

THE EIGENVALUE DECOMPOSITION OF NORMAL MATRICES BY THE SKEW-SYMMETRIC PART

SIMON MATAIGNE* AND KYLE A. GALLIVAN†

Abstract. We propose a new method for computing the eigenvalue decomposition of a dense real normal matrix A through the decomposition of its skew-symmetric part. The method relies on algorithms that are known to be efficiently implemented, such as the bidiagonal singular value decomposition and the symmetric eigenvalue decomposition. The advantages of this method stand for normal matrices with few real eigenvalues, such as random orthogonal matrices. We provide a stability and a complexity analysis of the method. The numerical performance is compared with existing algorithms. In most cases, the method has the same operation count as the Hessenberg factorization of a dense matrix. Finally, we provide experiments for the application of computing a Riemannian barycenter on the special orthogonal group.

Key words. Normal matrices, eigenvalue decomposition, orthogonal matrices, Schur decomposition, matrix logarithm, Riemannian barycenter, Karcher mean.

AMS subject classifications. 65F15, 15B10, 15B57, 62R30.

1. Introduction. Computing the eigenvalues and the eigenvectors of a matrix is one of the most fundamental problems in linear algebra. It is well-known that certain classes of matrices allow fast algorithms with convergence and accuracy guarantees. For example, Hermitian matrices ($A = \overline{A}^\top =: A^*$) enjoy a wide scope of applications and it is known that the QR algorithm [38, 39] converges globally [79, 80, 81]. Additionally, for Hermitian matrices, a divide and conquer method and a multisectioning method were designed [24, 31, 56]. Robust symmetric eigensolvers are available in **LAPACK** [6, 30]. All these methods consist of reducing the Hermitian matrix to a tridiagonal form under unitary ($UU^* = U^*U = I$) similarity transformations [58, 59, 63, 78], and then solving the tridiagonal eigenvalue problem (EVP). For general matrices, the QR algorithm remains a state of the art approach to obtain the Schur form (upper triangular matrix with eigenvalues on the diagonal).

Although not nearly as popular as the symmetric EVP, the real skew-symmetric EVP ($A = -A^\top$) is known to be even more efficiently solvable [60, 75] because, first, skew-symmetric matrices can also be reduced to a tridiagonal form [52, 75]. Secondly, the zero diagonal makes the EVP equivalent to a bidiagonal singular value problem with half the size of the original EVP. It can be solved both rapidly and accurately [27, 28, 36]. This procedure exploits that eigenvalues arise by pairs of conjugate purely imaginary numbers. Unfortunately, in the complex setting, skew-Hermitian matrices turn into Hermitian matrices when multiplied by the imaginary unit. This strategy cannot be extended; real-valued matrices are more structured.

The focus of this paper, normal matrices ($AA^* = A^*A$), includes the two previous classes. It was shown that the shifted QR algorithm converges for normal matrices [10]. However, normal matrices are

*ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium. (simon.mataigne@uclouvain.be). Simon Mataigne is a Research Fellow of the Fonds de la Recherche Scientifique - FNRS. This work was supported by the Fonds de la Recherche Scientifique - FNRS under Grant no T.0001.23.

†Department of Mathematics, Florida State University, Tallahassee, FL, USA. (kgallivan@fsu.edu). This work was initiated while Kyle A. Gallivan was a Visiting Professor at the ICTEAM Institute, UCLouvain.

diagonalizable under unitary similarity transformations, i.e., the Schur form is diagonal. Consequently, the costly QR iterations operate on a dense upper Hessenberg matrix while converging to a diagonal matrix without exploiting this property. The key observation for a better method is the following: the Hermitian part $H := \frac{A+A^*}{2}$ and the skew-Hermitian part $\Omega := \frac{A-A^*}{2}$ of a normal matrix A commute: $H\Omega = \Omega H$. It is known that two normal matrices H and Ω that commute are jointly diagonalizable, i.e., there exists a unitary matrix U such that $U^* H U = \Lambda_H$ and $U^* \Omega U = \Lambda_\Omega$ [51, Thm. 2.5.5], and thus $A = U(\Lambda_H + \Lambda_\Omega)U^*$. The latter decomposition can be obtained by the *simultaneous diagonalization* of two Hermitian matrices (here H and $i\Omega$) using the Jacobi method [19, 40]. These algorithms were notably developed for generalized EVPs [23, 77]. Another approach, called the *do-one-then-do-the-other* (DODO), considers the sequential diagonalization of H and Ω [66]. Moreover, a method leveraging random combinations of H and $i\Omega$ has been proposed recently [48]. Finally, the normal EVP can be transformed into a complex symmetric EVP, more details can be found in [9, 37, 71]. However, all previously cited methods consider *complex* matrices. For real matrices, a Jacobi method [83] was designed. The aim of this paper is to show that the additional structure of real matrices enables a specific and computationally efficient approach.

Unitary matrices, called *orthogonal matrices* when real-valued, are an important subset of normal matrices. A rich literature has emerged from the representation of $n \times n$ unitary Hessenberg matrices in their *Schur parametric form*, i.e., as the product of n elementary Householder matrices (or equivalently as Givens rotations):

$$H = H_1 H_2 \dots H_n \text{ where}$$

$$H_k = \text{diag} \left[I_{k-1}, \begin{bmatrix} -\alpha_k & \beta_k \\ \beta_k & \alpha_k^* \end{bmatrix}, I_{n-k-3} \right] \text{ with } |\alpha_k|^2 + |\beta_k|^2 = 1 \text{ for } k = 1, 2, \dots, n-1,$$

and $H_n = \text{diag} [I_{n-1}, -\alpha_n]$ with $|\alpha_n| = 1$. This compact storage format allows cheap and stable QR iterations, called UHQR, requiring only $\mathcal{O}(n)$ flops per iteration when only the eigenvalues are desired [44, 65]. This method can be implemented as a bulge-chasing algorithm [8] and it was shown that appropriate shifts produce quadratic convergence [73, 74]. The success of the Schur parametric form also led to the development of alternative algorithms such as divide-and-conquer [43, 46] and SVD-based methods [4, 21]. The common step for all previously cited methods is the initial reduction to Hessenberg form. Our method avoids this computationally expensive step.

Contributions. In this paper, we show that real normal matrices, with orthogonal matrices as a special case, admit a fast factorization to the eigenvalue decomposition (EVD) by exploiting the structure of their invariant spaces. This structure only exists in the real setting. In most cases,¹ if a normal matrix $A \in \mathbb{R}^{n \times n}$ has r real eigenvalues (r does not need to be known in advance), its EVD can be obtained by solving the rank $n - r$ EVP of the skew-symmetric part of A , plus a $r \times r$ symmetric EVP. The method is designed for matrices where $r \ll n$, a frequent property of random orthogonal matrices (e.g., Haar-distributed). We provide a stability and a complexity analysis of the method and compare its numerical performance with existing algorithms. When all eigenvalues have distinct imaginary parts, i.e., well-separated in floating point arithmetic, the method has the same operation count as the Hessenberg factorization of a dense matrix. When the imaginary parts of the eigenvalues are clustered, we propose refinement steps based on theoretical bounds that allow to maintain a user-specified accuracy of the decomposition. We corroborate our results in a suite of experiments. The method facilitates efficiently computing the matrix logarithm of orthogonal

¹The precise meaning of “most cases” is described in [section 6](#).

matrices, the main computational bottleneck for calculating Riemannian logarithms and barycenters on matrix manifolds. In particular in these applications, the EVP is an intermediate computation used in iterative algorithms where speed is more important than accuracy to machine precision. In addition, the normal EVP also finds applications in generalized EVPs [20, 23] and computation of Pisarenko frequency estimates [5, 25].

Overview. Section 2 introduces the essential properties of the normal EVP and related concepts. Section 3 presents theoretical results that lead to the exact arithmetic formulation of our method for the normal EVP: Algorithm 3.1. In section 4, we analyze the impact of floating-point arithmetic through a sensitivity analysis of the invariant spaces of the skew-symmetric part with respect to perturbations of the original matrix; this yields Algorithm 4.1. Section 5 corroborates the perturbation analysis through a series of experiments. In section 6, we examine the computational complexity of the algorithm and validate the analysis with timing experiments. Finally, in section 7, we apply Algorithm 4.1 to enhance the performance of Riemannian gradient descent for computing the Riemannian barycenter on $\text{SO}(n)$.

Reproducibility. An implementation of the method, as well as the codes to perform all experiments presented in this paper, are available at <https://github.com/smataigne/NormalEVP.jl>.

Notation. We write $A = [a_{ij}]_{i,j=1}^n$ to denote matrix entries. We use $\text{sym}(A) := \frac{A+A^\top}{2}$ and $\text{skew}(A) := \frac{A-A^\top}{2}$ for, respectively, the symmetric part and the skew-symmetric part of the matrix A . $A_{a:b}$ denotes the selection of *columns* indexed by a to b included and $A_{a:b,c:d}$ refers to the submatrix of *rows* a to b and *columns* c to d included. The Frobenius norm and the spectral norm are denoted $\|\cdot\|_F$ and $\|\cdot\|_2$, respectively. The identity matrix of size $n \times n$ is denoted I_n . Since orthogonal transformations are central to normal EVPs, we recall the notation for the orthogonal group $\text{O}(n)$ defined as

$$\text{O}(n) := \{Q \in \mathbb{R}^{n \times n} \mid Q^\top Q = QQ^\top = I_n\}.$$

$\text{O}(n)$ admits two connected components: matrices with determinant 1 and matrices with determinant -1 , sometimes respectively referred to as the set of rotations and reflections. Rotations are best known as the special orthogonal group $\text{SO}(n) := \{Q \in \text{O}(n) \mid \det(Q) = 1\}$. Finally, the set of $n \times p$ matrices, $n \geq p$, with orthonormal columns is called the Stiefel manifold $\text{St}(n, p)$, namely

$$\text{St}(n, p) := \{V \in \mathbb{R}^{n \times p} \mid V^\top V = I_p\}.$$

In the special case $n = p$, we have $\text{St}(n, n) = \text{O}(n)$.

2. Preliminaries.

2.1. The normal eigenvalue problem. In the context of *real* normal matrices, eigenvalues and eigenvectors are, in general, complex-valued. Nonetheless, it is preferred for most eigenvalue algorithms to compute with real arithmetic. For a normal matrix A , obtaining the real Schur decomposition (RSD), i.e., $A = QSQ^\top$ with $Q \in \text{O}(n)$, is equivalent to obtaining the EVD. Indeed, the real Schur form S is block-diagonal with 2×2 and/or 1×1 blocks [51, Thm 2.5.8]. Each 2×2 block has the form $\lambda \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$, $\lambda > 0$, $\theta \in (0, \pi)$, and is associated with the conjugate pair of eigenvalues $\{\lambda e^{\pm i\theta}\}$. The transition from the

RSD to the EVD uses for each 2×2 block the relation

$$(2.1) \quad \lambda \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix} \begin{bmatrix} \lambda e^{-i\theta} & 0 \\ 0 & \lambda e^{i\theta} \end{bmatrix} \begin{bmatrix} 1 & -i \\ 1 & i \end{bmatrix}.$$

In this paper, we consider a permuted version of the real Schur form S , introduced in [Property 2.1](#). The latter permuted version allows writing the matrix transformations easily and succinctly. Therefore, we will refer to [\(2.2\)](#) as the “real Schur decomposition”, although it differs from the usual block-diagonal definition by a permutation.

PROPERTY 2.1 ([51, Thm 2.5.8]). *Every normal matrix $A \in \mathbb{R}^{n \times n}$ is similar to a real Schur form $S \in \mathbb{R}^{n \times n}$ under orthogonal transformation by a matrix $Q \in O(n)$ such that*

$$(2.2) \quad A = QSQ^\top = Q \begin{bmatrix} \Lambda \cos(\Theta) & 0 & -\Lambda \sin(\Theta) \\ 0 & \check{\Lambda} & 0 \\ \Lambda \sin(\Theta) & 0 & \Lambda \cos(\Theta) \end{bmatrix} Q^\top,$$

where $\check{\Lambda} = \text{diag}(\check{\lambda}_1, \dots, \check{\lambda}_r)$ contains the r real eigenvalues of A . Let $p := \frac{n-r}{2}$,² then $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_p)$ and $\Theta = \text{diag}(\theta_1, \dots, \theta_p)$ are real, can be assumed positive with $\theta_j \in (0, \pi)$ w.l.o.g. and are associated with the pairs of eigenvalues $\{\lambda_j e^{\pm i\theta_j}\}$ for $j = 1, \dots, p$. The columns of the matrix Q are called Schur vectors. *Special case 1: if A is symmetric, $r = n$. Special case 2: if A is skew-symmetric, $\cos(\Theta) = 0$ and $\check{\Lambda} = 0$, [51, Cor. 2.5.11].*

REMARK 2.2. The Schur vectors of A associated with real eigenvalues are also eigenvectors of A . Moreover, in view of [\(2.1\)](#), if $q_1, q_2 \in \text{St}(n, 1)$ is a pair of orthogonal Schur vectors associated with the same (λ, θ) -block, then $\text{span}([q_1 \ q_2])$ is an invariant subspace of A and $\frac{\sqrt{2}}{2}(q_1 + iq_2)$ and $\frac{\sqrt{2}}{2}(q_1 - iq_2)$ are eigenvectors of A .

2.2. The orthogonal groups. The Lie groups $O(n)$ and $SO(n)$ are particularly important subsets of normal matrices for applications [1, 22, 53]. Computing the RSD is a key primitive to obtain the matrix logarithm, which is required for computing Riemannian logarithms [57, 61, 67, 84, 85] and Riemannian barycenters [15, 16, 54, 55, 82, 85], notably on the Stiefel manifold. The principal matrix logarithm can be efficiently computed via the relation

$$(2.3) \quad \log(A) = Q \log(S) Q^\top = Q \begin{bmatrix} \log(\Lambda) & 0 & -\Theta \\ 0 & \log(\check{\Lambda}) & 0 \\ \Theta & 0 & \log(\Lambda) \end{bmatrix} Q^\top.$$

For both $O(n)$ and $SO(n)$, the property that eigenvalues lie on the unit circle yields $\Lambda = I_p$ and $\check{\Lambda} = \text{diag}(\pm 1, \dots, \pm 1)$. In particular, we have $\log(\Lambda) = 0$. Note that $\log(\check{\Lambda})$ admits a real element if and only if negative eigenvalues arise by pairs since $\begin{bmatrix} 0 & -\pi \\ \pi & 0 \end{bmatrix} \in \log\left(\begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}\right)$.

2.3. Invariance groups. The decomposition from [\(2.2\)](#) is often not unique. It is key to this paper to describe the admissible transformations of the Schur vectors Q such that [\(2.2\)](#) remains satisfied. To this end, we introduce the notion of the *invariance group* of a matrix in [Definition 2.3](#).

DEFINITION 2.3. *Given a matrix $A \in \mathbb{R}^{n \times n}$, the invariance group of A , written $\text{ig}(A)$, is defined as the set*

$$(2.4) \quad \text{ig}(A) := \{Q \in O(n) \mid QAQ^\top = A\}.$$

²Note that p is always an integer since r is odd if and only if n is odd.

An equivalent definition of $\text{ig}(A)$ is the set of orthogonal matrices that commute with A . It is easily verified that $\text{ig}(A)$ is a group under matrix multiplication since for all $Q_1, Q_2 \in \text{ig}(A)$, we have $Q_1 Q_2 \in \text{ig}(A)$. The set $\text{ig}\left(\begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix}\right)$ is well-known as the *ortho-symplectic group* $\text{OSp}(2m)$, see, e.g., [32].

Let A be a normal matrix and $A = QSQ^\top$ be a RSD. For every $R \in \text{ig}(S)$, we have

$$A = (QR)(R^\top SR)(QR)^\top = (QR)S(QR)^\top.$$

Therefore, QR is also a matrix of Schur vectors. If $r < n$ (A is not symmetric) or if $\check{\lambda}$ has repeated diagonal entries, then $\text{ig}(S)$ is uncountable. Indeed, for each 2×2 (λ, θ) -block of S and each real eigenvalue $\check{\lambda}$ of multiplicity $m > 1$, we respectively have

$$\text{ig}\left(\lambda \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}\right) = \text{SO}(2) \quad \text{and,} \quad \text{ig}(\check{\lambda} I_m) = \text{O}(m).$$

Finally, we highlight a property of invariance groups that is important in [section 3](#).

LEMMA 2.4. *For all pairs of matrices $A, B \in \mathbb{R}^{n \times n}$, we have*

$$\text{ig}(A) \cap \text{ig}(B) \subseteq \text{ig}(A + B).$$

Proof. Let $Q \in \text{ig}(A) \cap \text{ig}(B)$. Therefore, we have $Q A Q^\top = A$ and $Q B Q^\top = B$, and by summation, $Q(A + B)Q^\top = A + B$. Hence, $Q \in \text{ig}(A + B)$. $\square$

2.4. The even-odd permutation. To conclude the preliminaries, we introduce the notions of even-odd permutation and bidiagonalization of a tridiagonal skew-symmetric matrix. These notions are essential to the skew-symmetric EVP and are used in [section 3](#).

DEFINITION 2.5. *An even-odd permutation matrix $P_{\text{eo}} \in \mathbb{R}^{n \times n}$ is a permutation such that for every vector $[a_1 \ a_2 \ \dots \ a_n] \in \mathbb{R}^{1 \times n}$, we have*

$$[a_1 \ a_2 \ \dots \ a_n] P_{\text{eo}} = [a_1 \ a_3 \ \dots \ a_{\lceil \frac{n}{2} \rceil} \ a_2 \ a_4 \ \dots \ a_{\lfloor \frac{n}{2} \rfloor}].$$

The even-odd permutation is well-known in the field of discrete mathematics and it is related to the notion of Red-Black ordering [68]. This permutation allows the reduction of the skew-symmetric EVP to a bidiagonal SVD problem by a similarity transformation, described in [Property 2.6](#).

PROPERTY 2.6 ([75]). *Let $T \in \mathbb{R}^{n \times n}$ be a skew-symmetric tridiagonal matrix and $P_{\text{eo}} \in \mathbb{R}^{n \times n}$ be an even-odd permutation matrix, then*

$$(2.5) \quad P_{\text{eo}}^\top T P_{\text{eo}} = \begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix},$$

where $\tilde{B} \in \mathbb{R}^{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil}$ is bidiagonal.

3. The eigenvalue decomposition of normal matrices. In this section, we describe our method in exact arithmetic for computing the RSD of a normal matrix A . We show how the skew-symmetric part of a normal matrix reveals the structure of its Schur vectors and how these computations can be organized to design a fast RSD/EVD algorithm. In [section 4](#), we analyze the effect of floating point arithmetic. In particular, the method relies on the assumption that the three following routines are available:

- Routine 1: returns the SVD of a real bidiagonal matrix.
- Routine 2: returns the RSD/EVD of a real symmetric matrix.
- Routine 3: returns the RSD of a (small) real normal matrix.

In the context of exact arithmetic, these routines can be viewed as oracles, whereas in practice, they are efficiently implemented, for example, in **LAPACK**. Indeed, **Routine 1** and **Routine 2** are respectively available as **bdsqr/bdsdc** and **syev/syevd/syevr** routines. For the coherence of this discussion, the calls to **Routine 3** should be exceptional and only applied on small-size problems. **Routine 3** may be chosen to be **LAPACK's gees** routine since the shifted QR algorithm is known to converge for normal matrices [10]. **Routine 3** may also be any of the previously cited methods for normal or orthogonal EVPs, e.g., [4, 74, 83]. All of these **LAPACK's** routines are stable, see [6].

3.1. The skew-symmetric eigenvalue problem. The method starts by obtaining the RSD of $\Omega := \text{skew}(A)$. It is known to be efficiently done in three steps [75].

First, we compute $\Omega = \tilde{Q}T\tilde{Q}^\top$ with $\tilde{Q} \in O(n)$ and T skew-symmetric and tridiagonal. This decomposition can be obtained in at least two ways, both adapted versions of symmetric tridiagonalization. Either by using Householder reflectors [78], see, e.g., **SkewLinearAlgebra.jl**, or by using the Lanczos algorithm [52]. Householder tridiagonalization is preferred for dense matrices for its stability and use of level 3 BLAS operations [17, 62, 42]. Lanczos algorithm finds applications for large sparse matrices.

Secondly, an even-odd permutation matrix P_{eo} yields

$$\Omega = \tilde{Q}T\tilde{Q}^\top = \tilde{Q}P_{\text{eo}} \begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix} P_{\text{eo}}^\top \tilde{Q}^\top,$$

where $\tilde{B} \in \mathbb{R}^{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil}$ is bidiagonal. In the context of exact arithmetic, a transformation $G \in O(n)$ can isolate the central square matrix 0_r in (3.1). G is obtained by a finite process described in <https://github.com/smataigne/NormalEVP.jl>. In practice, we compute the SVD of $\tilde{B}$ directly. For theoretical purposes, we consider the transformation G in this section. This leads to

$$(3.1) \quad \Omega = \tilde{Q}P_{\text{eo}}G \begin{bmatrix} 0 & 0 & -B^\top \\ 0 & 0_r & 0 \\ B & 0 & 0 \end{bmatrix} G^\top P_{\text{eo}}^\top \tilde{Q}^\top,$$

where $B \in \mathbb{R}^{p \times p}$ is bidiagonal.

