

A Jacobi–Davidson type method for the generalized singular value problem

M. E. Hochstenbach ^{a,*}

^a*Department of Mathematics and Computing Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands.*

Abstract

We discuss a new method for the iterative computation of some of the generalized singular values and vectors of a large sparse matrix. Our starting point is the augmented matrix formulation of the GSVD. The subspace expansion is performed by (approximately) solving a Jacobi–Davidson type correction equation, while we give several alternatives for the subspace extraction. Numerical experiments illustrate the performance of the method.

Key words: Generalized singular value decomposition (GSVD), partial GSVD, Jacobi–Davidson, subspace method, augmented matrix, correction equation, (inexact) accelerated Newton, refined extraction, harmonic extraction.

1 Introduction

The *generalized singular value decomposition* (GSVD) was introduced by Van Loan [21] and further developed by Paige and Saunders [15]. Let $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{p \times n}$ be given. The generalized singular values of the pair (A, B) are [21, Def. 1]

$$\Sigma(A, B) = \{\sigma \geq 0 : A^T A - \sigma^2 B^T B \text{ singular}\}.$$

The (diagonal form of the) GSVD of A and B is given by [21, Th. 2], [15, p. 399]

$$U^T A X = \Sigma_1 = \text{diag}(\alpha_j), \quad V^T B X = \Sigma_2 = \text{diag}(\beta_j), \quad (1.1)$$

* <http://www.win.tue.nl/~hochsten>.

This research was supported in part by NSF grant DMS-0405387.

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{p \times p}$ are orthogonal, $X \in \mathbb{R}^{n \times n}$ is nonsingular, $\Sigma_1 \in \mathbb{R}^{m \times n}$, $\Sigma_2 \in \mathbb{R}^{p \times n}$, $\alpha_j^2 + \beta_j^2 = 1$ and $\alpha_j, \beta_j \geq 0$. We can assume that the α_j and β_j are ordered such that

$$\alpha_1 \geq \cdots \geq \alpha_n, \quad \beta_1 \leq \cdots \leq \beta_n;$$

in this case the generalized singular values $\sigma_j = \alpha_j/\beta_j$ are nonincreasing. There is also a triangular form of the GSVD; see [15].

If B is square and nonsingular, the GSVD of (A, B) gives the SVD of AB^{-1} : $AB^{-1} = U\Sigma_1\Sigma_2^{-1}V^T$. In the special case $B = I_n$, the $n \times n$ identity matrix, we get the singular values of A . The GSVD has many applications, such as the computation of the Kronecker form of the matrix pencil $A - \lambda B$ [14], solving linear matrix equations [4], weighted least squares [3, 21], constrained least squares [6, pp. 580ff.], the common null space of two matrices [6, pp. 602ff.], regularization of ill-posed problems [7], information retrieval [12], and linear discriminant analysis [16] and is also useful for the method of particular solutions [2]. For an elaborate and interesting review, with relations with the cosine-sine decomposition, we refer the reader to work by Bai [1].

We denote the j th column of U, V, X by u_j, v_j, x_j , respectively. The GSVD is closely connected to two different generalized eigenvalue problems. In the first place, the pencil

$$(A^T A, B^T B) \tag{1.2}$$

has eigenvectors x_j with corresponding eigenvalues $\alpha_j^2/\beta_j^2 = \sigma_j^2$. In this paper, we will pursue the second form: the generalized eigenvalue problem

$$\left(\begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}, \begin{bmatrix} I & 0 \\ 0 & B^T B \end{bmatrix} \right) \tag{1.3}$$

has eigenvalues $\lambda_j = \pm\alpha_j/\beta_j = \pm\sigma_j$ corresponding to eigenvectors

$$\begin{bmatrix} u_j \\ \pm x_j / \beta_j \end{bmatrix} \tag{1.4}$$

(see [1]). This form expresses the GSVD as a structured eigenvalue problem; in this paper we will exploit this specific structure. An important issue in the computation of the GSVD is the question how to cope with the cross-product matrices $A^T A$ and $B^T B$. Of course we will never form these matrices explicitly. One potential advantage of the form (1.3) over (1.2) is that the action of multiplication by $A^T A$ is effectively split up in two separate actions. This may for instance be favorable when A has a large condition number.

When not A , but B has a problematic condition number, we may interchange the roles of A and B in (1.3); see Section 8 for more comments on this matter.

In several applications, such as the generalized total least squares problem, the matrices A and B are large and sparse, and one is interested in a *partial GSVD*: only a few of the generalized singular vectors corresponding to the smallest or largest generalized singular values are needed; see [22] and the references therein. There seems to be only one earlier paper concerning a partial GSVD for large sparse matrices. Zha [22] proposed a method for finding generalized singular pairs using the Cosine-Sine decomposition and Lanczos bidiagonalization. On the one hand, his method is attractive for working with $[A^T \ B^T]^T$, hence avoiding the cross-product matrices $A^T A$ and $B^T B$. On the other hand, a difficulty is that full-dimension orthogonal projections have to be computed in every step; inaccuracies in the projections limit the accuracy of the computed generalized singular pairs.

In this paper, we examine a Jacobi–Davidson type subspace method, which is related to the Jacobi–Davidson method for the singular value problem (JDSVD, [8, 9]), which in turn is inspired by the Jacobi–Davidson method for the eigenvalue problem [19]. We will discuss similarities and differences between the proposed method and JDSVD in Section 8.

The generalized singular value problem may have special types of generalized singular values: zero values ($\alpha/\beta = 0/1$), infinite values ($\alpha/\beta = 1/0$), and undefined values ($\alpha/\beta = 0/0$); these values are called *trivial*. Zero generalized singular values correspond to zero singular values of A , infinite generalized singular values to zero singular values of B , and undefined generalized singular values to zero singular values of $[A^T \ B^T]^T$. Because of these facts it is not very sensible to use the GSVD for the computation of trivial values; we will therefore concentrate on the computation of nontrivial values. Let $\rho = \text{rank}([A^T \ B^T]^T)$. If $\rho > \min\{m, p\}$ or $n > \rho$, then the presence of trivial values is guaranteed. However, in most applications $m \geq n$, $p \geq n$, and $\rho = n$ [1]. In this paper we assume that $p \geq n$ and that B is of full rank. If is not the case, but $m \geq n$ and A is of full rank, we may interchange the role of A and B (cf. comments in Sections 7 and 8). Although we need this assumption in several of the theoretical results, we remark that the resulting method may still be tried if this assumption is not satisfied.

We will now introduce some notational conventions for later use. We will write $\|\cdot\|$ for the Euclidean norm, and $\kappa(X)$ for the associated (two-norm) condition number of a matrix X . Since by assumption B has full rank, $(x, y)_{(B^T B)^{-1}} := y^T (B^T B)^{-1} x$ is an inner product. The corresponding norm satisfies $\|x\|_{(B^T B)^{-1}}^2 = (x, x)_{(B^T B)^{-1}}$. Inspired by the equality $\|Z\|_F^2 = \text{trace}(Z^T Z)$ for a real matrix

Z , we define the $(B^T B)^{-1}$ -Frobenius norm of Z by

$$\|Z\|_{(B^T B)^{-1}, F}^2 = \text{trace}(Z^T (B^T B)^{-1} Z). \quad (1.5)$$

The rest of the paper is organized as follows. In Section 2 and Section 3, we focus on the subspace expansion and subspace extraction of the new subspace method. Section 4 concentrates on the computation of the smallest generalized singular pairs, while Section 5 examines the convergence of the method and the relation with an inexact accelerated Newton method. We will discuss various properties of the method in Section 6. After numerical experiments in Section 7, we will present our conclusions in Section 8.

2 Subspace extraction

Our starting point is given by (1.3) and (1.4). Let

$$w_j = x_j / \beta_j;$$

we will come back to the relation between the w_j and x_j in Section 6.3. Inspired by (1.4), we will work with two search spaces, $\tilde{\mathcal{U}}$ for the generalized singular vectors u_j , and $\tilde{\mathcal{W}}$ for the scaled generalized singular vectors w_j .

Suppose we have k -dimensional search spaces $\tilde{\mathcal{U}}$ and $\tilde{\mathcal{W}}$, and look for an approximation

$$(\theta, \tilde{u}, \tilde{w}) \approx (\sigma, u, w), \quad \text{with} \quad \theta \geq 0, \quad \tilde{u} \in \tilde{\mathcal{U}}, \quad \tilde{w} \in \tilde{\mathcal{W}}.$$

As usual in subspace methods, \tilde{U} and \tilde{W} will be $m \times k$ and $n \times k$ search matrices whose columns form bases for $\tilde{\mathcal{U}}$ and $\tilde{\mathcal{W}}$, respectively.

