A consistently oriented basis for eigenanalysis

Jay Damask¹

Received: 2 December 2019 / Accepted: 31 May 2020
© The Author(s) 2020, corrected publication 2021

Abstract

Repeated application of eigen-centric methods to an evolving dataset reveals that eigenvectors calculated by well-established computer implementations are not stable along an evolving sequence. This is because the sign of any one eigenvector may point along either the positive or negative direction of its associated eigenaxis, and for any one eigenfunction call, the sign does not matter when calculating a solution. This work reports a model-free algorithm that creates a consistently oriented basis of eigenvectors. The algorithm postprocesses any well-established eigen call and is therefore agnostic to the particular implementation of the latter. Once consistently oriented, directional statistics can be applied to the eigenvectors in order to track their motion and summarize their dispersion. When a consistently oriented eigensystem is applied to methods of machine learning, the time series of training weights becomes interpretable in the context of the machine-learning model. Ordinary linear regression is used to demonstrate such interpretability. A reference implementation of the algorithm reported herein has been written in Python and is freely available, both as source code and through the `thucyd` Python package.

Keywords Eigenvectors · Eigenvalues · Singular-value decomposition · Directional statistics

Mathematics Subject Classification 15A18 · 65F15 · 62H11

1 Overview

Eigenanalysis plays a central role in many aspects of data science because the natural coordinates and modal strengths of a multifactor dataset are revealed [1,6,10,20]. Singular-value decomposition and eigendecomposition are two of the most common instances of eigenanalysis [5,24], yet there are a wide range of adaptations, such as eigenface [14]. The eigenanalysis ecosystem is so important that highly optimized implementations are available in core numerical libraries such as LAPACK [2]. This article does not seek to improve the existing theoretical or numerical study of eigenanalysis as it applies to a stand-alone problem. Instead, the focus of this article is on the improvement to eigenanalysis when applied sequentially to an *evolving* dataset.

In the domain of neural networks, the study of data streams is well known. For instance, the *evolving granular neural network* was introduced by Leite et al in 2009 (see [12] and references therein) to “tackle classification problems in continuously changing environments.” Hou et al subsequently

proposed *feature-evolvable streaming learning* to manage real-world situations in which changes in the features themselves, or the way in which they are measured, are handled in a consistent manner [7]. Although developed independently, the work in this paper applies the spirit of handling data streams to the problem of eigenanalysis.

To perform eigenanalysis on a stream of data, the eigenvectors and -values are calculated upon each update. The researcher may then expect to track the pointing direction of the vectors. He or she will be disappointed to find, however, that the *orientation* of the eigenvector basis produced by current eigenanalysis algorithms is not consistent along the stream. Rather, the eigenvectors flip between pointing along the positive or negative direction of the associated eigenaxis, thereby leading to an unconnected evolution.

This paper demonstrates a method for the *consistent orientation* of an eigensystem, in which the orientation is made consistent after an initial call to a canonical eigenanalysis routine. As a consequence of the method detailed below, the evolution of the eigenvectors themselves can be tracked. Eigenvectors that span $\mathbb{R}^n$ lie on a hypersphere $\mathbb{S}^{n-1} \subseteq \mathbb{R}^n$, and the mean pointing direction on the sphere and dispersion about the mean are statistics of interest, as is any statistically

✉ Jay Damask
jaydamask@buell-lane-press.co

¹ Buell Lane Press, New York, NY, USA

Fig. 1 An ellipsoid-shaped point cloud in $\mathbb{R}^3$ typical of centered, variance-bound data having a stationary or slowly changing underlying random process. The point cloud is quoted in a constituent basis (π_1, π_2, π_3) , while the eigenvectors lie along the natural axes of the data. Here, the ambiguity of the direction of v_2 is unresolved. *Left*, a left-handed basis of eigenvectors. *Right*, a right-handed basis of eigenvectors

significant drift. Basic results from directional statistics (see [9,13,15]) are used to study these directional features.

2 Inconsistent orientation

Given a square, symmetric matrix, $M^T = M$, with real-valued entries, $M \in \mathbb{R}^{n \times n}$, the eigendecomposition is

$$M = V \Lambda V^T, \quad (1)$$

where the columns of $V \in \mathbb{R}^{n \times n}$ are the eigenvectors of M , the diagonal entries along $\Lambda \in \mathbb{R}^{n \times n}$ are the eigenvalues, and entries in both are purely real. Moreover, when properly normalized, $V^T V = I$, indicating that the eigenvectors form an orthonormal basis.

To inspect how the eigenvectors might be oriented, write out V as $(v_1, v_2, \dots, v_n)$, where v_k is a column vector, and take the inner product two ways:

$$(v_1, v_2, \dots, v_n)^T (v_1, v_2, \dots, v_n) = I \quad \text{and} \\ (v_1, -v_2, \dots, v_n)^T (v_1, -v_2, \dots, v_n) = I.$$

A choice of $+v_2$ is indistinguishable from a choice of $-v_2$ because in either case the equality in (1) holds. In practice, the eigenvectors take an *inconsistent orientation*, even though they will have a particular orientation once returned from a call to an eigenanalysis routine.