The third and final step is a call to **Routine 1** to compute the SVD of the bidiagonal block $B = U\Sigma V^\top$. This yields

$$(3.2) \quad \Omega = \tilde{Q}P_{\text{eo}}G \begin{bmatrix} V & 0 & 0 \\ 0 & I_r & 0 \\ 0 & 0 & U \end{bmatrix} \begin{bmatrix} 0 & 0 & -\Sigma \\ 0 & 0_r & 0 \\ \Sigma & 0 & 0 \end{bmatrix} \begin{bmatrix} V^\top & 0 & 0 \\ 0 & I_r & 0 \\ 0 & 0 & U^\top \end{bmatrix} G^\top P_{\text{eo}}^\top \tilde{Q}^\top.$$

By defining the matrix $\hat{Q} := \tilde{Q}P_{\text{eo}}G \begin{bmatrix} V & 0 & 0 \\ 0 & I_r & 0 \\ 0 & 0 & U \end{bmatrix}$, the decomposition (3.2) becomes a RSD of Ω . This decomposition can be compared with the skew-symmetric part of the (unknown) RSD of $A = QSQ^\top$ from (2.2):

$$(3.3) \quad \Omega = Q\text{skew}(S)Q^\top = Q \begin{bmatrix} 0 & 0 & -\Lambda \sin(\Theta) \\ 0 & 0_r & 0 \\ \Lambda \sin(\Theta) & 0 & 0 \end{bmatrix} Q^\top.$$

By comparing (3.2) and (3.3), it must hold that $\Sigma = \Lambda \sin(\Theta)$ if we assume that both are sorted, e.g., in decreasing order. Unfortunately, it does not hold true for the Schur vectors: $Q \neq \widehat{Q}$ in general. However, by equating (3.2) and (3.3), it holds that

$$\widehat{Q}^\top Q \text{skew}(S) Q^\top \widehat{Q} = \text{skew}(S) \implies \widehat{Q}^\top Q \in \text{ig}(\text{skew}(S)).$$

We recall that $\widehat{Q}$ is a matrix of Schur vectors of A if and only if $\widehat{Q}^\top Q \in \text{ig}(S)$. We have thus not yet computed the Schur vectors of A since, in general, $\text{ig}(\text{skew}(S)) \not\subseteq \text{ig}(S)$. Designing the next steps of the algorithm requires finding the relation between $\text{ig}(\text{skew}(S))$ and $\text{ig}(S)$. We show in the next subsection that it depends on the number r of real eigenvalues and the diagonal entries of $\Lambda \sin(\Theta)$.

3.2. Invariant spaces of the skew-symmetric part: preparatory lemmas. The rest of section 3 is dedicated to the description of the relation between the invariance groups $\text{ig}(\text{skew}(S))$ and $\text{ig}(S)$. By completing this analysis, we will be able to find a matrix $R \in O(n)$ such that $(\widehat{Q}R)^\top Q \in \text{ig}(S)$, and thus $A = (\widehat{Q}R)S(\widehat{Q}R)^\top$. Let us start by dividing the eigenvalues of A in three sets:

- Set 1: The eigenvalues with distinct nonzero imaginary parts.
- Set 2: The eigenvalues with repeated nonzero imaginary parts.
- Set 3: The real eigenvalues.

A preliminary result is showing that R features a block structure related to the aforementioned sets of eigenvalues of A . This is done in Lemmas 3.1 and 3.2. To this end, we partition Q and $\widehat{Q}$ column-wise to match the blocks of the decomposition (2.2):

$$Q := [Q_{c,1} \ Q_r \ Q_{c,2}] \quad \text{and} \quad \widehat{Q} := [\widehat{Q}_{c,1} \ \widehat{Q}_r \ \widehat{Q}_{c,2}],$$

where the notation follows because, by definition, Q_r contains the Schur vectors associated with the *real* eigenvalues and $[Q_{c,1} \ Q_{c,2}]$ contains the Schur vectors associated with *complex* eigenvalues. Lemma 3.1 shows the equivalence between the real eigenspace of A and the nullspace of $\text{skew}(A)$.

LEMMA 3.1. *Let $A \in \mathbb{R}^{n \times n}$ be a normal matrix with exactly $r > 0$ real eigenvalues and $\check{\Lambda} \in \mathbb{R}^{r \times r}$ be a diagonal matrix with these eigenvalues as diagonal entries. If a matrix $V \in \text{St}(n, r)$ satisfies $AV = V\check{\Lambda}$, then $\text{skew}(A)V = 0$. Conversely, if $\text{skew}(A)V = 0$, then there exists a matrix $\check{R} \in O(r)$ such that $A(V\check{R}) = (V\check{R})\check{\Lambda}$.*

Proof. ($\implies$) Assume $AV = V\check{\Lambda}$. Then, by definition, there is a RSD $A = QSQ^\top$ such that $V^\top Q = [0_{r \times p} \ I_{r \times r} \ 0_{r \times p}]$. It follows that $\text{skew}(A)V = Q\text{skew}(S)Q^\top V = Q0 = 0$.

($\impliedby$) Assume $\text{skew}(A)V = 0$ and let $A = QSQ^\top$ be a RSD. Then it follows that $\ker(\text{skew}(A)) = \text{span}(Q_r)$. Moreover, since $\text{skew}(A)V = 0$, $\text{span}(V) \subseteq \ker(\text{skew}(A))$. Since both Q_r and V have column rank r , their columns span the same subspace of $\mathbb{R}^n$ and there is $\check{R} \in \mathbb{R}^{r \times r}$ such that $Q_r = V\check{R}$. Finally, since $Q_r \in \text{St}(n, r)$, $I_r = Q_r^\top Q_r = \check{R}^\top V^\top V \check{R} = \check{R}^\top \check{R}$. Hence, $\check{R} \in O(r)$. $\square$

Since $\widehat{Q}$ satisfies $\text{skew}(A)\widehat{Q} = \widehat{Q}\text{skew}(S)$ by (3.2), it holds in particular that $\widehat{Q}_r$ spans the kernel of $\text{skew}(A)$: $\text{skew}(A)\widehat{Q}_r = 0$. Therefore, Lemma 3.1 ensures the existence of $\check{R} \in O(r)$ such that $Q_r = \widehat{Q}_r \check{R}$. As a corollary of Lemma 3.1, it also holds that

$$Q_r^\top [\widehat{Q}_{c,1} \ \widehat{Q}_{c,2}] = 0 \quad \text{and} \quad \widehat{Q}_r^\top [Q_{c,1} \ Q_{c,2}] = 0.$$

Therefore, we can conclude that the matrix R such that $(\widehat{Q}R)^\top Q \in \text{ig}(S)$ has the block structure

$$(3.4) \quad R = \begin{bmatrix} R_a & 0 & R_b \\ 0 & \check{R} & 0 \\ R_c & 0 & R_b \end{bmatrix} \in O(n),$$

where $\begin{bmatrix} R_a & R_b \\ R_c & R_d \end{bmatrix} \in O(2p)$, $R_a \in \mathbb{R}^{p \times p}$ and $\check{R} \in O(r)$.

We can show that R is even more structured by considering [Set 1](#) and [Set 2](#) separately. [Lemma 3.2](#) shows that, for each set, R acts on independent sets of columns of $\widehat{Q}$.

LEMMA 3.2. *For every normal matrix $A \in \mathbb{R}^{n \times n}$ that has $m \geq 1$ eigenvalues with imaginary part $\sigma > 0$, it holds for every $V \in \text{St}(n, 2m)$ satisfying $\text{skew}(A)V = V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}$ that there exists a matrix $R \in O(2m)$ such that*

$$A(VR) = (VR) \begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix},$$

where $D_m = \text{diag}(d_1, \dots, d_m)$ and $d_j \pm i\sigma$ is an eigenvalue of A for $j = 1, \dots, m$.

Proof. Let $k := 2(p - m)$ and $A = \check{Q}\check{S}\check{Q}^\top$ be a RSD with

$$(3.5) \quad \check{S} = \begin{bmatrix} D_m & -\sigma I_m & 0 & 0 & 0 \\ \sigma I_m & D_m & 0 & 0 & 0 \\ 0 & 0 & \check{\Lambda} & 0 & 0 \\ 0 & 0 & 0 & \Lambda_k \cos(\Theta_k) & -\Lambda_k \sin(\Theta_k) \\ 0 & 0 & 0 & \Lambda_k \sin(\Theta_k) & \Lambda_k \cos(\Theta_k) \end{bmatrix},$$

where Λ_k, Θ_k are $k \times k$ diagonal matrices.³ It follows that

$$(3.6) \quad \text{skew}(A)V = V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \iff \text{skew}(\check{S})\check{Q}^\top V = \check{Q}^\top V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}.$$

Let us divide $\check{Q}^\top V$ row-wise in three blocks such that

$$\check{Q}^\top V := \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \text{ with } X \in \mathbb{R}^{2m \times 2m}, Y \in \mathbb{R}^{r \times 2m} \text{ and } Z \in \mathbb{R}^{2(p-m) \times 2m}.$$

Equation (3.6) implies that $0 = Y \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}$ and therefore $Y = 0$. Moreover, letting $Z = \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix}$ with appropriate blocks, we have

$$(3.7) \quad \begin{aligned} & \begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} Z = Z \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \\ \iff & \begin{cases} \Lambda_k \sin(\Theta_k) Z_{21} = -\sigma Z_{12}, \\ \Lambda_k \sin(\Theta_k) Z_{12} = -\sigma Z_{21}, \end{cases} \quad \text{and,} \quad \begin{cases} \Lambda_k \sin(\Theta_k) Z_{11} = \sigma Z_{22}, \\ \Lambda_k \sin(\Theta_k) Z_{22} = \sigma Z_{11}. \end{cases} \end{aligned}$$

Since all diagonal entries of $\Lambda_k \sin(\Theta_k)$ are nonzero and distinct from σ , (3.7) is possible if and only if $Z = 0$. Finally, since $\check{Q} \in O(n)$ and $V \in \text{St}(n, 2m)$, we have $\check{Q}^\top V \in \text{St}(n, 2m)$ and thus $X \in O(2m)$. Letting $R = X^\top$, we have $\check{Q}_{1:2m} = VR$ and this concludes the proof. $\square$

³There is a permutation matrix P such that if $A = QSQ^\top$ is a RSD from (2.2), then $\check{Q} = QP$ and $\check{S} = P^\top SP$.

Lemmas 3.1 and 3.2 show that the matrix R such that $\widehat{Q}R$ is a matrix of Schur vectors acts independently on three sets of columns of $\widehat{Q}$. This fact is the basis for the main structure of the algorithm: the separate computation of the three sets of eigenvalues and eigenvectors, as detailed in the next subsections.

3.3. The eigenvalues with distinct nonzero imaginary parts. In this subsection, we show in Theorem 3.3 that the Schur vectors of A for eigenvalues with distinct nonzero imaginary parts can be obtained by computing the Schur vectors of the skew-symmetric part Ω . Intermediate lemmas are given in Appendix A.

THEOREM 3.3. *Let $A \in \mathbb{R}^{n \times n}$ be a normal matrix. If $\Lambda_k \sin(\Theta_k)$ is diagonal and contains k distinct non-zero imaginary parts of the non-repeated eigenvalues of $\text{skew}(A)$, then, for all $V \in \text{St}(n, 2k)$ such that*

$$\text{skew}(A)V = V \begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix},$$

we have

$$AV = V \begin{bmatrix} \Lambda_k \cos(\Theta_k) & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & \Lambda_k \cos(\Theta_k) \end{bmatrix} =: VS_{2k}.$$

If, additionally, $k = \frac{n}{2}$, then $A = VS_n V^\top$ is a real Schur decomposition (2.2).

Proof. Let $A = QSQ^\top$ be a RSD and take a permutation matrix P such that $(P^\top SP)_{1:2k, 1:2k} = S_{2k}$. By Lemma 3.2, it follows that $QP = [VR \ (QP)_{(k+1):n}]$ for some $R \in O(2k)$ and this yields

$$\begin{aligned} \begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} &= V^\top \text{skew}(A)V && \text{By assumption.} \\ &= V^\top Q \text{skew}(S) Q^\top V \\ &= V^\top (QP)(P^\top \text{skew}(S)P)(QP)^\top V \\ &= [R \ 0_{2k \times (n-2k)}](P^\top \text{skew}(S)P)[R \ 0_{2k \times (n-2k)}]^\top \\ &= R \begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} R^\top. \end{aligned}$$

Therefore, we have $R \in \text{ig} \left(\begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} \right)$. Leveraging Lemma A.2, there is a diagonal matrix $\Phi \in \mathbb{R}^{k \times k}$ such that $R = \begin{bmatrix} \cos(\Phi) & -\sin(\Phi) \\ \sin(\Phi) & \cos(\Phi) \end{bmatrix}$. Moreover, it is easily verified that every such matrix R also belongs to $\text{ig} \left(\begin{bmatrix} \Lambda_k \cos(\Theta_k) & 0 \\ 0 & \Lambda_k \cos(\Theta_k) \end{bmatrix} \right)$. This yields

$$R \in \begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} \subset \begin{bmatrix} \Lambda_k \cos(\Theta_k) & 0 \\ 0 & \Lambda_k \cos(\Theta_k) \end{bmatrix}.$$

Finally, it follows by Lemma 2.4 that

$$R \in \text{ig} \left(\begin{bmatrix} \Lambda_k \cos(\Theta_k) & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & \Lambda_k \cos(\Theta_k) \end{bmatrix} \right) =: \text{ig}(S_{2k}).$$

Consequently, since $QP = [VR \ (QP)_{(k+1):n}]$, we have

$$AV = (QP)(P^\top SP)(QP)^\top V = [VR \ (QP)_{(k+1):n}] \begin{bmatrix} S_{2k} & 0 \\ 0 & * \end{bmatrix} \begin{bmatrix} R^\top \\ 0_{(n-2k) \times 2k} \end{bmatrix} = VRS_{2k}R^\top = VS_{2k}.$$

When $2k = n$, $V \in O(n)$ and the real Schur decomposition is readily obtained since $S_n = S$ and $A = VSV^\top$. $\square$

The conclusion of [Theorem 3.3](#) is that for all eigenvalues that have distinct, non-zero imaginary parts, the invariant subspaces of $\text{skew}(A)$ are always invariant subspaces of A . However, this does not hold true for [Set 2](#) and [Set 3](#) and this implies additional computations, as shown next.

3.4. The eigenvalues with repeated nonzero imaginary parts. Assume there is a diagonal entry σ in $\Lambda \sin(\Theta)$ that has multiplicity $m > 1$. Then, the associated $2m \times 2m$ block in S can be isolated using a permutation P such that

$$(3.8) \quad (P^\top SP)_{1:2m, 1:2m} = \begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix},$$

where D_m is a diagonal submatrix of $\Lambda \cos(\Theta)$. The proof of [Theorem 3.3](#) relies on the following inclusion of invariant groups:

$$\text{ig} \left(\begin{bmatrix} 0 & -\Lambda_k \sin(\Theta_k) \\ \Lambda_k \sin(\Theta_k) & 0 \end{bmatrix} \right) \subset \text{ig} \left(\begin{bmatrix} \Lambda_k \cos(\Theta_k) & 0 \\ 0 & \Lambda_k \cos(\Theta_k) \end{bmatrix} \right).$$

Unfortunately, [Lemma A.3](#) shows that the block $\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}$ admits a larger invariance group than when the eigenvalues have distinct imaginary parts, i.e., if Σ_m is an $m \times m$ diagonal matrix with distinct non-zero entries, then $\text{ig} \left(\begin{bmatrix} 0 & -\Sigma_m \\ \Sigma_m & 0 \end{bmatrix} \right) \subset \text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right)$. However, if $D_m \neq dI_m$, then $\text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right) \not\subset \text{ig} \left(\begin{bmatrix} D_m & 0 \\ 0 & D_m \end{bmatrix} \right)$. This relation is illustrated in [Figure 1](#).




FIGURE 1. Venn diagram of the sets $\text{ig} \left(\begin{bmatrix} D_m & 0 \\ 0 & D_m \end{bmatrix} \right)$, $\text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right)$ and $\text{ig} \left(\begin{bmatrix} 0 & -\Sigma_m \\ \Sigma_m & 0 \end{bmatrix} \right)$ where D_m and Σ_m are diagonal with non-zero diagonal entries and $\sigma > 0$.

As shown in [Theorem 3.4](#), it follows from [Figure 1](#) that if $V \in \text{St}(n, 2m)$ spans the invariant subspace of $\text{skew}(A)$ associated to σ , then $V^\top AV$ is not a real Schur form. Instead, $V^\top AV$ reveals a very structured *symmetric skew-Hamiltonian* matrix.

THEOREM 3.4. *For every normal matrix $A \in \mathbb{R}^{n \times n}$ that has $m > 1$ eigenvalues with imaginary part $\sigma > 0$, for all $V \in \text{St}(n, 2m)$ such that*

$$\text{skew}(A)V = V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix},$$

we have

$$AV = V \begin{bmatrix} \tilde{H} & -\tilde{\Omega} - \sigma I_m \\ \tilde{\Omega} + \sigma I_m & \tilde{H} \end{bmatrix},$$

where $\tilde{H} \in \mathbb{R}^{m \times m}$ is symmetric and $\tilde{\Omega} \in \mathbb{R}^{m \times m}$ is skew-symmetric.

Proof. Let $A = \check{Q}\check{S}\check{Q}^\top$ be a RSD with the structure of (3.5). By Lemma 3.2, there is $R \in O(2m)$ such that $\check{Q}_{1:2m} = VR$. Since $\check{Q}\text{skew}(\check{S})\check{Q}^\top V = V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}$, it follows that

$$(3.9) \quad R \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} R^\top = \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}.$$

Hence, $R \in \text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right)$. In Lemma A.3, we recall the following equivalence:

$$(3.10) \quad R \in \text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right) \iff R \in O(2m) \text{ and } R = \begin{bmatrix} E & -F \\ F & E \end{bmatrix},$$

where $E, F \in \mathbb{R}^{m \times m}$. Recall that such matrices are called *ortho-symplectic*, see, e.g., [32, Eq. 5.1]. In view of (3.10), it holds that

$$(3.11) \quad \begin{aligned} AV &= \text{sym}(A)V + \text{skew}(A)V \\ &= \check{Q}\text{sym}(\check{S})\check{Q}^\top V + \text{skew}(A)V \\ &= VR \begin{bmatrix} D_m & 0 \\ 0 & D_m \end{bmatrix} R^\top + V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \\ (3.12) \quad &= V \begin{bmatrix} ED_mE^\top + FD_mF^\top & (FD_mE^\top - ED_mF^\top)^\top - \sigma I_m \\ FD_mE^\top - ED_mF^\top + \sigma I_m & ED_mE^\top + FD_mF^\top \end{bmatrix}. \end{aligned}$$

Defining $\tilde{H} := ED_mE^\top + FD_mF^\top$ and $\tilde{\Omega} := FD_mE^\top - ED_mF^\top$ concludes the proof. $\square$

Let us consider again Q , the Schur vectors of A , $\hat{Q}$, the Schur vectors of $\text{skew}(A)$, and the permutation matrix P from (3.8). In view of Lemma 3.2, $(QP)_{1:2m} = (\hat{Q}P)_{1:2m}R$ with $R \in O(2m)$ and Theorem 3.4 holds with $V = (\hat{Q}P)_{1:2m}$. Therefore, for each diagonal entry σ in $\Lambda \sin(\Theta)$ that has multiplicity $m > 1$, it follows from Theorem 3.4 that we must compute $V^\top AV$ and solve an additional $2m \times 2m$ EVP given by

$$(3.13) \quad V^\top AV = \begin{bmatrix} \tilde{H} & -\tilde{\Omega} - \sigma I_m \\ \tilde{\Omega} + \sigma I_m & \tilde{H} \end{bmatrix} = R \begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix} R^\top,$$

where $R \in O(2m)$ and $D_m \in \mathbb{R}^{m \times m}$ are to be computed. Solving (3.13) is enough to obtain the correct Schur vectors by $(QP)_{1:2m} := VR$.

1. If $D_m = dI_m$, $d \in \mathbb{R}$, then it is the unique case where $\text{ig} \left(\begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix} \right) \subset \text{ig} \left(\begin{bmatrix} dI_m & 0 \\ 0 & dI_m \end{bmatrix} \right)$. Consequently, $\tilde{H} = dI_m$ and $\tilde{\Omega} = 0$. No additional work is needed, one can assume $R = I_{2m}$ without loss of generality.
2. If D_m is not a multiple of the identity, we can recover an ortho-symplectic eigenvector matrix R from the *symmetric skew-Hamiltonian* EVP of $\begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix}$. Several methods have been designed for Hamiltonian and skew-Hamiltonian EVPs, see, e.g., [12, 76, 13, 70, 35], notably a QR algorithm relying exclusively on ortho-symplectic similarity transformations [70]. For completeness, we briefly discuss in Appendix B how the m first steps of Lanczos tridiagonalization, with restarting and reorthogonalization for stability, are enough to construct an ortho-symplectic matrix Y such that

$$(3.14) \quad \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix} = Y \begin{bmatrix} T_m & 0 \\ 0 & T_m \end{bmatrix} Y^\top = R \begin{bmatrix} D_m & 0 \\ 0 & D_m \end{bmatrix} R^\top,$$

where T_m is an $m \times m$ symmetric tridiagonal matrix. Consequently, once the EVP of $T_m = ZD_mZ^\top$ is solved with a call to [Routine 2](#), we obtain $R := Y \begin{bmatrix} Z & 0 \\ 0 & Z \end{bmatrix}$.