First we define the *residual* of the triple $(\theta, \tilde{u}, \tilde{w})$ by

$$r = \begin{bmatrix} A\tilde{w} - \theta\tilde{u} \\ A^T\tilde{u} - \theta B^T B\tilde{w} \end{bmatrix}. \quad (2.1)$$

Since $\tilde{u} \in \tilde{\mathcal{U}}$ and $\tilde{w} \in \tilde{\mathcal{W}}$, we can write $\tilde{u} = \tilde{U}c$ and $\tilde{w} = \tilde{W}d$ for (low-dimensional) k -vectors c and d . To derive approximate triples, we impose two Galerkin conditions

$$\begin{aligned} A\tilde{W}d - \theta\tilde{U}c &\perp \tilde{\mathcal{U}}, \\ A^T\tilde{U}c - \theta B^T B\tilde{W}d &\perp \tilde{\mathcal{W}}. \end{aligned}$$

This is equivalent to the low-dimensional, projected system

$$\begin{aligned}\tilde{U}^T A \tilde{W} d &= \theta \tilde{U}^T \tilde{U} c, \\ \tilde{W}^T A^T \tilde{U} c &= \theta \tilde{W}^T B^T B \tilde{W} d.\end{aligned}\tag{2.2}$$

This suggests that it is convenient to let \tilde{U} have orthonormal columns and \tilde{W} have $B^T B$ -orthonormal columns; that is, $\tilde{W}^T B^T B \tilde{W} = I_k$. Then the subspace extraction amounts to computing singular values θ with corresponding right and left singular vectors d and c (of unit norm) of the projected matrix

$$H = \tilde{U}^T A \tilde{W},$$

where the columns of \tilde{U} are orthonormal and the columns of \tilde{W} are $B^T B$ -orthonormal. In line with the terminology for the standard eigenvalue problem, the names *Ritz value* for θ and *Ritz vectors* for \tilde{u} and \tilde{w} suggest themselves. Since $\|c\| = \|d\| = 1$, it follows that $\|\tilde{u}\| = \|B\tilde{w}\| = 1$. The Ritz value is equal to the *Rayleigh quotient*

$$\theta(\tilde{u}, \tilde{w}) = \tilde{u}^T A \tilde{w} \quad \left(= \frac{\tilde{u}^T A \tilde{w}}{\tilde{w}^T B^T B \tilde{w}} = \frac{\tilde{u}^T A \tilde{w}}{\tilde{u}^T \tilde{u}} \right)\tag{2.3}$$

of the vectors \tilde{u} and \tilde{w} . This Rayleigh quotient has the property that it minimizes the first part of the residual (2.1):

$$\theta = \underset{\gamma}{\operatorname{argmin}} \|A\tilde{w} - \gamma\tilde{u}\|.$$

However, it does not minimize the second part of the residual $\|A^T\tilde{u} - \gamma B^T B \tilde{w}\|$; in general, we cannot choose a value that minimizes the norms of both parts of the residual simultaneously. Indeed, the norm of the second part is minimized by $\gamma = \tilde{w}^T B^T B A^T \tilde{u} / \|B^T B \tilde{w}\|^2$, since for that γ value we have $A^T\tilde{u} - \gamma B^T B \tilde{w} \perp B^T B \tilde{w}$. Given \tilde{u} and \tilde{w} , we can also choose the θ that minimizes $\|r\|$. By setting the derivative of $\|r\|^2$ with respect to θ to zero we get that the minimizing value is

$$\frac{\tilde{u}^T A (\tilde{w} + B^T B \tilde{w})}{1 + \|B^T B \tilde{w}\|^2}.$$

We will not use this approximation in the method; it may, however, use it as a post-processing (refinement) step.

The following result, a generalization of [8, Th. 4.1], shows that given \tilde{U} and \tilde{W} , with orthonormal and $B^T B$ -orthonormal columns, respectively, H minimizes the *residual matrices*

$$R_1(K) := A\tilde{W} - \tilde{U}K \quad \text{and} \quad R_2(L) := A^T\tilde{U} - B^T B \tilde{W}L^T.$$

Theorem 2.1 For given $m \times k$ matrix \tilde{U} and $n \times k$ matrix \tilde{W} , let $H = \tilde{U}^T A \tilde{W}$.

- (a) If the columns of the given matrix \tilde{U} are orthonormal, then for all $k \times k$ matrices K we have $\|R_1(H)\| \leq \|R_1(K)\|$. Moreover, H is the unique minimizer 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 the given matrix \tilde{W} are $B^T B$ -orthonormal, then H minimizes $\|R_2(L)\|_{(B^T B)^{-1}}$ over all $k \times k$ matrices L , and H is the unique minimizer with respect to the $(B^T B)^{-1}$ -Frobenius norm (1.5).

Proof: Part (a) was proved in [8, Th. 4.1]. For (b), suppose that the columns of \tilde{W} are $B^T B$ -orthonormal; then

$$\begin{aligned} R_2(L)^T (B^T B)^{-1} R_2(L) &= \tilde{U}^T A (B^T B)^{-1} A^T \tilde{U} + LL^T - LH^T - HL^T \\ &= \tilde{U}^T A (B^T B)^{-1} A^T \tilde{U} - HH^T + (L - H)(L - H)^T \\ &= R_2(H)^T (B^T B)^{-1} R_2(H) + (L - H)(L - H)^T. \end{aligned}$$

Since $(L - H)(L - H)^T$ is positive semidefinite, it follows that (see, e.g., [20, p. 42])

$$\begin{aligned} \|R_2(L)\|_{(B^T B)^{-1}}^2 &= \lambda_{\max}(R_2(L)^T (B^T B)^{-1} R_2(L)) \\ &\geq \lambda_{\max}(R_2(H)^T (B^T B)^{-1} R_2(H)) \\ &= \|R_2(H)\|_{(B^T B)^{-1}}^2. \end{aligned}$$

For uniqueness, we realize, using (1.5), that for $L \neq H$

$$\begin{aligned} \|R_2(L)\|_{(B^T B)^{-1}, F}^2 &= \text{trace}(R_2(L)^T (B^T B)^{-1} R_2(L)) \\ &= \text{trace}(R_2(H)^T (B^T B)^{-1} R_2(H)) + \|L - H\|_F^2 \\ &> \|R_2(H)\|_{(B^T B)^{-1}, F}^2. \end{aligned}$$

The following theorem, a generalization of [8, Th. 4.3], ensures monotonic convergence to the largest generalized singular values.

Theorem 2.2 Let $\tilde{U}_k = [\tilde{u}_1 \dots \tilde{u}_k]$, $k = 1, 2, \dots$, be a sequence of $m \times k$ matrices with orthonormal columns and let $\tilde{W}_k = [\tilde{w}_1 \dots \tilde{w}_k]$, $k = 1, 2, \dots$, be a sequence of $n \times k$ matrices with $B^T B$ -orthonormal columns. Let $\theta_k^{(k)} \leq \dots \leq \theta_1^{(k)}$ be the ordered singular values of $H_k := \tilde{U}_k^T A \tilde{W}_k$. Then, for all fixed j and increasing k , the $\theta_j^{(k)}$ converge monotonically increasing to the σ_j .

Proof: H_k is a submatrix of H_{k+1} , so according to [11, (3.1.4)] $\theta_j^{(k+1)} \geq \theta_j^{(k)}$

for $1 \leq j \leq k$. Because of the orthogonality conditions on the columns of \tilde{U}_k and \tilde{W}_k , the $\theta_j^{(k)}$ converge to the σ_j .

This monotonicity is often of great value in practice: not only do we have a lower bound on the largest generalized singular value(s) during the process, but also a more rapid convergence. Similar to the convergence for the smallest singular values in JDSVD [8], the convergence to the smallest generalized singular values is not monotonic in general. The smallest generalized singular values correspond to the interior eigenvalues of the augmented matrix formulation (1.3) and are often much harder to find than the largest generalized singular values. Section 4 will be devoted to this subject.

For any approximation (θ, \tilde{w}) we have the following generalization of the Bauer–Fike theorem (see, e.g., [17, Th. 3.6]).

Theorem 2.3 *Let (θ, \tilde{w}) be an approximate eigenpair of the pencil $(A^T A, B^T B)$ with corresponding residual $A^T A \tilde{w} - \theta^2 B^T B \tilde{w}$. Then there is a generalized singular value σ such that*

$$|\theta^2 - \sigma^2| \leq \frac{\kappa(X) \|(A^T A - \theta^2 B^T B) \tilde{w}\|}{\sigma_{\min}^2(B)},$$

where the columns of X are the generalized singular vectors satisfying $A^T A X = B^T B X \Sigma^2$.