A similar situation holds for singular-value decomposition (SVD). For a centered panel of data $P \in \mathbb{R}^{m \times n}$ where $m > n$, SVD gives

$$P = U \Lambda^{1/2} V^T, \quad (2)$$

where the columns in panel $U \in \mathbb{R}^{m \times n}$ are the projections of the columns of P onto the eigenvectors V , scaled to have unit variance. The eigenvectors have captured the natural

orientation of the data. Unlike the example above, the eigensystem here is positive semidefinite, so all eigenvalues in Λ are nonnegative, and in turn a scatter plot of data in P typically appears as an ellipsoid in $\mathbb{R}^n$ (see Fig. 1). Following the thought experiment above, a sign flip of an eigenvector in V flips the sign of the corresponding column in U , thereby preserving the equality in (2). Consequently, the eigenvector orientations produced by SVD also have inconsistent orientation, since either sign satisfies the equations.

In isolation, such inconsistent eigenvector orientation does not matter. The decomposition is correct and is generally produced efficiently. However, in subsequent analysis of the eigensystem—principal component analysis (PCA) being a well-known example [6]—and for an evolving environment, the inconsistency leads to an inability to *interpret* subsequent results and the inability to track the eigenvector evolution.

For instance, a simple linear regression performed on the eigensystem (whether with reduced dimension or not) with dependent variable $y \in \mathbb{R}^m$ reads

$$y = U \Lambda^{1/2} \beta + \epsilon, \quad (3)$$

where the weight vector is $\beta \in \mathbb{R}^n$, $\beta = (\beta_1, \dots, \beta_n)^T$. Each entry in β corresponds to a column in U , which in turn corresponds to an eigenvector in V . If the sign of one of the eigenvectors is flipped (compare *left* and *right* in Fig. 1), the corresponding column in U has its sign flipped, *which flips the sign of the corresponding β entry*. In an evolving environment, a particular regression might be denoted by index k , and the k th regression reads

$$y_k = U_k \Lambda_k^{1/2} \beta_k + \epsilon_k. \quad (4)$$

When tracking any particular entry in the β vector, say β_i , it will generally be found that $\beta_{i,k}$ flips sign throughout the evolution of system (see Fig. 2). There are two consequences: First, for a single regression, the *interpretation* of how a change of the i th factor corresponds to a change in y

Fig. 2 Left, in $\mathbb{R}^2$ a regression along the eigenaxes yields the tuple (β_1, β_2) . Right, top, because the eigenvector direction can flip from one eigenanalysis to the next, the signs of the β -tuple also flip, defeating any attempt at interpretability. Right, bottom, ideally there are no sign flips of the eigenvectors and therefore no flips of the β -tuple along a sequence

is unknown because the i th eigenvector can take either sign. Second, as the data evolves, the sign of $\beta_{i,k}$ will generally change, even if the unsigned β_i value remains fixed, so that the time series of β cannot be meaningfully analyzed. If these sources of sign flips can be eliminated, then a time series of weights β can be interpreted as we might have expected (see Fig. 2 lower right).

3 Toward a consistent orientation

To construct a consistent eigenvector basis, we must begin with a reference, and SVD offers a suitable start. The data in matrix P of (2) are presented with n columns, where, typically, each column refers to a distinct observable or feature. Provided that the data are well formed, such as having zero mean, bound variance, and negligible autocorrelation along each column, the row space of P forms the *constituent basis* $\Pi \in \mathbb{R}^{n \times n}$ with basis vectors $\pi \in \mathbb{R}^n$ such that $\Pi = (\pi_1, \pi_2, \dots, \pi_n)$ (see Fig. 1). If we label the columns in P by $(p_1, \dots, p_n)$, then the basis Π when materialized onto the constituent axes is simply

$$\Pi = \begin{pmatrix} p_1 & p_2 & \cdots & p_n \\ 1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix} \begin{pmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \end{pmatrix}$$

or, $\Pi = I$.

Now, when plotted as a scatter plot in $\mathbb{R}^n$, the data in the rows of P typically form an ellipsoid in the space. When the columns of P have zero linear correlation, the axes of the

ellipsoid align to the constituent axes. As correlation is introduced across the columns of P , the ellipsoid rotates away. Given such a case, we might seek to rotate the ellipsoid back onto the constituent axes. The constituent axes can therefore act as our reference.

This idea has good flavor, but is incomplete. An axis is different from a vector, for a vector may point along either the positive or negative direction of a particular axis. Therefore, ambiguity remains. To capture this mathematically, let us construct a rotation matrix $R \in \mathbb{R}^{n \times n}$ such that R^T rotates the data ellipse onto the constituent axes. A modified SVD will read

$$PR = U \Lambda^{1/2} V^T R.$$

If it were the case that $R^T V = I$, then the goal would be achieved. However, $R^T V = I$ is not true. Instead,

$$R^T V \sim \text{diag}(\pm 1, \pm 1, \dots, \pm 1), \quad (5)$$

is the general case, where the sign on the k th entry depends on the corresponding eigenvector orientation in V . While the following is an oversimplification of the broader concept of orientation, it is instructive to consider that a basis in $\mathbb{R}^3$ may be oriented in a right- or left-handed way. It is not possible to rotate one handedness onto the other because of the embedded reflection.