As studied in [section 4](#), in practice, imaginary parts of the eigenvalues will not be exactly repeated but rather clustered and therefore, the symmetric skew-Hamiltonian structure cannot be guaranteed. This jeopardizes the attractive computational route based on (3.14). In order to maintain the accuracy of the global RSD, it is more robust to call [Routine 3](#) to solve (3.13) directly. This choice has no influence on the computational complexity if m remains small, which is expected.

3.5. The real eigenvalues. We conclude this section by considering the computation of the r real eigenvalues of A and the associated eigenvectors. In [Theorem 3.5](#), we show that an additional $r \times r$ symmetric EVP needs to be solved to obtain the eigenvectors of A .

THEOREM 3.5. *For every normal matrix $A \in \mathbb{R}^{n \times n}$ that has $r > 0$ real eigenvalues and for all $V \in \text{St}(n, r)$ such that $\text{skew}(A)V = 0$, we have*

$$AV = VH,$$

where the matrix $H \in \mathbb{R}^{r \times r}$ is symmetric.

Proof. For every RSD $A = QSQ^\top$ where S follows the structure of (2.2), we know by [Lemma 3.1](#) that $Q = [Q_{1:p} \mid V\check{R} \mid Q_{(p+r+1):n}]$ for some $\check{R} \in \text{O}(r)$. Therefore, it holds that

$$\begin{aligned} AV &= QSQ^\top V \\ &= Q \begin{bmatrix} \Lambda \cos(\Theta) & 0 & -\Lambda \sin(\Theta) \\ 0 & \check{\Lambda} & 0 \\ \Lambda \sin(\Theta) & 0 & \Lambda \cos(\Theta) \end{bmatrix} \begin{bmatrix} 0_{p \times r} \\ \check{R}^\top \\ 0_{p \times r} \end{bmatrix} \\ &= Q \begin{bmatrix} 0_{p \times r} \\ \check{\Lambda} \check{R}^\top \\ 0_{p \times r} \end{bmatrix} \\ &= V\check{R}\check{\Lambda}\check{R}^\top. \end{aligned}$$

Defining the symmetric matrix $H := \check{R}\check{\Lambda}\check{R}^\top$ completes the proof. $\square$

3.6. The algorithm in exact arithmetic. We have gathered all pieces to design an algorithm in exact arithmetic. First, we obtain $\hat{Q}$ from RSD of $\Omega = \text{skew}(A)$ (3.2). By [Theorem 4.4](#), the invariant spaces associated to [Set 1](#) need no further action. For each invariant space of Ω associated to [Set 2](#), i.e, a complex eigenvalue with a repeated imaginary part, an additional symmetric skew-Hamiltonian EVP needs to be solved. This gives the transformation $[Q_{c,1} \ Q_{c,2}] = [\hat{Q}_{c,1} \ \hat{Q}_{c,2}] \begin{bmatrix} R_a & R_b \\ R_c & R_d \end{bmatrix}$ from (3.4). Finally, by [Theorem 4.7](#), the decomposition $H := \hat{Q}_r^\top A \hat{Q}_r = \check{R}\check{\Lambda}\check{R}^\top$ using a call to [Routine 2](#) completes the determination of the real Schur decomposition

$$A = \hat{Q} \begin{bmatrix} R_a & 0 & R_b \\ 0 & \check{R} & 0 \\ R_c & 0 & R_d \end{bmatrix} \begin{bmatrix} \Lambda \cos(\Theta) & 0 & -\Lambda \sin(\Theta) \\ 0 & \check{\Lambda} & 0 \\ \Lambda \sin(\Theta) & 0 & \Lambda \cos(\Theta) \end{bmatrix} \begin{bmatrix} R_a & 0 & R_b \\ 0 & \check{R} & 0 \\ R_c & 0 & R_d \end{bmatrix}^\top \hat{Q}^\top.$$

The complete method in exact arithmetic is summarized in [Algorithm 3.1](#). In a broad set of problems, steps 5 and 7 of [Algorithm 3.1](#) are either not necessary or very simple. Indeed, for Haar-distributed random

orthogonal matrices [7, 64], it holds with probability 1 that $r \in \{0, 1\}$ if $\det(A) = 1$ and $r \in \{1, 2\}$ if $\det(A) = -1$ [34]. In particular, if $r = 1$, $H = \pm 1$ and $\check{R} = 1$. Step 7 of Algorithm 3.1 is thus trivial. Moreover, all complex eigenvalues have distinct imaginary parts with probability 1 and thus step 5 is skipped. For general normal matrices, the distribution of the eigenvalues, and in particular the value of r , are specific to the context of the problem. Nonetheless, the case where all or a large majority of the eigenvalues belong to Set 1 is common.

Algorithm 3.1 Exact arithmetic algorithm for the real Schur decomposition of a normal matrix

Input: A normal matrix $A \in \mathbb{R}^{n \times n}$.

Output: Q and S where $Q \in O(n)$ and $A = QSQ^\top$ is a RSD.

step 1: Compute the skew-symmetric part $\Omega := \frac{1}{2}(A - A^\top)$.

step 2.1: Compute a tridiagonal reduction $\Omega = \tilde{Q}T\tilde{Q}^\top$.

step 2.2: Apply the permutation $\begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix} := P_{\text{eo}}^\top T P_{\text{eo}}$ where $\tilde{B} \in \mathbb{R}^{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil}$ is bidiagonal.

step 3: Compute the (permuted) SVD of $\tilde{B} = U \begin{bmatrix} 0 & 0_{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil} \\ \Sigma & 0 \end{bmatrix} V^\top$ where $\text{diag}(\Sigma) > 0$.

step 4: Define $\hat{Q} := \tilde{Q}P_{\text{eo}} \begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix}$, $\Lambda \sin(\Theta) := \Sigma$ and $p := \frac{n-r}{2}$.

for each non-zero diagonal entry σ of Σ with multiplicity $m > 1$ **do**

step 5.1: Define V as the $2m$ columns of $\hat{Q}$ such that $\Omega V = V \begin{bmatrix} 0 & -\sigma I_m \\ \sigma I_m & 0 \end{bmatrix}$.

step 5.2: Compute the matrix $V^\top AV =: \begin{bmatrix} \tilde{H} & -\tilde{\Omega} - \sigma I_m \\ \tilde{\Omega} + \sigma I_m & \tilde{H} \end{bmatrix}$ where $\tilde{H} = \tilde{H}^\top$ and $\tilde{\Omega} = -\tilde{\Omega}^\top$.

step 5.3: Compute the RSD $\begin{bmatrix} \tilde{H} & -\tilde{\Omega} - \sigma I_m \\ \tilde{\Omega} + \sigma I_m & \tilde{H} \end{bmatrix} = R \begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix} R^\top$.

step 5.4: Replace the columns of $\hat{Q}$ selected at step 5.1 by VR .

end for

step 6: Define the diagonal matrix $\Lambda \cos(\Theta) := \hat{Q}_{1:p}^\top A \hat{Q}_{1:p}$.

if $r > 0$ **then**

step 7.1: Define $V := \hat{Q}_{(p+1):(p+r)}$ such that $\Omega V = 0$.

step 7.2: Compute the symmetric matrix $H := V^\top AV$.

step 7.3: Compute the EVD of $H = \check{R}\check{\Lambda}\check{R}^\top$.

step 7.4: Replace $\hat{Q}_{(p+1):(p+r)}$ by $V\check{R}$.

end if

return $Q := \hat{Q}$ and $S := \begin{bmatrix} \Lambda \cos(\Theta) & 0 & -\Lambda \sin(\Theta) \\ 0 & \check{\Lambda} & 0 \\ \Lambda \sin(\Theta) & 0 & \Lambda \cos(\Theta) \end{bmatrix}$.

3.7. The eigenvalue decomposition by the symmetric part. Algorithm 3.1 shows that the real Schur decomposition of real normal matrix can be obtained from its skew-symmetric part. It is natural to ask if this approach can be adapted to start from the *symmetric part*. The answer is positive but several drawbacks arise compared to the skew-symmetric approach.

- Drawback 1: The symmetric part does not allow separating real from complex eigenvalues easily. Indeed, given an EVD of $\text{sym}(A) = V\Lambda V^\top$, one can not distinguish a real eigenvalue of A from the real part of the complex eigenvalue, except if its multiplicity is 1. Making this separation requires computing the product $V^\top AV$. In Algorithm 3.1, there is a direct correspondence between the real eigenspace of A and the kernel of $\text{skew}(A)$.
- Drawback 2: The aforementioned separation between real and complex eigenvalues can not be obtained if some real parts of complex eigenvalues are equal to real eigenvalues of A . Then, the

associated sub-block of $V^\top AV$ is a normal matrix with no additional structure.

- Drawback 3: The symmetric eigensolver for $\text{sym}(A)$ does not take advantage of the even multiplicity of the real parts of complex eigenvalues, while the skew-symmetric part uses this advantage by computing a half-the-size bidiagonal SVD.
- Drawback 4: For every two-dimensional invariant subspace of A , the problem arises that

$$\text{ig} \left(\begin{bmatrix} \lambda \cos(\theta) & -\lambda \sin(\theta) \\ \lambda \sin(\theta) & \lambda \cos(\theta) \end{bmatrix} \right) = \text{SO}(2) \subset \text{ig} \left(\begin{bmatrix} \lambda \cos(\theta) & 0 \\ 0 & \lambda \cos(\theta) \end{bmatrix} \right) = \text{O}(2).$$

Since an element of $\text{O}(2)$ belongs to $\text{SO}(2)$ with probability $\frac{1}{2}$, there is a probability $\frac{1}{2}$ for each base to be incorrectly oriented.

The aforementioned drawbacks matter if a significant proportion of the eigenvalues are complex. If most eigenvalues are real, then, the approach by the symmetric part can be efficient. However, applications, e.g., with orthogonal matrices, are usually dominated by complex eigenvalues. This makes the approach by the skew-symmetric-part more relevant for applications to the best of the authors' knowledge.

3.8. Exact versus floating point arithmetic: an example. We have not yet adressed an essential algorithmic question: how accurate is [Algorithm 3.1](#) in floating point arithmetic and how should it be implemented in order to always obtain a desired accuracy? Indeed, floating point arithmetic brings an additional challenge. In some cases, the invariant subspaces of Ω are less sensitive to perturbations of A than the invariant subspaces of A themselves, and, if not analyzed and addressed in the design of the algorithm, this can be a source of inaccuracies. We show next an illustrative example.

Consider a variable $\tau \geq 0$, a $p \times p$ random symmetric matrix E , $\|E\|_2 = 1$, and the normal matrix $A(\tau)$ defined by

$$A(\tau) = \begin{bmatrix} \sin(\tau E) & -\cos(\tau E) \\ \cos(\tau E) & \sin(\tau E) \end{bmatrix} \quad \text{such that,} \quad A(0) = \begin{bmatrix} 0 & -I_p \\ I_p & 0 \end{bmatrix}.$$

Notice that $A(0) = \Omega(0)$. For τ small enough, Taylor series yield $\cos(\tau E) \approx I_p - \frac{\tau^2}{2} E^2$ and $\sin(\tau E) \approx \tau E$ such that, using the spectral norm $\|\cdot\|_2$, we have

$$\|A(\tau) - A(0)\|_2 \approx \tau \quad \text{and,} \quad \|\Omega(\tau) - \Omega(0)\|_2 \approx \frac{\tau^2}{2}.$$

Around $\tau = 0$, it follows that $\Omega(\tau)$ is less sensitive to small variations of τ than $A(\tau)$. Let ε_m denote the machine precision. Then, applying [Algorithm 3.1](#) on $A(\sqrt{\varepsilon_m})$ or $A(0)$ will yield similar Schur vectors in output since, in floating point arithmetic, $\Omega(\sqrt{\varepsilon_m}) \approx \Omega(0)$. Therefore, $\mathcal{O}(\sqrt{\varepsilon_m})$ -errors will be made in the computation of the Schur vectors of $A(\sqrt{\varepsilon_m})$. This is verified in practice, here for $p = 100$ and $\varepsilon_m = 2.22 \cdot 10^{-16}$, we have

$$\text{Naive implementation of } \text{Algorithm 3.1:} \quad \|A(\sqrt{\varepsilon_m})\widehat{Q} - \widehat{Q}\widehat{S}\|_F = 6.52 \cdot 10^{-9} \|A(\sqrt{\varepsilon_m})\|_F.$$

For methods relying on the diagonalization of commuting matrices, it is well known that sensitivity analyses depend on eigenvalue gaps. In [\[47\]](#), which studies the simultaneous diagonalization of symmetric matrices, a probabilistic accuracy bound is derived in terms of the smallest eigenvalue gap among the family of matrices. Similarly, the do-one-then-do-the-other algorithm of [\[66, Alg. 2\]](#) requires a refinement of the invariant subspaces that is governed by the eigenvalue gaps. As shown in [section 4](#), the specific setting

of real normal matrices, together with [Algorithm 4.1](#), allows for more precise statements, again involving eigenvalue gaps.

In the next section, we study how the loss of accuracy occurs and how [Algorithm 3.1](#) should be implemented in order to obtain a desired accuracy μ .

4. Analysis of numerical stability and accuracy. In this section, we analyze the effect of floating point arithmetic on the accuracy of [Algorithm 3.1](#). We rely on standard assumptions of floating point arithmetic. In particular, ε_m denotes the machine precision and, for every scalar $\alpha \in \mathbb{C}$, $\text{fl}(\alpha) = \alpha(1 + \Delta\alpha)$ with $|\Delta\alpha| \leq \varepsilon_m$ models its float representation. The *absolute gap* between two numbers $\alpha, \hat{\alpha} \in \mathbb{C}$ is defined as $|\alpha - \hat{\alpha}|$ while the *relative gap* refers to $\frac{|\alpha - \hat{\alpha}|}{|\alpha| + |\hat{\alpha}|}$. The notation $\alpha \lesssim \beta$ means that there is a moderately growing function $\varphi(n)$, e.g., n or n^2 , independent of the values of α and β , such that $\alpha \leq \varphi(n)\beta$.

We will analyze step by step the numerical properties of the algorithm. This will determine how to implement [Algorithm 3.1](#) to obtain high accuracy on both the eigenvalues and the Schur vectors/invariant spaces. The floating point arithmetic version of the method is given in [Algorithm 4.1](#). The algorithm requires to identify clusters of eigenvalues. A definition for a cluster of eigenvalues, justified by [Theorems 4.4, 4.6](#) and [4.7](#) is given [Definition 4.1](#).

DEFINITION 4.1. *Given $\delta > 0$, a matrix A , a set of $2m > 0$ eigenvalues $\{\pm i\sigma_j\}_{j=1}^m$ of $\Omega = \text{skew}(A)$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m \geq 0$ is a δ -cluster of Ω if and only if*

$$|\sigma_j - \sigma_{j+1}| \leq \delta \|A\|_F \quad \text{and} \quad |\sigma_k - \sigma| > \delta \|A\|_F \quad \text{for } j = 1, \dots, m-1, \text{ and, } k = 1, \dots, m,$$

where $i\sigma$ is any eigenvalue of Ω not in $\{\pm i\sigma_j\}_{j=1}^m$.

[Definition 4.1](#) is similar to the clustering criterion used in [\[66, Alg. 2\]](#). Instead of the Frobenius norm $\|A\|_F$, the results of [\[66\]](#) employ the *spread* of the eigenvalues, i.e., the maximum distance between the eigenvalues. This difference in definitions reflects the distinct hypotheses and analytical frameworks underlying the respective results.

4.1. Analysis of steps 1 to 4. **Step 1** already entails some subtleties since it consists of n^2 subtractions. Some rare pathological cases exist where this operation might be an initial source of inaccuracy. Indeed, letting $\Omega = \text{skew}(A)$ and $\hat{\Omega} = \text{skew}(\text{fl}(A))$ such that $\hat{\omega}_{ij} = \frac{\text{fl}(a_{ij}) - \text{fl}(a_{ji})}{2}$, we obtain

$$(4.1) \quad \left| \frac{\omega_{ij} - \hat{\omega}_{ij}}{\omega_{ij}} \right| = \left| \frac{a_{ij}\Delta a_{ij} - a_{ji}\Delta a_{ji}}{a_{ij} - a_{ji}} \right| \leq \varepsilon_m \frac{|a_{ij}| + |a_{ji}|}{|a_{ij} - a_{ji}|}.$$

Based on the elementwise relation [\(4.1\)](#) and acknowledging that $\text{Tr}(|A|^2) \leq \|A\|_F^2$ by Cauchy-Schwarz inequality, we can also deduce the more practical matrix-relation

$$(4.2) \quad \frac{\|\Omega - \hat{\Omega}\|_F}{\|\Omega\|_F} \leq \varepsilon_m \frac{2\|A\|_F}{\|\Omega\|_F}.$$

In view of [\(4.1\)](#), whenever symmetrically positioned entries of A are *relatively* close, the numerical approximation $\hat{\omega}_{ij}$ may become inaccurate. However, relative inaccuracies only occur for ω_{ij} close to zero. Therefore, if $\|\Omega\|_F \gg \varepsilon_m \|A\|_F$, then $\hat{\Omega}$ is accurate in the sense of [\(4.2\)](#), even if some entries are not. An example of a matrix set where such entrywise inaccuracies occur in general with negligible effect is

$$(4.3) \quad \{U \in \text{O}(2m) \mid U = \exp \begin{bmatrix} 0 & -M^\top \\ M & 0 \end{bmatrix}, \quad M \in \mathbb{R}^{m \times m}\}.$$

Algorithm 4.1 Floating point arithmetic algorithm for the real Schur decomposition of a normal matrix (nrmschur)

Input: A normal matrix $A \in \mathbb{R}^{n \times n}$, $\delta, \delta_r = \sqrt{\varepsilon_m}$ and $\mu = \varepsilon_m t$, $t \geq 1$.

Output: $\hat{Q}$ and $\hat{S}$ where $d_{O(n)}(\hat{Q}) \lesssim \varepsilon_m$ and $\|A\hat{Q} - \hat{Q}\hat{S}\|_F \lesssim \mu\|A\|_F$.

step 1: Compute the skew-symmetric part $\hat{\Omega} \approx \frac{1}{2}(A - A^\top)$.

step 2.1: Tridiagonalize $\hat{\Omega} \approx \tilde{Q}T\tilde{Q}^\top$.

step 2.2: Apply the permutation $\begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix} := P_{\text{eo}}^\top T P_{\text{eo}}$ where $\tilde{B} \in \mathbb{R}^{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil}$ is bidiagonal.

step 3: ([Routine 1](#)) Compute the (permuted) SVD of $\tilde{B} \approx \hat{U} \begin{bmatrix} 0 & 0 & \hat{\Sigma} \end{bmatrix} \hat{V}^\top$ where $\text{diag}(\hat{\Sigma}) > 0$ and $\hat{r}$ is the size of the δ_r -cluster of eigenvalues of $\hat{\Omega}$ around 0.

step 4: Compute $\hat{Q} \approx \tilde{Q}P_{\text{eo}} \begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix}$, set $\Lambda \sin(\Theta) := \hat{\Sigma}$ and $\hat{p} := \frac{n-\hat{r}}{2}$.

for each δ -cluster of $2m$ (non-zero) eigenvalues of $\hat{\Omega}$ **do**

step 5.1: Set V as the $2m$ columns of $\hat{Q}$ associated to the cluster and compute $M \approx V^\top A V$.

step 5.2: ([Routine 3](#)) Compute the RSD $M \approx R\tilde{S}_M R^\top$.

step 5.3: Replace the $2m$ columns of $\hat{Q}$ from step 5.1 by VR .

end for

step 6: Set $\Lambda \cos(\Theta)$ as the diagonal of $\hat{Q}_{1:\hat{p}}^\top \text{sym}(A) \hat{Q}_{1:\hat{p}}$.

if there is a δ_r -cluster of $\hat{r}$ eigenvalues of Ω around 0, **then**

step 7.1: Compute the matrix $H \approx \hat{Q}_{\hat{p}+1:\hat{p}+\hat{r}}^\top A \hat{Q}_{\hat{p}+1:\hat{p}+\hat{r}}$.

if $\|\text{skew}(H)\|_F \leq \varepsilon_m \|H\|_F$, **then**

step 7.2a: ([Routine 2](#)) Compute the EVD of $\text{sym}(H) \approx \check{R}\check{\Lambda}\check{R}^\top$.

else

step 7.2b: ([Routine 3](#)) Compute the RSD of $H \approx \check{R}\check{S}_H\check{R}^\top$.

end if

step 7.3: Set $\hat{Q}_{\hat{p}+1:\hat{p}+\hat{r}} \approx \hat{Q}_{\hat{p}+1:\hat{p}+\hat{r}}\check{R}$.

end if

step 8: (Optional, see [subsection 4.5](#)) Subspaces correction steps. For each $\frac{1}{t}$ -cluster of eigenvalues of $\hat{\Omega}$, repeat step 5 with $\delta = \frac{1}{t}$ and consider a Jacobi algorithm [83] for [Routine 3](#).

return $\hat{Q}$ and $\hat{S} := \begin{bmatrix} \Lambda \cos(\Theta) & 0 & -\Lambda \sin(\Theta) \\ 0 & \check{\Lambda} & 0 \\ \Lambda \sin(\Theta) & 0 & \Lambda \cos(\Theta) \end{bmatrix}$.

The matrix set (4.3) notably arises in the fundamental Riemannian computation of geodesics on the Grassmann manifold [11]. Entrywise inaccuracies can also be neglected if $\|\Omega\|_F \leq \varepsilon_m \|A\|_F$. In this case, the matrix A can be numerically considered as symmetric. In conclusion, computing a skew-symmetric part is dubious when $\varepsilon_m \|A\|_F \ll \|\Omega\|_F \ll \|A\|_F$. However, it is shown in [section 6](#) that [Algorithm 4.1](#) can handle small skew-symmetric parts.

