatrix} \begin{bmatrix} \widehat{c} \\ \widehat{d} \end{bmatrix} \right\|, \tag{6.2}$$

which amounts to the SVD of a tall skinny $(m+n) \times 2k$ matrix. For $\tau = 0$, we get back the refined approach of Section 4.

We note that both approaches may fail to produce approximations to both the left and right singular vectors, since it is not guaranteed that the computed $\widetilde{c}, \widetilde{d}, \widehat{c}$, and $\widehat{d}$ are nonzero. Suppose, for example, that $\widetilde{c} = 0$ and $\widetilde{d} \neq 0$. Then a possibility is to solve $\widetilde{c}$ from $H^T\widetilde{c} = \widetilde{d}$ (as in the $\mathcal{V}$-harmonic method). See the next section for the case of a nonsquare or singular H.

7 Nonsquare or singular matrices

Most characterizations of the extraction processes given above use A^{-1} and A^{-T}. Moreover, in the $\mathcal{U}$- and $\mathcal{V}$-harmonic approach, we have to solve a system with H or H^T. We now show that no difficulties arise from a nonsquare or singular A, or from a singular H. We first show that such a nonsquare or singular A gives no problems for the double-harmonic approach, using characterizations (iv) or (v) of Section 3.2. Let

$$U' = (I - P_{\mathcal{N}(A^T)})U \quad \text{and} \quad V' = (I - P_{\mathcal{N}(A)})V. \tag{7.1}$$

Notice that $A' := (I - P_{\mathcal{N}(A^T)})A(I - P_{\mathcal{N}(A)}) = A$. By using characterization (iv) of the double-harmonic method we obtain that

$$\begin{cases} AV\widetilde{d} - \widetilde{\eta}U\widetilde{c} & \perp \quad A\mathcal{V}, \\ A^TU\widetilde{c} - \widetilde{\theta}V\widetilde{d} & \perp \quad A^T\mathcal{U}, \end{cases}$$

if and only if

$$\begin{cases} A'V'\widetilde{d} - \widetilde{\eta}U'\widetilde{c} & \perp \quad A'\mathcal{V}', \\ (A')^TU'\widetilde{c} - \widetilde{\theta}V'\widetilde{d} & \perp \quad (A')^T\mathcal{U}'. \end{cases}$$

So for the double-harmonic approach we may assume without loss of generality that A is nonsingular. Another way to see this is via item (v) in Section 3.2: since we assumed that AV and A^TU are of full rank, we see that nonsquare or singular A form no difficulty.

Though characterization (ii) of the refined approach (see Section 4) is suitable only for nonsingular matrices, it can be seen from the items (i) and (iii) that nonsquare or singular A give no problems.

In the $\mathcal{U}$- and $\mathcal{V}$-harmonic method, we have to solve a system of the form $H\widetilde{d} = \widetilde{c}$ or $H^T\widetilde{c} = \widetilde{d}$. Independent of the properties of A, the matrix H may be singular. In this circumstance, the systems involving H or H^T do not have a unique solution. This is precisely the situation where infinite harmonic Ritz values may occur in the $\mathcal{U}$- and $\mathcal{V}$-harmonic method, see (iii) in Section 3.1.

A possibility is then to solve the small system in a least squares sense, that is, take the pseudoinverse: $\widetilde{d} = H^+\widetilde{c}$ or $\widetilde{c} = (H^T)^+\widetilde{d}$. However, these vectors may be zero, in that case the $\mathcal{U}$- and $\mathcal{V}$-harmonic approaches fail to give an approximation to one of the two singular vectors. In numerical experiments, however, we have not encountered this situation.

We mention that for nonsquare or singular matrices, the theorems in Sections 3 and 4 still hold, as far as the nonzero singular values are concerned.

The following examples illustrate some properties of the extraction methods.

Example 7.1 Take $A = \text{diag}(1, 2, 3)$, and suppose that we try to find an approximation to the smallest singular triple of A from the search spaces $\mathcal{U} = \text{span}(e_1, e_3)$, and $\mathcal{V} = \text{span}(e_1, e_2)$. First we consider the standard extraction: the singular triples of $H = \text{diag}(1, 0)$ lead to approximate singular triples $(\theta, u, v) = (0, e_3, e_2)$ and $(1, e_1, e_1)$. So, although the standard extraction finds the smallest triple of A, it is not safe to take the smallest $(0, e_3, e_2)$ as an approximation to the smallest triple of A. Both the $\mathcal{U}$-harmonic and $\mathcal{V}$-harmonic approach have to deal with a singular H. As discussed, we could enlarge $\mathcal{U}$ and $\mathcal{V}$ (by for instance the residual vectors r_1 and r_2) to avoid the singularity, or take the pseudoinverse for the systems involving H or H^T. The latter option gives $(\theta, \eta, u, v) = (1, 1, e_1, e_1)$ for both the $\mathcal{U}$- and $\mathcal{V}$-harmonic approach. The double-harmonic approach finds the tuples $(\theta, \eta, u, v) = (1, 1, e_1, e_1)$ and $(\infty, \infty, e_3, e_2)$. Here it is safe to take the smallest tuple. (We can get rid of the infinite harmonic value by taking the Rayleigh quotient of e_3 and e_2: this gives the approximate value 0. This is an example where it can be seen that the Rayleigh quotients often make more sense than the harmonic values.) Finally, the refined approach gives the correct solution $(\theta, u, v) = (1, e_1, e_1)$. The conclusion is that in this example, the standard approach is the only one having difficulties to determine the "smallest" singular vectors that are present in the search spaces.

