

A Structure-Preserving LOBPCG Method for Computing Several of the Largest Generalized Singular Triplets

Xinyu Shan

xyshan@simis.cn

Shanghai Institute for Mathematics and Interdisciplinary Sciences

Research Article

Keywords: Generalized singular value decomposition, Structure-preserving LOBPCG algorithm, Jordan–Wielandt pencil, StructuredGalerkin condition, Improved Hetmaniuk–Lehoucq trick

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11171518/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Structure-Preserving LOBPCG Method for Computing Several of the Largest Generalized Singular Triplets

Xinyu Shan^{1,2*}

^{1*}Shanghai Institute for Mathematics and Interdisciplinary Sciences (SIMIS), Shanghai, 200433, China.

²Research Institute of Intelligent Complex Systems, Fudan University, Shanghai, 200433, China.

Corresponding author(s). E-mail(s): xyshan@simis.cn;

Abstract

We develop a structure-preserving locally optimal block preconditioned conjugate gradient (LOBPCG) method for computing several of the largest positive generalized singular values of a matrix pair $(\mathbf{A}, \mathbf{B})$, where $\mathbf{B}$ has full column rank, together with the associated generalized singular vectors. The method is derived from the Hermitian-definite Jordan–Wielandt pencil and therefore avoids explicitly forming the normal-equation matrices. By exploiting the block structure and spectral symmetry of this pencil, we construct separate search subspaces for the left and right generalized singular vectors. The resulting Rayleigh–Ritz procedure requires only a singular value decomposition of a small projected cross matrix. We adapt the improved Hetmaniuk–Lehoucq (IHL) trick to construct stable conjugate search directions in the two component spaces. We also establish an explicit correspondence between the coefficient matrices produced by the componentwise and augmented Rayleigh–Ritz procedures. This correspondence permits the componentwise IHL construction to be used in an augmented, structure-preserving LOBPCG formulation without explicitly assembling the augmented basis. Numerical experiments demonstrate the computational advantages of the proposed method over unstructured and Gram-matrix-based LOBPCG implementations.

Keywords: Generalized singular value decomposition; Structure-preserving LOBPCG algorithm; Jordan–Wielandt pencil; Structured Galerkin condition; Improved Hetmaniuk–Lehoucq trick

MSC Classification: 65F15 , 65F25 , 15A18

1 Introduction

The generalized singular value decomposition (GSVD) was introduced by Van Loan [1] and subsequently developed by Paige and Saunders [2] and others [3, 4]. It appears in a wide range of numerical linear algebra and data analysis problems, including eigenvalue problems [5], least-squares problems [6], information retrieval [7], and real-time signal processing [8].

Several numerical algorithms based on dense matrix factorizations have been developed for computing the GSVD [9, 10], and recent advances in the stable computation of the CS decomposition provide an alternative approach [11]. For large-scale sparse matrix pairs, iterative methods are generally more suitable when only a few generalized singular triplets are required. Zha’s algorithm [12], which combines the CS decomposition with the Lanczos bidiagonalization process, computes a few extremal generalized singular values and their associated vectors. Jacobi–Davidson-type methods [13, 14] can target generalized singular values at arbitrary locations and are particularly well suited for computing interior ones. More recently, Liu et al. [15] proposed a contour-integral-based algorithm that exploits the structure of the associated Jordan–Wielandt matrix pencil to compute generalized singular values in a prescribed interval. The LOBPCG method has also been applied to the GSVD through the normal-equation generalized eigenvalue problem [16].

Originally proposed by Knyazev [17], the LOBPCG method computes extremal eigenpairs of symmetric-definite generalized eigenvalue problems. At each iteration, the LOBPCG method performs a Rayleigh–Ritz procedure in the trial subspace spanned by the current approximations, preconditioned residuals, and previous search directions. Hetmaniuk and Lehoucq observed that a poorly chosen basis can make the Gram matrix in the Rayleigh–Ritz procedure ill-conditioned, thereby slowing or even preventing convergence [18]. They proposed a basis-updating strategy that maintains orthogonality, commonly called the IHL trick. An improved variant (IHL), together with additional strategies for numerical robustness, was analyzed in [19]. The IHL trick has also been employed in a structure-preserving LOBPCG method for the Bethe–Salpeter eigenvalue problem [20].

In this paper, we develop a structure-preserving LOBPCG method for computing several of the largest nontrivial generalized singular values of a matrix pair (A, B) and their associated generalized singular vectors. Unlike the existing LOBPCG method based on the normal-equation formulation [16], our method is derived from the generalized Hermitian-definite eigenvalue problem associated with the Jordan–Wielandt matrix pencil [21]. Instead of treating the resulting augmented eigenvalue problem as an unstructured problem, we exploit its block structure and construct separate search subspaces for the left and right generalized singular vectors. Under a simplified Galerkin condition, the Rayleigh–Ritz procedure reduces to computing a small SVD of the projected cross matrix. To update the two search subspaces, we adapt the IHL trick to the componentwise setting. We further establish an explicit correspondence between the coefficient representations arising from the simplified and structured Galerkin conditions. This correspondence enables the componentwise IHL construction to be incorporated into the augmented structure-preserving LOBPCG framework.

The paper is organized as follows. In Section 2, we review the GSVD, the Jordan–Wielandt pencil, and the structured Galerkin condition. Section 3 develops the structure-preserving LOBPCG method. Numerical experiments are reported in Section 4. Conclusions are drawn in Section 5.

2 Preliminaries

Throughout this paper, let $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{p \times n}$, with B of full column rank, i.e., $\text{rank}(B) = n$. We use *nontrivial* to mean strictly positive and write $q = \text{rank}(A) > 0$. The derivation of the method uses k as the working block size. A^H denotes the conjugate transpose of A . A^T represents the transpose of A . I_d represents the identity matrix of order d . $\text{span}(M)$ denotes the subspace spanned by the columns of M . Unless otherwise stated, orthogonality is Euclidean. The weighted inner product on the right component space is $\langle x, y \rangle_{B^H B} = x^H B^H B y$.

2.1 Generalized singular value decomposition

Definition 1 (Generalized singular value decomposition) Let $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{p \times n}$, and assume that $\text{rank}(B) = n$. Then there exist unitary matrices $U \in \mathbb{C}^{m \times m}$ and $V \in \mathbb{C}^{p \times p}$, together with a nonsingular matrix $X \in \mathbb{C}^{n \times n}$, such that

$$U^H A X = \Sigma_A, \quad V^H B X = \Sigma_B,$$

where

$$\Sigma_A = \begin{bmatrix} C & 0 \\ 0 & 0 \end{bmatrix}, \quad \Sigma_B = \begin{bmatrix} S & 0 \\ 0 & 0 \\ 0 & I_{n-q} \end{bmatrix},$$

$C = \text{diag}(\alpha_1, \dots, \alpha_q)$, $S = \text{diag}(\beta_1, \dots, \beta_q)$, with $\alpha_i > 0$, $\beta_i > 0$, $\alpha_i^2 + \beta_i^2 = 1$, $i = 1, \dots, q$. The quantities

$$\sigma_i = \frac{\alpha_i}{\beta_i}, \quad i = 1, \dots, q,$$

are called the *nontrivial generalized singular values* of the pair (A, B) , ordered as $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_q > 0$. Let x_i denote the column of X associated with the pair (α_i, β_i) , and define $w_i = x_i / \beta_i$. Let u_i and v_i be the corresponding columns of U and V . Then the left and scaled right generalized singular vectors u_i and w_i satisfy

$$A w_i = \sigma_i u_i, \quad A^H u_i = \sigma_i B^H B w_i, \quad u_i^H u_j = \delta_{ij}, \quad w_i^H B^H B w_j = \delta_{ij}.$$

The triple (σ_i, u_i, w_i) is referred to as a *generalized singular triplet* of the pair (A, B) . The normalized GSVD components are recovered as

$$\beta_i = (1 + \sigma_i^2)^{-1/2}, \quad \alpha_i = \sigma_i \beta_i, \quad x_i = \beta_i w_i, \quad v_i = B w_i.$$

The nontrivial GSVD components of the matrix pair (A, B) can be characterized by two equivalent Hermitian generalized eigenvalue problems. In particular, the right generalized singular vectors in the original GSVD scaling x_i satisfy

$$A^H A x_i = B^H B x_i \sigma_i^2. \quad (1)$$

Alternatively, the generalized singular triplet (σ_i, u_i, w_i) satisfies the augmented generalized eigenvalue problem induced by the Jordan–Wielandt pencil:

$$\check{A} \begin{bmatrix} u_i \\ w_i \end{bmatrix} = \check{B} \begin{bmatrix} u_i \\ w_i \end{bmatrix} \sigma_i, \quad i = 1, \dots, q, \quad (2)$$

where

$$\check{A} = \begin{bmatrix} 0 & A \\ A^H & 0 \end{bmatrix}, \quad \check{B} = \begin{bmatrix} I_m & 0 \\ 0 & B^H B \end{bmatrix}.$$

Compared with (1), formulation (2) avoids explicitly forming $A^H A$ and avoids squaring the generalized singular values. It is therefore often preferable in terms of numerical stability, particularly when A is ill-conditioned.