Step 2 is the tridiagonal reduction using Householder reflectors. It is known to be backward stable [41, p. 416], i.e.,

$$(4.4) \quad \hat{\Omega} + \Delta\hat{\Omega} = \tilde{Q}P_{\text{eo}} \begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix} P_{\text{eo}}^\top \tilde{Q}^\top \text{ with } \|\Delta\hat{\Omega}\|_F \lesssim \varepsilon_m \|\hat{\Omega}\|_F,$$

where $\Delta\hat{\Omega}$ is exactly skew-symmetric, $\tilde{Q} \in O(n)$ is the exact product of the Householder reflectors numerically computed during the reduction and P_{eo} is an even-odd permutation (2.5). Moreover, Householder reflectors can be assembled very accurately as $\tilde{Q}_{\text{asm}}$ such that $\|\tilde{Q}_{\text{asm}} - \tilde{Q}\|_F \lesssim \varepsilon_m$ [49, Eq. 19.13]. Defining the distance

to orthogonality of a matrix V by

$$(4.5) \quad d_{O(n)}(V) := \min_{Q \in O(n)} \|Q - V\|_F,$$

we can write $d_{O(n)}(\tilde{Q}_{\text{asm}}) \lesssim \varepsilon_m$.

Step 3 is the bidiagonal SVD. Algorithms such as **bdsqr** provide a mixed forward-backward stable SVD of $\tilde{B} \in \mathbb{R}^{\lfloor \frac{n}{2} \rfloor \times \lceil \frac{n}{2} \rceil}$:

$$\tilde{B} + \Delta\tilde{B} = \hat{U}\hat{\Sigma}\hat{V}^\top \quad \text{and} \quad \|\Delta\tilde{B}\|_F \lesssim \varepsilon_m \|\tilde{B}\|_F,$$

such that $\max\{d_{O(\lfloor \frac{n}{2} \rfloor)}(\hat{U}), d_{O(\lceil \frac{n}{2} \rceil)}(\hat{V})\} \lesssim \varepsilon_m$ [6, 28].

REMARK 4.2. The singular values of $\tilde{B}$ are computed with high *relative* accuracy by algorithms such as **bdsqr** [28, Thm. 6 and 7]. This holds because relatively small perturbations of $\tilde{B}$ yield relatively perturbed singular values [30, Chap. 4.3]. Moreover, it was conjectured [28] and proved [26, Thm. 6.1] that left and right singular vectors of $\tilde{B}$ are also very accurately computed. However, the Bauer-Fike theorem [51, Thm. 6.3.2] and its corollary [51, Cor. 6.3.4] indicate that, due to the perturbation $\Delta\hat{\Omega}$ in (4.4), relative accuracy of the eigenvalues of Ω cannot be achieved. Consequently, the relative accuracy of the singular values of $\tilde{B}$ is not relevant.

Step 4 is a product of matrices. Define $\hat{Q}_{\text{exact}}$ as the *exact* matrix product of $\tilde{Q}_{\text{asm}}P_{\text{eo}}$ by $\begin{bmatrix} \hat{V} & 0 \\ 0 & \hat{U} \end{bmatrix}$. Since the permutation is performed exactly, w.l.o.g., we assume $P_{\text{eo}} = I_n$. We obtain that

$$\begin{aligned} d_{O(n)}(\hat{Q}_{\text{exact}}) &\leq \|\tilde{Q}_{\text{asm}} \begin{bmatrix} \hat{V} & 0 \\ 0 & \hat{U} \end{bmatrix} - \tilde{Q} \begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix}\|_F \\ &= \|\tilde{Q}_{\text{asm}} \begin{bmatrix} \hat{V} - V + V & 0 \\ 0 & \hat{U} - U + U \end{bmatrix} - \tilde{Q} \begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix}\|_F \\ &\leq \|\tilde{Q}_{\text{asm}} - \tilde{Q}\|_F + \|\tilde{Q}_{\text{asm}} \begin{bmatrix} \hat{V} - V & 0 \\ 0 & \hat{U} - U \end{bmatrix}\|_F \\ &\lesssim \varepsilon_m. \end{aligned}$$

Finally, it holds by [49, Eq. 3.13] that $\|\hat{Q} - \hat{Q}_{\text{exact}}\|_F \lesssim \varepsilon_m \|\tilde{Q}_{\text{asm}}\|_F \|\begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix}\|_F \lesssim \varepsilon_m$. Therefore, by the triangular inequality, we have $d_{O(n)}(\hat{Q}) \lesssim \varepsilon_m$. By summing all errors from steps 1 to 4, we obtain the skew-symmetric error matrix $\Delta_2\hat{\Omega}$:

$$\begin{aligned} \Delta_2\hat{\Omega} &:= \Delta\hat{\Omega} + \tilde{Q} \begin{bmatrix} 0 & -\Delta\tilde{B}^\top \\ \Delta\tilde{B} & 0 \end{bmatrix} \tilde{Q}^\top \\ &\quad + (\tilde{Q}_{\text{asm}} - \tilde{Q}) \begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix} \tilde{Q}^\top - \tilde{Q} \begin{bmatrix} 0 & -\tilde{B}^\top \\ \tilde{B} & 0 \end{bmatrix}^\top (\tilde{Q}_{\text{asm}} - \tilde{Q})^\top \\ &\quad + (\hat{Q} - \hat{Q}_{\text{exact}}) \begin{bmatrix} 0 & -\hat{\Sigma} \\ \hat{\Sigma} & 0 \end{bmatrix} \hat{Q}^\top - \hat{Q} \begin{bmatrix} 0 & -\hat{\Sigma} \\ \hat{\Sigma} & 0 \end{bmatrix}^\top (\hat{Q} - \hat{Q}_{\text{exact}})^\top + \mathcal{O}(\varepsilon_m^2). \end{aligned}$$

$\mathcal{O}(\varepsilon_m^2)$ gathers terms where, e.g., $\tilde{Q}_{\text{asm}} - \tilde{Q}$ appears twice. If we consider $(\hat{Q}, \hat{\Sigma})$ as the output of the skew-symmetric RSD, this yields the mixed forward-backward stable decomposition of the skew-symmetric part $\hat{\Omega}$:

$$(4.6) \quad \hat{\Omega} + \Delta_2\hat{\Omega} = \hat{Q} \begin{bmatrix} 0 & -\hat{\Sigma} \\ \hat{\Sigma} & 0 \end{bmatrix} \hat{Q}^\top \quad \text{with} \quad \|\Delta_2\hat{\Omega}\|_F \lesssim \varepsilon_m \|\hat{\Omega}\|_F \quad \text{and} \quad d_{O(n)}(\hat{Q}) \lesssim \varepsilon_m.$$

Mixed forward-backward stability refers to the fact that $\hat{Q}$ is near orthogonal. In this context, the decomposition would be said *backward* stable if $\hat{Q}$ was exactly orthogonal. As shown in section 5, $d_{O(n)}(\hat{Q})$ is very close to machine precision in practice.

4.2. Sensitivity of the invariant subspaces: Set 1. Let A be a normal matrix with an exact RSD $A = QSQ^\top$. We first assume that the matrix A has no real eigenvalues ($r = 0, n = 2p$) and that all eigenvalues have distinct imaginary parts. Given a backward stable RSD of $\hat{\Omega}$, i.e., $\hat{\Omega}\hat{Q} - \hat{Q}\text{skew}(\hat{S}) = \hat{E}$ such that $\|\hat{E}\|_F \lesssim \varepsilon_m \|\hat{\Omega}\|_F$, we quantify $\|A\hat{Q} - \hat{Q}S\|_F$ in [Theorem 4.4](#). We first assume that $\hat{Q}$ is *exactly* orthogonal such that [Theorem 4.4](#) focuses on the important conclusions. We address deviation from orthogonality in [Corollary 4.5](#). This quantification is key for two points: first, characterizing how accurately eigenvalues are computed at step 6 by the Rayleigh quotients and second, designing correction steps to achieve a specified accuracy on the Schur vectors, if needed.

We start by showing in [Lemma 4.3](#) that it is equivalent to consider the numerical output $\hat{\Omega}\hat{Q} - \hat{Q}\text{skew}(\hat{S}) = \hat{E}$ with $\|\hat{E}\|_F \leq \hat{\tau}\varepsilon_m \|\hat{\Omega}\|_F$ from (4.6) and the more practical relation $\Omega\hat{Q} - \hat{Q}\text{skew}(S) = E$ with $\|E\|_F \leq \tau\|A\|_F$ and $\tau > 0$. In the next lemma, we assume that S and $\hat{S}$ are organized such that $\text{skew}(S)$ and $\text{skew}(\hat{S})$ have their eigenvalues sorted, e.g., in decreasing order.

LEMMA 4.3. *Let $A = QSQ^\top$ be a normal matrix, $\Omega := \text{skew}(A)$ and $\hat{\Omega} := \text{skew}(\text{fl}(A))$. If the orthogonal matrix $\hat{Q} \in O(2p)$ is such that $\hat{\Omega}\hat{Q} - \hat{Q}\text{skew}(\hat{S}) = \hat{E}$ is a backward stable RSD of $\hat{\Omega}$, namely, if $\|\hat{E}\|_F \leq \hat{\tau}\varepsilon_m \|\hat{\Omega}\|_F$ for $\hat{\tau} > 0$, then it follows that*

$$\|\Omega\hat{Q} - \hat{Q}\text{skew}(S)\|_F \leq (2\hat{\tau} + 4)\varepsilon_m \|A\|_F + o(\varepsilon_m).$$

Proof. Let us relate $\hat{Q}$ to Ω instead of $\hat{\Omega}$.

$$\begin{aligned} \hat{\Omega}\hat{Q} - \hat{Q}\text{skew}(\hat{S}) &= \hat{E} \\ \iff \Omega\hat{Q} - \hat{Q}\text{skew}(S) &= \hat{E} - (\hat{\Omega} - \Omega)\hat{Q} - \hat{Q}\text{skew}(S - \hat{S}). \end{aligned}$$

Now let $\text{skew}(\hat{S}_2)$ be the *exact* real Schur form of $\hat{\Omega}$. Then, we have

$$\begin{aligned} \|\Omega\hat{Q} - \hat{Q}\text{skew}(S)\|_F &\leq \|\hat{E}\|_F + \|(\hat{\Omega} - \Omega)\hat{Q}\|_F + \|\hat{Q}\text{skew}(S - \hat{S})\|_F \\ &= \|\hat{E}\|_F + \|\hat{\Omega} - \Omega\|_F + \|\text{skew}(S - \hat{S})\|_F \\ &\leq \|\hat{E}\|_F + \|\hat{\Omega} - \Omega\|_F + \|\text{skew}(S - \hat{S}_2)\|_F + \|\text{skew}(\hat{S}_2 - \hat{S})\|_F. \end{aligned}$$

By [50, Thm. 1], we have $\|\text{skew}(S - \hat{S}_2)\|_F \leq \|\hat{\Omega} - \Omega\|_F$. Moreover, since $\hat{\Omega} = \hat{Q}\text{skew}(\hat{S})\hat{Q}^\top + \hat{E}\hat{Q}^\top$, again by [50, Thm. 1], we have $\|\text{skew}(\hat{S}_2 - \hat{S})\|_F \leq \|\hat{E}\hat{Q}^\top\|_F = \|\hat{E}\|_F$. By (4.2), we have $\|\hat{\Omega} - \Omega\|_F \leq 2\varepsilon_m \|A\|_F$ and $\|\hat{E}\|_F \leq \hat{\tau}\varepsilon_m \|\hat{\Omega}\|_F \leq \hat{\tau}\varepsilon_m (\|\Omega\|_F + \varepsilon_m 2\|A\|_F)$. This yields

$$\begin{aligned} \|\Omega\hat{Q} - \hat{Q}\text{skew}(S)\|_F &\leq 2\|\hat{E}\|_F + 2\|\hat{\Omega} - \Omega\|_F \\ &\leq 2\hat{\tau}\varepsilon_m (\|\Omega\|_F + \varepsilon_m 2\|A\|_F) + 2(2\varepsilon_m \|A\|_F) \\ &\leq (2\hat{\tau} + 4)\varepsilon_m \|A\|_F + 4\hat{\tau}\varepsilon_m^2 \|A\|_F. \end{aligned} \quad \square$$

We can now prove the stability result of [Theorem 4.4](#).

THEOREM 4.4. *Let $A = QSQ^\top \in \mathbb{R}^{2p \times 2p}$ be a normal matrix with eigenvalues $\lambda_j e^{\pm i\theta_j}$, $\lambda_j > 0$, $\theta_j \in (0, \pi)$ for $j = 1, \dots, p$ and $\Omega := \text{skew}(A)$. Assume moreover that the imaginary parts of the eigenvalues are all distinct. Then, if the matrix $\hat{Q} \in O(2p)$ satisfies $\hat{\Omega}\hat{Q} - \hat{Q}\text{skew}(S) = E$ with $\|E\|_F \leq \tau\varepsilon_m \|A\|_F$, we have*

$$(4.7) \quad \|A\hat{Q} - \hat{Q}S\|_F \leq \tau\varepsilon_m \|A\|_F \left(1 + \max_{\substack{i,j=1,\dots,p \\ i \neq j}} \left| \frac{\lambda_i \cos(\theta_i) - \lambda_j \cos(\theta_j)}{\lambda_i \sin(\theta_i) - \lambda_j \sin(\theta_j)} \right| \right).$$

Moreover, if A is orthogonal, then

$$(4.8) \quad \max_{\substack{i,j=1,\dots,p \\ i \neq j}} \left| \frac{\lambda_i \cos(\theta_i) - \lambda_j \cos(\theta_j)}{\lambda_i \sin(\theta_i) - \lambda_j \sin(\theta_j)} \right| = \max_{\substack{i,j=1,\dots,p \\ i \neq j}} \left| \tan \left(\frac{\theta_i + \theta_j}{2} \right) \right|.$$

Proof. Let us introduce that $\Omega = Q \text{skew}(S) Q^\top$. This leads to

$$(4.9) \quad \Omega \hat{Q} - \hat{Q} \text{skew}(S) = E \iff \text{skew}(S)X - X \text{skew}(S) = \tilde{E},$$

where we defined $X := Q^\top \hat{Q}$ and $\tilde{E} := Q^\top E$. Recall that $S = \begin{bmatrix} \Lambda \cos(\Theta) & -\Lambda \sin(\Theta) \\ \Lambda \sin(\Theta) & \Lambda \cos(\Theta) \end{bmatrix}$ such that (4.9) is the concatenation of $p^2 = \frac{n^2}{4}$ small 2×2 Sylvester equations. Indeed, for every pair $i, j = 1, \dots, p$, using the short-hand notation $s_i := \sin(\theta_i)$ and $c_i := \cos(\theta_i)$, we have

$$(4.10) \quad \begin{bmatrix} 0 & -\lambda_i s_i \\ \lambda_i s_i & 0 \end{bmatrix} X_{ij} - X_{ij} \begin{bmatrix} 0 & -\lambda_j s_j \\ \lambda_j s_j & 0 \end{bmatrix} = \tilde{E}_{ij},$$

where X_{ij} and $\tilde{E}_{ij}$ are non-overlapping 2×2 blocks of X and $\tilde{E}$ respectively. If $i \neq j$, by [14, Thm. VII.2.8], we have $\|X_{ij}\|_F \leq \frac{1}{|\lambda_i s_i - \lambda_j s_j|} \|\tilde{E}_{ij}\|_F$. Therefore, acknowledging that $\|A\hat{Q} - \hat{Q}S\|_F = \|SX - XS\|_F \leq \|\text{sym}(S)X - X\text{sym}(S)\|_F + \|\tilde{E}\|_F$, we obtain

$$\begin{aligned} \|A\hat{Q} - \hat{Q}S\|_F^2 &= \sum_{i,j=1}^p \left\| \lambda_i \begin{bmatrix} c_i & -s_i \\ s_i & c_i \end{bmatrix} X_{ij} - X_{ij} \lambda_j \begin{bmatrix} c_j & -s_j \\ s_j & c_j \end{bmatrix} \right\|_F^2 \\ &\leq \sum_{i,j=1}^p \left(|\lambda_i c_i - \lambda_j c_j| \|X_{ij}\|_F + \|\tilde{E}_{ij}\|_F \right)^2 \\ &\leq \left(1 + \max_{\substack{i,j=1,\dots,p \\ i \neq j}} \left| \frac{\lambda_i c_i - \lambda_j c_j}{\lambda_i s_i - \lambda_j s_j} \right| \right)^2 \sum_{i,j=1}^p \|\tilde{E}_{ij}\|_F^2. \end{aligned}$$

The claim (4.7) follows directly by observing that $\sqrt{\sum_{i,j=1}^p \|\tilde{E}_{ij}\|_F^2} = \|\tilde{E}\|_F = \|E\|_F \leq \tau \varepsilon_m \|A\|_F$. If A is orthogonal, then $\lambda_j = 1$ for $j = 1, \dots, p$. (4.8) follows from simple trigonometric identities. $\square$

Generalizing Theorem 4.4 when the matrix $\hat{Q}$ is not exactly orthogonal is simple and does not change its conclusion. This is shown in Corollary 4.5.

COROLLARY 4.5. *Assume in Theorem 4.4 that $\hat{Q} \notin \text{O}(2p)$ but $d_{\text{O}(n)}(\hat{Q}) \leq \kappa \varepsilon_m$ with $\kappa > 0$. Then the matrix $\hat{Q}_* \in \arg \min_{Q \in \text{O}(n)} \|Q - \hat{Q}\|_F$ satisfies the hypotheses of Theorem 4.4 and $\|A\hat{Q} - \hat{Q}S\|_F \leq 2\kappa \varepsilon_m \|A\|_F + \|A\hat{Q}_* - \hat{Q}_*S\|_F$.*

Proof. First, we have

$$\begin{aligned} \|A\hat{Q} - \hat{Q}S\|_F &\leq \|A(\hat{Q} - \hat{Q}_*) - (\hat{Q} - \hat{Q}_*)S\|_F + \|A\hat{Q}_* - \hat{Q}_*S\|_F \\ &\leq 2\|\hat{Q} - \hat{Q}_*\|_F \|A\|_F + \|A\hat{Q}_* - \hat{Q}_*S\|_F \\ &\leq 2\kappa \varepsilon_m \|A\|_F + \|A\hat{Q}_* - \hat{Q}_*S\|_F. \end{aligned}$$

Moreover, $\hat{Q}_*$ satisfies the hypotheses of Theorem 4.4:

$$\begin{aligned} \|\Omega \hat{Q}_* - \hat{Q}_* \text{skew}(S)\|_F &\leq \|\hat{\Omega}(\hat{Q}_* - \hat{Q}) - (\hat{Q}_* - \hat{Q}) \text{skew}(S)\|_F + \|\hat{\Omega} \hat{Q} - \hat{Q} \text{skew}(S)\|_F \\ &\leq \|\hat{Q}_* - \hat{Q}\|_F (\|\Omega\|_F + \|\text{skew}(S)\|_F) + \tau \varepsilon_m \|A\|_F \\ &\leq (2\kappa + \tau) \varepsilon_m \|A\|_F. \end{aligned}$$

$\square$




FIGURE 2. Accuracy in double precision of [Algorithm 4.1](#) (without any correction step 8) for 100×100 orthogonal (left plot) and normal (right plot) matrices when the average phase of the eigenvalues varies from $\theta = 0$ to $\theta = \frac{\pi}{2}$. The Schur vectors are sampled with Haar distribution on $O(100)$. The eigenvalues of the matrices are sampled with phase uniformly distributed in $[\theta - \Delta\theta, \theta + \Delta\theta]$, with $\Delta\theta = 0.005$. For normal matrices, the moduli of the eigenvalues are uniformly sampled in $[0, 2]$. Results are averaged over 50 runs. The accuracy follows the behavior predicted by [Theorem 4.4](#).

[Theorem 4.4](#) is very informative on the accuracy that [Algorithm 4.1](#) will achieve when it is allowed to take its fastest route by skipping steps 5 and 7. We recall that this setting occurs with very high probability for Haar-distributed orthogonal matrices, and more generally, for random orthogonal matrices. In particular, if the bound in (4.7) is not satisfactory to the user, i.e., if $\max_{\substack{i,j=1,\dots,p \\ i \neq j}} \left| \frac{\lambda_i \cos(\theta_i) - \lambda_j \cos(\theta_j)}{\lambda_i \sin(\theta_i) - \lambda_j \sin(\theta_j)} \right| \geq t$, the sets of eigenvalues responsible for the loss of accuracy can be identified and the associated invariant subspaces can be corrected. We detail this approach in [subsection 4.5](#). In [Figure 2](#), we present the results of a numerical experiment corroborating the bounds of [Theorem 4.4](#) and assessing their relevance.