Proof: We have

$$\begin{aligned} \|r\| &\geq \sigma_{\min}(A^T A - \theta^2 B^T B) \\ &\geq \sigma_{\min}(A^T A X - \theta^2 B^T B X) \sigma_{\min}(X^{-1}) \\ &\geq \sigma_{\min}(\Sigma^2 - \theta^2 I) \sigma_{\min}(B^T B) \sigma_{\min}(X) \sigma_{\min}(X^{-1}) \\ &\geq \min|\sigma^2 - \theta^2| \sigma_{\min}(B^T B) \sigma_{\min}(X) \sigma_{\max}^{-1}(X), \end{aligned}$$

from which the conclusion follows.

As is not unusual in this type of result, the columns of X in the proposition above can still be scaled to tighten the upper bound; the problem of finding the scaling of X that minimizes this upper bound is highly nontrivial.

3 Subspace expansion

We now come to the genuine Jacobi–Davidson part of the method: the subspace expansion. Suppose we have an approximate triple $(\theta, \tilde{u}, \tilde{w})$, $\|\tilde{u}\| = \|B\tilde{w}\| = 1$, and we would like to enlarge the search spaces $\tilde{\mathcal{U}}$ and $\tilde{\mathcal{W}}$ to further

improve the approximation. Then we look for orthogonal updates $s \perp \tilde{u}$ and $t \perp \tilde{w}$ such that the updated vectors are generalized singular vectors in the sense that

$$\begin{aligned} A(\tilde{w} + t) &= \mu_1(\tilde{u} + s), \\ A^T(\tilde{u} + s) &= \mu_2 B^T B(\tilde{w} + t), \end{aligned} \quad (3.1)$$

for certain μ_1 and μ_2 . We note that we choose $t \perp \tilde{w}$ instead of $t \perp B^T B \tilde{w}$ to get an orthogonal subspace expansion which we hope will lead to fast convergence. Rewriting these equations to an equation involving the residual (2.1) gives

$$\begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = -r + \begin{bmatrix} (\mu_1 - \theta)\tilde{u} \\ (\mu_2 - \theta)B^T B \tilde{w} \end{bmatrix} + \begin{bmatrix} (\mu_1 - \theta)s \\ (\mu_2 - \theta)B^T B t \end{bmatrix}. \quad (3.2)$$

Since we are not interested in approximating trivial generalized singular values, we assume that $B(\tilde{w} + t) \neq 0$. We have

$$\begin{aligned} \mu_1 &= \tilde{u}^T A(\tilde{w} + t) &&= \theta + \mathcal{O}(\|t\|), \\ \mu_2 &= (\tilde{w} + t)^T A^T(\tilde{u} + s) / \|B(\tilde{w} + t)\|^2 &&= \theta + \mathcal{O}(\|s\| + \|t\|), \end{aligned} \quad (3.3)$$

so

$$\sigma^2 = \mu_1 \mu_2 = \theta^2 + \mathcal{O}(\|s\| + \|t\|).$$

Therefore, the last term on the right-hand side in (3.2) is of second order, that is, $\mathcal{O}((\|s\| + \|t\|)^2)$; we will neglect this term in the following. The idea is not to discard the (first-order) second term on the right-hand side, but to project the equation such that this term cancels. This forms the essence of asymptotically quadratic convergence; see also Section 5.¹ Apart for projecting out the last term on the right-hand side in (3.2), we want to fix the residual r to preserve the available information. Since the first component of the residual is orthogonal to \tilde{u} , and the second to \tilde{w} , this suggests using the projection

$$P = \begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - B^T B \tilde{w}\tilde{w}^T \end{bmatrix}, \quad (3.4)$$

¹ In the context of subspace methods, *asymptotic convergence* means the convergence behavior of the approximate quantities when they are sufficiently close to the true quantities.

which combines an orthogonal projection with an oblique projection and satisfies our two requirements. The resulting *correction equation* is

$$P \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = -r, \quad s \perp \tilde{u}, \quad t \perp \tilde{w}. \quad (3.5)$$

As usual in Jacobi–Davidson type methods, in practice we will often solve this equation approximately. Since the operator in (3.4) is not symmetric in general, GMRES is a reasonable solver. One of the advantages of Jacobi–Davidson type methods is that we may use a preconditioner, if available, to speed up the linear solve; see also Section 6.4.

Since the projected operator in (3.5) maps $\text{span}(\tilde{u})^\perp \times \text{span}(\tilde{w})^\perp$ onto itself, it can easily be repeated in the context of a Krylov subspace method. As an alternative correction equation, we may consider

$$\begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - B^T B \tilde{w}\tilde{w}^T \end{bmatrix} \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - \tilde{w}\tilde{w}^T B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = -r,$$

for $s \perp \tilde{u}$ and $t \perp B^T B \tilde{w}$. This formulation has the advantage that the operator is symmetric, but on the other hand, since it maps $\text{span}(\tilde{u})^\perp \times \text{span}(B^T B \tilde{w})^\perp$ to $\text{span}(\tilde{u})^\perp \times \text{span}(\tilde{w})^\perp$, we need a projected preconditioner of the form

$$\begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - B^T B \tilde{w}\tilde{w}^T \end{bmatrix} M \begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - \tilde{w}\tilde{w}^T B^T B \end{bmatrix}$$

to solve these equations by a Krylov subspace method; see also the remarks in [18] and [10].

4 Alternative extraction methods

In this section, we develop alternative extraction methods that are often more suitable for small generalized singular values than the standard extraction from Section 2. The alternatives, harmonic and refined extraction processes, are generalizations of those proposed in [9] for small singular triples. It turns out that some of these extractions can also be used for large generalized singular values and for generalized singular values close to a *target* $\tau \neq 0$.

4.1 Refined extractions

A refined Rayleigh–Ritz extraction for the standard eigenvalue problem was advocated in [13], see also [20], and was proposed for the singular value problem in [9]. A *refined extraction* process is possible for the generalized singular value problem for a target $0 \leq \tau \leq \infty$. To minimize the residual (2.1), we solve

$$\begin{bmatrix} \hat{c} \\ \hat{d} \end{bmatrix} = \underset{\substack{c, d \in \mathbb{R}^k \\ \|c\| = \|d\| = 1}}{\operatorname{argmin}} \left\| \begin{bmatrix} -\tau \tilde{U} & A\tilde{W} \\ A^T \tilde{U} & -\tau B^T B \tilde{W} \end{bmatrix} \begin{bmatrix} c \\ d \end{bmatrix} \right\|$$

and take $\hat{u} = \tilde{U}\hat{c}$ and $\hat{w} = \tilde{W}\hat{d}$.² Here τ can be either a fixed target, or a varying approximate generalized singular value, for instance, the Rayleigh quotient.

When we are interested in the smallest generalized singular value(s), we may focus on target $\tau = 0$ giving

$$\hat{d} = \underset{d \in \mathbb{R}^k, \|d\|=1}{\operatorname{argmin}} \|A\tilde{W}d\|, \quad \hat{c} = \underset{c \in \mathbb{R}^k, \|c\|=1}{\operatorname{argmin}} \|A^T \tilde{U}c\|.$$

This approach amounts to computing a small SVD of the thin tall matrices $A\tilde{W} \in \mathbb{R}^{m \times k}$ and $A^T \tilde{U} \in \mathbb{R}^{n \times k}$. However, since B does not play a role in this extraction process, we may not expect good results unless the minimal generalized singular value is very small.

When trying to find (very) large generalized singular values, with $\tau \rightarrow \infty$, the refined approach reduces to

$$\hat{d} = \underset{d \in \mathbb{R}^k, \|d\|=1}{\operatorname{argmin}} \|B\tilde{W}d\|, \quad \hat{w} = \tilde{W}\hat{d},$$

after which we can take

$$\hat{u} = \tilde{U}H\hat{d} / \|\tilde{U}H\hat{d}\|$$

as in the standard extraction. After the vector extraction, we can take a Rayleigh quotient (2.3) to get a (new) approximation to the generalized singular value. Although, because of the properties of the SVD, the smallest and largest singular values of $A\tilde{W}$ and $B\tilde{W}$ converge monotonically (decreasing

² In the literature, results of standard extraction are usually denoted with plain letters, those of harmonic extraction with a tilde, and those of refined extraction with a hat. Here, we have used a tilde to denote both the standard and harmonic extraction, but will still use a hat for the refined extraction.

and increasing, respectively), in general they will not converge to a generalized singular value, since the extraction only involves A (for $\tau = 0$) or B (for $\tau = \infty$). To ensure convergence in practice, we may have to let the target converge to the wanted eigenvalue, for instance by setting τ equal to the current Rayleigh quotient every now and then. (We note that this is also necessary in the refined approach for the standard eigenvalue problem.) See also the numerical experiments in Section 7.