For a positive generalized singular value, the corresponding triplet also satisfies the dual augmented relation

$$\begin{bmatrix} B^H & B \end{bmatrix} \begin{bmatrix} v_i \\ x_i/\alpha_i \end{bmatrix} = \begin{bmatrix} I & A^H A \end{bmatrix} \begin{bmatrix} v_i \\ x_i/\alpha_i \end{bmatrix} \sigma_i^{-1}.$$

This formulation requires $\alpha_i > 0$, ensuring that the scaling x_i/α_i is well defined. Moreover, the right-hand coefficient matrix $\text{diag}(I, A^H A)$ is positive definite if and only if A has full column rank. Under this additional assumption, the dual relation defines a Hermitian-definite generalized eigenvalue problem and provides a useful alternative, particularly when $B^H B$ is ill-conditioned.

The method developed below is based exclusively on (2). A key structured property of the augmented pencil is its spectral symmetry about the origin: if $\sigma_i \neq 0$ is an eigenvalue with eigenvector $[u_i^H, w_i^H]^H$, then $-\sigma_i$ is also an eigenvalue, with the corresponding eigenvector $[u_i^H, -w_i^H]^H$. These paired eigenvalue relations can be written compactly as

$$\begin{bmatrix} A & \\ A^H & \end{bmatrix} \begin{bmatrix} u_i & u_i \\ w_i & -w_i \end{bmatrix} = \begin{bmatrix} I & \\ & B^H B \end{bmatrix} \begin{bmatrix} u_i & u_i \\ w_i & -w_i \end{bmatrix} \begin{bmatrix} \sigma_i \\ -\sigma_i \end{bmatrix},$$

Thus, the nonzero eigenvalues occur in opposite-sign pairs $\pm\sigma_i$, with the associated eigenvectors exhibiting the paired structure shown above.

2.2 Trace variational characterization

Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k$ denote the k largest generalized singular values of the matrix pair (A, B) . Applying the Ky Fan trace variational principle to the associated Hermitian-definite pencil yields

$$\sum_{i=1}^k \sigma_i = \max_{Z^H \check{B} Z = I_k} \text{trace}(Z^H \check{A} Z).$$

The maximum is attained at

$$Z = \begin{bmatrix} U_{\max} \\ W_{\max} \end{bmatrix} \in \mathbb{C}^{(m+n) \times k},$$

where the columns of $\sqrt{2}U_{\max}$ and $\sqrt{2}W_{\max}$ are the left and right generalized singular vectors associated with $\Sigma_k = \text{diag}(\sigma_1, \dots, \sigma_k)$, respectively.

Similarly,

$$-\sum_{i=1}^k \sigma_i = \min_{Z^H \tilde{B} Z = I_k} \text{trace}(Z^H \tilde{A} Z).$$

The minimum is attained by

$$Z = \begin{bmatrix} U_{\min} \\ W_{\min} \end{bmatrix} \in \mathbb{C}^{(m+n) \times k},$$

whose columns are generalized eigenvectors associated with the k smallest eigenvalues $-\Sigma_k = \text{diag}(-\sigma_1, \dots, -\sigma_k)$. The blocks $\sqrt{2}U_{\min}$ and $-\sqrt{2}W_{\min}$ correspond to the left and right generalized singular vectors associated with the largest generalized singular values Σ_k .

Consequently, the partial GSVD of the matrix pair (A, B) associated with its k largest generalized singular values can be recovered equivalently from either of the two extremal trace problems for the associated Hermitian-definite pencil $(\tilde{A}, \tilde{B})$. Specifically, maximizing the trace of $Z^H \tilde{A} Z$ over k -dimensional $\tilde{B}$ -orthonormal subspaces yields the positive eigenvalue branch, whereas minimizing it yields the corresponding negative branch.

2.3 LOBPCG method

The LOBPCG method introduced by Knyazev [17] is a block iterative eigensolver for computing a few eigenpairs associated with either the smallest or the largest eigenvalues of standard or generalized Hermitian eigenvalue problems induced by a Hermitian-definite pencil (A_1, A_2) (i.e., $A_1^H = A_1 \in \mathbb{C}^{N \times N}$, $A_2^H = A_2 \in \mathbb{C}^{N \times N}$, and A_2 is positive definite). At iteration i , the next approximation $Z^{(i+1)}$ is sought in the $3k$ -dimensional space spanned by the current approximation $Z^{(i)}$, the (preconditioned) residual $F^{(i)}$, and the previous approximation $Z^{(i-1)}$, i.e., in $\text{span}(Z^{(i)}, Z^{(i-1)}, F^{(i)})$. Assuming that $V_{1,1}$ is nonsingular, replacing the previous approximation $Z^{(i-1)}$ by the stabilized search direction $P^{(i)} = Z^{(i)} - Z^{(i-1)}V_{1,1}$ preserves the search subspace unchanged, since $\text{span}(P^{(i)}, Z^{(i)}) = \text{span}(Z^{(i-1)}, Z^{(i)})$. This reformulation improves the numerical stability of the orthogonalization as the iterates converge [18, 19].

A key practical issue is to construct $P^{(i+1)}$ while maintaining a well-conditioned A_2 -orthonormal basis without repeatedly orthogonalizing large vectors. The IHL trick [19] achieves this by transferring the essential orthogonalization to a small coefficient matrix. For the full-rank case, assume that $[Z^{(i)}, Z_{\perp}^{(i)}]$ is an A_2 -orthonormal search basis of dimension $3k$, with its first k columns equal to $Z^{(i)}$. The coefficient

matrix produced by the Rayleigh–Ritz procedure is unitary and can be partitioned as

$$V = \begin{bmatrix} V_{1,1} & V_{1,2} \\ V_{2,1} & V_{2,2} \end{bmatrix}, \quad V_{1,1} \in \mathbb{C}^{k \times k}, \quad V_{1,2} \in \mathbb{C}^{k \times 2k},$$

where $V_{2,1} \in \mathbb{C}^{2k \times k}$ and $V_{2,2} \in \mathbb{C}^{2k \times 2k}$. Update $Z^{(i+1)}$ and $P^{(i+1)}$ as

$$Z^{(i+1)} = [Z^{(i)}, Z_{\perp}^{(i)}] \begin{bmatrix} V_{1,1} \\ V_{2,1} \end{bmatrix}, \quad P^{(i+1)} = Z^{(i+1)} - Z^{(i)}V_{1,1} = [Z^{(i)}, Z_{\perp}^{(i)}] \begin{bmatrix} 0 \\ V_{2,1} \end{bmatrix},$$

where $Z_{\perp}^{(i)} = [P^{(i)}, F^{(i)}]$. To obtain a direction orthogonal to $Z^{(i+1)}$, consider the orthogonal projection of $[0, V_{2,1}^H]^H$ onto the orthogonal complement of $\text{span}([V_{1,1}^H, V_{2,1}^H]^H)$:

$$\begin{bmatrix} 0 \\ V_{2,1} \end{bmatrix} - \begin{bmatrix} V_{1,1} \\ V_{2,1} \end{bmatrix} \begin{bmatrix} V_{1,1} \\ V_{2,1} \end{bmatrix}^H \begin{bmatrix} 0 \\ V_{2,1} \end{bmatrix} = \begin{bmatrix} V_{1,2} \\ V_{2,2} \end{bmatrix} \begin{bmatrix} V_{1,2} \\ V_{2,2} \end{bmatrix}^H \begin{bmatrix} 0 \\ V_{2,1} \end{bmatrix} = - \begin{bmatrix} V_{1,2} \\ V_{2,2} \end{bmatrix} V_{1,2}^H V_{1,1}.$$