4.3. Sensitivity of the invariant subspaces: Set 2. From now on, we use the notation c, s for $\cos(\theta)$ and $\sin(\theta)$. Assume still that $A = QSQ^\top \in \mathbb{R}^{2p \times 2p}$ ($r = 0$) is a normal matrix with eigenvalues $\lambda_j e^{\pm i\theta_j}$, $\lambda_j > 0$, $\theta_j \in (0, \pi)$ for $j = 1, \dots, p$. We next consider the case of eigenvalues that have repeated imaginary

parts (Set 2). Assume a set of indices $\mathcal{J} \subset \{1, 2, \dots, p\}$, $|\mathcal{J}| = m$, such that for all $j \in \mathcal{J}$, we have $\lambda_j s_j = \sigma > 0$. In other words, we can numerically compute $\hat{V}$ such that $\hat{\Omega}\hat{V} - \hat{V}\hat{\sigma} \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} = \hat{E}$ with $\|\hat{E}\|_F \lesssim \varepsilon_m \|\hat{\Omega}\|_F$. In view of Corollary 4.5, w.l.o.g., we can consider $\hat{V} \in \text{St}(n, 2m)$. We set $\mathcal{I} := \{1, 2, \dots, p\} \setminus \mathcal{J}$. The proof of Theorem 4.4 fails when [14, Thm. VII.2.8] is used after (4.10) because, for $j_1, j_2 \in \mathcal{J}$, $\lambda_{j_1} s_{j_1} = \lambda_{j_2} s_{j_2}$.

As expected, the accuracy of the invariant subspace computation depends on the separation between $\pm i\sigma$ and the other eigenvalues values of Ω . In subsection 4.5 we use Theorem 4.6 to design an algorithm that satisfies a desired accuracy.

THEOREM 4.6. *Let $A = QSQ^\top \in \mathbb{R}^{2p \times 2p}$ be a normal matrix with eigenvalues $\lambda_j e^{\pm i\theta_j}$, $\lambda_j > 0$, $\theta_j \in (0, \pi)$ for $j = 1, \dots, p$ and $\Omega := \text{skew}(A)$. Assume there is $\mathcal{J} \subset \{1, \dots, p\}$, $|\mathcal{J}| = m$, such that for all $j \in \mathcal{J}$, we have $\lambda_j s_j = \sigma$. Set $\mathcal{I} := \{1, \dots, p\} \setminus \mathcal{J}$. If the matrix $\hat{V} \in \text{St}(2p, 2m)$ satisfies $\hat{\Omega}\hat{V} - \hat{V}\sigma \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} = E$ with $\|E\|_F \leq \tau \varepsilon_m \|A\|_F$, then we have*

$$\begin{aligned} & \min_{R \in \text{O}(2m)} \|A(\hat{V}R) - (\hat{V}R) \begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix}\|_F \\ & \leq \tau \varepsilon_m \|A\|_F \left(1 + \frac{\max_{i \in \mathcal{I}, j \in \mathcal{J}} |\lambda_i c_i - \lambda_j c_j|}{\min_{i \in \mathcal{I}} |\lambda_i s_i - \sigma|} + \varepsilon_m \tau \|A\|_F \frac{\max_{j_1, j_2 \in \mathcal{J}} |\lambda_{j_1} c_{j_1} - \lambda_{j_2} c_{j_2}|}{\min_{i \in \mathcal{I}} |\lambda_i s_i - \sigma|^2} \right). \end{aligned}$$

Proof. Define $X := Q^\top \hat{V}$, then it follows that

$$\text{skew}(S)X - X\sigma \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} = Q^\top E =: \tilde{E}.$$

There exist two permutations $P_l \in \text{O}(n)$, $P_r \in \text{O}(2m)$ such that $P_l^\top X P_r = \begin{bmatrix} X_{\mathcal{I}, \mathcal{J}} \\ X_{\mathcal{J}, \mathcal{J}} \end{bmatrix}$ where the block $X_{\mathcal{I}, \mathcal{J}}$ comprises all 2×2 blocks X_{ij} such that for $i \in \mathcal{I}, j \in \mathcal{J}$. We have

$$(4.11) \quad \begin{bmatrix} 0 & -\lambda_i s_i \\ \lambda_i s_i & 0 \end{bmatrix} X_{ij} - X_{ij} \begin{bmatrix} 0 & -\sigma \\ \sigma & 0 \end{bmatrix} = (P_l^\top \tilde{E} P_r)_{ij}.$$

If it hold that $\tilde{E} = 0$, then we would have $X_{\mathcal{I}, \mathcal{J}} = 0$ and $X_{\mathcal{J}, \mathcal{J}} \in \text{O}(2m)$. However, since $\tilde{E} \neq 0$ and $X \in \text{St}(n, 2m)$, we have

$$\begin{aligned} & X_{\mathcal{I}, \mathcal{J}}^\top X_{\mathcal{I}, \mathcal{J}} + X_{\mathcal{J}, \mathcal{J}}^\top X_{\mathcal{J}, \mathcal{J}} = I_{2m} \implies \|X_{\mathcal{J}, \mathcal{J}}^\top X_{\mathcal{J}, \mathcal{J}} - I_{2m}\|_F \leq \|X_{\mathcal{I}, \mathcal{J}}\|_F^2 \\ & \text{and by (4.11), } \|X_{\mathcal{I}, \mathcal{J}}\|_F^2 \leq \sum_{i \in \mathcal{I}, j \in \mathcal{J}} \frac{1}{|\lambda_i s_i - \sigma|^2} \|(P_l^\top \tilde{E} P_r)_{ij}\|_F^2 \leq \frac{\|E\|_F^2}{\min_{i \in \mathcal{I}} |\lambda_i s_i - \sigma|^2}. \end{aligned}$$

The last inequality is obtained using [14, Thm. VII.2.8]. The deviation of $X_{\mathcal{J}, \mathcal{J}}$ from orthogonality can thus be quantified. In addition, by taking the SVD $X_{\mathcal{J}, \mathcal{J}} = U_X \Sigma_X V_X^\top$, we can write $\|X_{\mathcal{J}, \mathcal{J}}^\top X_{\mathcal{J}, \mathcal{J}} - I_{2m}\|_F = \|\Sigma_X^2 - I_{2m}\|_F \geq \|\Sigma_X - I_{2m}\|_F$, where the last inequality stands because $|\alpha^2 - 1| \geq |\alpha - 1|$ for all $\alpha \geq 0$. Define $\begin{bmatrix} \tilde{S}_{2(p-m)} & 0 \\ 0 & \tilde{S}_{2m} \end{bmatrix} := P_l^\top S P_l$, then for all $R \in \text{O}(2m)$, we have

$$\begin{aligned} (4.12) \quad & \|A\hat{V}R - \hat{V}R\tilde{S}_{2m}\|_F \leq \|\text{sym}(S)XR - XR\text{sym}(\tilde{S}_{2m})\|_F + \|E\|_F \\ & \leq \|\text{sym}(\tilde{S}_{2(p-m)})X_{\mathcal{I}, \mathcal{J}}R - X_{\mathcal{I}, \mathcal{J}}R\text{sym}(\tilde{S}_{2m})\|_F \\ & \quad + \|\text{sym}(\tilde{S}_{2m})X_{\mathcal{J}, \mathcal{J}}R - X_{\mathcal{J}, \mathcal{J}}R\text{sym}(\tilde{S}_{2m})\|_F + \|E\|_F. \end{aligned}$$

The first term of (4.12) can be bounded by

$$\begin{aligned}
\|\text{sym}(\tilde{S}_{2(p-m)})X_{\mathcal{I},\mathcal{J}}R - X_{\mathcal{I},\mathcal{J}}R\text{sym}(\tilde{S}_{2m})\|_{\text{F}} &\leq \max_{i \in \mathcal{I}, j \in \mathcal{J}} |\lambda_i c_i - \lambda_j c_j| \|X_{\mathcal{I},\mathcal{J}}R\|_{\text{F}} \\
(4.13) \qquad \qquad \qquad &\leq \frac{\max_{i \in \mathcal{I}, j \in \mathcal{J}} |\lambda_i c_i - \lambda_j c_j|}{\min_{i \in \mathcal{I}} |\lambda_i s_i - \sigma|} \|E\|_{\text{F}}.
\end{aligned}$$

Moreover, the second term can also be bounded by

$$\begin{aligned}
&\|\text{sym}(\tilde{S}_{2m})X_{\mathcal{J},\mathcal{J}}R - X_{\mathcal{J},\mathcal{J}}R\text{sym}(\tilde{S}_{2m})\|_{\text{F}} \\
&= \|\text{sym}(\tilde{S}_{2m})(X_{\mathcal{J},\mathcal{J}}R - I_{2m}) - (X_{\mathcal{J},\mathcal{J}}R - I_{2m})\text{sym}(\tilde{S}_{2m})\|_{\text{F}} \\
(4.14) \qquad \qquad \qquad &\leq \max_{j_1, j_2 \in \mathcal{J}} |\lambda_{j_1} c_{j_1} - \lambda_{j_2} c_{j_2}| \|X_{\mathcal{J},\mathcal{J}}R - I_{2m}\|_{\text{F}}.
\end{aligned}$$

In particular, for $R = V_X U_X^\top$ which minimizes $\|X_{\mathcal{J},\mathcal{J}}R - I_{2m}\|_{\text{F}}$, we have

$$(4.15) \qquad \|X_{\mathcal{J},\mathcal{J}}R - I_{2m}\|_{\text{F}} = \|\Sigma_X - I_{2m}\|_{\text{F}} \leq \frac{\|E\|_{\text{F}}^2}{\min_{i \in \mathcal{I}} |\sigma - \lambda_i s_i|^2}.$$

The theorem follows directly by inserting (4.13), (4.14) and (4.15) in (4.12) and acknowledging that $\|E\|_{\text{F}} \leq \tau \varepsilon_{\text{m}} \|A\|_{\text{F}}$. $\square$

4.4. Sensitivity of the invariant subspaces: Set 3. The case of real eigenvalues (Set 3) can be handled in a similarly to the case of repeated imaginary parts. Indeed, the proof of Theorem 4.7 is very close to that of Theorem 4.6, if $\begin{bmatrix} D_m & -\sigma I_m \\ \sigma I_m & D_m \end{bmatrix}$ is replaced by $\check{\Lambda}$. For this reason, the proof is given in Appendix C.

THEOREM 4.7. *Let $A = QSQ^\top \in \mathbb{R}^{n \times n}$ be a normal matrix with real eigenvalues $\check{\lambda}_k$ for $k \in \mathcal{K} := \{1, \dots, r\}$ and $\Omega := \text{skew}(A)$. Moreover, assume $\lambda_j e^{\pm i\theta_j}$ is an eigenvalue of A for $j \in \mathcal{I}$. If the matrix $\hat{V} \in \text{St}(n, r)$ satisfies $\Omega \hat{V} = \hat{E}$ with $\|E\|_{\text{F}} \leq \tau \varepsilon_{\text{m}} \|A\|_{\text{F}}$, then we have*

$$\begin{aligned}
&\min_{\check{R} \in \text{O}(2m)} \|\hat{A} \hat{V} \check{R} - \hat{V} \check{R} \check{\Lambda}\|_{\text{F}} \\
&\leq \tau \varepsilon_{\text{m}} \|A\|_{\text{F}} \left(1 + \frac{\max_{i \in \mathcal{I}, k \in \mathcal{K}} |\lambda_i c_i - \check{\lambda}_k|}{\min_{i \in \mathcal{I}} |\lambda_i s_i|} + \varepsilon_{\text{m}} \tau \|A\|_{\text{F}} \frac{\max_{k_1, k_2 \in \mathcal{K}} |\check{\lambda}_{k_1} - \check{\lambda}_{k_2}|}{\min_{i \in \mathcal{I}} |\lambda_i s_i|^2} \right).
\end{aligned}$$

Proof. See Appendix C. $\square$

4.5. Analysis of steps 5 to 8. The previous subsections determined the sensitivity of the computation of the Schur vectors of A by the Schur vectors of Ω . We can now provide conditions on steps 5 to 8 of Algorithm 4.1 to achieve a desired accuracy.

Step 5 requires identifying every approximately invariant subspace of Ω spanned by $\hat{V}$, i.e., $\hat{V}$ such that $\|\Omega \hat{V} - \hat{V} \sigma J_{2m}\|_{\text{F}} \lesssim \varepsilon_{\text{m}} \|A\|_{\text{F}}$, $m > 1$. In view of Corollary 4.5, we assume $\hat{V} \in \text{St}(n, 2m)$. To recover the Schur vectors of A , we need to compute $M := \hat{V}^\top A \hat{V}$. In exact arithmetic, $2m$ is the multiplicity of the singular value σ of Ω . In floating point arithmetic, $2\hat{m}$ must be the size of the cluster of singular values of $\hat{\Omega}$. If we take $\hat{V}$ as the $2\hat{m}$ computed Schur vectors associated to a δ -cluster of $\hat{\Omega}$, in view of Theorem 4.6, it holds that $\|\hat{A} \hat{V} - \hat{V} M\|_{\text{F}} \lesssim \frac{\varepsilon_{\text{m}}}{\delta} \|A\|_{\text{F}}$. If δ is chosen such that the δ -cluster is well-separated from the others and if the RSD $M \approx R \tilde{S}_{2m} R^\top$ is such that $\|M - R \tilde{S}_{2m} R^\top\|_{\text{F}} \lesssim \varepsilon_{\text{m}} \|A\|_{\text{F}}$, then after step 5, we have $\|\hat{A} \hat{V} R - \hat{V} R \tilde{S}_{2m}\|_{\text{F}} \lesssim \frac{\varepsilon_{\text{m}}}{\delta} \|A\|_{\text{F}}$ and $d_{\text{St}(n, 2m)}(\hat{V} R) \lesssim \varepsilon_{\text{m}}$. Notice here that the effective accuracy that is

obtained is in fact function of the largest δ such that the δ -cluster of $\widehat{\Omega}$ remains the same, rather than the user specified value δ . As mentioned in [section 3](#), several stable algorithms exist for [Routine 3](#) at step 5.2.

Step 6 obtains the real parts of the eigenvalues by computing Rayleigh quotients. It is known that the Rayleigh quotient offers the square of the accuracy of the eigenvector, see, e.g., [69, Lecture 27]. Therefore, by [Theorem 4.4](#), step 6 provides a worst-case $\mathcal{O}(\varepsilon_m \|A\|_F)$ -absolute accuracy on the real parts of the eigenvalues if $\max_{i,j} \left| \frac{\lambda_i c_i - \lambda_j c_j}{\lambda_i s_i - \lambda_j s_j} \right| \lesssim \frac{1}{\sqrt{\varepsilon_m}}$. It suggests considering δ -clusters of Ω with $\delta \leq \sqrt{\varepsilon_m}$ to guarantee high absolute accuracy on the eigenvalues.

Step 7 computes the real eigenvalues and the eigenvectors of A . This starts by estimating r by $\widehat{r}$. By [Theorem 4.7](#), considering a δ_r -cluster of the eigenvalues of $\widehat{\Omega}$ around zero will provide a worst-case $\mathcal{O}(\frac{\varepsilon_m}{\delta_r} \|A\|_F)$ -absolute accuracy on the real eigenspace of A . It is essential not to underestimate r , i.e., to have $\widehat{r} \geq r$. Therefore, δ_r should always be increased until the δ_r -cluster around zero is well separated from the other eigenvalues, if needed at the cost of the symmetry of $H := \widehat{Q}_{\widehat{p}+1:\widehat{p}+\widehat{r}}^\top A \widehat{Q}_{\widehat{p}+1:\widehat{p}+\widehat{r}}$. If $\|\text{skew}(H)\|_F \leq \varepsilon_m \|H\|_F$, then H can be considered exactly symmetric without loss of accuracy since the skew-symmetric part can be considered as a negligible perturbation. If H is symmetric, the decomposition at step 7.2a using, e.g., `syevr`, yields

$$H + \Delta H = \check{R} \check{\Lambda} \check{R}^\top \quad \text{with} \quad \|\Delta H\|_F \lesssim \varepsilon_m \|H\|_F,$$

with $d_{O(\widehat{r})}(\check{R}) \lesssim \varepsilon_m$ [6]. This yields $\|A \widehat{Q}_{\widehat{p}+1:\widehat{p}+\widehat{r}} \check{R} - \widehat{Q}_{\widehat{p}+1:\widehat{p}+\widehat{r}} \check{R} \check{\Lambda}\|_F \lesssim \frac{\varepsilon_m}{\delta_r} \|A\|_F$ with $d_{\text{St}(n,\widehat{r})}(\widehat{Q}_{\widehat{p}+1:\widehat{p}+\widehat{r}} \check{R}) \lesssim \varepsilon_m$. Only in certain circumstances will symmetric eigensolvers achieve relative accuracy [29, 30, 33, 72]. If H is not numerically symmetric, it means that $\widehat{r}$ overestimated r , [Routine 3](#) should then be preferred to [Routine 2](#) to maintain accuracy.

Step 8 can finally be considered to improve the accuracy of the RSD to a desired tolerance $\mu = \varepsilon_m t$, $t \geq 1$ if this accuracy has not yet been attained. By combining [Theorems 4.4, 4.6](#) and [4.7](#), we deduce that $\|A \widehat{Q} - \widehat{Q} \widehat{S}\|_F \lesssim \left(1 + \max_{i \neq j,k} \left\{ \frac{|\lambda_i c_i - \lambda_j c_j|}{|\lambda_i s_i - \lambda_j s_j|}, \frac{|\check{\lambda}_k - \lambda_j c_j|}{|\lambda_j s_j|} \right\}\right) \varepsilon_m \|A\|_F$. If $\max_{i \neq j,k} \left\{ \frac{|\lambda_i c_i - \lambda_j c_j|}{|\lambda_i s_i - \lambda_j s_j|}, \frac{|\check{\lambda}_k - \lambda_j c_j|}{|\lambda_j s_j|} \right\} \geq t$, the groups of eigenvalues responsible for making the constant larger than t can be identified. Then, to restore the desired accuracy, one should repeat step 5 with $\delta = \frac{1}{t}$ on the invariant subspaces associated with the $\frac{1}{t}$ -clusters of $\widehat{\Omega}$. Because the matrix $\widehat{Q}$ is already close to the Schur vectors of A , the matrix M at step 5 is a small perturbation of the Schur form. For such matrices, a Jacobi algorithm such as [83] is efficient since it converges quadratically. One sweep is enough in practice to obtain the desired accuracy. However, for specific matrices and t too small, the $\frac{1}{t}$ -clusters may include many eigenvalues. The user should be aware that step 8 may then require a significant additional computational work. In double precision, the authors recommend using $t \geq 100$ based on empirical experience. If the output of [Routine 3](#) satisfies $\|M - R \widetilde{S}_{2m} R^\top\|_F \lesssim \varepsilon_m \|A\|_F$, then finally, we obtain the decomposition $\|A \widehat{Q} - \widehat{Q} \widehat{S}\|_F \lesssim \mu \|A\|_F$ with $d_{O(n)}(\widehat{Q}) \lesssim \varepsilon_m$.

5. Numerical experiments on the accuracy. In this section, we present the results of several numerical experiments on the accuracy of [Algorithm 4.1](#) for a diverse set of $n \times n$ problems. The results are given in [Table 1](#) and [Figure 3](#). For five different distributions of the eigenvalues, we report the average accuracy of [Algorithm 4.1](#) over 100 runs. Codes reproducing the numerical experiments can be found at <https://github.com/smataigne/NormalEVP.jl>.

- Experiment E1 verifies that in the setting of [Theorem 4.4](#), if the bound of (4.7) remains small, then [Algorithm 4.1](#) achieves high accuracy on the RSD.
- Experiment E2 tests a random normal scenario representing an “average normal case”. In this case,

n	$\frac{\ A\hat{Q}-\hat{Q}\hat{S}\ _F}{\ A\ _F}$		$\frac{\ \hat{Q}^\top \hat{Q}-I\ _F}{\sqrt{n}}$		$\frac{\ \text{diag}(S-\hat{S})\ _F}{1+\ \text{diag}(S)\ _F}$	
	nrmschur	schur	nrmschur	schur	nrmschur	schur
E1: $\lambda = 1$, $\theta \sim \mathcal{U}(0, \frac{\pi}{4})$: Best accuracy scenario on $\text{SO}(n)$.						
10	8.0e-16	1.1e-15	6.7e-16	1.2e-15	3.8e-16	4.6e-16
32	1.3e-15	1.9e-15	1.2e-15	2.0e-15	6.6e-16	8.1e-16
100	1.5e-15	3.5e-15	1.6e-15	3.7e-15	6.6e-16	1.6e-15
316	1.5e-15	7.2e-15	2.0e-15	7.8e-15	5.8e-16	4.3e-15
1000	1.7e-15	2.3e-14	2.8e-15	2.3e-14	6.4e-16	1.9e-14
E2: $\lambda \sim \mathcal{U}(0, 2)$, $\theta \sim \mathcal{U}(0, \pi)$: Random normal scenario.						
10	2.7e-15	1.4e-15	6.2e-16	1.3e-15	3.4e-16	5.1e-16
32	1.6e-14	2.6e-15	1.1e-15	2.2e-15	6.2e-16	1.0e-15
100	1.2e-13	4.2e-15	1.5e-15	3.8e-15	7.2e-16	1.9e-15
316	4.0e-13	7.4e-15	1.9e-15	7.0e-15	6.0e-16	3.9e-15
1000	1.6e-12	1.3e-14	2.6e-15	1.2e-14	6.2e-16	8.4e-15
E3: $\lambda, \bar{\lambda} \sim \mathcal{U}(0, 2)$, $\theta \sim \mathcal{U}(0, \pi)$: Random normal scenario, 20% real eigenvalues.						
10	2.3e-15	1.4e-15	5.7e-16	1.3e-15	2.9e-16	4.9e-16
32	1.3e-14	2.6e-15	1.2e-15	2.2e-15	5.8e-16	1.0e-15
100	6.7e-14	4.1e-15	3.8e-15	3.7e-15	6.5e-16	1.9e-15
316	2.7e-13	6.9e-15	1.2e-14	6.5e-15	5.8e-16	3.7e-15
1000	8.7e-13	1.1e-14	2.9e-14	1.1e-14	5.7e-16	7.3e-15
E4: $\lambda \sim \mathcal{U}(0, 2)$, $\theta \sim \mathcal{U}(0, \pi)$: Random normal scen., 20% repeated imaginary parts.						
10	3.5e-15	1.3e-15	6.6e-16	1.3e-15	3.4e-16	4.6e-16
32	2.4e-14	2.0e-15	1.1e-15	1.9e-15	5.8e-16	7.5e-16
100	1.2e-13	3.6e-15	1.5e-15	3.5e-15	6.9e-16	1.4e-15
316	4.1e-13	7.4e-15	2.1e-15	7.9e-15	8.4e-16	4.2e-15
1000	1.5e-12	2.3e-14	3.1e-15	2.4e-14	1.2e-15	1.9e-14
E5: $\lambda \sim \mathcal{U}(0, 2)$, $\theta \sim \pi\sqrt{\varepsilon_m}\mathcal{N}(1, 1)$: Worst-case normal scenario.						
10	5.1e-10 $\xrightarrow{\text{step 8}}$ 1.2e-15	1.1e-15	9.9e-16	1.1e-15	5.3e-16	5.3e-16
32	2.1e-9 $\rightarrow$ 2.9e-15	1.8e-15	2.2e-15	1.9e-15	1.2e-15	6.1e-16
100	5.6e-9 $\rightarrow$ 5.8e-15	3.7e-15	4.1e-15	3.6e-15	2.9e-15	1.7e-15
316	9.7e-9 $\rightarrow$ 1.4e-14	6.1e-15	6.8e-15	6.5e-15	3.5e-15	2.5e-15
1000	1.1e-8 $\rightarrow$ 8.4e-14	9.7e-15	1.1e-14	1.0e-14	5.8e-15	3.8e-15

TABLE 1

Numerical experiment on the accuracy in double precision of *nrmschur* (Algorithm 4.1) and *schur* (LAPACK's *gees*) to compute $A \approx \hat{Q}\hat{S}\hat{Q}^\top$. Step 8 is considered only for the last case where exactly one Jacobi sweep is performed. We chose $\delta, \delta_r, \frac{1}{t}$ equal to $\sqrt{\varepsilon_m}$ to define clusters. S is the ground truth real Schur form. Each reported accuracy is averaged over 100 runs. The normal matrices A are randomly sampled with eigenvalues of different distributions specified in the table. The Schur vectors are always sampled with a Haar distribution on $\text{O}(n)$. For all $a, b \in \mathbb{R}$, $a\text{-}b$ denotes $a \cdot 10^{-b}$.

a moderate loss of accuracy of the Schur decomposition occurs as n grows because the separation between the imaginary parts of the eigenvalues decreases with n . However, the absolute accuracy on the eigenvalues remains close to machine precision thanks to the Rayleigh quotients.