It is worth noting that an objection might be raised at this point because the eigenvector matrix V is itself a unitary matrix with $\det(V) = +1$ and therefore serves as the sought-after rotation. Indeed, $V^T V = I$. The problem is that $V^T V$ is quadratic, so any embedded reflections cancel. In order to identify the vector orientation within V , we must move away from the quadratic form.

4 A representation for consistent orientation

The representation detailed herein uses the constituent basis $\Pi = I$ as the reference basis and seeks to make (5) an equality for a particular V . To do so, first define the diagonal sign matrix $S_k \in \mathbb{R}^{n \times n}$ as

$$S_k = \text{diag}(1, 1, \dots, s_k, \dots, 1), \quad \text{where } s_k = \pm 1. \quad (6)$$

The product $S = \prod S_k$ is $S = \text{diag}(s_1, s_2, \dots, s_n)$, and clearly $S^2 = I$. Given an entry $s_k = -1$, the VS_k product reflects the k th eigenvector such that $v_k \rightarrow -v_k$ while leaving the other eigenvectors unaltered.

An equality can now be written with the representation

$$R^T VS = I, \quad (7)$$

given suitable rotation and reflection matrices R and S .

This representation says that a rotation matrix R can be found to rotate VS onto the constituent basis Π , provided that a diagonal reflection matrix S is found to selectively flip the sign of eigenvectors in V to ensure complete alignment in the end. Once (7) is satisfied, the oriented eigenvector basis $\mathcal{V}$ can be calculated in two ways:

$$\mathcal{V} = VS = R. \quad (8)$$

To achieve this goal, the first step is to order the eigenvalues in descending order and reorder the eigenvectors in V accordingly. (Since the representation (7) holds for real symmetric matrices $M = M^T$ and not solely for positive semidefinite systems, the absolute value of the eigenvalues should be ordered in descending order.) Indexing into V below expects that this ordering has occurred.

Representation (7) is then expanded so that each eigenvector is inspected for a reflection, and a rotation matrix is constructed for its alignment to the constituent. The expansion reads

$$R_n^T \dots R_2^T R_1^T VS_1 S_2 \dots S_n = I. \quad (9)$$

There is a good separation of concerns here: Reflections are imparted by the S_k operators to the right of V , while rotations are imparted by the R_k^T operators to the left. Reflections *do not* preserve orientation, whereas rotations *do* preserve orientation.

The algorithm created to attain a consistent orientation follows the pattern

$$V \rightarrow VS_1 \rightarrow R_1^T VS_1 \rightarrow R_1^T VS_1 S_2 \rightarrow \dots \quad (10)$$

For each eigenvector, a possible reflection is computed first, followed by a rotation. The result of the rotation is to align

the target eigenvector with its associated constituent basis vector. The first rotation is performed on the entire space $\mathbb{R}^n$, the second is performed on the subspace $\mathbb{R}^{n-1}$ so that the alignment $v_1^T \pi_1 = 1$ is preserved, the third rotation operates on $\mathbb{R}^{n-2}$ to preserve the preceding alignments, and so on. The solution to (7) will include n reflection elements s_k and $n(n-1)/2$ rotation angles embedded in the n rotation matrices R_k .

Before working through the solution, it will be helpful to walk through an example of the attainment algorithm in $\mathbb{R}^3$.

5 Algorithm Example in $\mathbb{R}^3$

This example illustrates a sequence of reflections and rotations in $\mathbb{R}^3$ to attain alignment between V and Π , and along the way shows cases in which a reflection or rotation is simply an identity operator.

An example eigenvector V is created with a left-hand basis, such that an alignment to the constituent basis Π would yield $R^T V = \text{diag}(1, -1, 1)$. The precondition of ordered eigenvalues $\lambda_1 > \lambda_2 > \lambda_3$ has already been met. Figure 3a shows the relative orientation of V to Π .

The first step is to consider whether or not to reflect v_1 , the purpose being to bring v_1 into the nonnegative side of the half-space created by the hyperplane perpendicular to π_1 . To do so, the relative orientation of (v_1, π_1) is measured with an inner product. The reflection entry s_1 is then determined from

$$s_1 = \text{sign}(v_1^T \pi_1)$$

(with $\text{sign}(0) = 1$). From this, VS_1 is formed. A consequence of the resultant orientation is that the angle subtended between v_1 and π_1 is $\leq |\pi/2|$.

The next step is to align v_1 with π_1 through a rotation R_1 , and the result is shown in Fig. 3b. The principal result of $R_1^T VS_1$ is that $v_1^T \pi_1 = 1$: The basis vectors are aligned. Significantly, the remaining vectors (v_2, v_3) are perpendicular to π_1 , and therefore the subsequent rotation needs only to act in the $(\pi_2, \pi_3) \in \mathbb{R}^2$ subspace.