Since the columns of $[V_{1,2}^H, V_{2,2}^H]^H$ are already orthonormal, the remaining task is reduced to constructing an orthonormal basis for the column space of $V_{1,2}^H$, which is a small $2k \times k$ matrix. Omitting the right factor $V_{1,1}$ does not alter the range of the projected history block when $V_{1,1}$ is nonsingular. Assuming also that $V_{1,2}$ has full row rank, let $V_{1,2} = LQ^H$ be a thin LQ decomposition, where $L \in \mathbb{C}^{k \times k}$ and $Q \in \mathbb{C}^{2k \times k}$ satisfies $Q^H Q = I_k$. The IHL update then defines

$$[Z^{(i+1)}, P^{(i+1)}] = [Z^{(i)}, Z_{\perp}^{(i)}] V \begin{bmatrix} I_k \\ Q \end{bmatrix}.$$

Hence, the IHL trick replaces the orthogonalization of large-scale vectors by an orthogonalization involving only the small matrix $V_{1,2}$, thereby reducing the orthogonalization cost while improving the numerical robustness of the LOBPCG method.

2.4 Structured Galerkin condition

We next recall the Galerkin condition associated with the augmented Hermitian-definite formulation of the GSVD. Owing to the spectral symmetry of the pencil $(\tilde{A}, \tilde{B})$, an approximation to a positive generalized singular value σ_i is naturally accompanied by an approximation to its negative counterpart. Accordingly, let $[\tilde{U}^H, \tilde{W}^H]^H$ and $[\tilde{U}^H, -\tilde{W}^H]^H$ with $\tilde{U} \in \mathbb{C}^{m \times k}$, $\tilde{W} \in \mathbb{C}^{n \times k}$ be approximate eigenvector blocks associated with the approximate eigenvalues $\tilde{\Sigma}$ and $-\tilde{\Sigma}$, respectively, where $\tilde{\Sigma} = \text{diag}(\tilde{\sigma}_1, \dots, \tilde{\sigma}_k)$ is positive diagonal. A structure-preserving approximation is obtained by enforcing the Galerkin condition simultaneously for both eigenvalues in each symmetric pair, namely,

$$\text{span} \left(\tilde{A} \begin{bmatrix} \tilde{U} & \tilde{U} \\ \tilde{W} & -\tilde{W} \end{bmatrix} - \tilde{B} \begin{bmatrix} \tilde{U} & \tilde{U} \\ \tilde{W} & -\tilde{W} \end{bmatrix} \begin{bmatrix} \tilde{\Sigma} & \\ & -\tilde{\Sigma} \end{bmatrix} \right) \perp \text{span} \left(\begin{bmatrix} \tilde{U} & \tilde{U} \\ \tilde{W} & -\tilde{W} \end{bmatrix} \right). \quad (3)$$

Following [15], we refer to (3) as the *structured Galerkin condition*, which is equivalent to

$$\text{Cassesspan}(A\tilde{W} - \tilde{U}\tilde{\Sigma}) \perp \text{span}(\tilde{U}), \quad \text{span}(A^H\tilde{U} - B^HB\tilde{W}\tilde{\Sigma}) \perp \text{span}(\tilde{W}).$$

Suppose further that the approximate left and right generalized singular vectors satisfy $\tilde{U}^H\tilde{U} = I_k$, and $\tilde{W}^HB^HB\tilde{W} = I_k$. Then (3) reduces to the *simplified Galerkin condition* in [13]:

$$\tilde{U}^HA\tilde{W} = \tilde{\Sigma}.$$

We next describe how to construct approximations satisfying the simplified Galerkin condition from prescribed left and right trial subspaces. Let $S_U \in \mathbb{C}^{m \times 3k}$ and $S_W \in \mathbb{C}^{n \times 3k}$ be bases for the search subspaces, normalized such that $S_U^HS_U = I_{3k}$ and $S_W^HS_W = I_{3k}$. We seek approximate left and right generalized singular vectors of the form $\tilde{U} = S_U V_U$ and $\tilde{W} = S_W V_W$, where $V_U, V_W \in \mathbb{C}^{3k \times k}$. Enforcing the residuals to be orthogonal to the corresponding trial subspaces gives

$$S_U^H(A\tilde{W} - \tilde{U}\tilde{\Sigma}) = 0, \quad S_W^H(A^H\tilde{U} - B^HB\tilde{W}\tilde{\Sigma}) = 0.$$

Defining the projected cross matrix $M = S_U^H A S_W$, the above conditions reduce to

$$M V_W = V_U \tilde{\Sigma}, \quad M^H V_U = V_W \tilde{\Sigma}.$$

Thus, V_U and V_W are left and right singular vectors, and $\tilde{\Sigma}$ contains the corresponding singular values of M . Consequently, the desired approximations $\tilde{U} = S_U V_U$ and $\tilde{W} = S_W V_W$ are obtained by computing the selected singular triplets of the small projected matrix M .

3 A structure-preserving LOBPCG method

We now develop an efficient structure-preserving LOBPCG method for computing the k largest nontrivial generalized singular values of the matrix pair (A, B) . The method is motivated by the Hermitian-definite pencil $(\check{A}, \check{B})$, where $\check{B}$ is positive definite because B has full column rank. Rather than explicitly assembling augmented vectors, the method works with separate left and right component spaces. Its central structural requirement is to preserve, throughout the iteration, the pairing between the components associated with the positive and negative eigenvalue branches.

The k largest positive approximate eigenpairs are denoted by $(\Theta, [U^H, W^H]^H)$. Their negative counterparts, $(-\Theta, [U^H, -W^H]^H)$, represent the k smallest approximate eigenpairs. The residuals for the eigenproblem (2) are

$$R_+ = \check{A} \begin{bmatrix} U \\ W \end{bmatrix} - \check{B} \begin{bmatrix} U \\ W \end{bmatrix} \Theta = \begin{bmatrix} R_U \\ R_W \end{bmatrix}, \quad R_- = \check{A} \begin{bmatrix} U \\ -W \end{bmatrix} + \check{B} \begin{bmatrix} U \\ -W \end{bmatrix} \Theta = \begin{bmatrix} -R_U \\ R_W \end{bmatrix},$$

where

$$R_U = A W - U \Theta, \quad R_W = A^H U - B^H B W \Theta.$$

The residuals R_+ and R_- inherit the block symmetry of the paired approximate eigenvectors, up to an irrelevant overall sign. Let T_+ and T_- be the preconditioners for R_+ and R_- , respectively. Let

$$\begin{bmatrix} F_U \\ F_W \end{bmatrix} = T_+ \begin{bmatrix} R_U \\ R_W \end{bmatrix}, \quad \begin{bmatrix} F_U \\ -F_W \end{bmatrix} = T_- \begin{bmatrix} -R_U \\ R_W \end{bmatrix},$$

where $T_- = -JT_+J$ with $J = \text{diag}(-I_m, I_n)$. The preconditioned residuals still inherit the block structure.

3.1 The IHL trick

At initialization, the history blocks are empty, $P_U = P_W = []$. Consequently, the first Rayleigh–Ritz extraction uses the $2k$ -column bases $[U, F_U]$ and $[W, F_W]$. This first extraction constructs P_U and P_W ; all subsequent extractions use the $3k$ -column bases $[U, P_U, F_U]$ and $[W, P_W, F_W]$. For clarity, the analysis below is written for this steady-state $3k$ -dimensional case. The first extraction follows from the same argument after replacing the dimensions $3k$ and $2k$ of the full basis and its complementary block by $2k$ and k , respectively.