4.2 Harmonic extractions

The largest generalized singular values are exterior eigenvalues of (1.3), the smallest are interior eigenvalues. Galerkin conditions usually work favorably for exterior values; in addition, we were able to show monotonic convergence for large generalized singular values in Section 2. This motivates us to look at the inverted problem.

Assume for the moment that $A \in \mathbb{R}^{n \times n}$ is square and nonsingular; as before $B \in \mathbb{R}^{p \times n}$ with full rank. The smallest generalized singular value(s) are the largest generalized singular value(s) of the problem involving σ^{-1} :

$$\begin{aligned} A^{-1}u &= \sigma^{-1}w, \\ A^{-T}(B^TB)w &= \sigma^{-1}u. \end{aligned}$$

If we write $B^TB = GG^T$ (for instance, a Cholesky decomposition), then we get

$$\begin{aligned} G^T A^{-1}u &= \sigma^{-1}y, \\ A^{-T}Gy &= \sigma^{-1}u. \end{aligned}$$

with $y = G^T w$. Therefore, this problem is also symmetric, and we expect monotonic convergence to the smallest generalized singular values for suitable Galerkin conditions, which is indeed the case; see Theorem 4.1.

This suggests to consider *Petrov–Galerkin conditions* on the residuals:

$$\begin{aligned} A^{-1}\tilde{u} - \tilde{\theta}^{-1}\tilde{w} &\perp \tilde{\mathcal{X}}, \\ A^{-T}(B^TB)\tilde{w} - \tilde{\theta}^{-1}\tilde{u} &\perp \tilde{\mathcal{Y}}, \end{aligned}$$

for certain *test spaces* $\tilde{\mathcal{X}}$ and $\tilde{\mathcal{Y}}$. To avoid working with the inverse of large sparse matrices, we want to make a suitable choice for these test spaces. For the first equation, we may choose $\tilde{\mathcal{X}} = A^T\tilde{\mathcal{U}}$ leading to the standard Galerkin

condition $A\tilde{w} - \tilde{\theta}\tilde{u} \perp \tilde{\mathcal{U}}$, or $\tilde{\mathcal{X}} = A^T A \tilde{\mathcal{W}}$ yielding the requirement $A\tilde{w} - \tilde{\theta}\tilde{u} \perp A\tilde{\mathcal{W}}$. For the second condition, we may choose $\tilde{\mathcal{Y}} = A\tilde{\mathcal{W}}$, leading to the standard Galerkin condition $A^T\tilde{u} - \tilde{\theta}B^T B\tilde{w} \perp \tilde{\mathcal{W}}$, or $\tilde{\mathcal{Y}} = AA^T\tilde{\mathcal{U}}$ giving $A^T\tilde{u} - \tilde{\theta}B^T B\tilde{w} \perp A^T\tilde{\mathcal{W}}$. We summarize possible Galerkin conditions in Table 1.

Table 1

Possible Galerkin conditions to obtain harmonic extractions.

First condition	Second condition
(1a) $A\tilde{w} - \tilde{\theta}\tilde{u} \perp \tilde{\mathcal{U}}$	(2a) $A^T\tilde{u} - \tilde{\theta}B^T B\tilde{w} \perp \tilde{\mathcal{W}}$
(1b) $A\tilde{w} - \tilde{\theta}\tilde{u} \perp A\tilde{\mathcal{W}}$	(2b) $A^T\tilde{u} - \tilde{\theta}B^T B\tilde{w} \perp A^T\tilde{\mathcal{U}}$

Note that the combination (1a) and (2a) gives the standard extraction of Section 2. The combination (1a) and (2b) is a generalization of the $\tilde{\mathcal{U}}$ -harmonic approach in [9]; a good way to implement constraint (2b) would be using a QR-decomposition of $A^T\tilde{\mathcal{U}}$. Here, we will focus on the combination of (1b) and (2a), which we will call the *harmonic approach* for $\tau = 0$. An important reason for this choice is that the resulting extraction method converges monotonically to the smallest generalized singular values; see Theorem 4.1.

This harmonic approach is characterized by the equations

$$\begin{aligned}\tilde{W}^T A^T A \tilde{W} \tilde{d} &= \tilde{\theta} \tilde{W}^T A^T \tilde{U} \tilde{c}, \\ \tilde{W}^T A^T \tilde{U} \tilde{c} &= \tilde{\theta} \tilde{W}^T B^T B \tilde{W} \tilde{d}.\end{aligned}$$

In particular, \tilde{d} solves a projected GSVD equation:

$$\tilde{W}^T A^T A \tilde{W} \tilde{d} = \tilde{\theta}^2 \tilde{W}^T B^T B \tilde{W} \tilde{d}, \quad \tilde{c} = \tilde{\theta} (\tilde{W}^T A^T \tilde{U})^{-1} \tilde{W}^T B^T B \tilde{W} \tilde{d}.$$

This is a generalization of what was called the \mathcal{V} -harmonic approach in [9]. Although our derivation assumed that A is square and invertible, for the resulting equations this is no longer needed. It seems necessary that $H = \tilde{U}^T A \tilde{W}$ is invertible, but this is not important in practice: if H is not invertible, we may expand the search spaces $\tilde{\mathcal{U}}$ and $\tilde{\mathcal{W}}$ by random vectors, or restart the method. Also, we may take the pseudoinverse, setting $\tilde{c} = \tilde{\theta} (\tilde{W}^T A^T \tilde{U})^+ \tilde{W}^T B^T B \tilde{W} \tilde{d}$. Note that this harmonic approach resembles an approach based on (1.2) in the extraction phase. However, it determines an extra vector \tilde{u} , and splits up the action of $A^T A$ in the expansion phase. In a practical implementation, we may again take $B^T B$ -orthonormal \tilde{W} .

The approximations $\tilde{\theta}$ to the generalized singular values have the property of monotonic convergence—also for the smallest values, as is shown in the following theorem. Denote the approximate generalized singular values in step k of the harmonic approach by

$$\tilde{\theta}_k^{(k)} \leq \dots \leq \tilde{\theta}_1^{(k)}.$$

Theorem 4.1 *In the harmonic approach, the approximate generalized singular values $\tilde{\theta}_j^{(k)}$ converge monotonically to both the largest and the smallest generalized singular values:*

$$\sigma_{\min} \leq \tilde{\theta}_k^{(k)} \leq \tilde{\theta}_{k-1}^{(k-1)}, \quad \tilde{\theta}_1^{(k-1)} \leq \tilde{\theta}_1^{(k)} \leq \sigma_{\max}.$$

Proof: With \tilde{W} a $B^T B$ -orthonormal basis for $\tilde{\mathcal{W}}$, the $\tilde{\theta}_j^{(k)}$ are the singular values of $A\tilde{W}_k$. Since $A\tilde{W}_k$ is a submatrix of $A\tilde{W}_{k+1}$, the result now follows from [11, (3.1.4)].

A harmonic approach is also possible for a target $0 < \tau < \infty$. Denote

$$\mathbf{A} = \begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} I & 0 \\ 0 & B^T B \end{bmatrix}, \quad Z = \begin{bmatrix} \tilde{U} & 0 \\ 0 & \tilde{W} \end{bmatrix}. \quad (4.1)$$

The (usual) harmonic approach on the pencil (\mathbf{A}, \mathbf{B}) for the target τ and search space $\text{span}(Z)$ consists of determining the eigenpair(s) $(\xi, [\tilde{c}^T \tilde{d}^T]^T)$ of the generalized eigenvalue problem

$$Z^T(\mathbf{A} - \tau\mathbf{B})^T(\mathbf{A} - \tau\mathbf{B})Z \begin{bmatrix} \tilde{c} \\ \tilde{d} \end{bmatrix} = \xi Z^T(\mathbf{A} - \tau\mathbf{B})^T\mathbf{B}Z \begin{bmatrix} \tilde{c} \\ \tilde{d} \end{bmatrix}$$

with the smallest $|\xi|$ [20, p. 296] (see also [5]). With the QR -decomposition

$$(\mathbf{A} - \tau\mathbf{B})Z = \begin{bmatrix} -\tau\tilde{U} & A\tilde{W} \\ A^T\tilde{U} & -\tau B^T B\tilde{W} \end{bmatrix} = QR,$$

this amounts to solving the generalized eigenvalue problem

$$Q^T \begin{bmatrix} \tilde{U} & 0 \\ 0 & B^T B\tilde{W} \end{bmatrix} \begin{bmatrix} \tilde{c} \\ \tilde{d} \end{bmatrix} = \xi^{-1} R \begin{bmatrix} \tilde{c} \\ \tilde{d} \end{bmatrix},$$

for which $|\xi^{-1}|$ is maximal. In this approach, convergence is not monotonic, which is not surprising since interior generalized singular values are approximated. It may be checked that for $\tau = 0$ this approach amounts to a combination of (1b) and (2b) in Table 1.