With this sequence of operations in the full space $\mathbb{R}^3$ now complete, the subspace $\mathbb{R}^2$ is treated next. As before, the first step is to inspect the orientation of v_2 with π_2 . However, care must be taken because v_2 is no longer in its original orientation; instead, it was rotated with the whole basis during $R_1^T V$, as shown in Fig. 3b. It is the *new* orientation of v_2 that needs inspection and that orientation is $R_1^T v_2$. As seen in Fig. (b), $R_1^T v_2$ and π_2 point in opposite directions, so eigenvector v_2 needs to be reflected. We have

$$s_2 = \text{sign}((R_1^T v_2)^T \pi_2),$$

Fig. 3 Example of the orientation of a left-handed eigenvector basis V to the constituent basis Π . **a** Orthonormal eigenvectors (v_1, v_2, v_3) as they are oriented with respect to the constituent basis vectors (π_1, π_2, π_3) . **b** Rotation of V so that $v_1^T \pi_1 = 1$. **c** Reflection of v_2 so that $(R_1^T v_2)^T \pi_2 \geq 0$. **d** Rotation of $R_1^T V S_1$ by R_2 to align V with Π

which in this example is $s_2 = -1$. The updated representation of the eigenbasis is $R_1^T V S_1 S_2$, as seen in Fig. 3c.

To complete work in $\mathbb{R}^2$, rotation R_2 is found to bring $R_1^T v_2 s_2$ into alignment with π_2 . The sequence $R_2^T R_1^T V S_1 S_2$ is illustrated in Fig. 3d.

The solution needs to work through the last subspace $\mathbb{R}$ to align v_3 . The reflection entry s_3 is evaluated from

$$s_3 = \text{sign} \left((R_2^T R_1^T v_3)^T \pi_3 \right),$$

where $s_3 = +1$ in this example. From this, we arrive at $R_2^T R_1^T V S_1 S_2 S_3$. Canonically, a final rotation R_3 will carry $R_2^T R_1^T v_3 s_3$ into π_3 , but since we are in the last subspace of $\mathbb{R}$, the operation is a tautology. Still, R_3 is included in representation (9) for symmetry.

Lastly, the final, oriented eigenvector basis can be calculated two ways,

$$\mathcal{V} = V S_1 S_2 S_3 = R_1 R_2 R_3; \quad (11)$$

compare with (8).

In summary, the algorithm walks down from the full space $\mathbb{R}^3$ into two subspaces in order to perform rotations that do not upset previous alignments. As this example shows, $S_1 = S_3 = I$, so there are in fact two reflection operators that are simply identities. Moreover, $R_3 = I$, since the prior rotations automatically align that last basis vector to within a reflection.

6 Algorithm details

The algorithm that implements (9) is now generalized to $\mathbb{R}^n$. There are three building blocks to the algorithm: sorting of the eigenvalues in descending order, calculation of the reflection entries, and construction of the rotation matrices. Of the three, construction of the rotation matrices requires the most work and has not been previously addressed in detail.

6.1 Sorting of eigenvalues

The purpose of sorting the eigenvalues in descending order and rearranging the corresponding eigenvectors accordingly is to ensure that the first rotation operates on the axis of maximum variance, the second rotation on the axis of next-largest variance, and so on, the reason being that the numerical certainty of the principal ellipsoid axis is higher than that of any other axis. Since the rotations are concatenated in this method, an outside error introduced by an early rotation persists throughout the calculation. By sorting the eigenvalues and vectors first, any accumulated errors will be minimized.

Since the algorithm works just as well for real-valued symmetric matrices, it is reasonable to presort the eigenvalues in descending order of their absolute value. From a data-science perspective, we generally expect positive semidefinite systems; therefore, a system with negative eigenvalues is unlikely, and if one is encountered, there may be larger upstream issues to address.

6.2 Reflection calculation

Reflection entries are calculated by inspecting the inner product between a rotated eigenvector and its associated constituent basis (several examples are given in the previous section). The formal reflection-value expression for the k th basis vector is

$$s_k = \text{sign} \left(\left(\mathcal{R}_k^T v_k \right)^T \pi_k \right) \quad (12)$$

where

$$\mathcal{R}_k^T \equiv R_{k-1}^T \cdots R_2^T R_1^T. \quad (13)$$

Despite the apparent complexity of these expressions, the practical implementation is simply to inspect the sign of the (k, k) entry in the $R_{k-1}^T \cdots R_1^T V$ matrix.

6.3 Rotation matrix construction

One way to state the purpose of a particular rotation is that the rotation reduces the dimension of the subspace by one. For an eigenvector matrix V in $\mathbb{R}^n$, denoted by $V_{(n)}$, the rotation acts such that

$$R_1^T (V_{(n)} S_1) = \begin{pmatrix} 1 & \\ & V_{(n-1)} \end{pmatrix}. \quad (14)$$

Here, S_1 ensures that the $(1, 1)$ entry on the right is $+1$. Next, the sign of the $(2, 2)$ matrix entry is inspected and s_2 set accordingly. The next rotation decrements the dimension again such that

$$R_2^T (R_1^T V_{(n)} S_1 S_2) = \begin{pmatrix} 1 & & \\ & 1 & \\ & & V_{(n-2)} \end{pmatrix}. \quad (15)$$

As before, S_2 ensures that the $(2, 2)$ entry on the right is $+1$. The 1 entries along the diagonal reflect the fact that a eigenvector has been aligned with its corresponding constituent basis vector. Subsequent rotations preserve the alignment by operating in an orthogonal subspace.

6.3.1 Alignment to one constituent basis vector

To simplify the notation for the analysis below, let us collapse V and all operations prior to the k th rotation into a single working matrix, denoted by W_k . The left-hand side of the two equations above then deals with $R_1^T W_1$ and $R_2^T W_2$. The ℓ th column of the k th working matrix is $w_{k,\ell}$, and the m th column entry is $w_{k,\ell,m}$.