Assume $3k \leq \min(m, n)$. Let

$$S_U = [U, P_U, F_U] \in \mathbb{C}^{m \times 3k}, \quad S_W = [W, P_W, F_W] \in \mathbb{C}^{n \times 3k}$$

be bases for the search subspaces so that $S_U^H S_U = I_{3k}$ and $S_W^H B^H B S_W = I_{3k}$. The simplified Galerkin condition leads to the following small SVD in the Rayleigh–Ritz procedure:

$$S_U^H A S_W = U^P \Theta (W^P)^H, \quad (4)$$

where $(U^P)^H U^P = I_{3k}$ and $(W^P)^H W^P = I_{3k}$, and $\Theta = \text{diag}(\theta_1, \dots, \theta_{3k}) \geq 0$.

At the next step, the updated basis matrices $[U, P_U]$ and $[W, P_W]$ are required to satisfy the orthonormality conditions $[U, P_U]^H [U, P_U] = I_{2k}$, and $[W, P_W]^H B^H B [W, P_W] = I_{2k}$. These conditions could be enforced by explicitly orthogonalizing P_U against U in the Euclidean inner product and P_W against W in the $B^H B$ -inner product, followed by appropriate normalization. Such operations, however, are performed on m - and n -dimensional vectors and may therefore be computationally expensive. We instead employ the IHL trick [18, 19] to construct P_U and P_W through low-dimensional coefficient transformations, while preserving the required orthogonality of the corresponding high-dimensional basis vectors. In the present GSVD setting, the left and right coefficient transformations must preserve the same number of columns and their columnwise correspondence, so that the i th left and right components can be paired to form a well-defined approximate eigenvector of the augmented problem.

Partition the unitary matrices $U^P, W^P \in \mathbb{C}^{3k \times 3k}$ obtained from (4) into 2×2 block form

$$U^P = [U_1^P, U_2^P] = \begin{bmatrix} U_{1,1}^P & U_{1,2}^P \\ U_{2,1}^P & U_{2,2}^P \end{bmatrix}, \quad W^P = [W_1^P, W_2^P] = \begin{bmatrix} W_{1,1}^P & W_{1,2}^P \\ W_{2,1}^P & W_{2,2}^P \end{bmatrix}.$$

Here,

$$U_1^p, W_1^p \in \mathbb{C}^{3k \times k}, \quad U_2^p, W_2^p \in \mathbb{C}^{3k \times 2k}, \quad U_{1,2}^p, W_{1,2}^p \in \mathbb{C}^{k \times 2k}.$$

The first k columns U_1^p and W_1^p determine the updated approximations

$$U = S_U U_1^p, \quad W = S_W W_1^p,$$

whereas U_2^p and W_2^p span their orthogonal complements. The projected relationship is

$$\begin{bmatrix} 0 \\ U_{2,1}^p \end{bmatrix} - \begin{bmatrix} U_{1,1}^p \\ U_{2,1}^p \end{bmatrix} \begin{bmatrix} U_{1,1}^p \\ U_{2,1}^p \end{bmatrix}^H \begin{bmatrix} 0 \\ U_{2,1}^p \end{bmatrix} = -U_2^p (U_{1,2}^p)^H U_{1,1}^p.$$

Similarly,

$$\begin{bmatrix} 0 \\ W_{2,1}^p \end{bmatrix} - \begin{bmatrix} W_{1,1}^p \\ W_{2,1}^p \end{bmatrix} \begin{bmatrix} W_{1,1}^p \\ W_{2,1}^p \end{bmatrix}^H \begin{bmatrix} 0 \\ W_{2,1}^p \end{bmatrix} = -W_2^p (W_{1,2}^p)^H W_{1,1}^p.$$

If $U_{1,1}^p$ is nonsingular, then

$$\text{span}(U_2^p (U_{1,2}^p)^H U_{1,1}^p) = \text{span}(U_2^p (U_{1,2}^p)^H).$$

Thus the right factor $U_{1,1}^p$ can be omitted when constructing a basis for the conjugate direction. Since U_2^p has orthonormal columns, only the small matrix $(U_{1,2}^p)^H$ needs to be orthogonalized. The same argument applies to the right component when $W_{1,1}^p$ is nonsingular. The necessary orthogonalization can therefore be carried out through small LQ factorizations of $U_{1,2}^p$ and $W_{1,2}^p$. Let $U_{1,2}^p = L_U Q_U^H$ and $W_{1,2}^p = L_W Q_W^H$ be the thin LQ factorizations, where $L_U, L_W \in \mathbb{C}^{k \times k}$ are nonsingular, and $Q_U, Q_W \in \mathbb{C}^{2k \times k}$ have orthonormal columns, i.e., $Q_U^H Q_U = I_k$, and $Q_W^H Q_W = I_k$. Let

$$Z_U = U_2^p Q_U, \quad Z_W = W_2^p Q_W,$$

The new P_U and P_W are then updated by

$$P_U = S_U Z_U, \quad P_W = S_W Z_W. \quad (5)$$

The IHL update (5) preserves the required orthogonality and normalization conditions: $P_U^H P_U = I_k$, $P_U^H U = 0$, $P_W^H B^H B P_W = I_k$, and $P_W^H B^H B W = 0$. The range-preservation argument follows directly from the projection identities above.

In finite precision, the paired blocks have equal numbers of columns, but the independent orthogonalizations in the left and right component spaces may reveal different numerical ranks. This mismatch can occur when F_U and F_W are orthogonalized in their respective inner products and when $(U_{1,2}^p)^H$ and $(W_{1,2}^p)^H$ are orthogonalized to form Q_U and Q_W . The approximation blocks U and W , by contrast, retain the same active column count by construction. We identify nearly dependent directions using a prescribed rank-revealing tolerance and retain only the paired directions whose column indices

are accepted on both sides. The same common-index rule is applied when constructing Q_U and Q_W . Thus, coordinated deflation preserves the one-to-one correspondence and equal column counts of all retained left and right directions. Such deflation is a finite-precision safeguard and is not commonly triggered in practice. To keep the notation and analysis concise, we assume throughout the remainder of this section that no deflation occurs; accordingly, all subsequent derivations use the nominal steady-state dimension $3k$. The implementation nevertheless applies the common-index rule above whenever a rank deficiency is detected. Algorithm 1 summarizes the LOBPCG method with the IHL trick.

Algorithm 1 LOBPCG–GSVD solver

Require: $A \in \mathbb{C}^{m \times n}$, $B \in \mathbb{C}^{p \times n}$ with $\text{rank}(B) = n$; $1 \leq l \leq k \leq \text{rank}(A)$, $3k \leq \min(m, n)$; initial U , W with $U^H U = W^H B^H B W = I_k$; a rule for applying T_+ ; maximum iteration count $N_{\max}$.

Ensure: The leading l approximate triplets.

- 1: Compute $U^H A W = U^P \Theta (W^P)^H$, with descending singular values.
- 2: $U \leftarrow U U^P$, $W \leftarrow W W^P$, $P_U \leftarrow []$, $P_W \leftarrow []$.
- 3: **for** $i = 1, 2, \dots, N_{\max}$ **do**
- 4: $R_U \leftarrow A W - U \Theta$, $R_W \leftarrow A^H U - B^H B W \Theta$.
- 5: **if** the leading l approximate triplets satisfy the stopping criterion **then**
- 6: **return** $\Theta(1 : l, 1 : l)$, $U(:, 1 : l)$, $W(:, 1 : l)$.
- 7: **end if**
- 8: Apply $[F_U^H, F_W^H]^H \leftarrow T_+[R_U^H, R_W^H]^H$.
- 9: Orthogonalize F_U against $[U, P_U]$, and $B^H B$ -orthogonalize F_W against $[W, P_W]$.
- 10: Remove unmatched columns from F_U and F_W to preserve their columnwise correspondence.
- 11: Form $S_U = [U, P_U, F_U]$, $S_W = [W, P_W, F_W]$.
- 12: Compute the SVD: $S_U^H A S_W = U^P \Theta (W^P)^H$, with descending singular values.
- 13: Apply rank-revealing orthogonalization to $(U_{1,2}^P)^H$ and $(W_{1,2}^P)^H$.
- 14: Form Q_U , Q_W using the common accepted indices.
- 15: $U \leftarrow S_U U_1^P$, $W \leftarrow S_W W_1^P$, $\Theta \leftarrow \Theta(1 : k, 1 : k)$.
- 16: $Z_U = U_2^P Q_U$, $Z_W = W_2^P Q_W$.
- 17: $P_U \leftarrow S_U Z_U$, $P_W \leftarrow S_W Z_W$.
- 18: **end for**
- 19: **return** $\Theta(1 : l, 1 : l)$, $U(:, 1 : l)$, $W(:, 1 : l)$.