Example 7.2 Let A be as in Example 7.1, and let $\mathcal{U} = \mathcal{V} = \mathrm{span}(e_2, (e_1 + e_3)/\sqrt{2})$. Then $H = \mathrm{diag}(2, 2)$ has a double singular value and the standard extraction does not know which approximate vectors to take. The other methods do know how to decide: for target $\tau = 0$, all three harmonic approaches and the refined approach take $(u, v) = (e_2, e_2)$ as approximate vectors with approximate singular value $\theta = \eta = 2$. Also for target $\tau = 2$, the double-harmonic and refined method (see Section 6) yield $(2, e_2, e_2)$ as approximate triple.

8 Lanczos bidiagonalization

We now study the different extraction methods in the context of Lanczos bidiagonalization. After k steps of Lanczos bidiagonalization with starting vector v_1 we have the relations [2, p. 495]:

$$\left\{ \begin{array}{lcl} AV_k & = & U_k B_{k,k}, \\ A^T U_k & = & V_{k+1} B_{k+1,k}^T, \end{array} \right.$$

where $B_{k,k}$ and $B_{k+1,k}^T$ are a $k \times k$ upper, and a $(k+1) \times k$ lower bidiagonal matrix, respectively. This implies that the standard extraction process (see Section 2) takes the singular triples of

$$U_k^T A V_k = B_{k,k}$$

as approximations to the singular triples of A. The $\mathcal{V}$-harmonic method (see Section 3.2, characterization (iii)) reduces to

$$\left\{ \begin{array}{lcl} B_{k,k}^T B_{k,k} \widetilde{d} & = & \widetilde{\theta} B_{k,k}^T \widetilde{c}, \\ B_{k,k}^T \widetilde{c} & = & \widetilde{\eta} \widetilde{d}. \end{array} \right.$$

Here $B_{k,k}$ is nonsingular due to the assumption that AV and $A^T U$ are of full rank. Hence, one may check that the $\mathcal{V}$-harmonic approach also takes the singular triples of $B_{k,k}$ as approximate singular triples. Similar remarks can be made for the $\mathcal{U}$-harmonic method.

Item (v) of the double-harmonic approach (Section 4) deduces to

$$\left\{ \begin{array}{lcl} \widetilde{\eta} B_{k,k} \widetilde{d} & = & B_{k,k+1} B_{k+1,k}^T \widetilde{c}, \\ \widetilde{\theta} B_{k,k}^T \widetilde{c} & = & B_{k,k}^T B_{k,k} \widetilde{d}. \end{array} \right.$$

Again we may assume that $B_{k,k}$ is nonsingular. Then $B_{k,k}\widetilde{d} = \widetilde{\theta}\widetilde{c}$ and $B_{k,k+1} B_{k+1,k}^T \widetilde{c} = (\widetilde{\theta}\widetilde{\eta})\widetilde{c}$. Therefore, $\widetilde{c}$ is a left singular vector of $B_{k,k+1}$. When $B_{k,k+1}$ and $B_{k,k}$ do not differ much (as will be true when $\mathcal{U}$ and $\mathcal{V}$ are nearly invariant singular subspaces), $\widetilde{c}$ is close to a left singular vector of $B_{k,k}$, and hence $\widetilde{d}$ is close to a right left singular vector of $B_{k,k}$.

The refined approach considers

$$\left\{ \begin{array}{lcl} \min_{\substack{c \in \mathbb{R}^k \\ \|c\| = 1}} \|A^T U_k c\| & = & \min_{\substack{c \in \mathbb{R}^k \\ \|c\| = 1}} \|B_{k+1,k}^T c\|, \\ \min_{\substack{d \in \mathbb{R}^k \\ \|d\| = 1}} \|AV_k d\| & = & \min_{\substack{d \in \mathbb{R}^k \\ \|d\| = 1}} \|B_{k,k} d\|, \end{array} \right.$$

to which the smallest (left and right, respectively) singular vectors of $B_{k,k+1}$ and $B_{k,k}$ are the solutions.

We conclude that in Lanczos bidiagonalization (a two-sided subspace method where we choose a specific subspace expansion), all extraction processes do essentially the same: approximating singular triples of A by those of $B_{k,k}$ or $B_{k,k+1}$. Since the new extraction processes in this paper are often good for the minimal singular triple, the standard extraction is also fine in this case. This may be seen as an explanation why Lanczos bidiagonalization is, besides for the largest singular values, also successful for the approximation of the smallest singular triples. For other two-sided subspace methods, such as JDSVD, the extraction processes may differ much, see also the numerical experiments.

Finally, for completeness, we give the extraction processes of Section 6 for interior singular values in the case of Lanczos bidiagonalization. The double harmonic approach attempts to determine (the "smallest") eigenpairs of

$$\begin{bmatrix} B_{k,k+1}B_{k+1,k}^T & -\tau B_{k,k} \\ -\tau B_{k,k} & B_{k,k}^T B_{k,k} \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix} = \widetilde{\theta} \begin{bmatrix} -\tau I_k & B_{k,k} \\ B_{k,k}^T & -\tau I_k \end{bmatrix} \begin{bmatrix} \widetilde{c} \\ \widetilde{d} \end{bmatrix}.$$

The refined extraction for target $\tau \geq 0$ considers

$$\min_{\substack{c,\, d \in \mathbb{R}^k \\ \|c\| = \|d\| = 1}} \left\| \begin{bmatrix} AV_k d - \tau U_k c \\ AU_k^T c - \tau V_k d \end{bmatrix} \right\| = \min_{\substack{c,\, d \in \mathbb{R}^k \\ \|c\| = \|d\| = 1}} \left\| \begin{bmatrix} B_{k,k} d - \tau c \\ B_{k+1,k}^T c - \tau I_{k+1,k} d \end{bmatrix} \right\|,$$

where $I_{k+1,k}$ is the identity with an extra $(k+1)$th zero row. We conclude that in Lanczos bidiagonalization, the extraction methods for the smallest singular triples are all more or less equivalent, but the extraction methods for interior singular triples differ.