5 Convergence

We will now prove that the method has asymptotically quadratic convergence to generalized singular values when the correction equations are solved exactly,

and linear convergence when they are solved with a sufficiently small residual reduction.

Definition 5.1 We call a generalized singular value σ_j simple if $\sigma_i \neq \sigma_j$, for all $i \neq j$.

Lemma 5.2 Let (σ, u, w) be a generalized singular triple: $Aw = \sigma u$ and $A^T u = \sigma B^T B w$, where σ is a simple nontrivial generalized singular value, and $\|u\| = \|Bw\| = 1$. Write

$$P_\infty = \begin{bmatrix} I - uu^T & 0 \\ 0 & I - B^T B w w^T \end{bmatrix}. \quad (5.1)$$

Then the map

$$P_\infty \begin{bmatrix} -\sigma I & A \\ A^T & -\sigma B^T B \end{bmatrix}$$

is a bijection from $u^\perp \times w^\perp$ onto itself.

Proof: Suppose

$$P_\infty \begin{bmatrix} -\sigma I & A \\ A^T & -\sigma B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = 0$$

for $s \perp u, t \perp w$. We will show that $s = t = 0$; this proves that the operator in the lemma is injective and therefore also a bijection from $u^\perp \times w^\perp$ onto itself. Omitting the projector, there must be scalars α, β such that

$$\begin{bmatrix} -\sigma I & A \\ A^T & -\sigma B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = \begin{bmatrix} \alpha u \\ \beta B^T B w \end{bmatrix}. \quad (5.2)$$

Multiplying the first equation by A^T gives

$$\sigma A^T s = A^T A t - \alpha A^T u = A^T A t - \alpha \sigma B^T B w.$$

Substituting this in σ times the second equation in (5.2),

$$\sigma A^T s - \sigma^2 B^T B t = \beta \sigma B^T B w,$$

we get

$$(A^T A - \sigma^2 B^T B) t = \sigma(\alpha + \beta) B^T B w.$$

Left-multiplying this last equation by w^T gives $0 = w^T (A^T A - \sigma^2 B^T B) t = \sigma(\alpha + \beta)$, hence $A^T A t = \sigma^2 B^T B t$. But since $t \perp w$ and σ is a simple generalized singular value, we must have $t = 0$. In that case we get from the first

equation in (5.2) that $\alpha u + \sigma s = 0$. Since $s \perp u$, left-multiplication by s^T renders $\sigma \|s\|^2 = 0$. Because of the assumption that σ is a nontrivial generalized singular value we have $\sigma \neq 0$ hence $s = 0$ completing the proof.

Theorem 5.3 *With the assumptions of Lemma 5.2, assume that the correction equations (3.5) are solved exactly in every step. If the initial vectors (\tilde{u}, \tilde{w}) are close enough to (u, w) , then the sequence of approximations (\tilde{u}, \tilde{w}) converges quadratically to (u, w) .*

Proof: Let P , \mathbf{A} , and \mathbf{B} be as in (3.4) and (4.1). Let $[s_1^T t_1^T]^T$ with $s_1 \perp \tilde{u}$ and $t_1 \perp \tilde{w}$ be the exact solution to the correction equation

$$P(\mathbf{A} - \theta \mathbf{B}) \begin{bmatrix} s_1 \\ t_1 \end{bmatrix} = -r. \quad (5.3)$$

Moreover, let $\alpha u = \tilde{u} + s$, $s \perp \tilde{u}$, and $\beta w = \tilde{w} + t$, $t \perp \tilde{w}$, for certain scalars α and β , satisfy (3.1); note that these decompositions are possible since $u^T \tilde{u} \neq 0$ and $w^T \tilde{w} \neq 0$ because of the assumption that the vectors (\tilde{u}, \tilde{w}) are close to (u, w) . Projecting (3.2) yields

$$P(\mathbf{A} - \theta \mathbf{B}) \begin{bmatrix} s \\ t \end{bmatrix} = -r + P \begin{bmatrix} (\mu_1 - \theta)s \\ (\mu_2 - \theta)B^T B t \end{bmatrix}. \quad (5.4)$$

Subtracting (5.3) from (5.4) gives

$$P(\mathbf{A} - \theta \mathbf{B}) \begin{bmatrix} s - s_1 \\ t - t_1 \end{bmatrix} = P \begin{bmatrix} (\mu_1 - \theta)s \\ (\mu_2 - \theta)B^T B t \end{bmatrix}.$$

Lemma 5.2 implies that for (\tilde{u}, \tilde{w}) close enough to (u, w) , $P(\mathbf{A} - \theta \mathbf{B})$ is a bijection from $\tilde{u}^\perp \times \tilde{w}^\perp$ onto itself. Together with (3.3) this implies asymptotic quadratic convergence:

$$\left\| \begin{bmatrix} \alpha u - (\tilde{u} + s_1) \\ \beta w - (\tilde{w} + t_1) \end{bmatrix} \right\| = \left\| \begin{bmatrix} s - s_1 \\ t - t_1 \end{bmatrix} \right\| = \mathcal{O} \left(\left\| \begin{bmatrix} s \\ t \end{bmatrix} \right\|^2 \right).$$

Remark If we solve the correction equation exactly for the expansion $t \perp B^T B \tilde{w}$ (instead of $t \perp w$), we even get cubic asymptotic convergence. The reason for this is that in this case (cf. (3.3))

$$\begin{aligned} \mu_1 &= \tilde{u}^T A (\tilde{w} + t) = \theta + (\tilde{u} + s)^T A t + \mathcal{O}(\|s\| \|t\|) = \theta + \mathcal{O}(\|s\| \|t\|), \\ \mu_2 &= \tilde{w}^T A^T (\tilde{u} + s) = \theta + (\tilde{w} + t)^T A^T s + \mathcal{O}(\|s\| \|t\|) = \theta + \mathcal{O}(\|s\| \|t\|), \end{aligned}$$

since, with the notation of the previous proof,

$$(\tilde{u} + s)^T A = \alpha \sigma w^T B^T B \quad \text{and} \quad (\tilde{w} + t)^T A^T = \beta \sigma u^T.$$

This reduces the discarded term in the correction equation by an extra order of magnitude. However, since in practice we do not solve the correction equations exactly, this faster asymptotic convergence rate may not be very relevant. Instead, we choose an orthogonal subspace expansion $t \perp \tilde{w}$ which may be more important for fast overall convergence.

Theorem 5.4 *With the assumptions of Lemma 5.2, assume that the correction (3.5) are solved inexactly in every step with residual reduction*

$$\eta < (\kappa(P_\infty(\mathbf{A} - \sigma\mathbf{B})))^{-1},$$

where P_∞ , \mathbf{A} , and \mathbf{B} are as in (5.1) and (4.1), and the operator is seen as bijection from $u^\perp \times w^\perp$ onto itself as in Lemma 5.2. If the initial vectors (\tilde{u}, \tilde{w}) are close enough to (u, w) , then the sequence of approximations (\tilde{u}, \tilde{w}) converges linearly to (u, w) .