3.2 Coefficient correspondence between two Rayleigh–Ritz procedures

Assume that $S_U^H S_U = I_{3k}$, and $S_W^H B^H B S_W = I_{3k}$. Then

$$S = \frac{1}{\sqrt{2}} \begin{bmatrix} S_U & S_U \\ S_W & -S_W \end{bmatrix} \in \mathbb{C}^{(m+n) \times 6k}$$

is $\check{B}$ -orthonormal, i.e., $S^H \check{B} S = I_{6k}$. The structured Galerkin condition with respect to the augmented search subspace $\text{span}(S)$ leads to the reduced eigenvalue problem

$$S^H \check{A} S V_+ = V_+ \Theta, \quad (6)$$

where $\Theta \in \mathbb{R}^{3k \times 3k}$ is diagonal with positive diagonal entries, and $V_+ \in \mathbb{C}^{6k \times 3k}$ has orthonormal columns $V_+^H V_+ = I_{3k}$. The corresponding approximate eigenvector block of $(\check{A}, \check{B})$ is $X_+ = S V_+$.

The IHL trick in Section 3.1 is formulated in terms of U^p and W^p , the coefficient matrices associated with the basis matrices S_U and S_W . Their leading columns update U and W , while their remaining columns are used to construct P_U and P_W . In contrast, the augmented projection (6) directly yields only V_+ , the coefficient matrix of $X_+ = S V_+$ with respect to the augmented basis S . This motivates recovering U^p and W^p from V_+ ; the required relationship is established in the following proposition.

Proposition 1 *Let S_U, S_W, S be as above, and suppose that $M = S_U^H A S_W$ is nonsingular. Let $V_+ = [(V_+^{\text{top}})^H, (V_+^{\text{bot}})^H]^H \in \mathbb{C}^{6k \times 3k}$ have orthonormal columns satisfying $S^H \check{A} S V_+ = V_+ \Theta$, where Θ is positive diagonal, V_+^{top} and V_+^{bot} have equal size. Then the matrices*

$$U^p = V_+^{\text{top}} + V_+^{\text{bot}}, \quad W^p = V_+^{\text{top}} - V_+^{\text{bot}} \quad (7)$$

are unitary matrices satisfying

$$M = U^p \Theta (W^p)^H. \quad (8)$$

Furthermore, define $V_- = [(V_+^{\text{bot}})^H, (V_+^{\text{top}})^H]^H$. Then $S^H \check{A} S V_- = -V_- \Theta$, $V_-^H V_- = I_{3k}$ and $V_+^H V_- = 0$. Consequently,

$$[X_+, X_-] = [S V_+, S V_-] = \frac{1}{\sqrt{2}} \begin{bmatrix} S_U U^p & S_U U^p \\ S_W W^p & -S_W W^p \end{bmatrix}.$$

Proof Let

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} I_{3k} & I_{3k} \\ I_{3k} & -I_{3k} \end{bmatrix}, \quad D = \begin{bmatrix} 0 & M \\ M^H & 0 \end{bmatrix}.$$

Then $H^H = H$, $H^2 = I_{6k}$, and $S^H \check{A} S = H D H$. By the definitions of U^p and W^p , applying H to V_+ gives

$$H V_+ = \frac{1}{\sqrt{2}} \begin{bmatrix} U^p \\ W^p \end{bmatrix}.$$

Multiplying $S^H \check{A} S V_+ = V_+ \Theta$ from the left by H therefore gives

$$D \begin{bmatrix} U^P \\ W^P \end{bmatrix} = \begin{bmatrix} U^P \\ W^P \end{bmatrix} \Theta.$$

Consequently, $MW^P = U^P \Theta$, and $M^H U^P = W^P \Theta$.

Set $G_U = (U^P)^H U^P$, $G_W = (W^P)^H W^P$. Since $V_+^H V_+ = I_{3k}$, we have $G_U + G_W = 2I_{3k}$. Moreover, the preceding singular-vector relations imply $G_U \Theta = \Theta G_W$. Eliminating G_W gives

$$(G_U - I_{3k})\Theta + \Theta(G_U - I_{3k}) = 0.$$

Since Θ is positive diagonal, this homogeneous Sylvester equation has only the trivial solution. Hence $G_U = G_W = I_{3k}$. Thus, U^P and W^P are unitary. It follows from $MW^P = U^P \Theta$ that $M = U^P \Theta (W^P)^H$, which proves (8). According to the definition of V_- , we obtain

$$H V_- = \frac{1}{\sqrt{2}} \begin{bmatrix} U^P \\ -W^P \end{bmatrix}.$$

Using the two singular-vector relations above, we obtain

$$D \begin{bmatrix} U^P \\ -W^P \end{bmatrix} = - \begin{bmatrix} U^P \\ -W^P \end{bmatrix} \Theta.$$

Therefore, $S^H \check{A} S V_- = H D H V_- = -V_- \Theta$. Furthermore,

$$V_-^H V_- = I_{3k}, \quad V_+^H V_- = \frac{1}{2}(G_U - G_W) = 0.$$

Finally, since

$$S = \begin{bmatrix} S_U & 0 \\ 0 & S_W \end{bmatrix} H,$$

direct substitution gives

$$S V_+ = \frac{1}{\sqrt{2}} \begin{bmatrix} S_U U^P \\ S_W W^P \end{bmatrix}, \quad S V_- = \frac{1}{\sqrt{2}} \begin{bmatrix} S_U U^P \\ -S_W W^P \end{bmatrix}.$$

Hence,

$$[S V_+, S V_-] = \frac{1}{\sqrt{2}} \begin{bmatrix} S_U U^P & S_U U^P \\ S_W W^P & -S_W W^P \end{bmatrix}. \quad \square$$

3.3 The IHL trick from the structured Galerkin condition

Algorithm 1 constructs the IHL coefficients from the projected SVD, where U^P and W^P are directly available. In the augmented Rayleigh–Ritz procedure, however, the projected eigenproblem provides V_+ and Θ instead. Proposition 1 allows us to recover U^P and W^P from V_+ . The partition and orthogonalization described in Section 3.1 can then be applied to obtain

$$Z_U = U_2^P Q_U, \quad Z_W = W_2^P Q_W.$$

Thus, the IHL update for the augmented procedure can be constructed through $P_U = S_U Z_U$, $P_W = S_W Z_W$, $U = S_U U_1^P$, and $W = S_W W_1^P$.

Under the full-row-rank assumptions of Section 3.1, define the paired augmented coefficient matrix

$$C_P = \frac{1}{2} \begin{bmatrix} Z_U + Z_W & Z_U - Z_W \\ Z_U - Z_W & Z_U + Z_W \end{bmatrix} \in \mathbb{C}^{6k \times 2k}.$$

The matrix P can be computed by

$$P = SC_P = \frac{1}{\sqrt{2}} \begin{bmatrix} P_U & P_U \\ P_W & -P_W \end{bmatrix}.$$

Thus the augmented conjugate direction can be represented and updated through its two components; neither S nor the large matrix P needs to be explicitly assembled.