9 Applications

In this section, we study two applications where the new extraction processes can be used: the least squares problem and the truncated SVD.

9.1 The least squares problem

The *least squares problem*

$$\min_v \|b - Av\|,$$

with minimal norm solution

$$A^+ b = \sum_{\sigma_j \neq 0} \sigma_j^{-1} (x_j^T b) y_j, \tag{9.1}$$

has often been successfully attacked by methods based on Lanczos bidiagonalization. For example, LSQR [6] chooses $u_1 = b/\beta$, where $\beta = \|b\|$, and forms U_k and V_k such that $A^T U_k = V_k B_{k,k}^T$ and $AV_k = U_{k+1} B_{k+1,k}$. Suppose we look for a solution $v \in \mathcal{V}_k$, say $v = V_k d$. Then

$$\|b - Av\| = \|\beta U_{k+1} \underline{e}_1 - AV_k d\| = \|\beta U_{k+1} \underline{e}_1 - U_{k+1} B_{k+1,k} d\| = \|\beta \underline{e}_1 - B_{k+1,k} d\|,$$

where $\underline{e}_1$ is the first unit vector in $\mathbb{R}^{k+1}$. Now LSQR takes the approximation

$$v = \beta V_k B_{k+1,k}^+ \underline{e}_1. \tag{9.2}$$

For other two-sided SVD methods, such as JDSVD [3], we can use a similar idea, although in general, we will not have short recurrences as in LSQR. Let $\widetilde{\mathcal{U}}$ a test space, yet to be determined, and let $[\, \widetilde{U} \;\; \widetilde{U}_\perp \,]$ form an orthogonal basis. For ease omit the index k. Then, again with $\beta u_1 = b$ and $v = Vd$ we get

$$\|b - Av\| = \left\| \begin{bmatrix} \widetilde{U}^T (\beta U e_1 - AVd) \\ \widetilde{U}_\perp^T (\beta U e_1 - AVd) \end{bmatrix} \right\|.$$

Now we neglect the second part $\|\widetilde{U}_\perp^T(\beta U e_1 - AVd)\|$, this is equivalent to requiring

$$b - Av \perp \widetilde{\mathcal{U}}.$$

The minimal norm solution to $\min_d \|\beta \widetilde{U}^T U e_1 - \widetilde{U}^T AVd\|$ leads to

$$\widetilde{v} = \beta V(\widetilde{U}^T AV)^+ \widetilde{U}^T U e_1.$$

To get a good solution, $V(\widetilde{U}^T AV)^+\widetilde{U}^T$ should be a good approximation to A^+. Since a pseudoinverse is mainly determined by its smallest singular values and vectors, we realize that the extraction of those small singular triples from the search spaces $\mathcal{U}$ and $\mathcal{V}$ is crucial. We have already seen that the choice $\widetilde{U} = U$ (leading to $v = \beta V H^+ e_1$) is often not satisfactory for the smallest singular triples. The choice $\widetilde{U} = AV$, as in the $\mathcal{V}$-harmonic and double-harmonic approaches, is more promising. With this choice, we have $d = \beta(V^TA^TAV)^{-1}V^TA^TUe_1 = \beta(AV)^+Ue_1$, so d is the least squares solution to $\min_d \|\beta U e_1 - AVd\|$. Then $v = \beta V(AV)^+Ue_1$. For Lanczos bidiagonalization, this gives $v = \beta V_k B_{k,k}^+ e_1$, which resembles the LSQR solution (9.2).

Since AV and $H^T = V^TA^TU$ are already computed in the $\mathcal{V}$-harmonic and double-harmonic methods, these approaches can therefore also be useful for least square problems: they may give an approximate solution of the least squares problem at low additional costs during the process.

As already mentioned, methods such as JDSVD may need restarts and deflation from time to time. When we would like to use these methods for the least squares problem, special care has to be taken when the maximum dimension of the search space has been reached (restart), or when a singular triple has been found (deflation). With restarts, we have to ensure that $b = \beta u_1 \in \mathcal{U}$. Therefore, we restart with the span of the best (say) $l-1$ left vectors in $\mathcal{U}$, together with u_1 as the new left search space $\mathcal{U}$ (what "the best" means, depends on the extraction method). For the new right search space $\mathcal{V}$, we take the span of the best $l-1$ right vectors in $\mathcal{V}$, together with $V(AV)^+Ue_1$, the minimal norm solution to $\min_{v\in\mathcal{V}} \|b - Av\|$. Since we include the best approximation so far to the least squares problem in the new search space $\mathcal{V}$, we get monotonic convergence for the least squares solution, that is,

$$\|b - Av_{k+1}\| \le \|b - Av_k\|.$$

This is trivial when we expand the search space $\mathcal{V}$, but by restarting in the this way, it is also valid at restarts.