With this notation, the actions of R_1 and R_2 above on the principal working columns $w_{1,1}$ and $w_{2,2}$ are

$$R_1^T \begin{pmatrix} w_{1,1,1} \\ w_{1,1,2} \\ \vdots \\ w_{1,1,n} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad \text{and} \quad R_2^T \begin{pmatrix} 0 \\ w_{2,2,2} \\ \vdots \\ w_{2,2,n} \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix}. \quad (16)$$

Strictly speaking, the 1 entries in the column vectors on the right-hand sides are $w_{k,k}^T w_{k,k}$, but since V is expected to be orthonormal, the inner product is unity.

Equations like (16) are well known in the linear-algebra literature, for instance [5,24], and form the basis for the implementation of QR decomposition (see also [19]). There are two approaches to the solution: use of Householder reflection matrices and use of Givens rotation matrices. The Householder approach is reviewed in Sect. 7.1, but in the end it is not suitable in this context. The Givens approach works well, and the solution here has an advantage over the classical formulation because we can rely on $w_{k,k}^T w_{k,k} = 1$.

To simplify the explanation once again, let us focus on the four-dimensional space $\mathbb{R}^4$. In this case, $R_1^T w_{1,1}$ needs to zero out three entries below the pivot, and likewise $R_2^T w_{2,2}$ needs to zero out two entries below its pivot. Since one Givens rotation matrix can zero out one column entry at most, a cascade of Givens matrices is required to realize the pattern in (16) and the length of the cascade depends on the dimension of the current subspace.

A Givens rotation imparts a rotation within an $\mathbb{R}^2$ plane embedded in a higher-dimensional space. While a simple two-dimensional counterclockwise rotation matrix is

$$R(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix},$$

an example Givens rotation matrix in $\mathbb{R}^4$ is

$$R_{.,3}(\theta_{.,3}) = \begin{pmatrix} c_3 & -s_3 & & \\ & 1 & & \\ s_3 & c_3 & & \\ & & & 1 \end{pmatrix} \rightarrow \begin{pmatrix} \cdot & \cdot & & \\ & \cdot & & \\ \cdot & \cdot & & \\ & & & \cdot \end{pmatrix},$$

where the dot notation on the right represents the pattern of nonzero entries. The full rotation R_1 is then decomposed into a cascade of Givens rotations:

$$R_1(\theta_{1,2}, \theta_{1,3}, \theta_{1,4}) = R_{1,2}(\theta_{1,2}) R_{1,3}(\theta_{1,3}) R_{1,4}(\theta_{1,4}). \quad (17)$$

The order of the component rotations is arbitrary and is selected for best analytic advantage. That said, a different rotation order yields different Givens rotation angles, so there is no uniqueness to the angles calculated below. The only

material concern is that the rotations must be applied in a consistent order.

An advantageous cascade order of Givens rotations to expand the left-hand equation in (16) follows the pattern

$$\underbrace{\begin{pmatrix} \cdot & \cdot \\ \cdot & \cdot \\ & \cdot \end{pmatrix}}_{R_{1,2}} \underbrace{\begin{pmatrix} \cdot & \cdot \\ & \cdot \\ \cdot & \cdot \end{pmatrix}}_{R_{1,3}} \underbrace{\begin{pmatrix} \cdot & \cdot \\ & \cdot \\ \cdot & \cdot \end{pmatrix}}_{R_{1,4}} \underbrace{\begin{pmatrix} \cdot \\ & \cdot \\ & \cdot \end{pmatrix}}_{\delta_1} = \underbrace{\begin{pmatrix} \cdot \\ \cdot \\ \cdot \\ \cdot \end{pmatrix}}_{w_{1,1}}, \quad (18)$$

where the rotation R_1 has been moved to the other side of the original equation, and δ_k denotes a vector in $\mathbb{R}^{n \times 1}$, in which all entries are zero except for a unit entry in the k th row. Each Givens matrix moves some of the weight from the first row of δ_1 into a specific row below, and otherwise the matrices are not coupled.

In the same vein, Givens rotations to represent R_2 follow the pattern

$$\underbrace{\begin{pmatrix} \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \end{pmatrix}}_{R_{2,3}} \underbrace{\begin{pmatrix} \cdot & \cdot \\ & \cdot \\ \cdot & \cdot \end{pmatrix}}_{R_{2,4}} \underbrace{\begin{pmatrix} \cdot \\ & \cdot \\ & \cdot \end{pmatrix}}_{\delta_2} = \underbrace{\begin{pmatrix} \cdot \\ \cdot \\ \cdot \\ \cdot \end{pmatrix}}_{w_{2,2}}.$$

Similar to before, weight in the $(2, 1)$ entry of δ_2 is shifted into the third and fourth rows of the final vector. It is evident that R_3 requires only one Givens rotation and that for R_4 is simply the identity matrix.