- Experiments E3 and E4 are similar to E2, but enforce nonempty sets [Set 2](#) and [Set 3](#), respectively. The accuracy of the RSD remains equivalent to E2, as well as the absolute accuracy on the eigenvalues.
- Finally, experiment E5 tests the worst case for [Algorithm 4.1](#) in terms of accuracy. The skew-




FIGURE 3. Evolution of the accuracy in double precision of *nrmschur* (Algorithm 4.1) without correction step 8. The matrix size n varies from 10 to 1000. The plots shows results for experiments E1 to E4 from section 5. The results are averaged over 100 runs.

symmetric part Ω is small such that the eigenvalues of Ω are relatively close and induce sensitivity of the invariant subspaces. Moreover, Ω is large enough such that eigenvalues of Ω are not considered as $\sqrt{\varepsilon_m}$ -clusters and the accuracy is not restored by step 5. However, the high accuracy on the eigenvalues is maintained by the Rayleigh quotients. We show that one sweep of the Jacobi method [83] allows to restore high accuracy of the RSD.

In Table 1, the deviation from orthogonality is measured by the simple criterion $\|\hat{Q}^\top \hat{Q} - I\|_F$. We recall in Lemma A.4 that this measure is equivalent to $d_{O(n)}(\hat{Q})$ when it is small. The results of Figure 3 show that for the experiments E2 to E4, the accuracy of the RSD grows proportionally to the matrix size n . Moreover, for the experiment E1, the accuracy is almost insensitive to n .

6. Complexity analysis and running time performance. We analyze the best and worst case computational complexities of Algorithm 4.1. We exclude the cases where the matrix A is either purely skew-symmetric or symmetric since Algorithm 4.1 reduces to obtaining, respectively, the skew-symmetric RSD with Routine 1 or the symmetric EVD with Routine 2. These ideal cases are not relevant to our discussion. By analogy with algorithms such as quicksort, we will show that worst-case complexity analysis is not the most relevant here, as worst cases are very unlikely to occur.

Since the cost of the method is $\mathcal{O}(n^3)$, the $\mathcal{O}(n^2)$ costs are neglected. Following the notation from [69, Lec. 8], we say that a method requires $\sim g(n, p, r)$ flops if

$$(6.1) \quad \lim_{n, p, r \rightarrow \infty} \frac{\text{number of flops}}{g(n, p, r)} = 1.$$

The numerical experiments presented in this section are performed on a processor *Intel(R) Core(TM) i7-8750H CPU @ 2.20GHz* in single-threaded framework. Codes reproducing the numerical experiments can be found at <https://github.com/smataigne/NormalEVP.jl>.

6.1. The intuitive (but not) best-case complexity. Intuitively, the best case for Algorithm 4.1 is when all eigenvalues are complex and none of the imaginary parts are clustered (relatively to ε_m). The

assumptions of [Theorem 3.3](#) are then matched with $k = \lfloor \frac{n}{2} \rfloor = p$ ($r = 1$ for n odd). We recall that for Haar-distributed special orthogonal matrices [\[7, 64\]](#), this case occurs with probability 1 [\[34\]](#). Step 1 is negligible since it requires $\mathcal{O}(n^2)$ flops.

The cost of step 2 is $\sim \frac{4}{3}n^3$ for the skew-symmetric tridiagonal reduction [\[69, Lec. 26\]](#), plus an additional cost of $\sim \frac{4}{3}n^3$ to assemble the Householder reflectors [\[41, Sec. 5.1.6\]](#). Indeed, given $M \in \mathbb{R}^{k \times k}$ and $v \in \mathbb{R}^k$, performing $(I - vv^\top)M$ requires $\sim 4k^2$ flops. Therefore, assembling all reflectors of size $k = 1, \dots, n-1$ has a cost $\sim \sum_{k=1}^{n-1} 4k^2 \sim \frac{4}{3}n^3$ flops.

At step 3, the call to [Routine 1](#) can be neglected since an $\mathcal{O}(p^2)$ implementation exists [\[45\]](#).

Step 4 requires $\sim 4np^2$ flops, step 5 is skipped and step 6 costs $\sim 2n^2p$ flops. Step 7 is skipped (or costs $\mathcal{O}(n^2)$ if $r = 1$). The total cost is thus

$$(6.2) \quad \sim \frac{8}{3}n^3 + 4np^2 + 2n^2p \text{ flops.}$$

Given $p = \lfloor \frac{n}{2} \rfloor$, the total cost is $\sim \frac{14}{3}n^3$, among which $\sim \frac{6}{3}n^3$ are level 3 BLAS operations [\[18\]](#). Moreover, the $\sim \frac{8}{3}n^3$ flops from the tridiagonalization also admit blocked implementations, maximizing the use of level 3 BLAS (see, e.g., [SkewLinearAlgebra.jl](#)).

In the best case, the total cost of [Algorithm 4.1](#) is thus equivalent to a Hessenberg factorization, which is known to cost $\frac{10}{3}n^3$ flops to obtain the H factor [\[69, Lec. 26\]](#) and $\sim \frac{4}{3}n^3$ to assemble the Householder reflectors, hence $\sim \frac{14}{3}n^3$ flops in total. This result is [Algorithm 4.1](#)'s best advantage over methods that include an initial reduction of the matrix A to Hessenberg form [\[4, 43, 73, 74\]](#), while the latter methods usually offer better accuracy on the eigenvectors in general.

[Figure 4](#) illustrates the performance of [nrmschur](#) ([Algorithm 4.1](#)) compared to Julia calls to LAPACK's Hessenberg factorization ([gehrd+orghr](#)) and [schur](#) routine ([gees](#)). We also include the performance of our Julia implementation of the UHQR algorithm [\[44\]](#). UHQR's computational complexity for the eigenvalues is $\sim \frac{10}{3}n^3$ for the Hessenberg reduction plus $\mathcal{O}(n^2)$ for QR iterations, thus faster than LAPACK's eigenvalue routine ([hseqr](#)). However, if the Schur/eigen-vectors are desired, the complexity of UHQR is $\sim \frac{14}{3}n^3$ for the Hessenberg reduction plus $\mathcal{O}(n^3)$ because Givens rotations of the QR are accumulated.

With due caution in interpreting complexity from timing experiments, the numerical results in [Figure 4](#) corroborate that the cost of [Algorithm 4.1](#) is comparable to that of a Hessenberg factorization. Moreover, it outperforms [gees](#) by a factor 10 for $n \approx 100$ and a factor 2 for $n = 10000$. We can observe that UHQR (only for eigenvalues) lies above [Algorithm 4.1](#) however, a FORTRAN implementation of UHQR would be faster. Nonetheless, this highlights that because [Algorithm 4.1](#) relies exclusively on high performance software, namely LAPACK, it can be efficiently implemented in any computer language where an interface to this software is available.

6.2. The true best-case complexity. If we consider $r > 1$ real eigenvalues, then there is an additional cost of $\sim 2n^2r + 2nr^2$ flops at step 7.1. The call to [Routine 2](#) at step 7.2a requires $\sim (\frac{4}{3} + \frac{4}{3})r^3$ flops for tridiagonalization plus $\mathcal{O}(r^2)$ flops to solve the tridiagonal EVP [\[30\]](#). Finally, step 7.3 demands $\sim 2nr^2$ flops. However, notice that $p := \frac{n-r}{2}$ decreases when r rises. Hence, all costs from [\(6.2\)](#) depending on p decrease accordingly. For r sufficiently small, we show that the decrease of p has a larger effect on the operation count than the increase of r . For $r > 0$, a pessimistic (but realistic, see [Remark 6.1](#)) total cost of [Algorithm 4.1](#) is

$$(6.3) \quad \sim \frac{8}{3}(n^3 + r^3) + n(n-r)^2 + n^2(n-r) + 2n^2r + 4nr^2 \text{ flops.}$$




FIGURE 4. Top: Benchmark of `nrmschur` (Algorithm 4.1) on Haar-distributed matrices of $\text{SO}(n)$ w.r.t. the Julia routines `schur` (LAPACK’s `gees`) and `hessenberg` (LAPACK’s `gehrd` and `orghr`). The performance of the UHQR algorithm for eigenvalues [44] as well as a representative $\mathcal{O}(n^3)$ are also drawn. The plot shows experiments for $n \in \{10, 32, 100, 316, 1000, 3162, 10000\}$. Bottom: the same results are displayed but divided by the computational time of `hessenberg`.

By setting $r = \alpha n$ with $\alpha \in [0, 1]$, we obtain the complexity

$$(6.4) \quad \sim \left(\frac{8}{3} \alpha^3 + 5\alpha^2 - \alpha + \frac{14}{3} \right) n^3 \text{ flops.}$$

The polynomial in α achieves a local minimum at $\alpha = \frac{\sqrt{5^2+8}-5}{8} \approx 0.093$ for an optimal cost of $4.61n^3$ flops, instead of $\sim \frac{14}{3}n^3 \approx 4.66n^3$ when $\alpha = 0$. Although, the variation of the constant is too small to yield significant gain in performance, the take-home conclusion is that “a few” real eigenvalues (up to 20%) do not impact the performance negatively. The numerical experiment of Figure 5 confirms the relevance of this operation count. The average runtime fits the expected behavior when α varies.

REMARK 6.1. When the number r of real eigenvalues grows, exact arithmetic could use only $n - r$ Householder reflectors to tridiagonalize $\text{skew}(A)$ instead of $n - 1$. Indeed, since $\text{skew}(A)$ has rank $n - r$




FIGURE 5. Numerical experiments over 600×600 random normal matrices with proportion α of real eigenvalues. For each value α ($\Delta\alpha = 0.02$), 200 matrices are sampled. The Schur vectors are Haar-distributed on $O(n)$, the real eigenvalues are normally distributed with unit variance and pairs of conjugate eigenvalues are sampled with uniformly distributed phases in $[0, \pi]$ and radii in $[0, 2]$. The theoretical complexity from (6.4) is represented by the dotted green line. The polynomial $\frac{8}{3}\alpha^3 + 5\alpha^2 - \alpha + \frac{14}{3}$ is scaled by a constant factor to match the intercept ($\alpha = 0$) of the numerical experiments.

and since each column of a tridiagonal matrix with a non-zero subdiagonal element is independent from all previous columns, the (exact) tridiagonal form can not have more than $n - r$ non-zero subdiagonal elements. In short, tridiagonalization is a rank-revealing method. The most expensive reflectors being the $n - k$ ones of largest size, the worst case complexity in exact arithmetic of the tridiagonalization is $\sim \frac{8}{3}(n^3 - r^3)$ instead of $\sim \frac{8}{3}n^3$. However, due to numerical roundoff, logical zeros are perturbed. Given random matrices $U, V \in \mathbb{R}^{n \times l}$, $l \ll n$, it can be verified that LAPACK's methods `sytrd` on $UU^\top$ and `gehrd` on $UV^\top$ produce $n - 1$ non-trivial reflectors. This validates in practice the cost $\sim \frac{8}{3}n^3$ considered in (6.3).

6.3. The worst-case complexity. The worst case for Algorithm 4.1 is when all eigenvalues are complex (all but one if n is odd) and their imaginary parts form two conjugate clusters. This case maximizes the additional work of Routine 3 at step 5. For simplicity, assume n even such that $r = 0$ and at step 5, $\Sigma = \sigma I_p$ such that $V = \hat{Q}$. Essentially, all steps but step 5.2 are a waste a computational resources.

Assuming we ignore the (potentially inaccurate) symmetric skew-Hamiltonian structure, step 5.1 requires $\sim 4n^3$ flops. Step 5.2 requires $\sim \frac{14}{3}n^3$ and $\mathcal{O}(n^3)$ flops for, respectively, the Hessenberg reduction and Routine 3. Step 5.3 demands $\sim 2n^3$ flops. Algorithm 4.1 returns after step 5 since $\Lambda \cos(\Theta)$ was computed at step 5.2. Taking into account the $\sim \frac{8}{3}n^3 + \mathcal{O}(n^2) + n^3$ flops of steps 2 to 4, the total cost is $\sim \frac{43}{3}n^3 + \mathcal{O}(n^3)$, thus a large increase of the computational cost.

Figure 6 reports an experiment that is similar to the best case (Figure 4) but for the worst case. Without surprise, the average runtime of Algorithm 4.1 exceeds that of `gees` since `gees` becomes a subroutine of Algorithm 4.1. The practical advantage of Algorithm 4.1 stems from the unlikelihood of having $\Sigma = \sigma I_p$. Most likely, the multiplicity of the entries of Σ , if any, remains small compared to n . Then, the cost of step 5 remains negligible and the performance does not significantly differ from the $\sim \frac{14}{3}n^3$ flops of the best case.

7. Application: Riemannian barycenter on manifolds. The Riemannian barycenter is a classical problem from statistics on manifolds, in particular on $SO(n)$ and $St(n, p)$, see, e.g., [15, 54, 55, 82, 85]. For




FIGURE 6. Worst-case benchmark, as in Figure 4, with $n \times n$ random normal matrices such that $\Sigma = \sigma I_p$, $\sigma \in \mathbb{R}$ and $D_p \sim \mathcal{N}(0, 1)$. The Schur vectors are sampled with Haar-distribution on $O(n)$ [64]. The plot shows $n \in \{10, 32, 100, 316, 1000\}$.

both manifolds, intensive computations of matrix logarithms of orthogonal matrices are required, as detailed next for $SO(n)$. Given a collection of N matrices $\{X_i\}_{i=1}^N$ where $X_i \in SO(n)$ for $i = 1, \dots, N$, one is interested in finding a Riemannian center of mass $X_c \in SO(n)$ of the collection, i.e., a matrix satisfying

$$(7.1) \quad X_c \in \operatorname{argmin}_{X \in SO(n)} \frac{1}{2} \sum_{i=1}^N d(X_i, X)^2,$$

where $d(X_a, X_b) := \|\log(X_a^\top X_b)\|_F$ is the Riemannian distance function on $SO(n)$. Well-posedness of the problem, existence and uniqueness of X_c are discussed in [2]. The definition (7.1) extends the usual notion of center of mass in the Euclidean space. To solve (7.1), we consider the Riemannian gradient descent [3, 54], i.e., for $k = 1, 2, \dots$ we compute

$$(7.2) \quad X_c^{k+1} = X_c^k \exp \left(-\frac{1}{N} \sum_{i=1}^N \log(X_i^\top X_c^k) \right).$$

To test the performance of Algorithm 4.1, we perform 100 iterations of Riemannian gradient descent using (7.2). We perform a fixed number of iterations instead of taking a convergence criterion to evaluate more clearly the impact of the RSD routine. The Riemannian gradient step is computed using the relation $\log(A) = Q \log(S) Q^\top$ (see (2.3)). The RSD is computed using either `schur` (LAPACK's `gees`) or `nrmschur` (Algorithm 4.1). We vary the size n of $SO(n)$, and the number N of samples in the collection. The results are displayed in Table 2. The larger n and N , the larger the improvement. Indeed, the factor of improvement in running time varies from 1.72 in the smallest configuration to 4.38 for $(N = 64, n = 200)$.

8. Conclusion. In this paper, we proposed a new method for computing the real Schur decomposition of a dense real normal matrix, or equivalently, its eigenvalue decomposition. We showed that the skew-symmetric part of the matrix is a valuable tool for obtaining this decomposition, as its two-dimensional invariant subspaces coincide with those of the original matrix.

Our approach was the following. First, we developed the matrix theory results that motivate the algorithm. This naturally led to an exact arithmetic version of the method, which we detail in Algorithm 3.1.

$N \backslash n$	25		50		100		200	
	schur	nrmschur	schur	nrmschur	schur	nrmschur	schur	nrmschur
16	0.307	0.178	1.31	0.663	10.1	2.56	42.9	11.6
32	0.570	0.327	2.90	1.35	21.0	4.97	77.4	17.8
64	1.02	0.527	5.87	2.23	43.2	9.56	185	42.2

TABLE 2

Time in seconds to perform 100 steps of Riemannian Gradient Descent (RGD) for the Riemannian barycenter problem. The sizes n of $\text{SO}(n)$ varies from 25 to 200 and the number of samples from 16 to 64. Orthogonal EVDs are respectively computed with *schur* and *nrmschur*.

Next, we examined the method's robustness under floating point arithmetic, resulting in the floating point variant [Algorithm 4.1](#).

By analyzing the sensitivity of the invariant subspaces of the skew-symmetric part, we theoretically quantified in [section 4](#) and numerically corroborated in [section 5](#) how an unfavorable distribution of eigenvalues can reduce the numerical accuracy of the real Schur decomposition if no additional computations are performed. While this loss of accuracy may often be negligible in applications such as Riemannian computations, we demonstrated how it can be mitigated by applying correction steps to selected inaccurate invariant subspaces, using the bounds provided in [Theorems 4.4, 4.6](#) and [4.7](#). The numerical accuracy of [Algorithm 4.1](#) was verified in [section 5](#).

We then analyzed the computational complexity of the method. In particular, we showed that in scenarios such as Haar-distributed orthogonal matrices, the cost of [Algorithm 4.1](#) is comparable to that of performing a Hessenberg reduction followed by assembling the Householder reflectors, i.e., $\sim \frac{14}{3}n^3$.

Finally, we highlighted the relevance of this method in the context of statistics on manifolds, and computations on the Stiefel manifold and the orthogonal group, where the Riemannian logarithm plays a central role and requires an efficient eigenvalue decomposition.

Appendix A. Special invariance groups. In this section, we provide important lemmas about invariance groups and we characterize two of them in [Lemmas A.2](#) and [A.3](#).

LEMMA A.1. *For every $A \in \mathbb{R}^{n \times n}$, all $Q \in \text{ig}(A)$ and all $V \in \text{O}(n)$, we have $VQV^\top \in \text{ig}(VAV^\top)$. In particular, V can be a permutation matrix.*

Proof. Since $QAQ^\top = A$ and $V^\top V = I_n$, we have $(VQV^\top)(VAV^\top)(VQV^\top) = (VAV^\top)$. $\square$

LEMMA A.2. *Let $\Sigma \in \mathbb{R}^{p \times p}$ be a diagonal matrix with distinct non-zero diagonal values. An orthogonal matrix $Q \in \text{O}(2p)$ is such that*

$$(A.1) \quad Q \begin{bmatrix} 0 & -\Sigma \\ \Sigma & 0 \end{bmatrix} Q^\top = \begin{bmatrix} 0 & -\Sigma \\ \Sigma & 0 \end{bmatrix},$$

if and only if $Q = \begin{bmatrix} \cos(\Phi) & -\sin(\Phi) \\ \sin(\Phi) & \cos(\Phi) \end{bmatrix}$ for some diagonal matrix $\Phi \in \mathbb{R}^{p \times p}$.