Now consider deflation. Suppose we have detected a singular triple (σ, x, y), where $\sigma \neq 0$. By decomposing

$$A = (I - xx^T)A(I - yy^T) + \sigma xy^T \quad \text{and} \quad b = (I - xx^T)b + (xx^T)b,$$

we get

$$\min_v \|b - Av\|^2 = \min_v \|(I - xx^T)(b - A(I - yy^T)v)\|^2 + |x^Tb - \sigma y^T v|^2.$$

So we may conclude that the (minimal norm) solution v has a component $\sigma^{-1}x^Tb$ in the direction of y. This may also be seen from (9.1). With $\widetilde{v} = (I - yy^T)v$, we are left with a *deflated least squares problem*

$$\min_{\widetilde{v}\perp y} \|(I - xx^T)(b - A\widetilde{v})\|.$$

Hence, if a triple has been found, we restart with the best $l-1$ left vectors in $\mathcal{U}$ and $(I - xx^T)b$ as the new $\mathcal{U}$. (In this case, "the best" means the best vectors to find the next singular triple.) For the new $\mathcal{V}$, we take the best $l-1$ right vectors in $\mathcal{V}$, together with $(I - yy^T)v$, where, as before, $v = V(AV)^+Ue_1$ is the current best approximation to the least squares problem. Of course, this procedure can be repeated when more singular triples are found. See Section 10 for numerical experiments.

9.2 The truncated SVD

We may also use the standard and double-harmonic methods to give an approximation to the *truncated SVD* of A. The solution to

$$\min_{\mathrm{rank}(B)=k} \|A - B\|$$

is given by $B = A_k := \sum_{j=1}^{k} \sigma_j x_j y_j^T$ (unique if σ_k is simple). Analogously, the solution to

$$\min_{\mathrm{rank}(B)=k} \|A^+ - B\|$$

is given by $B = A_{-k} := \sum_{j=1}^{k} \sigma_{-j} x_{-j} y_{-j}^T$, where the sum is over nonzero singular values of A (unique if σ_{-k} is simple).

In view of the discussed extraction processes, we expect that when $\mathcal{U}_k$ and $\mathcal{V}_k$ are search spaces for the largest triples,

$$P_{\mathcal{U}_k} A P_{\mathcal{V}_k} = U_k U_k^T A V_k V_k^T = U_k H_k V_k^T$$

may be a reasonable approximation to A_k. In [9], the authors use the Lanczos bidiagonalization to approximate A_k by $U_k B_{k,k} V_k^T$. Although the Lanczos process has in principle short recurrences, (some) reorthogonalization of the vectors appears to be necessary [9].

Now consider the situation that $\mathcal{U}_k$ and $\mathcal{V}_k$ are search spaces for the smallest triples. Then from (3.1), we know that $\widetilde{H} = G_U H^{-1} G_V^T$ can be viewed as a projected approximation to A, which attempts to approximate the smallest portion of the singular spectrum well. Hence, as a reasonable approximation to A_{-k}, we may take

$$U_k \widetilde{H}_k V_k^T = U_k G_U H_k^{-1} G_V^T V_k^T.$$

See Section 10 for numerical experiments.

10 Numerical experiments

The following experiments were carried out in MATLAB. For all experiments where the random generator is used, we first put the "seed" to 0 by the command "`rand('seed',0)`", so that our results are reproducible.

Experiment 10.1 Up to rounding errors, it is not a loss of generality to consider diagonal matrices [3]. For the first example we take $A = \mathrm{diag}(1:100)$. We build up four-dimensional search spaces $\mathcal{U}$ and $\mathcal{V}$ to find the minimal singular triple. The first basis vector of the left search space $\mathcal{U}$ is $u_1 = e_1 + \varepsilon_U w_U$, where w_U is a random vector of unit length. We complement $\mathcal{U}$ by three random vectors. The right search space $\mathcal{V}$ is formed in a similar way: take $v_1 = e_1 + \varepsilon_V w_V$, and add three random vectors.

We consider two cases. For the first we take $\varepsilon_U = 10^{-3}$ and $\varepsilon_V = 10^{-1}$. This means that the left search space is good, while the right search space is not very accurate. It appears that

$$\angle(\mathcal{U}, e_1) \approx 3.5 \cdot 10^{-3} \quad \text{and} \quad \angle(\mathcal{V}, e_1) \approx 3.0 \cdot 10^{-1},$$

these angles also give the best possible approximate vectors in $\mathcal{U}$ and $\mathcal{V}$. For the second case we take $\varepsilon_U = \varepsilon_V = 10^{-3}$, in other words: left and right search spaces of good quality. In this case

$$\angle(\mathcal{U}, e_1) \approx 3.5 \cdot 10^{-3} \quad \text{and} \quad \angle(\mathcal{V}, e_1) \approx 3.2 \cdot 10^{-3}.$$

Table 3 gives the results of the 5 different extraction processes. We display the error in the approximate vectors u and v, and the error in the approximate value ρ, the Rayleigh quotient of u and v.

TABLE 3: The 5 different extraction processes for the minimal singular triple of $A = \mathrm{diag}(1:100)$. Column 2 to 4 are for $\varepsilon_U = 10^{-3}$ and $\varepsilon_V = 10^{-1}$, while column 5 to 7 represent $\varepsilon_U = \varepsilon_V = 10^{-3}$.