6.3.2 Solution for Givens rotation angles

Now, to solve for the Givens angles, the cascades are multiplied through. For the R_1 cascade in (18), multiplying through yields

$$\begin{pmatrix} c_2 & c_3 & c_4 \\ s_2 & c_3 & c_4 \\ s_3 & c_4 \\ s_4 \end{pmatrix} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix}, \quad (19)$$

where entries a_k denote row entries in $w_{1,1}$ and are used only to further simplify the notation. This is the central equation, and it has a straightforward solution. Yet first, there are several important properties to note:

1. The L_2 norms of the two column vectors are both unity.
2. While there are four equations, the three angles ($\theta_2, \theta_3, \theta_4$) are the only unknowns. The fourth equation is constrained by the L_2 norm.
3. The a_1 entry (or specifically the $w_{1,1,1}$ entry) is non-negative: $a_1 \geq 0$. This holds by construction because the

associated reflection matrix, applied previously, ensures the nonnegative value of this leading entry.

4. In order to uniquely take the arcsine, the rotation angles are restricted to the domain $\theta \in [-\pi/2, \pi/2]$. As a consequence, the sign of each row is solely determined by the sine-function values, the cosine-function values being nonnegative on this domain: $c_k \geq 0$. This leads to a global constraint on the solution since $c_2 c_3 c_4 \geq 0$ could otherwise hold for pairs of negatively signed cosine-function values.

With these properties in mind, a solution to (19) uses the arcsine method, which starts from the bottom row and works upward. The sequence in this example reads

$$\begin{aligned} \theta_4 &= \sin^{-1}(a_4) \\ \theta_3 &= \sin^{-1}(a_3 / \cos \theta_4) \\ \theta_2 &= \sin^{-1}(a_2 / (\cos \theta_4 \cos \theta_3)), \end{aligned} \quad (20)$$

where $\theta_k \in [-\pi/2, \pi/2]$. The sign of each angle is solely governed by the sign of the corresponding a_k entry. Moreover, other than the first angle, the Givens angles are coupled in the sense that no one angle can be calculated from the entries of working matrix W_k alone, and certainly not prior to the rotation of W_{k-1} into W_k . This shows once again that there is no shortcut to working down from the full space through each subspace until $\mathbb{R}^1$ is resolved.

Past the first equation, the arguments to the arcsine functions are quotients, and it will be reassuring to verify that these quotients never exceed unity in absolute value. For θ_3 , we can write

$$a_3^2 = 1 - a_4^2 - a_2^2 - a_1^2 = \cos^2 \theta_4 - (a_1^2 + a_2^2),$$

and thus $a_3^2 \leq \cos^2 \theta_4$, so the arcsine function has a defined value. Similarly, for θ_2 , we have

$$a_2^2 = 1 - a_4^2 - a_3^2 - a_1^2 = \cos^2 \theta_4 - \sin^2 \theta_3 \cos^2 \theta_4 - a_1^2,$$

and therefore $a_2^2 \leq \cos^2 \theta_3 \cos^2 \theta_4$. Again, the arcsine function has a defined value. This pattern holds in any dimension.

The preceding analysis is carried out for each eigenvector in V . For each subspace k , there are $k - 1$ angles to resolve. For convenience, the Givens angles can be organized into the upper right triangle of a square matrix. For instance, in $\mathbb{R}^4$,

$$\theta \leftarrow \begin{pmatrix} 0 & \theta_{1,2} & \theta_{1,3} & \theta_{1,4} \\ & 0 & \theta_{2,3} & \theta_{2,4} \\ & & 0 & \theta_{3,4} \\ & & & 0 \end{pmatrix}. \quad (21)$$

Lastly, it is worth noting that Dash constructed a storage matrix of angles similar to (21) when calculating the embed-

ded angles in a correlation matrix [4]. His interest, shared here, is to depart from the Cartesian form of a matrix—whose entries in his case are correlation values and in the present case are eigenvector entries—and cast them into polar form. Polar form captures the natural representation of the system.

6.3.3 Summary

This section has shown that each eigenvector rotation matrix R_k is constructed from a concatenation of elemental rotation matrices called Givens rotations. Each Givens rotation has an angle, and provided that the rotation order is always preserved, the angles are a meaningful way to represent the orientation of the eigenbasis within the constituent basis.

The representation for θ in (21) highlights that there are $n(n-1)/2$ angles to calculate to complete the full basis rotation in $\mathbb{R}^n$. It will be helpful to connect the angles to a rotation matrix, so let us define a generator function $\mathcal{G}(\cdot)$ such that

$$R = \mathcal{G}(\theta) \quad \text{and} \quad R_k = \mathcal{G}(\theta, k). \quad (22)$$

The job of the generator is to concatenate Givens rotations in the correct order, using the recorded angles, to correctly produce a rotation matrix, either a matrix R_k for one vector or the matrix R for the entire basis.

7 Consideration of alternative approaches

The preceding analysis passes over two choices made during the development of the algorithm. These alternatives are detailed in this section, together with the reasoning for not selecting them.

7.1 Householder reflections

Equations nearly identical to (16) appear in well-known linear algebra texts, such as [5, 23, 24], typically in the context of QR decomposition, along with the advice that Householder reflections are preferred over Givens rotations because all entries below a pivot can be zeroed out in one step. This advice is accurate in the context of QR decomposition, but does not hold in the current context.

Recall that in order to align eigenvector v_1 to constituent basis π_1 , $n-1$ Givens rotations are necessary for a space of dimension n . Arranged appropriately, each rotation has the effect of introducing a zero in the column vector below the pivot. A Householder reflection aligns v_1 to π_1 in a single operation by reflecting v_1 onto π_1 about an appropriately aligned hyperplane in $\mathbb{R}^n$. In general, this hyperplane is neither parallel nor perpendicular to v_1 , and therefore neither are any of the other eigenvectors in V .