Proof. We proceed by induction. Assume first $p = 1$, i.e., $\Sigma = \sigma \in \mathbb{R}$. Then, for every $Q \in \text{O}(2)$, there is $\phi \in \mathbb{R}$ such that either $Q = \begin{bmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{bmatrix}$ (Q is a rotation) or $Q = \begin{bmatrix} \cos(\phi) & \sin(\phi) \\ \sin(\phi) & -\cos(\phi) \end{bmatrix}$ (Q is a reflection). It is easily verified that only rotations belong to $\text{ig}\left(\begin{bmatrix} 0 & -\sigma \\ \sigma & 0 \end{bmatrix}\right)$. Thus, [Lemma A.2](#) holds for $p = 1$.

Assume [Lemma A.2](#) holds for $p \geq 1$ and let $\Sigma \in \mathbb{R}^{p \times p}$, $S_1 = \begin{bmatrix} 0 & -\Sigma \\ \Sigma & 0 \end{bmatrix}$ and $S_2 = \begin{bmatrix} 0 & -\sigma \\ \sigma & 0 \end{bmatrix}$ with $\sigma \in \mathbb{R}$ be non-zero and distinct from all entries of Σ . Let $Q = \begin{bmatrix} Q_1 & Q_2 \\ Q_3 & Q_4 \end{bmatrix} \in O(2p+2)$ with $Q_1 \in \mathbb{R}^{2p \times 2p}$ and $Q_4 \in \mathbb{R}^{2 \times 2}$. Then, if one imposes

$$Q \begin{bmatrix} S_1 & 0 \\ 0 & S_2 \end{bmatrix} Q^\top = \begin{bmatrix} S_1 & 0 \\ 0 & S_2 \end{bmatrix} \text{ and } Q \in O(2p+2),$$

it must hold in particular that

$$\begin{cases} Q_3 S_1 Q_3^\top + Q_4 S_2 Q_4^\top = S_2, \\ Q_3 S_1 Q_1^\top + Q_4 S_2 Q_2^\top = 0, \end{cases} \quad \text{and} \quad \begin{cases} Q_1^\top Q_1 + Q_3^\top Q_3 = I_{2p}, \\ Q_1^\top Q_2 + Q_3^\top Q_4 = 0. \end{cases}$$

By multiplying the second equation by Q_1 on the right, we obtain $Q_3 S_1 Q_1^\top Q_1 + Q_4 S_2 Q_2^\top Q_1 = 0$. Leveraging both orthogonality constraints, we obtain $Q_3 S_1 (I_{2p} - Q_3^\top Q_3) - Q_4 S_2 Q_4^\top Q_3 = 0$. Introducing the first equation, we obtain $Q_3 S_1 = S_2 Q_3$. Let $Q_3 := \begin{bmatrix} Q_{31} & Q_{32} \end{bmatrix}$ such that

$$\begin{aligned} Q_3 S_1 = S_2 Q_3 &\iff Q_{32} \Sigma = S_2 Q_{31} \text{ and } -Q_{31} \Sigma = S_2 Q_{32}. \\ &\iff Q_{31} \Sigma^2 = \sigma^2 Q_{31} \text{ and } Q_{32} \Sigma^2 = \sigma^2 Q_{32}. \end{aligned}$$

Since σ is distinct from all entries of Σ , $Q_3 = 0$. Hence, $Q_4 \in O(2)$ and $Q_2 = 0$. Finally, $Q \in \text{ig} \left(\begin{bmatrix} S_1 & 0 \\ 0 & S_2 \end{bmatrix} \right)$ if and only if it is block-diagonal, $Q_1 \in \text{ig}(S_1)$ and $Q_4 \in \text{ig}(S_2) = \text{SO}(2)$.

To conclude, we can define $\tilde{\Sigma} := \text{diag}(\Sigma, \sigma)$ and take a permutation matrix P such that $P \begin{bmatrix} S_1 & 0 \\ 0 & S_2 \end{bmatrix} P^\top = \begin{bmatrix} 0 & -\tilde{\Sigma} \\ \tilde{\Sigma} & 0 \end{bmatrix}$. Applying [Lemma A.1](#) is enough to conclude that [Lemma A.2](#) holds for $p+1$. Therefore, it holds for all $p \geq 1$. $\square$

LEMMA A.3. An orthogonal matrix $Q = \begin{bmatrix} A & C \\ B & D \end{bmatrix} \in O(2m)$, $A \in \mathbb{R}^{m \times m}$, is such that

$$(A.2) \quad Q \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} Q^\top = \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix},$$

if and only if $A = D$ and $B = -C$. The matrix Q is called ortho-symplectic [\[32\]](#).

Proof. It is enough to notice that [\(A.2\)](#) is equivalent to

$$(A.3) \quad Q \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} = \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} Q \iff \begin{bmatrix} C & -A \\ D & -B \end{bmatrix} = \begin{bmatrix} -B & -D \\ A & C \end{bmatrix}.$$

Therefore $Q \in O(2m)$, $A = D$ and $B = -C$ are necessary and sufficient conditions. $\square$

Finally, we briefly recall the relation between the distance to the orthogonal group $d_{O(n)}(\hat{Q})$ from [\(4.5\)](#) and $\|Q^\top Q - I\|_F$. These measures are equivalent for $d_{O(n)}(\hat{Q})$ small.

LEMMA A.4. For all matrices $\hat{Q} \in \mathbb{R}^{n \times n}$, it holds that

$$d_{O(n)}(\hat{Q}) \leq \|\hat{Q}^\top \hat{Q} - I\|_F \leq 2d_{O(n)}(\hat{Q}) + d_{O(n)}(\hat{Q})^2.$$

Proof. Let Σ be the diagonal matrix of the singular values of $\hat{Q}$. Then, it is well-known that $d_{O(n)}(\hat{Q}) = \|\Sigma - I\|_F$. Moreover, $\|\hat{Q}^\top \hat{Q} - I\|_F = \|\Sigma^2 - I\|_F$. It follows that $\|\Sigma^2 - I\|_F = \|(\Sigma - I)(\Sigma + I)\|_F \geq \|\Sigma - I\|_F$.

Moreover,

$$\|\Sigma^2 - I\|_F = \|(\Sigma - I + I)^\top (\Sigma - I + I) - I\|_F \leq 2\|\Sigma - I\|_F + \|\Sigma - I\|_F^2.$$

This concludes the proof. $\square$

Appendix B. Symmetric skew-Hamiltonian matrices. In this section, we highlight two important properties of the symmetric skew-Hamiltonian matrix $\begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix}$ from [Theorem 3.4](#). This matrix appears when a normal matrix has complex eigenvalues with repeated imaginary parts.

LEMMA B.1. *A matrix $A \in \mathbb{R}^{2m \times 2m}$ is symmetric skew-Hamiltonian, i.e., $A = \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix}$ with $\tilde{H}$ symmetric and $\tilde{\Omega}$ skew-symmetric, if and only if there is an eigenvalue decomposition $A = Q \begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix} Q^\top$ where Q is ortho-symplectic and Λ is diagonal.*

Proof. ($\implies$) Assume $A = \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix}$. Since A is symmetric, all eigenvalues are real. Moreover, $\begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$ is a normed eigenvector of A if and only if $\begin{bmatrix} -v_2 \\ v_1 \end{bmatrix}$ is an eigenvector of A . Indeed,

$$(B.1) \quad \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \lambda \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \iff \begin{cases} (\tilde{H} - \lambda I)v_1 = \tilde{\Omega}v_2, \\ (\tilde{H} - \lambda I)v_2 = -\tilde{\Omega}v_1. \end{cases} \iff \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix} \begin{bmatrix} -v_2 \\ v_1 \end{bmatrix} = \lambda \begin{bmatrix} -v_2 \\ v_1 \end{bmatrix}.$$

Therefore, the multiplicity of every eigenvalue is even and greater or equal to two. Since (B.1) holds for every eigenvector, the eigenvector's matrix always admits a form $Q = \begin{bmatrix} E & -F \\ F & E \end{bmatrix}$ with $A = Q \begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix} Q^\top$. By [Lemma A.3](#), Q is ortho-symplectic.

($\impliedby$) Since Q is ortho-symplectic, $Q = \begin{bmatrix} E & -F \\ F & E \end{bmatrix}$ by [Lemma A.3](#). Multiplying $Q \begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix} Q^\top$ such as in (3.12) concludes the proof. $\square$

THEOREM B.2. *For every symmetric skew-Hamiltonian matrix $A \in \mathbb{R}^{2m \times 2m}$, i.e., $A = \begin{bmatrix} \tilde{H} & -\tilde{\Omega} \\ \tilde{\Omega} & \tilde{H} \end{bmatrix}$ with $\tilde{H}$ symmetric and $\tilde{\Omega}$ skew-symmetric, if $\text{rank}(A) = 2m$ and every eigenvalue of A has multiplicity 2, then the m first steps of Lanczos tridiagonalization are enough in exact arithmetic to build an ortho-symplectic matrix Y where $A = Y \begin{bmatrix} T_m & 0 \\ 0 & T_m \end{bmatrix} Y^\top$ and T_m is an $m \times m$ tridiagonal symmetric matrix.*

Proof. We denote $K = [b \ Ab \dots A^{2m-1}b]$ a Krylov basis of A generated by b . Choose any $v \in \mathbb{R}^n$ such that $b := Av \neq 0$ and $K_{1:m}$ has column rank m .⁴ By [Lemma B.1](#), we know that there is an ortho-symplectic matrix Q such that $A = Q \begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix} Q^\top$ and $\Lambda \in \mathbb{R}^{m \times m}$ is diagonal. Let k_j denote the j th column of K , it follows that

$$k_j = Q \begin{bmatrix} \Lambda^j & 0 \\ 0 & \Lambda^j \end{bmatrix} Q^\top v.$$

⁴Notice that $\text{rank}(K) \leq m$ since every eigenvalue has multiplicity 2 by [Lemma B.1](#).

Moreover, for any integer pair i, j , with $1 \leq i, j \leq n$, we have

$$\begin{aligned}
k_i^\top \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} k_j &= v^\top Q \begin{bmatrix} \Lambda^i & 0 \\ 0 & \Lambda^i \end{bmatrix} Q^\top \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} Q \begin{bmatrix} \Lambda^j & 0 \\ 0 & \Lambda^j \end{bmatrix} Q^\top v \\
&= v^\top Q \begin{bmatrix} \Lambda^i & 0 \\ 0 & \Lambda^i \end{bmatrix} \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix} \begin{bmatrix} \Lambda^j & 0 \\ 0 & \Lambda^j \end{bmatrix} Q^\top v \\
&= v^\top Q \begin{bmatrix} 0 & -\Lambda^{i+j} \\ \Lambda^{i+j} & 0 \end{bmatrix} Q^\top v \\
&= 0.
\end{aligned}
\tag{B.2}$$

The last equality stands because $v^\top \Omega v = 0$ for every skew-symmetric matrices Ω . Letting $J_{2m} := \begin{bmatrix} 0 & -I_m \\ I_m & 0 \end{bmatrix}$, it follows from (B.2) that $K_{1:m}^\top J_{2m} K_{1:m} = 0$. Therefore, the matrix $[K_m \ J_{2m} K_{1:m}]$ is full rank. Consequently, the Gram-Schmidt QR factorization of $K_{1:m}$, namely, $K_{1:m} = Y_{1:m} R$, generates $Y := [Y_{1:m} \ J_{2m} Y_{1:m}]$ where Y is ortho-symplectic. Moreover, we also have

$$k_i^\top A J_{2m} k_j = k_{i+1}^\top J_{2m} k_j = 0, \tag{B.3}$$

and thus $K_{1:m}^\top A J_{2m} K_{1:m} = 0$. Finally, it is the praised property of Lanczos algorithm that $Y_{1:m}^\top A Y_{1:m} = T_m$ where T_m is tridiagonal and symmetric [41, Thm. 9.1.1]. Combining (B.2) and (B.3), we obtain

$$\begin{aligned}
Y^\top A Y &= \begin{bmatrix} Y_{1:m}^\top \\ Y_{1:m}^\top J_{2m}^\top \end{bmatrix} A [Y_{1:m} \ J_{2m} Y_{1:m}] \\
&= \begin{bmatrix} Y_{1:m}^\top A Y_{1:m} & Y_{1:m}^\top A J_{2m} Y_{1:m} \\ Y_{1:m}^\top J_{2m}^\top A Y_{1:m} & Y_{1:m}^\top J_{2m}^\top A J_{2m} Y_{1:m} \end{bmatrix} \\
&= \begin{bmatrix} T_m & 0 \\ 0 & T_m \end{bmatrix}.
\end{aligned}$$

We used that $J_{2m}^\top A J_{2m} = A$. This concludes the proof. $\square$

Algorithm B.1 is a simple modification of the usual modified Gram-Schmidt to build a matrix that is not only orthogonal, but also symplectic. In floating point arithmetic, the matrix Y is approximately orthogonal but *exactly* symplectic by construction. Instead of orthogonalizing any matrix, **Algorithm B.1** orthogonalizes the Krylov basis generated by A symmetric skew-Hamiltonian, we directly obtain Lanczos's procedure with full re-orthogonalization.

In **Figure 7**, we evaluate how much **Theorem B.2** can improve the performance compared to LAPACK's symmetric eigenvalue routine for small EVP's. The method using only m steps of Lanczos tridiagonalization is called **wxeigen** and is available at **wxeigen.jl** in the repository. We observe that **wxeigen** with full re-orthogonalization achieves good running time performance and accuracy.

Appendix C. Proof of **Theorem 4.7**.

Proof. Let $A = Q S Q^\top \in \mathbb{R}^{n \times n}$ be a normal matrix with real eigenvalue $\check{\lambda}_k$ for $k \in \mathcal{K} := \{1, \dots, r\}$. Moreover, assume $\lambda_j e^{\pm i\theta_j}$ is an eigenvalue of A for $j \in \mathcal{I}$. Finally, assume we have computed $\hat{V} \in \text{St}(n, r)$ such that $\Omega \hat{V} = E$ with $\|E\|_F \leq \tau \varepsilon_m \|A\|_F$. We have

$$\Omega \hat{V} = E \iff \text{skew}(S) Q^\top \hat{V} = Q^\top E =: \tilde{E}. \tag{C.1}$$

Algorithm B.1 Re-orthogonalized Lanczos for symmetric skew-Hamiltonian matrix.

INPUT: $v \in \mathbb{R}^{2m}$, $A \in \mathbb{R}^{2m \times 2m}$.

$\beta_0 = 0$, $y_1 = Av/\|Av\|_2$ and $y_{m+1} = J_{2m}y_1$.

for $i = 2, \dots, m$ **do**

$y_i = Ay_{i-1}$ and $\alpha_{i-1} = y_{i-1}^\top y_i$.

$y_i = y_i - \alpha_{i-1}y_{i-1} - \beta_{i-2}y_{i-2}$.

Lanczos step.

for $j = 1, \dots, i-1$ **do**

$y_i = y_i - (y_i^\top y_j)y_j - (y_i^\top y_{m+j})y_{m+j}$.

Full re-orthogonalization.

end for

$\beta_i = \|y_i\|_2$.

$y_i = y_i/\beta_i$ and $y_{m+i} = J_{2m}y_i$.

Build symplectic basis.

end for

$\alpha_m = y_m A y_m$.

return $Y \in O(2m) \cap OSp(2m)$, $\{\alpha_i\}_{i=1}^m$ and $\{\beta_i\}_{i=1}^{m-1}$.




FIGURE 7. Numerical performance evaluation of `wxeigen`, without and with full re-orthogonalization of the Lanczos vectors, compared to `eigen` (LAPACK's `syevr`). The plot on the left shows the running time performance. The plot on the right shows the accuracy performance in spectral norm. Taking advantage of the symmetric skew-Hamiltonian structure allows to design faster methods than the generic symmetric eigensolver.

Defining $X := Q^\top \hat{V}$, there are two permutation $P_l \in O(n)$, $P_r \in O(r)$ such that $P_l^\top X P_r = \begin{bmatrix} X_{\mathcal{I}, \mathcal{K}} \\ X_{\mathcal{K}, \mathcal{K}} \end{bmatrix}$ where the block $X_{\mathcal{I}, \mathcal{K}}$ comprises all 2×1 blocks X_{ik} such that for $i \in \mathcal{I}$, $k \in \mathcal{K}$, by (C.1), we have

$$\begin{bmatrix} 0 & -\lambda_i s_i \\ \lambda_i s_i & 0 \end{bmatrix} X_{ik} = (P_l^\top \tilde{E} P_r)_{ik}.$$

Since $X \in \text{St}(n, r)$, we have

$$X_{\mathcal{I}, \mathcal{K}}^\top X_{\mathcal{I}, \mathcal{K}} + X_{\mathcal{K}, \mathcal{K}}^\top X_{\mathcal{K}, \mathcal{K}} = I_r \implies \|X_{\mathcal{K}, \mathcal{K}}^\top X_{\mathcal{K}, \mathcal{K}} - I_r\|_F \leq \|X_{\mathcal{I}, \mathcal{K}}\|_F^2$$

and, $\|X_{\mathcal{I}, \mathcal{K}}\|_F^2 = \sum_{i \in \mathcal{I}, k \in \mathcal{K}} \frac{1}{|\lambda_i s_i|^2} \|(P_l^\top \tilde{E} P_r)_{ik}\|_F^2 \leq \frac{\|E\|_F^2}{\min_{i \in \mathcal{I}} |\lambda_i s_i|^2},$

where [14, Thm. VII.2.8] is used. The deviation of $X_{\mathcal{K}, \mathcal{K}}$ from orthogonality is quantified by that of $X_{\mathcal{I}, \mathcal{K}}$ from zero. In addition, by taking the SVD $X_{\mathcal{K}, \mathcal{K}} = U_X \Sigma_X V_X^\top$, we can write $\|X_{\mathcal{K}, \mathcal{K}}^\top X_{\mathcal{K}, \mathcal{K}} - I_r\|_F = \|\Sigma_X^2 - I_r\|_F \geq$

$\|\Sigma_X - I_r\|_F$ by Lemma A.4. Define $\begin{bmatrix} \tilde{S}_{n-r} & 0 \\ 0 & \check{\Lambda} \end{bmatrix} := P_1^\top S P_1$. For all $\check{R} \in O(r)$, we have

$$\begin{aligned} \|A\hat{V}\check{R} - \hat{V}\check{R}\check{\Lambda}\|_F &\leq \|\text{sym}(S)XR - XR\check{\Lambda}\|_F + \|E\|_F \\ (C.2) \quad &\leq \|\text{sym}(\tilde{S}_{n-r})X_{\mathcal{I},\mathcal{K}}\check{R} - X_{\mathcal{I},\mathcal{K}}\check{R}\check{\Lambda}\|_F \\ &\quad + \|\check{\Lambda}X_{\mathcal{K},\mathcal{K}}\check{R} - X_{\mathcal{K},\mathcal{K}}\check{R}\check{\Lambda}\|_F + \|E\|_F. \end{aligned}$$

The first term of (4.12) can be bounded by

$$\begin{aligned} \|\text{sym}(\tilde{S}_{n-r})X_{\mathcal{I},\mathcal{K}}\check{R} - X_{\mathcal{I},\mathcal{K}}\check{R}\check{\Lambda}\|_F &\leq \max_{i \in \mathcal{I}, k \in \mathcal{K}} |\lambda_i c_i - \check{\lambda}_k| \|X_{\mathcal{I},\mathcal{K}}\check{R}\|_F \\ (C.3) \quad &\leq \frac{\max_{i \in \mathcal{I}, k \in \mathcal{K}} |\lambda_i c_i - \check{\lambda}_k|}{\min_{i \in \mathcal{I}} |\lambda_i s_i|} \|E\|_F. \end{aligned}$$

Moreover, the second term can also be bounded by

$$\begin{aligned} \|\check{\Lambda}X_{\mathcal{K},\mathcal{K}}\check{R} - X_{\mathcal{K},\mathcal{K}}\check{R}\check{\Lambda}\|_F &= \|\check{\Lambda}(X_{\mathcal{K},\mathcal{K}}\check{R} - I_r) - (X_{\mathcal{K},\mathcal{K}}\check{R} - I_r)\check{\Lambda}\|_F \\ &\leq \max_{k_1, k_2 \in \mathcal{K}} |\check{\lambda}_{k_1} - \check{\lambda}_{k_2}| \|X_{\mathcal{K},\mathcal{K}}\check{R} - I_r\|_F. \end{aligned}$$

In particular, for $\check{R} = V_X U_X^\top$, which minimizes $\|X_{\mathcal{K},\mathcal{K}}\check{R} - I_r\|_F$, we have

$$(C.4) \quad \|X_{\mathcal{K},\mathcal{K}}\check{R} - I_r\|_F = \|\Sigma_X - I_r\|_F \leq \frac{\|E\|_F^2}{\min_{i \in \mathcal{I}} |\lambda_i s_i|^2}.$$

Inserting (C.3) and (C.4) in (C.2) results in Theorem 4.7. $\square$

Acknowledgment. The authors would like to thank Paul Van Dooren and Pierre-Antoine Absil for their invaluable advice during the preparation of this manuscript.