	$\varepsilon_U = 10^{-3}, \varepsilon_V = 10^{-1}$			$\varepsilon_U = 10^{-3}, \varepsilon_V = 10^{-3}$		
method	$\angle(u, e_1)$	$\angle(v, e_1)$	$\lvert\sigma_{\min} - \rho\rvert$	$\angle(u, e_1)$	$\angle(v, e_1)$	$\lvert\sigma_{\min} - \rho\rvert$
standard	$1.9e-1$	$8.0e-1$	$2.8e-1$	$4.0e-3$	$1.1e-2$	$2.5e-5$
$\mathcal{U}$-harmonic	$3.5e-3$	$7.6e-1$	$2.7e-1$	$3.5e-3$	$1.0e-2$	$2.6e-5$
$\mathcal{V}$-harmonic	$1.1e-2$	$3.1e-1$	$2.3e-2$	$4.4e-3$	$3.3e-3$	$6.8e-5$
double-harmonic	$3.6e-3$	$3.1e-1$	$4.4e-2$	$3.5e-3$	$3.3e-3$	$2.3e-7$
refined	$3.5e-3$	$3.1e-1$	$4.8e-2$	$3.5e-3$	$3.3e-3$	$8.1e-7$

Almost all errors of the new methods are smaller than those of the standard approach. Moreover, the new approaches are almost optimal in most cases, by which we mean that the extracted vectors are almost the best possible ones, given the search spaces. In view of the factors $\sqrt{1+2\frac{\widetilde{\gamma}^2}{\delta^2}} \approx 1.4$ for both the $\mathcal{U}$- and $\mathcal{V}$-harmonic approach (see Theorem 3.2), and $\sqrt{1+\frac{\gamma^2}{\delta^2}} \approx 11$ for the standard method (see Theorem 2.5), we already could suspect that the harmonic approaches would be superior. We mention that in the $\mathcal{V}$-harmonic method, the approximate left vector $\widetilde{u} = UH^{-T}\widetilde{d}$ is indeed much better that $A\widetilde{v}$ (cf. Lemma 3.1 and discussion). Similar remarks hold for the $\mathcal{U}$-harmonic method.

Experiment 10.2 For the remaining experiments, we use JDSVD, the Jacobi–Davidson type method for the singular value problem [3]. Unless mentioned otherwise, we set the following parameters for JDSVD. The dimension of the search spaces is at most 20, after which we restart with the best 10 vectors (remember that the meaning of "the best" depends on the extraction method). Recall from [3] that the JDSVD correction equation is of the form

$$\begin{bmatrix} I_m - \frac{u\widetilde{u}^T}{\widetilde{u}^Tu} & 0 \\ 0 & I_n - \frac{v\widetilde{v}^T}{\widetilde{v}^Tv} \end{bmatrix} \begin{bmatrix} -\zeta I_m & A \\ A^T & -\zeta I_n \end{bmatrix} \begin{bmatrix} I_m - uu^T & 0 \\ 0 & I_n - vv^T \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = -\begin{bmatrix} Av - \rho u \\ A^Tu - \rho v \end{bmatrix}.$$

Here u and v are the current approximate vectors, $\widetilde{u}$ and $\widetilde{v}$ are the test vectors (depending on the extraction method; for example, in the $\mathcal{V}$-harmonic approach we have $\widetilde{u} = Av$, $\widetilde{v} = v$), and we solve for orthogonal updates $s \perp u$ and $t \perp v$. On the right hand side of the correction equation, we take the Rayleigh quotient $\rho(u, v)$ as approximate singular value. In the beginning, the shift ζ in the left hand side of the correction equation is taken to be the target τ. The reason for this is that the Rayleigh quotient ρ is not likely to be accurate at that stage. So initially, the method behaves as an inexact inverse iteration with target τ. When we are close to convergence, in the sense that $\|r\| < 0.01$, we take the shift ζ equal to the Rayleigh quotient ρ. Then the methods works as an inexact Rayleigh quotient iteration. We solve the correction equation (the so-called "inner iteration") approximately, by 10 steps of unpreconditioned GMRES. (In practice, it advisable to use a preconditioner if one is available.) We continue with the method until $\|r\| < 1e-6$.

We take the 1850×712 matrix `well1850` from the Matrix Market [4], with $\sigma_{\min} \approx 1.6\cdot10^{-2}$. We perform 70 steps of JDSVD with double-harmonic extraction and target $\tau = 0$ to find the smallest singular triple. The starting vectors u_1 and v_1 are the vector of all ones. For the (20-dimensional) search spaces after 70 steps of the method, it appears that

$$\angle(\mathcal{U}, x) = 4.9 \cdot 10^{-3} \quad \text{and} \quad \angle(\mathcal{V}, y) = 3.8 \cdot 10^{-3}.$$

With these search spaces, we test the different extraction processes, see the first four columns of Table 4.

We see that the standard approach fails completely, apparently due to the selection of the wrong triple (a situation similar to that in Example 7.1). The new extraction methods perform reasonably well.

In the last three columns of Table 4, we give the number of outer iterations it takes before JDSVD with the specific extraction method has detected the smallest singular triple. For column

TABLE 4: The 5 different extraction processes for the minimal singular triple of `well1850`. Column 2 to 4 are the extraction results for the 20-dimensional search spaces produced after 70 steps of JDSVD. The last three columns give the number of outer steps needed for the computation of the smallest singular triple, with 10, 20, and 30 steps of GMRES to solve the correction equations, respectively.

method	$\angle(u, e_1)$	$\angle(v, e_1)$	$\lvert\sigma_{\min} - \rho\rvert$	GMRES_{10}	GMRES_{20}	GMRES_{30}
standard	$1.6e+0$	$1.6e+0$	$1.0e-2$	> 200	67	41
$\mathcal{U}$-harmonic	$1.0e-2$	$8.0e-3$	$8.7e-7$	155	72	41
$\mathcal{V}$-harmonic	$1.2e-2$	$6.7e-3$	$2.1e-7$	171	67	41
double-harmonic	$1.2e-2$	$1.1e-2$	$1.3e-6$	97	48	34
refined	$1.0e-2$	$6.7e-3$	$8.9e-7$	97	51	36

5, we use 10 steps of GMRES, for column 6 we perform 20 steps, and for the last column 30 steps to solve the correction equations. It appears that the less accurate we solve the correction equation (which is, of course, cheaper), the more advantageous the new extraction methods are, especially the double-harmonic and refined approach.