Proof: It follows from Lemma 5.2 that the condition number in the statement is finite. We employ the same notations as in Theorem 5.3, but here we are satisfied with approximate solutions $s_2 \perp \tilde{u}$, $t_2 \perp \tilde{w}$ to the correction equation such that $\|P(\mathbf{A} - \theta\mathbf{B}) \begin{bmatrix} s_2 \\ t_2 \end{bmatrix} + r\| \leq \eta \|r\|$. Then there are $0 \leq \eta_1 \leq \eta$, $f \perp \tilde{u}$, and $g \perp \tilde{w}$, $\|[f^T \ g^T]^T\| = 1$, such that

$$P(\mathbf{A} - \theta\mathbf{B}) \begin{bmatrix} s_2 \\ t_2 \end{bmatrix} = -r + \eta_1 \|r\| \begin{bmatrix} f \\ g \end{bmatrix}.$$

Subtracting this equation from (5.4) gives

$$P(\mathbf{A} - \theta\mathbf{B}) \begin{bmatrix} s - s_2 \\ t - t_2 \end{bmatrix} = -\eta_1 \|r\| \begin{bmatrix} f \\ g \end{bmatrix} + P \begin{bmatrix} (\mu_1 - \theta)s \\ (\mu_2 - \theta)B^T B t \end{bmatrix},$$

where the second term on the right-hand side is of second order as we have seen in the proof of Theorem 5.3. Furthermore,

$$\begin{aligned} r &= \begin{bmatrix} A\tilde{w} - \theta\tilde{u} \\ A^T\tilde{u} - \theta B^T B \tilde{w} \end{bmatrix} = - \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} + \begin{bmatrix} A(\tilde{w} + t) - \theta(\tilde{u} + s) \\ A^T(\tilde{u} + s) - \theta B^T B(\tilde{w} + t) \end{bmatrix} \\ &= - \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} + \begin{bmatrix} (\mu_1 - \theta)(\tilde{u} + s) \\ (\mu_2 - \theta)B^T B(\tilde{w} + t) \end{bmatrix}, \end{aligned}$$

so

$$r = Pr = -P \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} + P \begin{bmatrix} (\mu_1 - \theta)s \\ (\mu_2 - \theta)B^T B t \end{bmatrix}.$$

This means that

$$\|r\| \leq \|P(\mathbf{A} - \theta\mathbf{B})\| \left\| \begin{bmatrix} s \\ t \end{bmatrix} \right\| + \mathcal{O} \left(\left\| \begin{bmatrix} s \\ t \end{bmatrix} \right\|^2 \right)$$

and hence

$$\left\| \begin{bmatrix} s - s_2 \\ t - t_2 \end{bmatrix} \right\| \leq \eta_1 \kappa(P(\mathbf{A} - \theta \mathbf{B})) \left\| \begin{bmatrix} s \\ t \end{bmatrix} \right\| + \mathcal{O} \left(\left\| \begin{bmatrix} s \\ t \end{bmatrix} \right\|^2 \right).$$

We conclude that we get asymptotic linear convergence if the initial approximations are close enough to (u, w) and if η is as indicated in the theorem.

Finally, without giving further details, we mention that, similarly to [8], we can show that the new method can be interpreted as an inexact accelerated Newton scheme for the generalized singular value problem.

6 The algorithm

We now discuss various practical properties of the method. We first give pseudocode for the JDGSVD method in the following algorithm after which we discuss deflation and preconditioning.

6.1 Pseudocode

Algorithm: A Jacobi–Davidson type method for the GSVD

Input: Starting vectors u_1 and w_1 , a target τ , and a tolerance ε .

Output: An approximate triple (θ, u, w) for the generalized singular triple closest to the target τ satisfying $\left\| \begin{bmatrix} Aw - \theta u \\ A^T u - \theta B^T B w \end{bmatrix} \right\| \leq \varepsilon$.

- 1: $s = u_1, t = w_1, U_0 = [], W_0 = []$
for $k = 1, 2, \dots$ **do**
 - 2: $U_k = \text{RGS}(U_{k-1}, s)$
 $W_k = \text{RGS}_{B^T B}(W_{k-1}, t)$
 - 3: Compute k th column of $AW_k, A^T U_k,$ and $B^T B W_k$
Compute k th row and column of $H_k = U_k^T A W_k$
 - 4: Compute approximate generalized singular triple(s) $(\theta, c, d),$
with θ closest to $\tau,$ of the projected system
using standard, harmonic, or refined extraction techniques
 - 5: $u = U_k c, w = W_k d$
 - 6: $r = \begin{bmatrix} Aw - \theta u \\ A^T u - \theta B^T B w \end{bmatrix}$
 - 7: Stop if $\|r\| \leq \varepsilon$
 - 8: Solve (approximately) an $s \perp u, t \perp w$ from
$$\begin{bmatrix} I - uu^T & 0 \\ 0 & I - B^T B ww^T \end{bmatrix} \begin{bmatrix} -\theta I & A \\ A^T & -\theta B^T B \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = -r$$
-

In Step 2 of the algorithm, RGS and $\text{RGS}_{B^T B}$ stand for repeated Gram–Schmidt, a numerically stable way to form orthonormal, respectively $B^T B$ -orthonormal bases. In Step 4, we can choose between the different extraction techniques described in Sections 2 and 4, depending on the generalized singular values of interest. Every outer iteration costs four matrix vector products (MVs) with $A, A^T, B,$ and B^T (see Step 3). In addition, j steps of the inner iteration cost $4j$ MVs (the vector $B^T B w$ will be precomputed) and, if applicable, $j + 2$ applications of the preconditioner; see Section 6.4. Not included in this pseudocode, but included in our implementation for the experiments in Section 7, are deflation (see Section 6.2) and restarts. Also, we remark that in Step 8 of the JDGSVD algorithm, we may replace the shift θ by a given target $\tau,$ if applicable. This may be sensible in the beginning of the process if the Rayleigh quotient is not yet very accurate. This and other practical options are further discussed in Section 7.

6.2 Deflation

Deflation is the name of a technique that ensures that once we have detected a generalized singular value, we do not spend valuable effort to find it once again. We will need the following lemma.

Lemma 6.1 *If B is of full rank, then the w_j can be chosen to form an $A^T A$ -orthogonal and a $B^T B$ -orthonormal system. The u_j can be chosen to form an*

orthonormal system.

Proof: Since B is assumed to be of full rank, B^TB allows for a Cholesky decomposition $B^TB = GG^T$. With $y_j = G^T w_j$, the equation $A^T A w_j = \sigma_j^2 B^T B w_j$ becomes

$$G^{-1} A^T A G^{-T} y_j = \sigma_j^2 y_j.$$

Since the matrix on the left-hand side is symmetric, the eigenvectors y_j are orthogonal. Hence for $i \neq j$, $0 = y_i^T y_j = w_i^T B^T B w_j$. Moreover, $w_i^T A^T A w_j = \sigma_j^2 w_i^T B^T B w_j = 0$. We remark that in the case of multiple generalized singular values, the y_j can be chosen to be orthogonal, and likewise the w_j can also be chosen to be B^TB -orthogonal. Moreover, since for $i \neq j$

$$\sigma_i \sigma_j u_i^T u_j = w_i^T A^T A w_j = 0,$$

we see that the u_j form an orthogonal system: for $\sigma_i \sigma_j \neq 0$ we deduce $u_i^T u_j = 0$, while for $\sigma_i \sigma_j = 0$ the corresponding u -vectors can be chosen orthogonal. With the scaling $\|B w_j\| = 1$ for all j , we have that the w_j are $A^T A$ -orthogonal and B^TB -orthonormal, while the u_j are orthonormal.

Suppose we have already computed the generalized singular vectors $U_l = [u_1 \cdots u_l]$ and $W_l = [w_1 \cdots w_l]$, where BW_l has orthonormal columns. Using the preceding lemma, it can be checked that the pair of deflated matrices

$$\hat{A} := (I - U_l U_l^T) A (I - W_l W_l^T B^T B) \quad \text{and} \quad \hat{B} := B (I - W_l W_l^T B^T B) \quad (6.1)$$

has the same generalized singular values and vectors as the pair (A, B) , except that the computed values have been replaced by undefined ones (0/0).

6.3 A partial GSVD

If we compute l generalized singular values in combination with the deflation technique of the previous subsection, the result, in terms of the original undeflated A and B , is $AW_l = U_l R_1$, $A^T U_l = B^T B W_l R_2$, for upper triangular R_1 and R_2 . But we know even more. Since U_l has orthonormal columns, $R_1 = U_l^T A W_l$, and since W_l has B^TB -orthonormal columns, $R_2 = W_l^T A^T U_l$. Therefore, $R_2 = R_1^T$, and because both matrices are upper triangular we deduce that R_1 and R_2 are both diagonal and hence equal. Therefore, the JDGSVD method determines

$$AW_l = U_l S_l, \quad A^T U_l = B^T B W_l S_l, \quad (6.2)$$

where $S_l = \text{diag}(\sigma_1, \dots, \sigma_l)$ contains the computed generalized singular values.

If it is of interest to compute the full GSVD data: the α_j and β_j , as well as the vectors $V = [v_1 \cdots v_l]$ and $X = [x_1 \cdots x_l]$ (see Section 1), it is straightforward to compute a *partial GSVD* from (6.2).