The Householder operator is written as

$$H \equiv I - 2uu^T,$$

where u is the Householder vector that lies perpendicular to the reflecting hyperplane. The operator is Hermitian, unitary, involutory, and has $\det(H) = -1$. In fact, all eigenvalues are $+1$, except for one that is -1 .

Let us reconsider the eigenvector representation (7) and apply a general unitary transform U to V , as in

$$U^T V S = I. \quad (23)$$

(Unitary transform U is not to be confused with the matrix of projected data U in SVD equation (2).) The equality holds for any unitary operator, so Householder reflections are admitted. Let us then compare expanded representations in the form of (9) for both rotation and reflection operators to the left of V :

$$\begin{array}{ccc} \underbrace{R_n^T \dots R_2^T R_1^T}_{\text{basis rotations}} & V & \underbrace{S_1 S_2 \dots S_n}_{\text{eigenvector reflections}} = I \\ \underbrace{H_n^T \dots H_2^T H_1^T}_{\text{basis reflections}} & V & \underbrace{S_1 S_2 \dots S_n}_{\text{eigenvector reflections}} = I \end{array}$$

Even though (23) holds for Householder reflections, the equality itself *is not the goal*; rather, the *interpretability* of the eigenvector orientation is paramount. In the rotations method, reflections only reflect one eigenvector at a time, whereas rotations transform the basis as a whole. In contrast, the reflections method intermingles eigenvector reflections with basis reflections, thereby disrupting the interpretation of the basis orientation.

A comparison between the rotation and reflection methods is illustrated in Fig. 4. Here, the eigenvector basis spans $\mathbb{R}^2$ with an initial orientation of

$$V = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & -1 \\ -1 & +1 \end{pmatrix}. \quad (24)$$

Following the upper pane in the figure, where the reflection method is illustrated, it is found that $v_1^T \pi_1 < 0$ and therefore $s_1 = -1$ in order to reflect v_1 . Once $V S_1$ is completed, a clockwise rotation of 45° is imparted by R_1 to align v_1 with π_1 , thus completing $R_1^T V S_1$. Lastly, $v_2^T \pi_2 > 0$ (trivially so) and therefore $s_2 = +1$. The tuple of sign reflections for the eigenvectors in V is therefore $(s_1, s_2) = (-1, +1)$.

A parallel sequence is taken along the lower pane, where a Householder reflection is used in place of rotation. Expression $V S_1$ is as before, but to align v_1 with π_1 a reflection is used. To do so, Householder vector u_1 is oriented so that a reflection plane is inclined by 22.5° from the π_1 axis. Reflection of $V S_1$ about this plane indeed aligns v_1 to π_1 but has

Fig. 4 Rotation versus reflection transformation comparison. *Top*, the rotation method produces the sign tuple $(s_1, s_2) = (-1, +1)$. *Bottom*, the Householder reflection method produces the sign tuple $(-1, -1)$

the side effect of reflecting v_2 too. As a consequence, when it is time to inspect the orientation of v_2 with respect to π_2 , we find that $s_2 = -1$. Therefore, the tuple of eigenvector signs is now $(-1, -1)$.

Our focus here is not to discuss which tuple of signs is “correct,” but rather on how to interpret the representation. The choice made in this article is to use rotations instead of reflections for the basis transformations so that the basis orientation is preserved through this essential step. The use of Householder reflections, by contrast, scatters the basis orientation after each application.

7.2 Arctan calculation instead of arcsine

A solution to (19) is stated above by Equations (20) in terms of arcsine functions, and a requirement of this approach is that the range of angles be restricted to $\theta_k \in [-\pi/2, \pi/2]$. While this is the preferred solution, there is another approach that uses arctan functions instead of arcsine. The arctan solution to (19) starts at the top of the column vector and walks downward, the angles being

$$\begin{aligned} \theta_2 &= \arctan2(a_2, a_1 \csc \theta_1) \\ \theta_3 &= \arctan2(a_3, a_2 \csc \theta_2) \\ \theta_4 &= \arctan2(a_4, a_3 \csc \theta_3) \end{aligned} \quad (25)$$

The first angle is defined as $\theta_1 \equiv \pi/2$ (which is done solely for equation symmetry), and the arctan2 function is the four-quadrant form $\arctan2(y, x)$, in which the admissible angular domain is $\theta \in [-\pi, \pi]$. It would appear that the arctan method is a better choice.

The difficulty arises with edge cases. Let us consider the vector

$$a = (1/\sqrt{2}, 1/\sqrt{2}, 0, 0)^T.$$

The sequence of arctan angles is then

$$\begin{aligned} \theta_1 &= \pi/2 \\ \theta_2 &= \arctan2(1/\sqrt{2}, 1/\sqrt{2} \times \csc \theta_1) = \pi/4 \\ \theta_3 &= \arctan2(0, 1/\sqrt{2} \times \csc \theta_2) = 0 \\ \theta_4 &= \arctan2(0, 0 \times \csc \theta_3) \end{aligned}$$

The final arctan evaluation is not numerically stable. For computer systems that support signed zero, the arctan can either be $\arctan2(0, 0) \rightarrow 0$ or $\arctan2(0, -0) \rightarrow -\pi$, see [8,18]. The latter case is catastrophic because product $V S_1$ has forced the sign of $v_{1,1}$ to be positive, yet now $R_1^T V S_1$ flips the sign since $\cos \pi = -1$. After $R_1^T V S_1$ is constructed, the first eigenvector is not considered thereafter, so there is no chance, unless the algorithm is changed, to correct this spurious flip.