Experiment 10.3 Next, for $A = \mathrm{diag}(1 : 100)$, we perform 100 (outer) steps of JDSVD with each of the extraction processes, and count how many singular triples we find. See Table 5. For the second and third column we take target $\tau = 0$ (i.e., we look for the smallest singular triples), for the fourth and fifth column $\tau = 50.1$ (interior triples closest to 50.1), and for the last two columns $\tau = 105$ and $\tau = \infty$ (largest triples).

TABLE 5: The number of singular triples found within 100 outer iterations by the 5 different extraction processes for $A = \mathrm{diag}(1 : 100)$. Columns 2 and 3 are for $\tau = 0$ (minimal singular triples), respectively with no fix and fix $= 0.01$. Column 4 and 5 represent $\tau = 50.1$ (interior singular triples), with 10 and 20 steps of GMRES, respectively, while for the last two columns $\tau = 105$ and $\tau = \infty$ (largest singular triples).

	$\tau = 0$		$\tau = 50.1$		$\tau = 105$	$\tau = \infty$
method	no fix	fix	GMRES_{10}	GMRES_{20}		
standard	1	2	–	3	20	20
$\mathcal{U}$-harmonic	1	4	–	3	17	20
$\mathcal{V}$-harmonic	3	3	–	3	17	20
double-harmonic	3	7	1	4	20	–
refined	5	7	2	8	19	21

The "no fix" in the second column means that we take the shift ζ in the left hand side of the correction equation equal to the Rayleigh quotient from the beginning. For the results of column 3 through 6 we fix the target in the left hand side of the correction equation until $\|r\| < 0.01$. As can be seen from the second and third column, and as already has been suggested in Experiment 10.2, a "fix" gives better results than "no fix". Except for column 5, all correction equations are solved by 10 steps of unpreconditioned GMRES.

We see that for small and interior triples the $\mathcal{U}$- and $\mathcal{V}$-harmonic, and especially the double-harmonic and refined are superior compared with the standard approach. For the largest triples all methods are fine, with the exception that the double-harmonic needs a target $\tau < \infty$.

Experiment 10.4 In Figure 1 we compare the quality of extraction of the standard (a) and double-harmonic (b) method during the search for the smallest singular triple in Experiment 10.3. We plot $\angle(\mathcal{U}, e_1)$ (solid), $\angle(u, e_1)$ (dashed), $\angle(\mathcal{V}, e_1)$ (dots), and $\angle(v, e_1)$ (dash-dot).

In the double-harmonic approach (b), in every step $\angle(u, e_1) \approx \angle(\mathcal{U}, e_1)$ and $\angle(v, e_1) \approx \angle(\mathcal{V}, e_1)$, so this extraction process is almost optimal. We might say that the extraction is as good as the search spaces allow. In the standard extraction (a), the extraction is often good, but sometimes bad. In the expansion step (solving the correction equation), it is then unlikely to get a reasonable update, this can be regarded as the loss of one outer iteration. But what is worse: if we restart in a situation of bad extraction with few vectors, we may throw away the best part of the search space.

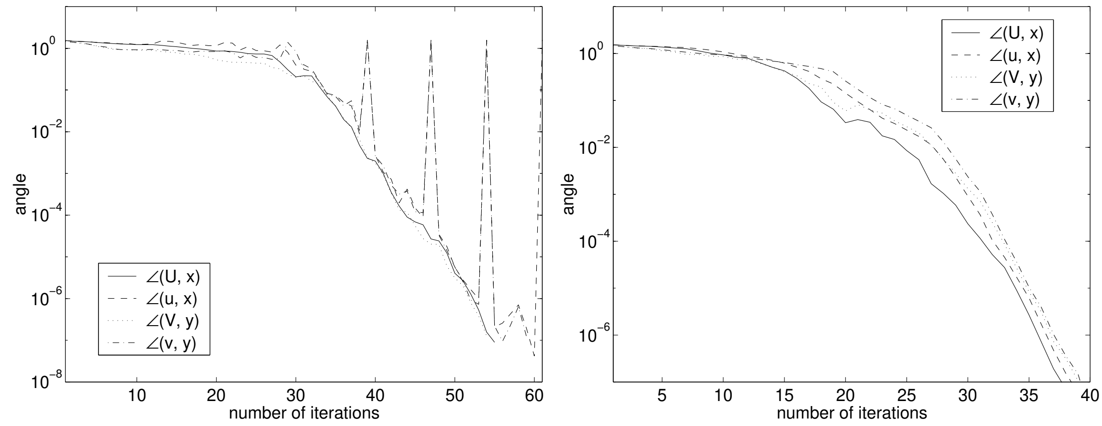

(a) (b)

FIGURE 1: The extraction results of the standard (a) and double-harmonic approach (b) for the computation of the smallest singular triple of $A = \text{diag}(1 : 100)$.