To describe the orthogonality and history-preserving properties of this construction, define

$$X_{\text{old}} = \frac{1}{\sqrt{2}} \begin{bmatrix} U_{\text{old}} & U_{\text{old}} \\ W_{\text{old}} & -W_{\text{old}} \end{bmatrix}, \quad X = \frac{1}{\sqrt{2}} \begin{bmatrix} U & U \\ W & -W \end{bmatrix}.$$

Here, U_{old} and W_{old} denote the previous approximations stored in the first k columns of S_U and S_W , respectively. The componentwise orthogonality relations satisfied by U , W , P_U , and P_W yield

$$P^H \check{B} P = I_{2k}, \quad P^H \check{B} X = 0.$$

Furthermore, direct calculation gives $XX^H \check{B} = \text{diag}(UU^H, WW^H B^H B)$. Set

$$Y = (I_{m+n} - XX^H \check{B})X_{\text{old}}.$$

Since $X^H \check{B} X = I_{2k}$, the matrix $XX^H \check{B}$ is the $\check{B}$ -orthogonal projector onto $\text{span}(X)$. Define

$$Y_U = (I_m - UU^H)U_{\text{old}}, \quad Y_W = (I_n - WW^H B^H B)W_{\text{old}}.$$

Then

$$Y = \frac{1}{\sqrt{2}} \begin{bmatrix} Y_U & Y_U \\ Y_W & -Y_W \end{bmatrix}.$$

To verify the componentwise range preservation, let $E = [I_k, 0_{k \times 2k}]^H$, so that $U_{\text{old}} = S_U E$ and $W_{\text{old}} = S_W E$. Using the unitarity of $[U_1^P, U_2^P]$ and the thin LQ factorization $U_{1,2}^P = L_U Q_U^H$, we obtain

$$Y_U = S_U (I_{3k} - U_1^P (U_1^P)^H) E = S_U U_2^P (U_{1,2}^P)^H = P_U L_U^H.$$

Likewise, $Y_W = P_W L_W^H$. Since L_U and L_W are nonsingular, it follows that

$$\text{span}(P_U) = \text{span}(Y_U), \quad \text{span}(P_W) = \text{span}(Y_W).$$

Moreover, for

$$H_k = \frac{1}{\sqrt{2}} \begin{bmatrix} I_k & I_k \\ I_k & -I_k \end{bmatrix},$$

which is unitary, we have

$$P H_k = \text{diag}(P_U, P_W), \quad Y H_k = \text{diag}(Y_U, Y_W).$$

Right multiplication by H_k does not change the column space. The two component-wise range equalities therefore imply $\text{span}(P) = \text{span}(Y)$. It remains to compare the two enlarged subspaces. Let $C = X^H \tilde{B} X_{\text{old}}$. From the definition of Y , we get $Y = X_{\text{old}} - XC$, which shows that $\text{span}(X, Y) = \text{span}(X, X_{\text{old}})$. Therefore, using $\text{span}(P) = \text{span}(Y)$, we obtain $\text{span}(X, X_{\text{old}}) = \text{span}(X, P)$. Thus, with X retained in the search space, replacing X_{old} with P leaves the augmented search subspace unchanged.

4 Numerical experiments

All experiments were performed in MATLAB R2022b on a Linux server with two 16-core Intel Xeon Gold 6226R processors at 2.90 GHz and 1024 GB of main memory. All reported CPU times are total runtimes and include every stage of the corresponding algorithm. We use the twelve sparse matrices A from the SuiteSparse Matrix Collection (formerly the University of Florida Sparse Matrix Collection) [22]; the matrix names are listed in Table 1. Here $\text{nnz}(A)$ denotes the number of nonzero entries in A , and a superscript $\top$ on a matrix name indicates that the transpose of the collection matrix is used. In the GSVD experiments, the constraint matrices are

$$B_1 = \begin{bmatrix} 5 & 1 & & & \\ 1 & 5 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & & 1 & \\ & & & & 1 & 5 \end{bmatrix} \in \mathbb{C}^{n \times n}, \quad B_2 = \begin{bmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \end{bmatrix}^\top \in \mathbb{C}^{(n+1) \times n},$$

i.e. B_1 is the tridiagonal matrix with diagonal 5 and off-diagonal 1, and B_2 is the transpose of the first-order difference operator on $n + 1$ grid values, mapping $\mathbb{R}^{n+1}$ to $\mathbb{R}^n$ before transposition. Both B_1 and B_2 have full column rank.

Table 1 List of test matrices.

ID	Matrix A	m	n	$\text{nnz}(A)$	$\ A\ _2$
1	plat1919	1,919	1,919	32,399	2.93
2	rosen10 $^\top$	6,152	2,056	68,302	20,200
3	GL7d12	8,899	1,019	37,519	14.4
4	3elt_dual	9,000	9,000	26,556	3
5	fv1	9,604	9,604	85,264	4.52
6	shuttle_eddy	10,429	10,429	103,599	16.2
7	nopoly	10,774	10,774	70,842	23.3
8	flower_5_4 $^\top$	14,721	5,226	43,942	5.53
9	barth5	15,606	15,606	61,484	4.23
10	L-9	17,983	17,983	71,192	4
11	crack_dual	20,141	20,141	60,086	3
12	rel8	345,688	12,347	821,839	18.3

Unless otherwise stated, we seek the $l = 35$ largest generalized singular values and the corresponding generalized singular vectors, and the block size is set to $k = \lceil 3l/2 \rceil$.

The initial blocks $U^{(0)} \in \mathbb{C}^{m \times k}$, $W^{(0)} \in \mathbb{C}^{n \times k}$ are randomly generated. A triplet $(\hat{\sigma}_i, \hat{u}_i, \hat{w}_i)$ is considered converged when

$$\begin{aligned} \|A\hat{w}_i - \hat{u}_i\hat{\sigma}_i\|_2 &\leq \text{tol} (\|A\|_2 \|\hat{w}_i\|_2 + |\hat{\sigma}_i|), \\ \|A^H\hat{u}_i - B^HB\hat{w}_i\hat{\sigma}_i\|_2 &\leq \text{tol} (\|A\|_2 + |\hat{\sigma}_i| \|B\|_2^2 \|\hat{w}_i\|_2). \end{aligned}$$

Unless stated otherwise, we set $\text{tol} = 10^{-14}\sqrt{m}$. Equivalently, the iteration stops when the largest of the two scaled residuals over the leading l approximate triplets satisfies

$$\max_{1 \leq i \leq l} \max \left\{ \frac{\|A\hat{w}_i - \hat{u}_i\hat{\sigma}_i\|_2}{\|A\|_2 \|\hat{w}_i\|_2 + |\hat{\sigma}_i|}, \frac{\|A^H\hat{u}_i - B^HB\hat{w}_i\hat{\sigma}_i\|_2}{\|A\|_2 + |\hat{\sigma}_i| \|B\|_2^2 \|\hat{w}_i\|_2} \right\} \leq 10^{-14}\sqrt{m},$$

or when the iteration count reaches 3000. The spectral norms used in the stopping criteria are estimated in MATLAB as `normA = normest(A, 1e-3)` and `normB = normest(B, 1e-3)`. In all convergence-history plots, the vertical axis shows the maximum of the scaled residuals over the leading l approximate triplets at each iteration. The relations $U^HU = I$ and $W^HB^HBW = I$ are maintained by the orthogonalization and IHL updates in Algorithm 1, up to roundoff and the tolerance used in the rank-revealing orthogonalization.

4.1 Unpreconditioned GSVD computation with $B = B_1$

We compare Algorithm 1 with two augmented implementations. The structured augmented variant, `LOBPCG-eig`, uses the structured Galerkin condition (3). It recovers the component coefficients by (7) and applies the IHL construction in Section 3.3. The third implementation, `LOBPCG`, is an augmented eigensolver that does not exploit the block structure of $(\tilde{A}, \tilde{B})$. The search subspaces of all three algorithms contain information about the k largest and the k smallest approximate generalized eigenpairs. All algorithms use the same initial guess.