Definition 6.2 $(\Sigma_1, \Sigma_2, U, V, X)$ is a *partial GSVD* of the pair (A, B) if, for $k \leq \min\{m, n, p\}$, $\Sigma_1, \Sigma_2 \in \mathbb{R}^{k \times k}$ are diagonal, $U \in \mathbb{R}^{m \times k}$ and $V \in \mathbb{R}^{p \times k}$ have orthonormal columns, $X \in \mathbb{R}^{n \times k}$ is of full rank, $AX = U\Sigma_1$, $BX = V\Sigma_2$, and $\Sigma_1^2 + \Sigma_2^2 = I_k$.

Since W_l has B^TB -orthonormal columns, BW_l has orthonormal columns, and $W_l^T A^T A W_l + W_l^T B^T B W_l = S_l^2 + I$. Therefore, we can cheaply compute the information Σ_1, Σ_2, U, V , and X from the computed triples (σ_j, u_j, w_j) as indicated in the following algorithm.

Algorithm: Computing a partial GSVD from the JDGSVD data

Input: $AW_l = U_l S_l$, $A^T U_l = B^T B W_l S_l$, with $S_l = \text{diag}(\sigma_1, \dots, \sigma_l)$;

U_l with orthonormal columns and W_l with B^TB -orthonormal columns.

Output: A partial GSVD $(\Sigma_1, \Sigma_2, U, V, X)$.

- 1: $V = BW_l$ (already computed in the JDGSVD algorithm)
 - 2: $\Sigma_2 = (I + S_l^2)^{-1/2}$
 - 3: $\Sigma_1 = S_l \Sigma_2$
 - 4: $X = W_l \Sigma_2$
-

6.4 Preconditioning

Given a target τ , we may try to use a preconditioner

$$M \approx \begin{bmatrix} -\tau I & A \\ A^T & -\tau B^T B \end{bmatrix}$$

to more efficiently solve the correction equations. For $\tau = 0$, which means that we are interested in the smallest generalized singular values, M could be an approximation to the augmented matrix $\begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}$. In the case that A is square, we may take

$$M = \begin{bmatrix} 0 & N \\ N^T & 0 \end{bmatrix}, \quad M^{-1} = \begin{bmatrix} 0 & N^{-T} \\ N^{-1} & 0 \end{bmatrix}, \quad (6.3)$$

where $N \approx A$ is a (relatively) cheaply invertible preconditioner for A .

Preconditioning the correction equation (3.5) means solving $s \perp \tilde{u}$ and $t \perp \tilde{w}$ from

$$\begin{bmatrix} I - \tilde{u}\tilde{u}^T & 0 \\ 0 & I - B^T B \tilde{w}\tilde{w}^T \end{bmatrix} M \begin{bmatrix} s \\ t \end{bmatrix} = b$$

for a right-hand side b . This means that

$$M \begin{bmatrix} s \\ t \end{bmatrix} = b + \begin{bmatrix} \alpha \tilde{u} \\ \beta B^T B \tilde{w} \end{bmatrix}$$

for certain α and β which are determined by the orthogonality conditions for s and t . From

$$\begin{bmatrix} s \\ t \end{bmatrix} = M^{-1}b + M^{-1} \begin{bmatrix} \tilde{u} & 0 \\ 0 & B^T B \tilde{w} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \tilde{u} & 0 \\ 0 & \tilde{w} \end{bmatrix}^T \begin{bmatrix} s \\ t \end{bmatrix} = 0$$

it can be verified that we have

$$\begin{bmatrix} s \\ t \end{bmatrix} = \left(I - M^{-1} \begin{bmatrix} \tilde{u} & 0 \\ 0 & B^T B \tilde{w} \end{bmatrix} \left(\begin{bmatrix} \tilde{u} & 0 \\ 0 & \tilde{w} \end{bmatrix}^T M^{-1} \begin{bmatrix} \tilde{u} & 0 \\ 0 & B^T B \tilde{w} \end{bmatrix} \right)^{-1} \begin{bmatrix} \tilde{u} & 0 \\ 0 & \tilde{w} \end{bmatrix}^T \right) M^{-1}b.$$

Since at the beginning of the inner iteration $M^{-1} \begin{bmatrix} \tilde{u} \\ 0 \end{bmatrix}$ and $M^{-1} \begin{bmatrix} 0 \\ B^T B \tilde{w} \end{bmatrix}$ may be precomputed, we need $j + 2$ actions with our preconditioner for j inner iterations.

7 Numerical experiments

We start with similar experiments as Zha [22], but we take the dimension of the matrices ten times as large ($n = 1000$ instead of $n = 100$). These test cases include matrices with high condition numbers and are therefore numerically challenging. In all experiments, we first initialize the Matlab's random generator by the command `rand('state', 0)`.

Experiment 7.1 We choose two diagonal matrices: for $j = 1, \dots, 1000$,

$$\begin{aligned} C &= \text{diag}(c_j), & c_j &= (n - j + 1)/2n, & S &= \sqrt{I - C^2}, \\ D &= \text{diag}(d_j), & d_j &= \lceil j/250 \rceil + r_j, \end{aligned}$$

where the r_j are random numbers chosen from a uniform distribution on the interval $(0, 1)$ and $\lceil \cdot \rceil$ denotes the `ceil` function. With $A = CD$ and $B = SD$, the condition numbers of A and B are modest. In separate runs of the algorithm, we look for the three largest and smallest generalized singular values, using the following default options of the proposed method:

Parameter	Meaning	Default value
<code>tol</code>	absolute tolerance for the outer iteration	10^{-6}
<code>maxit</code>	maximum number of outer iterations	1000
<code>mindim</code>	dimension of search spaces after restart	10
<code>maxdim</code>	maximum dimension of search spaces	30
<code>maxit_inner</code>	maximum iterations to solve correction equation	10
<code>inner_tol</code>	relative tolerance inner iteration	0
<code>fix</code>	fix target until $\ r\ \leq \text{fix}$ (see below)	0.01
<code>u1, w1</code>	initial search spaces	random
<code>krylov</code>	start-up with Krylov spaces	true
<code>M1, M2</code>	preconditioner $M = M_1 M_2$	$M_1 = M_2 = I$
<code>Mtype</code>	left or right preconditioning	left

The inner iterations are stopped once the inner tolerance `inner_tol` is met or if `maxit_inner` inner iterations have been carried out. The default choices for these parameters mean that the correction equations (3.5) are solved approximately by exactly ten steps of the GMRES method. If the `krylov` parameter is set, then the method first generates Krylov spaces of dimension `maxdim`; generated by $A^T A$ if we look for the largest generalized singular values, and by $B^T B$ if we look for the smallest generalized singular values. The only other parameter that may need some explanation is `fix`. On the left-hand side of the correction equation (3.5), we take θ equal to the target τ as long as the residual norm is larger than `fix`. If $\|r\|$ drops under this value, we assume that convergence has set in, and we take θ equal to the Rayleigh quotient in every step. We take the same starting vector for each of the extraction methods. This forms experiment (1) in Table 2.

For experiments (2), (3), and (4), we take

$$A = Q_1 C D Q_2, \quad B = Q_1 S D Q_2,$$

where Q_1 and Q_2 are two random orthogonal matrices. For experiment (2), we take D as in experiment (1). For experiments (3) and (4) we adjust D using

$$d_j = d_j - \min_{1 \leq i \leq 1000} d_i + 10^{-e}, \quad j = 1, \dots, 1000,$$

where we take $e = 6, 9$ for experiments (3) and (4), respectively. As can also be seen in Table 2, this choice affects the condition numbers of both A and

B . In [22] it is noted that ill-conditioning of the matrices limits the final attainable accuracy for Zha’s method. For JDGSVD, high condition numbers will generally also imply that the the correction equations are harder to solve.

Table 2

Results for experiments (1)–(4) (cf. [22]) for $n = 1000$. We give the number of outer iterations of the new method for computing the three largest and three smallest generalized singular values for different extraction techniques. Note the condition numbers in the two rightmost columns. For the cases indicated by an asterisk we used an inexact LU preconditioner with drop tolerance 10^{-3} .

#	σ_{\max}		σ_{\min}			$\kappa(A)$	$\kappa(B)$
	Standard	Harmonic	Standard	Harmonic	Refined		
(1)	81	157	784	334	280	4.4e2	5.7e0
(2)	93	153	703	393	277	4.4e2	5.7e0
(3)	662	597	16*	19*	–	1.9e6	4.3e6
(4)	577	517	59*	20*	–	1.9e9	4.3e9

The results for the computation of σ_{\max} are in line with the theory: the standard extraction is fine for the largest generalized singular values; the harmonic extraction ((1b) and (2a) in Table 1) is primarily meant for the smallest generalized singular values but also usable for the largest values. We remark that for experiments (3) and (4), computing σ_{\max} did not take much longer than for experiments (1) and (2); but computing the next two values did. The refined extraction (not in the table) with $\tau = \infty$ failed in all cases. This is natural in view of $\sigma_{\max} = 0.577 \cdots \ll \infty$. A suitable target might be helpful in this situation; how to cheaply find a first rough approximation to σ_{\max} is an interesting research question.