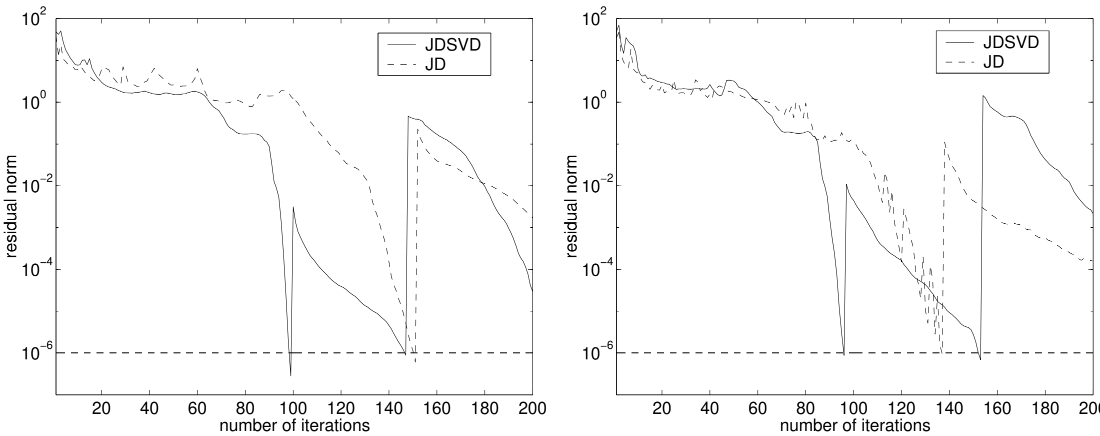

(a) (b)

FIGURE 2: JDSVD with double-harmonic approach versus JD with the harmonic Ritz approach (a), and JDSVD with refined extraction versus JD with refined Ritz (b) for the computation of the smallest singular triples of $A = \text{diag}(1 : 100)$.

Experiment 10.5 Next, we compare JDSVD with Jacobi–Davidson (JD) to compute the smallest singular triples. We compare these methods with two different extraction processes. For Figure 2(a), we use JD with the harmonic Ritz approach (see, for instance, [11, p. 292]) to compute the eigenpairs of the augmented matrix (2.4) closest to target $\tau = 0$. Recall (for instance, [3]) that finding an eigenpair of the augmented matrix gives full information on a singular triple of A, and vice versa. For JDSVD, we take the double-harmonic extraction method. All correction equations (JD and JDSVD) are solved approximately by 5 steps of GMRES. As the initial vectors we take v_1 random and $u_1 = Av_1$.

For Figure 2(b), JD uses the refined Ritz approach (see, for instance, [11, p. 289]), while JDSVD uses the refined Ritz approach of Section 4.

From Figure 2, it is clear that JDSVD easily beats JD with both extraction methods. A partly explanation of this fact could be that JD sees $\sigma_{\min}$ and $-\sigma_{\min}$ as two different eigenvalues of the augmented matrix. JDSVD avoids this "doubling".

Experiment 10.6 Now we illustrate the use the new extraction methods may have in producing approximate solutions to least squares problems with a simple example. Suppose we are interested

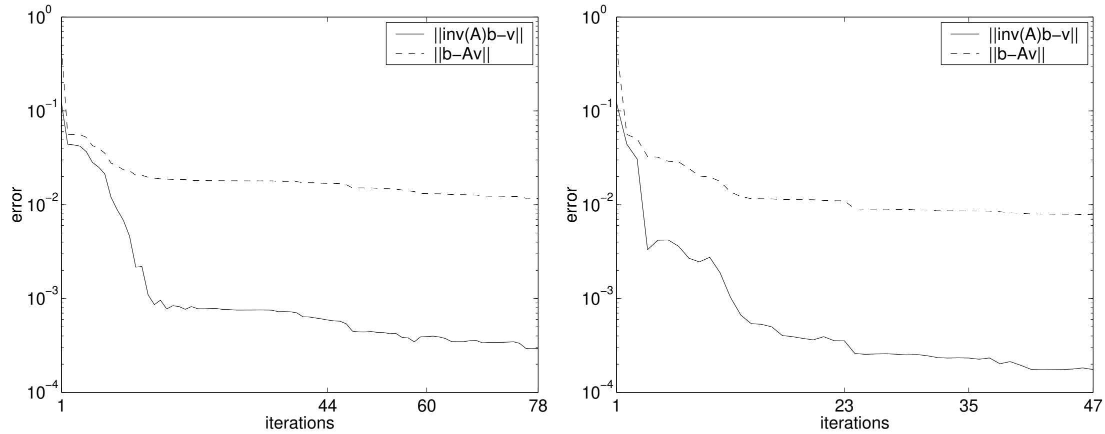

(a) (b)

FIGURE 3: The error and residual norm of the least squares solution produced by the $\mathcal{V}$-harmonic (a) and double-harmonic approach (b) during the computation of the three smallest singular triples of $A = \mathrm{diag}(1 : 100)$.

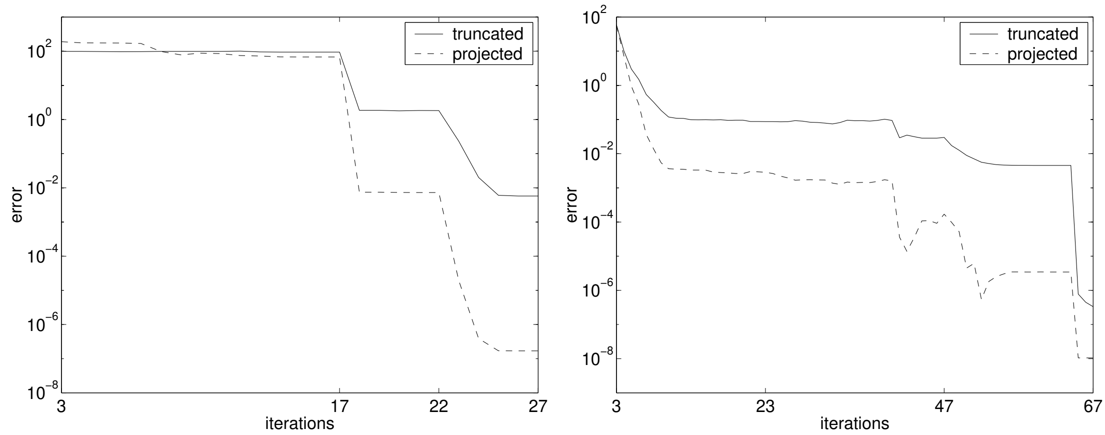

(a) (b)