We compute the 35 largest generalized singular triplets of (A, B_1) without preconditioning. Figure 1 presents the convergence histories, while Table 2 reports the corresponding CPU time. For the matrices in Table 1, all tests reach the prescribed tolerance within the iteration limit except those for $(\mathbf{fv1}, B_1)$. Reference generalized eigenvalues computed with MATLAB's `eigs` function for the pencil (A^HA, B^HB) indicate that the gap between the 35-th and the 54-th largest generalized singular values is 7.3889×10^{-5} , while Algorithm 1 reduces the maximum scaled residual below 10^{-7} at iteration 2136.

Table 2 lists the CPU times of the three algorithms. Benefiting from the IHL trick and the efficient exploitation of the structured Galerkin condition, Algorithm 1 achieves a substantial CPU time advantage. In contrast, the standard LOBPCG method, which does not exploit the underlying matrix structure, incurs the highest computational cost.




Fig. 1 Convergence histories for GSVD experiments with $B = B_1$, without preconditioning.

Table 2 CPU times (seconds) for the GSVD experiments with $B = B_1$, without preconditioning. A dagger indicates that the prescribed tolerance was not reached within 3000 iterations.

ID	1	2	3	4	5	6
Algorithm 1	1.373	3.476	3.345	12.22	209 [†]	4.893
LOBPCG-eig	1.757	4.476	4.198	14.58	250.4 [†]	5.758
LOBPCG	3.854	15.73	16.93	77.63	1507 [†]	34.79

ID	7	8	9	10	11	12
Algorithm 1	14.4	9.972	26.45	33.78	33.5	263.5
LOBPCG-eig	16.97	11.58	30.78	38.6	39.04	267.7
LOBPCG	126.2	77.52	268.9	362.2	351.3	2498

4.2 Preconditioned GSVD computation with $B = B_2$

For the pair (A, B_2) , we compute the 35 largest generalized singular triplets. For all three LOBPCG variants, the preconditioner is set to $T_+ = (\check{A} - \mu\check{B})^{-1}$, with the following adaptive shift rule. Let $(\hat{\sigma}_i, \hat{u}_i, \hat{w}_i)$ be the i th largest approximate generalized singular triplet. The preconditioner is switched on once the relative residual of the largest approximate triplet falls below 0.1. As long as the largest triplet has not converged, set $\mu = 2\hat{\sigma}_1$; once the first i triplets have converged but the $(i+1)$ -st has not, set $\mu = \hat{\sigma}_i$.

Table 3 CPU times (seconds) for the GSVD experiments with $B = B_2$, using direct or inexact MINRES solves for preconditioning.

ID	1	2	3	4	5	6
Direct solves						
Algorithm 1	2.976	2.747	5.461	10.37	21.59	9.953
LOBPCG-eig	3.405	2.927	5.587	10.6	22.61	10.1
LOBPCG	4.032	4.986	8.107	20.48	36.26	19.55
Inexact MINRES solves						
Algorithm 1	17.38	10.13	8.538	43.27	243.6	81.97
LOBPCG-eig	18.24	10.6	8.846	44.15	259.4	85.56
LOBPCG	21.28	12.84	11.88	67.63	380.4	132.3
ID	7	8	9	10	11	12
Direct solves						
Algorithm 1	20.04	52.27	23.79	20.38	29.9	1719
LOBPCG-eig	20.53	52.51	24.15	20.64	30.07	1701
LOBPCG	30.71	63.58	41.14	40.39	59.21	1925
Inexact MINRES solves						
Algorithm 1	93.22	31.69	135.6	131.6	194.7	1679
LOBPCG-eig	92.98	32.04	137.8	133.2	196.4	1701
LOBPCG	137.2	49.95	229.7	238.2	372.8	2713

Table 3 reports the total runtimes of the experiments using either direct or inexact MINRES solves for preconditioning (stopping tolerance 10^{-2} , at most 30 iterations). Each MINRES solve uses MATLAB’s default zero initial vector. Figure 2 presents the corresponding convergence histories for the MINRES-based implementation. For these sparse test problems, the solution of the preconditioning systems accounts for a substantial portion of the total computational cost. This common overhead reduces the relative benefit of the less expensive structured Rayleigh–Ritz procedure in Algorithm 1, compared with that observed in the unpreconditioned experiments. Nevertheless, Algorithm 1 remains consistently faster than the unstructured LOBPCG method and is generally competitive with, or faster than, the LOBPCG-eig variant.

The MINRES-based approximation preserves the overall convergence behavior of the eigensolver and provides a practical alternative when sparse direct factorization is prohibitively expensive in terms of computation or memory consumption. It is, however, not uniformly faster than a direct solve in the present experiments, since its efficiency depends on the convergence of the inner iterations.




Fig. 2 Convergence histories for the GSVD experiments with $B = B_2$, using inexact MINRES solves for preconditioning.

4.3 Comparison with the Gram-based solver

In addition, we compare our algorithm with a Gram-matrix-based algorithm. Recall that the squared generalized singular values of (A, B) coincide with the generalized eigenvalues of the Hermitian-definite Gram-type pencil

$$(A^H A)x = \lambda(B^H B)x, \quad \lambda = \sigma^2.$$

Consequently, the l largest generalized singular values can alternatively be computed by applying LOBPCG in the $B^H B$ inner product, equipped with the IHL trick, to the matrix pair $(A^H A, B^H B)$, and then recovering $\sigma_i = \sqrt{\lambda_i}$ from the l largest generalized eigenvalues.

Let $n = 1000$, and let $\hat{U}, \hat{V} \in \mathbb{R}^{n \times n}$ be random orthogonal matrices. For a prescribed sequence $\{\sigma_i\}_{i=1}^n$, define

$$\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n), \quad A = \hat{U} \Sigma \hat{V}^T B_1, \quad B = B_1.$$

The generalized singular values of the matrix pair (A, B_1) are exactly $\sigma_1, \dots, \sigma_n$. This construction provides a genuine GSVD problem with a nonidentity constraint matrix while allowing the exact generalized singular values to be prescribed explicitly. We consider the following two cases.

Case 1 ($l = 5$). The generalized singular values are prescribed as

$$\sigma_i = 10^{2 - \frac{7}{4}(i-1)}, \quad i = 1, \dots, 5, \quad \sigma_i = 10^{-8}, \quad i > 5.$$

Thus, the five desired generalized singular values range from $\sigma_1 = 10^2$ to $\sigma_5 = 10^{-5}$, spanning seven orders of magnitude. In the Gram formulation, these values are squared, so the corresponding generalized eigenvalues range from 10^4 to 10^{-10} , whereas the unwanted eigenvalues are of order 10^{-16} . We compute the five largest generalized singular values using a block size of eight.

Case 2 ($l = 50$). The first five generalized singular values are prescribed by

$$\sigma_i = 10^{2 - \frac{1}{4}(i-1)}, \quad i = 1, \dots, 5, \quad \sigma_i = 10^{-4 - \frac{i-6}{44}}, \quad i = 6, \dots, 50, \quad \sigma_i = 10^{-7}, \quad i > 50.$$

Consequently, the target set contains five large generalized singular values in $[10, 10^2]$ and forty-five much smaller values in $[10^{-5}, 10^{-4}]$. After passing to the Gram pencil $(A^T A, B_1^T B_1)$, the two groups lie in $[10^2, 10^4]$ and $[10^{-10}, 10^{-8}]$, respectively, while the unwanted eigenvalues are of order 10^{-14} . We compute the fifty largest generalized singular values using a block size of sixty.

For both cases, the stopping tolerance is $10^{-14} \sqrt{n}$, the maximum number of updates is 50, and no preconditioner is employed. Both methods use the same starting right subspace, the same IHL update, and the same deflation and stopping criteria. For the Gram-based method, positive Ritz pairs can be converted to triplets using $\hat{\sigma}_i = \sqrt{\hat{\lambda}_i}$ and $\hat{u}_i = A \hat{w}_i / \hat{\sigma}_i$, with $\|B \hat{w}_i\|_2 = 1$; the two triplet residuals should then be evaluated in the original equations. The relative error of the i th value is $|\hat{\sigma}_i - \sigma_i| / \sigma_i$, and the maximum is taken over $1 \leq i \leq l$.