For σ_{\min} , the harmonic extraction indeed does a better job than the standard extraction for most experiments. The refined extraction shows a good performance for modest-conditioned A and B thanks to the fact that the target $\tau = 0$ is rather accurate for $\sigma_{\min} \approx 5.0 \cdot 10^{-4}$. If we do not use a preconditioner in experiments (3) and (4), we have difficulties in computing the smallest singular value for ill-conditioned matrices A and B to the prescribed tolerance 10^{-6} . Only with a good preconditioner (inexact LU decomposition with drop tolerance 10^{-3}) we get a rapid convergence to the prescribed tolerance in various cases. The refined approach with the same preconditioner fails for more difficult test cases. We note that in experiments (3) and (4), due to the ill-conditioning of B , it turned out to be advantageous to turn off the `krylov` option.

Experiment 7.2 Next, we illustrate the use of interchanging the roles of A and B . We generate random sparse 1000×1000 A and B with a density of about 10% by the commands

```
n=1000; A = sprand(n,n,1e-1,1); B = sprand(n,n,1e-1,1e-2).
```

It turns out that $\kappa(A) \approx 9.4 \cdot 10^4$, $\kappa(B) \approx 6.0 \cdot 10^1$. Suppose we are interested in the largest generalized singular value. The convergence for σ_{\max} is slow; after 2000 outer iterations none of the three extraction processes has succeeded in finding this value. However, to find $\sigma_{\max}(A, B)$ we can also compute $\sigma_{\min}(B, A)$ instead. With target 0 and an ILU preconditioner with drop tolerance 0.001, both the standard and harmonic extraction find $\sigma_{\min}(B, A)$ in just 5 iterations. So in this case, since we have a good target and an appropriate preconditioner, the smallest generalized singular value is actually easier to detect than the largest, so that it is useful to interchange A and B .

8 Conclusions

We have proposed a new Jacobi–Davidson type method for the computation of some of the generalized singular values and corresponding vectors. The method is an accelerated (inexact) Newton method with asymptotically quadratic convergence if the correction equations are solved exactly. To accelerate the initial phase, we use subspace acceleration. While the convergence for the largest (exterior) values is naturally favorable (monotonic behavior), the method may also be used to compute the smallest generalized singular values. Preconditioners are relatively easy to obtain in the latter case, at least in the case of a square and invertible A . As there are no experiments in [22] for the smallest generalized singular values, it is not clear whether the Lanczos type method can be used for these values. Although the GSVD of complex matrices does not seem to be discussed in the literature, the method could be applied to these matrices with straightforward adaptations.

The described method may be seen as an adaptation of the JDSVD method for the singular value problem [8, 9], with a $B^T B$ -orthonormal basis \widetilde{W} . However, as the generalized singular value problem is mathematically more involved than the singular value problem, this is also true for the numerical solution in a number of aspects.

First, in JDGSVD we can choose between a $B^T B$ -orthogonal subspace expansion with cubic asymptotic convergence, and an orthogonal expansion with corresponding quadratic convergence. As the difference in asymptotic convergence rate may not be very important (both variants typically behave linearly if used inexactly), we preferred the orthogonal expansion.

Second, for the GSVD, the refined extraction is less attractive since the important cases $\tau = 0$ and $\tau = \infty$ are no longer as natural as for the singular value problem [9]. Two of the three different harmonic extractions in [9] also become more involved for the GSVD.

Third, per iteration, one needs four matrix vector products for the outer iteration, plus an additional four for every inner step; this is twice the number of JDSVD.

Fourth, both in the correction equation (3.5) and the deflation (6.1) an oblique projection (3.4) is present, which may affect the stability and efficiency. Although the cross-product matrix B^TB is never formed, it is applied. For ill-conditioned B , this may give numerical difficulties. The numerical experiments suggest that the effect of ill-conditioned matrices is that it takes longer to compute the largest generalized singular values, while for the the smallest generalized singular values good preconditioners and/or more modest tolerances are needed.

Finally, it is important to notice that we can interchange the role of A and B via the form $B^TBx = (1/\sigma^2)A^TAx$ if desired. This may be practical in the case that $\sigma_{\min}(B, A)$ is easier to compute than $\sigma_{\max}(A, B)$ (see Experiment 7.2), or in the case that the condition numbers of A and B differ greatly.

Acknowledgements

This paper is dedicated with pleasure to Henk van der Vorst. The author thanks Gerard Sleijpen for helpful discussions and comments. The excellent comments of the referees are greatly appreciated. In particular, one referee went far beyond the call of duty, which is very thankfully acknowledged.

References

- [1] Z. BAI, *The CSD, GSVD, their applications and computations*, Preprint Series 958, Institute for Mathematics and its Applications, University of Minnesota, April 1992.
- [2] T. BETCKE, *The generalized singular value decomposition and the method of particular solutions*, SIAM J. Sci. Comput., 30 (2008), pp. 1278–1295.
- [3] Å. BJÖRCK, *Numerical Methods for Least Squares Problems*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1996.
- [4] K.-w. E. CHU, *Singular value and generalized singular value decompositions and the solution of linear matrix equations*, Linear Algebra Appl., 88/89 (1987), pp. 83–98.
- [5] D. R. FOKKEMA, G. L. G. SLEIJPEN, AND H. A. VAN DER VORST, *Jacobi–Davidson style QR and QZ algorithms for the reduction of matrix pencils*, SIAM J. Sci. Comput., 20 (1998), pp. 94–125.

- [6] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The John Hopkins University Press, Baltimore, London, 3rd ed., 1996.
- [7] P. C. HANSEN, *Rank-deficient and Discrete Ill-posed Problems*, SIAM Monographs on Mathematical Modeling and Computation, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.
- [8] M. E. HOCHSTENBACH, *A Jacobi–Davidson type SVD method*, SIAM J. Sci. Comput., 23 (2001), pp. 606–628.
- [9] ———, *Harmonic and refined extraction methods for the singular value problem, with applications in least squares problems*, BIT, 44 (2004), pp. 721–754.
- [10] M. E. HOCHSTENBACH AND G. L. G. SLEIJPEN, *Two-sided and alternating Jacobi–Davidson*, Linear Algebra Appl., 358(1-3) (2003), pp. 145–172.
- [11] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.
- [12] P. HOWLAND, M. JEON, AND H. PARK, *Structure preserving dimension reduction for clustered text data based on the generalized singular value decomposition*, SIAM J. Matrix Anal. Appl., 25 (2003), pp. 165–179.
- [13] Z. JIA, *Refined iterative algorithms based on Arnoldi’s process for large unsymmetric eigenproblems*, Linear Algebra Appl., 259 (1997), pp. 1–23.
- [14] B. KÅGSTRÖM, *The generalized singular value decomposition and the general $(A - \lambda B)$ -problem*, BIT, 24 (1984), pp. 568–583.
- [15] C. C. PAIGE AND M. A. SAUNDERS, *Towards a generalized singular value decomposition*, SIAM J. Numer. Anal., 18 (1981), pp. 398–405.
- [16] C. H. PARK AND H. PARK, *A relationship between linear discriminant analysis and the generalized minimum squared error solution*, SIAM J. Matrix Anal. Appl., 27 (2005), pp. 474–492.
- [17] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, Manchester University Press, Manchester, UK, 1992.
- [18] G. L. G. SLEIJPEN, A. G. L. BOOTEN, D. R. FOKKEMA, AND H. A. VAN DER VORST, *Jacobi–Davidson type methods for generalized eigenproblems and polynomial eigenproblems*, BIT, 36 (1996), pp. 595–633.
- [19] G. L. G. SLEIJPEN AND H. A. VAN DER VORST, *A Jacobi–Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 401–425.
- [20] G. W. STEWART, *Matrix algorithms. Vol. II*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.
- [21] C. F. VAN LOAN, *Generalizing the singular value decomposition*, SIAM J. Numer. Anal., 13 (1976), pp. 76–83.
- [22] H. ZHA, *Computing the generalized singular values/vectors of large sparse or structured matrix pairs*, Numer. Math., 72 (1996), pp. 391–417.