FIGURE 4: JDSVD with standard (a) and double-harmonic (b) to approximate the truncated SVDs A_3 (a) and A_{-3} (b) of $A = \mathrm{diag}(1 : 100)$.

in the problem

$$\min_v \|b - Av\|,$$

where $A = \mathrm{diag}(1 : 100)$ and b is the vector of unit length with all entries equal. For Figure 3, we run JDSVD with $\mathcal{V}$-harmonic (a) and double-harmonic extraction (b) with target $\tau = 0$ and starting vectors $u_1 = v_1 = b$.

Depicted are the residual norm $\|b - Av\|$ and the error $\|A^{-1}b - v\|$. The numbers on the horizontal axis (except the number 1) indicate that at that iteration step, a singular triple is detected. We see that the errors decrease rapidly in the beginning; afterwards the convergence is slower but for the residual norm still monotonic, as predicted in Section 9.1. In particular, the true error seems to behave even more favorable than the residual norm.

Experiment 10.7 Finally, we illustrate the use of the extraction methods for the approximation of truncated SVDs. We run JDSVD to compute the three largest (a), respectively smallest (b) singular triples of $A = \mathrm{diag}(1 : 100)$. The starting vectors u_1 and v_1 are the vector of all ones, the target τ is 0 for (a) and ∞ for (b).

In Figure 4(a), we depict the error in the truncated SVD $\|P_{\mathcal{U}_3} A P_{\mathcal{V}_3} - A_3\|$ (here $A_3 = \sum_{j=98}^{100} j e_j e_j^T$, and $\mathcal{U}_3$ and $\mathcal{V}_3$ are the best three-dimensional part of $\mathcal{U}$ and $\mathcal{V}$), and the error in the "projected" truncated SVD $\|U_3^T AV_3 - \mathrm{diag}(98:100)\|$. For Figure 4(b), we give the error in the truncated SVD $\|P_{\mathcal{U}_3} A P_{\mathcal{V}_3} - A_{-3}\|$ and the error in the "projected" truncated SVD $\|U_3^T AV_3 H - \mathrm{diag}(1:3)\|$ (b). The numbers (except the number 3) on the horizontal axis in Figure 4 indicate the detection of a singular triple. For the "top" truncated SVD (a), the convergence corresponds, not surprisingly, with the detection of the singular triples.

11 Conclusions

For the accurate approximation of the minimal singular triple, we may use two separate subspaces. With respect to the *subspace expansion*, the Jacobi–Davidson SVD (inexact scaled RQI) is a competitor to Lanczos bidiagonalization when a (good) preconditioner is available, see [3].

With respect to the *subspace extraction*, the standard approach (via $H = U^T AV$) is fine for large singular triples, while for small and interior triples, the harmonic or refined approaches are more recommended. For the extraction of the smallest singular triple in Lanczos bidiagonalization we have seen that the standard, harmonic and refined approach are essentially equivalent.

Based on the theory and supported by numerical experiments, we can do the following recommendations:

- for the largest singular triples, we can choose any extraction method, except the double-harmonic (although this method may also perform well if we have a target $\tau < \infty$); the standard extraction is preferable because it is the cheapest;
- for interior singular triples we opt for the double-harmonic or refined approach;
- for the smallest singular triples we suggest any method except the standard; the double-harmonic and refined approach seem to be the most promising.

The $\mathcal{V}$-harmonic and double-harmonic method can also serve to give an approximate solution to a least squares problem; the standard and double-harmonic method can approximate the truncated SVDs.

MATLAB codes of JDSVD with incorporated extraction methods are available from the author on request.

References

[1] GOLUB, G. H., AND KAHAN, W. Calculating the singular values and pseudo-inverse of a matrix. *J. Soc. Indust. Appl. Math. Ser. B Numer. Anal. 2* (1965), 205–224.

[2] GOLUB, G. H., AND VAN LOAN, C. F. *Matrix Computations*, 3rd ed. The John Hopkins University Press, Baltimore, London, 1996.

[3] HOCHSTENBACH, M. E. A Jacobi–Davidson type SVD method. *SIAM J. Sci. Comput. 23*, 2 (2001), 606–628.

[4] The Matrix Market. `http://math.nist.gov/MatrixMarket`, a repository for test matrices.

[5] MORGAN, R. B. Computing interior eigenvalues of large matrices. *Linear Algebra Appl. 154/156* (1991), 289–309.

[6] PAIGE, C. C., AND SAUNDERS, M. A. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software 8*, 1 (1982), 43–71.

[7] PARLETT, B. N. *The Symmetric Eigenvalue Problem.* Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.

[8] SAAD, Y. *Numerical Methods for Large Eigenvalue Problems.* Manchester University Press, Manchester, UK, 1992.

[9] Simon, H. D., and Zha, H. Low-rank matrix approximation using the Lanczos bidiagonalization process with applications. *SIAM J. Sci. Comput. 21*, 6 (2000), 2257–2274 (electronic).

[10] Sleijpen, G. L. G., van der Vorst, H. A., and Meijerink, E. Efficient expansion of subspaces in the Jacobi–Davidson method for standard and generalized eigenproblems. *Electron. Trans. Numer. Anal. 7* (1998), 75–89.

[11] Stewart, G. W. *Matrix algorithms. Vol. II.* Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.