The numerical behavior is shown in Figure 3. In Case 1, Algorithm 1 converges in 32 iterations with a residual of 1.817×10^{-13} , whereas the Gram-based method fails to reach the prescribed tolerance within 50 iterations and attains only 1.061×10^{-10} . Their maximum relative errors in the computed generalized singular values are 1.128×10^{-11} and 1.083×10^{-5} , respectively. In Case 2, Algorithm 1 converges in 31 iterations with a residual of 1.686×10^{-13} , while the Gram-based method again reaches the iteration limit with a residual of 5.264×10^{-9} . The corresponding maximum relative errors are




Fig. 3 Convergence and accuracy comparison of Algorithm 1 and the Gram-based LOBPCG method for the GSVD with $B = B_1$. The top and bottom rows correspond to Cases 1 and 2, respectively. The plots in the left column show the convergence histories measured by the relative residuals, while those in the right column show the relative errors in the computed generalized singular values. In the right-column plots, index j denotes the j th largest desired generalized singular value σ_j .

1.681×10^{-11} and 1.732×10^{-2} . These results illustrate the loss of relative accuracy caused by explicitly forming the Gram matrices, particularly when very large and very small generalized singular values belong to the same target set.

5 Conclusions

In this paper, we developed a structure-preserving LOBPCG algorithm for computing several of the largest nontrivial generalized singular triplets of a large matrix pair (A, B) , where B has full column rank. By exploiting the positive-negative spectral symmetry of the associated Hermitian-definite Jordan–Wielandt pencil, the method represents the two spectral branches through paired left and right component spaces. When both component bases have full dimension, the resulting structured Galerkin condition replaces the Rayleigh–Ritz procedure on a $6k$ -dimensional augmented space by a complete SVD of the $3k \times 3k$ projected cross matrix $S_U^H A S_W$. The algorithm constructs its conjugate search directions by applying IHL to the left and right SVD coefficients, retaining the old approximation in the next search space. The established coefficient correspondence provides a simple way to construct the IHL update for the augmented LOBPCG algorithm from the projected eigenvector coefficients.

The present method assumes that B has full column rank and targets the largest nontrivial generalized singular values. Its convergence can deteriorate when the desired values are tightly clustered with the remaining spectrum, as observed for one of the

SuiteSparse test problems. Future work will therefore consider adaptive block-size and locking strategies for clustered spectra and extensions to rank-deficient matrix pairs and other parts of the GSVD spectrum.

Acknowledgements. The author thanks Meiyue Shao for helpful discussions.

- Funding: Funding not applicable
- Ethics approval and consent to participate: Not applicable
- Consent for publication: Not applicable
- Data availability: The data for this paper are available upon request.
- Materials availability: Not applicable
- Code availability: The code for this paper is available upon request.

Editorial Policies for:

Springer journals and proceedings: <https://www.springer.com/gp/editorial-policies>

Nature Portfolio journals: <https://www.nature.com/nature-research/editorial-policies>

Scientific Reports: <https://www.nature.com/srep/journal-policies/editorial-policies>

BMC journals: <https://www.biomedcentral.com/getpublished/editorial-policies>

References

- [1] Van Loan, C.F.: Generalizing the singular value decomposition. *SIAM J. Numer. Anal.* **13**(1), 76–83 (1976) <https://doi.org/10.1137/0713009>
- [2] Paige, C.C., Saunders, M.A.: Towards a generalized singular value decomposition. *SIAM J. Numer. Anal.* **18**(3), 398–405 (1981) <https://doi.org/10.1137/0718026>
- [3] Bai, Z., Zha, H.: A new preprocessing algorithm for the computation of the generalized singular value decomposition. *SIAM J. Sci. Comput.* **14**(4), 1007–1012 (1993) <https://doi.org/10.1137/0914060>
- [4] Park, H., Ewerbring, L.M.: An algorithm for the generalized singular value decomposition on massively parallel computers. *J. Parallel Distrib. Comput.* **17**(4), 267–276 (1993) <https://doi.org/10.1006/jpdc.1993.1026>
- [5] Betcke, T.: The generalized singular value decomposition and the method of particular solutions. *SIAM J. Sci. Comput.* **30**(3), 1278–1295 (2008) <https://doi.org/10.1137/060651057>
- [6] Hansen, P.C.: Regularization, GSVD and truncated GSVD. *BIT Numer. Math.* **29**, 491–504 (1989) <https://doi.org/10.1007/BF02219234>

- [7] Howland, P., Jeon, M., Park, H.: Structure preserving dimension reduction for clustered text data based on the generalized singular value decomposition. *SIAM J. Matrix Anal. Appl.* **25**(1), 165–179 (2003) <https://doi.org/10.1137/S0895479801393666>
- [8] Nakamura, K., Nakadai, K., Ince, G.: Real-time super-resolution sound source localization for robots. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 694–699 (2012). <https://doi.org/10.1109/IROS.2012.6385494> . <https://doi.org/10.1109/IROS.2012.6385494>
- [9] Bai, Z., Demmel, J.W.: Computing the generalized singular value decomposition. *SIAM J. Sci. Comput.* **14**(6), 1464–1486 (1993) <https://doi.org/10.1137/0914085>
- [10] Paige, C.C.: Computing the generalized singular value decomposition. *SIAM J. Sci. Comput.* **7**(4), 1126–1146 (1986) <https://doi.org/10.1137/0907077>
- [11] Sutton, B.D.: Stable computation of the CS decomposition: Simultaneous bidiagonalization. *SIAM J. Matrix Anal. Appl.* **33**(1), 1–21 (2012) <https://doi.org/10.1137/100813002>
- [12] Zha, H.: Computing the generalized singular values/vectors of large sparse or structured matrix pairs. *Numer. Math.* **72**, 391–417 (1996) <https://doi.org/10.1007/s002110050175>
- [13] Hochstenbach, M.E.: A Jacobi–Davidson type method for the generalized singular value problem. *Linear Algebra Appl.* **431**(3–4), 471–487 (2009) <https://doi.org/10.1016/j.laa.2009.03.003>
- [14] Huang, J., Jia, Z.: A cross-product free Jacobi–Davidson type method for computing a partial generalized singular value decomposition of a large matrix pair. *J. Sci. Comput.* **94**, 3 (2023) <https://doi.org/10.1007/s10915-022-02053-w>
- [15] Liu, Y., Shan, X., Shao, M.: A contour integral-based algorithm for computing generalized singular values. *J. Comput. Appl. Math.* **485**, 117583 (2026) <https://doi.org/10.1016/j.cam.2026.117583>
- [16] Deng, K.: Some problems in computing the pseudospectra of linear operators. Ph.d. thesis, University of Science and Technology of China, Hefei, China (2026)
- [17] Knyazev, A.V.: Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comput.* **23**(2), 517–541 (2001) <https://doi.org/10.1137/S1064827500366124>
- [18] Hetmaniuk, U., Lehoucq, R.: Basis selection in LOBPCG. *J. Comput. Phys.* **218**(1), 324–332 (2006) <https://doi.org/10.1016/j.jcp.2006.02.007>
- [19] Duersch, J.A., Shao, M., Yang, C., Gu, M.: A robust and efficient implementation

- of LOBPCG. SIAM J. Sci. Comput. **40**(5), 655–676 (2018) <https://doi.org/10.1137/17M1129830>
- [20] Shan, X., Shao, M.: A structure-preserving LOBPCG algorithm for the Bethe–Salpeter eigenvalue problem (2026). <https://doi.org/10.48550/arXiv.2603.07910> . <https://doi.org/10.48550/arXiv.2603.07910>
- [21] Huang, J., Jia, Z.: On choices of formulations of computing the generalized singular value decomposition of a large matrix pair. Numer. Algorithms **87**, 689–718 (2021) <https://doi.org/10.1007/s11075-020-00984-9>
- [22] Davis, T.A., Hu, Y.: The University of Florida Sparse Matrix Collection. ACM Trans. Math. Software **38**(1), 1–25 (2011) <https://doi.org/10.1145/2049662.2049663>