REFERENCES

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. [arXiv:http://press.princeton.edu/titles/8586.html](http://press.princeton.edu/titles/8586.html).
- [2] Bijan Afsari. Riemannian Lp center of mass: Existence, uniqueness, and convexity. *Proceedings of the American Mathematical Society*, 139(2):655–673, 2011. [arXiv:http://www.jstor.org/stable/41059320](http://www.jstor.org/stable/41059320).
- [3] Bijan Afsari, Roberto Tron, and René Vidal. On the convergence of gradient descent for finding the Riemannian center of mass. *SIAM Journal on Control and Optimization*, 51(3):2230–2260, 2013. URL: <https://doi.org/10.1137/12086282X>.
- [4] G. S. Ammar, W. B. Gragg, and L. Reichel. On the eigenproblem for orthogonal matrices. In *1986 25th IEEE Conference on Decision and Control*, pages 1963–1966, 1986. URL: <https://doi.org/10.1109/CDC.1986.267356>.
- [5] G. S. Ammar, W. B. Gragg, and L. Reichel. Determination Of Pisarenko Frequency Estimates As Eigenvalues Of An Orthogonal Matrix. In Franklin T. Luk, editor, *Advanced Algorithms and Architectures for Signal Processing II*, volume 0826, pages 143 – 145. International Society for Optics and Photonics, SPIE, 1988. URL: <https://doi.org/10.1117/12.942026>.
- [6] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [7] T. W. Anderson, I. Olkin, and L. G. Underhill. Generation of random orthogonal matrices. *SIAM Journal on Scientific and Statistical Computing*, 8(4):625–629, 1987. URL: <https://doi.org/10.1137/0908055>.
- [8] Jared L. Aurentz, Thomas Mach, Raf Vandebril, and David S. Watkins. Fast and stable unitary QR algorithm. *Electron. Trans. Numer. Anal.*, 44:327–341, 2015. [arXiv:https://etna.math.kent.edu/vol.44.2015/pp327-341.dir/pp327-341.pdf](https://etna.math.kent.edu/vol.44.2015/pp327-341.dir/pp327-341.pdf).

- [9] Jared L. Aurentz, Thomas Mach, Raf Vandebril, and David S. Watkins. Computing the eigenvalues of symmetric tridiagonal matrices via a Cayley transformation. *Electron. Trans. Numer. Anal.*, 46:447–459, 2017. [arXiv:https://etna.math.kent.edu/vol.46.2017/pp447-459.dir/pp447-459.pdf](#).
- [10] Steve Batterson. Convergence of the Francis shifted QR algorithm on normal matrices. *Linear Algebra and its Applications*, 207:181–195, 1994. URL: [https://doi.org/10.1016/0024-3795\(94\)90010-8](https://doi.org/10.1016/0024-3795(94)90010-8).
- [11] Thomas Bendokat, Ralf Zimmermann, and P.-A. Absil. A Grassmann manifold handbook: basic geometry and computational aspects. *Advances in Computational Mathematics*, 50, 01 2024. URL: <https://doi.org/10.1007/s10444-023-10090-8>.
- [12] Peter Benner and Heike Fassbender. An implicitly restarted symplectic Lanczos method for the Hamiltonian eigenvalue problem. *Linear Algebra and its Applications*, 263:75–111, 1997. URL: [https://doi.org/10.1016/S0024-3795\(96\)00524-1](https://doi.org/10.1016/S0024-3795(96)00524-1).
- [13] Peter Benner, Daniel Kressner, and Volker Mehrmann. Skew-Hamiltonian and Hamiltonian eigenvalue problems: Theory, algorithms and applications. In Zlatko Drmač, Miljenko Marušić, and Zvonimir Tutek, editors, *Proceedings of the Conference on Applied Mathematics and Scientific Computing*, pages 3–39, Dordrecht, 2005. Springer Netherlands. URL: https://link.springer.com/chapter/10.1007/1-4020-3197-1_1.
- [14] Rajendra Bhatia. *Matrix Analysis*, volume 169. Springer New York, NY, 1997. URL: <https://doi.org/10.1007/978-1-4612-0653-8>.
- [15] Dario A. Bini and Bruno Iannazzo. Computing the Karcher mean of symmetric positive definite matrices. *Linear Algebra and its Applications*, 438(4):1700–1710, 2013. 16th ILAS Conference Proceedings, Pisa 2010. URL: <https://doi.org/10.1016/j.laa.2011.08.052>.
- [16] Dario A Bini, Bruno Iannazzo, Ben Jeuris, and Raf Vandebril. Geometric means of structured matrices. *Bit Numerical Mathematics*, 54(1):55–83, 2014-03-01. URL: <https://doi.org/10.1007/s10543-013-0450-4>.
- [17] Christian Bischof and Charles Van Loan. The WY representation for products of Householder matrices. *SIAM Journal on Scientific and Statistical Computing*, 8(1):s2–s13, 1987. URL: <https://doi.org/10.1137/0908009>.
- [18] L Susan Blackford, Antoine Petitet, Roldan Pozo, Karin Remington, R Clint Whaley, James Demmel, Jack Dongarra, Iain Duff, Sven Hammarling, Greg Henry, et al. An updated set of basic linear algebra subprograms (BLAS). *ACM Transactions on Mathematical Software*, 28(2):135–151, 2002.
- [19] Angelika Bunse-Gerstner, Ralph Byers, and Volker Mehrmann. Numerical methods for simultaneous diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 14, 07 1998. URL: <https://doi.org/10.1137/0614062>.
- [20] Angelika Bunse-Gerstner and Ludwig Elsner. Schur parameter pencils for the solution of the unitary eigenproblem. *Linear Algebra and its Applications*, 154-156:741–778, 1991. URL: [https://doi.org/10.1016/0024-3795\(91\)90402-I](https://doi.org/10.1016/0024-3795(91)90402-I).
- [21] D. Calvetti, L. Reichel, and H. Xu. A CS decomposition for orthogonal matrices with application to eigenvalue computation. *Linear Algebra and its Applications*, 476:197–232, 2015. URL: <https://doi.org/10.1016/j.laa.2015.03.007>.
- [22] Rudrasis Chakraborty and Baba C. Vemuri. Statistics on the Stiefel manifold: Theory and applications. *The Annals of Statistics*, 2019. [arXiv:https://api.semanticscholar.org/CorpusID:13671005](https://api.semanticscholar.org/CorpusID:13671005).
- [23] J.-P. Charlier and P. van Dooren. A Jacobi-like algorithm for computing the generalized Schur form of a regular pencil. In Henk A. van der Vorst and Paul van Dooren, editors, *Parallel Algorithms for Numerical Linear Algebra*, volume 1 of *Advances in Parallel Computing*, pages 17–36. North-Holland, 1990. URL: <https://doi.org/10.1016/B978-0-444-88621-7.50007-4>.
- [24] Jan Cuppen. A divide and conquer method for the symmetric eigenproblem. *Numerische Mathematik*, 36:177–195, 06 1980. URL: <https://doi.org/10.1007/BF01396757>.
- [25] G. Cybenko. Computing Pisarenko frequency estimates. In *Proceedings of the Princeton Conference on Information Science*, Princeton, 1985.
- [26] Percy Deift, James Demmel, Luen-Chau Li, and Carlos Tomei. The bidiagonal singular value decomposition and hamiltonian mechanics. *SIAM Journal on Numerical Analysis*, 28(5):1463–1516, 1991. URL: <https://doi.org/10.1137/0728076>.
- [27] J Demmel and W Kahan. Computing small singular values of bidiagonal matrices with guaranteed high relative accuracy: LAPACK working note number 3. 2 1988. [arXiv:https://www.osti.gov/biblio/5039344](https://www.osti.gov/biblio/5039344).
- [28] James Demmel and W. Kahan. Accurate singular values of bidiagonal matrices. *SIAM Journal on Scientific and Statistical Computing*, 11(5):873–912, 1990. URL: <https://doi.org/10.1137/0911052>.
- [29] James Demmel and Krešimir Veselić. Jacobi’s method is more accurate than QR. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1204–1245, 1992. URL: <https://doi.org/10.1137/0613074>.
- [30] Inderjit Singh Dhillon. *A new $O(N^2)$ algorithm for the symmetric tridiagonal eigenvalue/eigenvector problem*. PhD thesis, 1997.
- [31] J. J. Dongarra and D. C. Sorensen. A fully parallel algorithm for the symmetric eigenvalue problem. *SIAM Journal on Scientific and Statistical Computing*, 8(2):s139–s154, 1987. URL: <https://doi.org/10.1137/0908018>.
- [32] Froilán M. Dopico and Charles R. Johnson. Parametrization of the matrix symplectic group and applications. *SIAM*

- Journal on Matrix Analysis and Applications*, 31(2):650–673, 2009. URL: <https://doi.org/10.1137/060678221>.
- [33] Stanley C. Eisenstat and Ilse C. F. Ipsen. Three absolute perturbation bounds for matrix eigenvalues imply relative bounds. *SIAM Journal on Matrix Analysis and Applications*, 20(1):149–158, 1998. URL: <https://doi.org/10.1137/S0895479897323282>.
 - [34] Massimiliano Fasi and Leonardo Robol. Sampling the eigenvalues of random orthogonal and unitary matrices. *Linear Algebra and its Applications*, 620:297–321, 2021. URL: <https://doi.org/10.1016/j.laa.2021.02.031>.
 - [35] Heike Fassbender, D. Steven Mackey, Niloufer Mackey, and Hongguo Xu. Hamiltonian square roots of skew-Hamiltonian matrices. *Linear Algebra and its Applications*, 287(1):125–159, 1999. URL: [https://doi.org/10.1016/S0024-3795\(98\)10137-4](https://doi.org/10.1016/S0024-3795(98)10137-4).
 - [36] Kurukula Fernando and Beresford Parlett. Accurate singular values and differential qd algorithms. *Numerische Mathematik*, 67:191–229, 03 1994. URL: <https://doi.org/10.1007/s002110050024>.
 - [37] Micol Ferranti and Raf Vandebril. Computing eigenvalues of normal matrices via complex symmetric matrices. *Journal of Computational and Applied Mathematics*, 259(PART A), 2013-03-01. URL: <https://doi.org/10.1016/j.cam.2013.08.036>.
 - [38] J. G. F. Francis. The QR transformation A unitary analogue to the LR transformation - Part 1. *Comput. J.*, 4:265–271, 1961. URL: <https://doi.org/10.1093/comjnl/4.3.265>.
 - [39] J. G. F. Francis. The QR transformation—Part 2. *The Computer Journal*, 4(4):332–345, 01 1962. URL: <https://doi.org/10.1093/comjnl/4.4.332>.
 - [40] H. H. Goldstine and L. P. Horwitz. A procedure for the diagonalization of normal matrices. *J. ACM*, 6(2):176–195, apr 1959. URL: <https://doi.org/10.1145/320964.320975>.
 - [41] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore MD, third edition, 1996.
 - [42] Kazushige Goto and Robert Van De Geijn. High-performance implementation of the level-3 BLAS. 35(1), jul 2008. URL: <https://doi.org/10.1145/1377603.1377607>.
 - [43] William Gragg and L. Reichel. A divide and conquer method for unitary and orthogonal eigenproblems. *Numerische Mathematik*, 57:695–718, 12 1990. URL: <https://doi.org/10.1007/BF01386438>.
 - [44] William B. Gragg. The QR algorithm for unitary Hessenberg matrices. *Journal of Computational and Applied Mathematics*, 16(1):1–8, 1986. URL: [https://doi.org/10.1016/0377-0427\(86\)90169-X](https://doi.org/10.1016/0377-0427(86)90169-X).
 - [45] Benedikt Groß and Bruno Lang. An $O(n^2)$ algorithm for the bidiagonal SVD. *Linear Algebra and its Applications*, 358(1):45–70, 2003. URL: [https://doi.org/10.1016/S0024-3795\(01\)00398-6](https://doi.org/10.1016/S0024-3795(01)00398-6).
 - [46] Ming Gu, Robert Guzzo, Xue-bin Chi, and Xing-Qin Cao. A stable divide and conquer algorithm for the unitary eigenproblem. *SIAM Journal on Matrix Analysis and Applications*, 25(2):385–404, 2003. URL: <https://doi.org/10.1137/S0895479899359539>.
 - [47] Haoze He and Daniel Kressner. Randomized joint diagonalization of symmetric matrices. *SIAM Journal on Matrix Analysis and Applications*, 45(1):661–684, 2024. URL: <https://doi.org/10.1137/22M1541265>.
 - [48] Haoze He and Daniel Kressner. A simple, randomized algorithm for diagonalizing normal matrices. *Calcolo*, 62(3):30, 2025. doi:10.1007/s10092-025-00654-z.
 - [49] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, second edition, 2002. URL: <https://doi.org/10.1137/1.9780898718027>.
 - [50] A. J. Hoffman and H. W. Wielandt. The variation of the spectrum of a normal matrix. *Duke Mathematical Journal*, 20(1):37 – 39, 1953. URL: <https://doi.org/10.1215/S0012-7094-53-02004-3>.
 - [51] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge; New York, 2nd edition, 2013.
 - [52] Jinzhi Huang and Zhongxiao Jia. A skew-symmetric Lanczos bidiagonalization method for computing several extremal eigenpairs of a large skew-symmetric matrix. *SIAM Journal on Matrix Analysis and Applications*, 45(2):1114–1147, 2024. URL: <https://doi.org/10.1137/23M1553029>.
 - [53] A. T. James. Normal Multivariate Analysis and the Orthogonal Group. *The Annals of Mathematical Statistics*, 25(1):40–75, 1954. URL: <https://doi.org/10.1214/aoms/1177728846>.
 - [54] Krzysztof Krakowski, Knut Hüper, and Jonathan Manton. On the computation of the Karcher mean on spheres and special orthogonal groups. 09 2007. arXiv:https://www.researchgate.net/publication/228753783_On_the_computation_of_the_Karcher_mean_on_spheres_and_special_orthogonal_groups.
 - [55] Yongdo Lim and Miklós Pálfi. Matrix power means and the Karcher mean. *Journal of Functional Analysis*, 262(4):1498–1514, 2012. URL: <https://doi.org/10.1016/j.jfa.2011.11.012>.
 - [56] Sy-Shin Lo, Bernard Philippe, and Ahmed Sameh. A multiprocessor algorithm for the symmetric tridiagonal eigenvalue problem. *SIAM Journal on Scientific and Statistical Computing*, 8(2):s155–s165, 1987. URL: <https://doi.org/10.1137/0908019>.
 - [57] Simon Mataigne, Ralf Zimmermann, and Nina Miolane. An Efficient Algorithm for the Riemannian Logarithm on the

- Stiefel Manifold for a Family of Riemannian Metrics. *SIAM Journal on Matrix Analysis and Applications*, 46(2):879–905, 2025. URL: <https://doi.org/10.1137/24M1647801>.
- [58] C. C. Paige. Computational Variants of the Lanczos Method for the Eigenproblem. *IMA Journal of Applied Mathematics*, 10(3):373–381, 12 1972. URL: <https://doi.org/10.1093/imamat/10.3.373>.
- [59] B. N. Parlett and D. S. Scott. The Lanczos algorithm with selective orthogonalization. *Mathematics of Computation*, 33(145):217–238, 1979. [arXiv:http://www.jstor.org/stable/2006037](http://www.jstor.org/stable/2006037).
- [60] Carolin Penke, Andreas Marek, Christian Vorwerk, Claudia Draxl, and Peter Benner. High performance solution of skew-symmetric eigenvalue problems with applications in solving the Bethe-Salpeter eigenvalue problem. *Parallel Computing*, 96:102639, 2020. URL: <https://doi.org/10.1016/j.parco.2020.102639>.
- [61] Quentin Rentmeesters. *Algorithms for data fitting on some common homogeneous spaces*. PhD thesis, Catholic University of Louvain (UCLouvain), 2013. [arXiv:https://dial.uclouvain.be/pr/boreal/fr/object/boreal/132587](https://dial.uclouvain.be/pr/boreal/fr/object/boreal/132587).
- [62] Robert Schreiber and Charles Van Loan. A storage-efficient WY representation for products of Householder transformations. *SIAM Journal on Scientific and Statistical Computing*, 10(1):53–57, 1989. URL: <https://doi.org/10.1137/0910005>.
- [63] Horst D. Simon. Analysis of the symmetric Lanczos algorithm with reorthogonalization methods. *Linear Algebra and its Applications*, 61:101–131, 1984. URL: [https://doi.org/10.1016/0024-3795\(84\)90025-9](https://doi.org/10.1016/0024-3795(84)90025-9).
- [64] G. W. Stewart. The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM Journal on Numerical Analysis*, 17(3):403–409, 1980. URL: <https://doi.org/10.1137/0717034>.
- [65] Michael Stewart. An error analysis of a unitary Hessenberg QR algorithm. *SIAM Journal on Matrix Analysis and Applications*, 28(1):40–67, 2006. URL: <https://doi.org/10.1137/04061948X>.
- [66] Brian D Sutton. Simultaneous diagonalization of nearly commuting Hermitian matrices: do-one-then-do-the-other. *IMA Journal of Numerical Analysis*, 44(2):1061–1089, 06 2023. URL: <https://doi.org/10.1093/imanum/drad033>.
- [67] Tom Szwagier and Xavier Pennec. Rethinking the Riemannian logarithm on flag manifolds as an orthogonal alignment problem. In Frank Nielsen and Frédéric Barbaresco, editors, *Geometric Science of Information*, pages 375–383, Cham, 2023. Springer Nature Switzerland.
- [68] Shunjiang Tao and Yunlin Xu. Parallel implementation of red-black ordering preconditioners for accelerating gmres solutions in neutron transport code pandas-moc. *Annals of Nuclear Energy*, 193:110008, 2023. URL: <https://doi.org/10.1016/j.anucene.2023.110008>.
- [69] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.
- [70] C. Van Loan. A symplectic method for approximating all the eigenvalues of a Hamiltonian matrix. *Linear Algebra and its Applications*, 61:233–251, 1984. URL: [https://doi.org/10.1016/0024-3795\(84\)90034-X](https://doi.org/10.1016/0024-3795(84)90034-X).
- [71] Raf Vandebril. A unitary similarity transform of a normal matrix to complex symmetric form. *Applied Mathematics Letters*, 24(2), 2011-02-01. URL: <https://doi.org/10.1016/j.aml.2010.08.037>.
- [72] Krešimir Veselić and Ivan Slapničar. Floating-point perturbations of Hermitian matrices. *Linear Algebra and its Applications*, 195:81–116, 1993. URL: [https://doi.org/10.1016/0024-3795\(93\)90258-P](https://doi.org/10.1016/0024-3795(93)90258-P).
- [73] Tai-Lin Wang and William B. Gragg. Convergence of the shifted QR algorithm for unitary Hessenberg matrices. *Mathematics of Computation*, 71(240):1473–1496, 2002. [arXiv:http://www.jstor.org/stable/4099960](http://www.jstor.org/stable/4099960).
- [74] Tai-Lin Wang and William B. Gragg. Convergence of the unitary QR algorithm with a unimodular Wilkinson shift. *Mathematics of Computation*, 72(241):375–385, 2003. [arXiv:http://www.jstor.org/stable/4099996](http://www.jstor.org/stable/4099996).
- [75] R. C. Ward and L. J. Gray. Eigensystem computation for skew-symmetric and a class of symmetric matrices. *ACM Trans. Math. Softw.*, 4(3):278–285, 1978. URL: <https://doi.org/10.1145/355791.355798>.
- [76] David S. Watkins. On Hamiltonian and symplectic Lanczos processes. *Linear Algebra and its Applications*, 385:23–45, 2004. Special Issue in honor of Peter Lancaster. URL: <https://doi.org/10.1016/j.laa.2002.11.001>.
- [77] J.F. Watters. Simultaneous quasi-diagonalization of normal matrices. *Linear Algebra and its Applications*, 9:103–117, 1974. URL: [https://doi.org/10.1016/0024-3795\(74\)90030-5](https://doi.org/10.1016/0024-3795(74)90030-5).
- [78] J. H. Wilkinson. Householder’s method for symmetric matrices. *Numer. Math.*, 4(1):354–361, dec 1962. URL: <https://doi.org/10.1007/BF01386332>.
- [79] J. H. Wilkinson. Convergence of the LR, QR, and Related Algorithms. *The Computer Journal*, 8(1):77–84, 04 1965. URL: <https://doi.org/10.1093/comjnl/8.1.77>.
- [80] J. H. Wilkinson. The QR Algorithm for Real Symmetric Matrices with Multiple Eigenvalues. *The Computer Journal*, 8(1):85–87, 04 1965. URL: <https://doi.org/10.1093/comjnl/8.1.85>.
- [81] J.H. Wilkinson. Global convergence of tridiagonal QR algorithm with origin shifts. *Linear Algebra and its Applications*, 1(3):409–420, 1968. URL: [https://doi.org/10.1016/0024-3795\(68\)90017-7](https://doi.org/10.1016/0024-3795(68)90017-7).
- [82] Shihui Ying, Han Qin, Yaxin Peng, and Zhijie Wen. Compute Karcher means on $SO(n)$ by the geometric conjugate gradient method. *Neurocomputing*, 215:169–174, 2016. SI: Stereo Data. URL: <https://doi.org/10.1016/j.neucom.2015.05.143>.
- [83] B.B Zhou and R.P Brent. An efficient method for computing eigenvalues of a real normal matrix. *Journal of Parallel and Distributed Computing*, 63(6):638–648, 2003. URL: [https://doi.org/10.1016/S0743-7315\(03\)00007-8](https://doi.org/10.1016/S0743-7315(03)00007-8).

- [84] Ralf Zimmermann. A matrix-algebraic algorithm for the Riemannian logarithm on the Stiefel manifold under the canonical metric. *SIAM Journal on Matrix Analysis and Applications*, 38(2):322–342, 2017. URL: <https://doi.org/10.1137/16M1074485>.
- [85] Ralf Zimmermann and Knut Hüper. Computing the Riemannian logarithm on the Stiefel manifold: Metrics, methods, and performance. *SIAM Journal on Matrix Analysis and Applications*, 43(2):953–980, 2022. URL: <https://doi.org/10.1137/21M1425426>.