Therefore, to avoid edge cases that may give false results, the arcsine method is preferred.

8 Application to regression

Returning now to the original motivation for a consistently oriented eigenvector basis, the workflow for regression and prediction on the eigenbasis becomes:

1. SVD: Find (V, Λ) from in-sample data P_{in} such that $P_{\text{in}} = U \Lambda^{1/2} V^T$.

2. Orient: Find (R, S) from V such that

$$R^T V S = I.$$

3. Regression: Find $(\hat{\beta}, \epsilon)$ from (y, U, S, Λ) such that

$$y = U S \Lambda^{1/2} \hat{\beta} + \epsilon. \quad (26)$$

4. Prediction: Calculate Ey from $(P_{\text{out}}, R, \hat{\beta})$, where P_{out} is out-of-sample data, with

$$Ey = P_{\text{out}} R \hat{\beta}. \quad (27)$$

The form of (26) here assumes that there is no exogenous correction to the direction of y in response to a change of the factors along the eigenbasis, as discussed in Sect. 1.

The workflow highlights that both elements of the orientation solution, rotation R and reflection S , are used, although for different purposes. Regression (26) requires the reflection matrix S to ensure that the signs of $\hat{\beta}$ faithfully align to the response of y . Predictive use with out-of-sample data requires the $\hat{\beta}$ estimate, as expected, but also the rotation matrix R . The rotation orients the constituent basis, which is observable, into the eigenbasis, which is not.

Without dimension reduction, the rotation in (27) is associative:

$$(P_{\text{out}} R) \hat{\beta} = P_{\text{out}} (R \hat{\beta}) \quad \text{for no dim reduction.} \quad (28)$$

However, associativity breaks with dimension reduction. Principal components analysis, for instance, discards all but the top few eigenvector components (as ranked by their corresponding eigenvalue) and uses the remaining factors in the regression. The number of entries in $\hat{\beta}$ equals the number of remaining components, not the number of constituent components. In this case, only $(P_{\text{out}} R)$ can be used. Typically, online calculation of this product is simple because the calculation is updated for each new observation. Rather than being an $\mathbb{R}^{m \times n}$ matrix, out-of-sample updates are $P_{\text{out}} \in \mathbb{R}^{1 \times n}$, which means that $P_{\text{out}} R$ requires only a BLAS level-2 (vector-matrix) call on an optimized system [3].

9 Treatment of evolving data

Eigenvectors and -values calculated from data are themselves sample estimates subject to uncertainties based on the particular sample at hand, statistical uncertainties based on the number of independent samples available in each dimension¹

¹ See [16], Section 4.4.2, Dispersion and Hidden Factors.

as well as on assumptions about the underlying distribution,² and scale-related numerical uncertainties based on the degree of colinearity among factors. All of these uncertainties exist for a single, static dataset.

Additional uncertainty becomes manifest when data evolve because the estimate of the eigensystem will vary at each point in time, this variation being due in part to sample noise, and possibly due to nonstationarity of the underlying random processes.

In the presence of evolving data, therefore, the eigensystem fluctuates, even when consistently oriented. It is natural, then, to seek statistics for the location and dispersion of the eigensystem. To do so, two orthogonal modes of variation are identified (see Fig. 5): stretch variation, which is tied to change of the eigenvalues, and wobble variation, which is tied to change of the eigenvectors. Each is treated in turn.

9.1 Stretch variation

Stretch variation is nominally simple because eigenvalues λ_i are scalar, real numbers. The mean and variance follow from the usual sample forms of the statistics. For an ensemble of N samples, the average eigenvalues are simply

$$\bar{\Lambda} = (1/N) \sum_{i=1}^N \Lambda[i]. \quad (29)$$

A challenge for the variance statistic is that eigenvalues are not independent since

$$\det(V) = \lambda_1 \lambda_2 \cdots \lambda_n = +1.$$

How the covariance manifests itself is specific to the dataset at hand.

9.2 Wobble variation

Quantification of wobble variation requires a different toolset because eigenvectors are *vectors*, and these orthonormal vectors point onto the surface of a hypersphere such that $v_k \in \mathbb{S}^{n-1} \subseteq \mathbb{R}^n$. Figure 6 left illustrates a case in point for $\mathbb{R}^3$: A stationary variation of the eigenvectors forms point constellations on the surface of the sphere $\mathbb{S}^2$, one constellation for each vector.³ The field of directional statistics

² See [16], Section 4.3, Maximum Likelihood Estimators.

³ Figure 6 is drawn using point constellations drawn from the von Mises–Fisher distribution [15] with concentration $\kappa = 100$. Sampling this distribution in $\mathbb{R}^2$ and $\mathbb{R}^3$ was done using the VONMISESFISHER class in the TENSORFLOW Probability package [26], which belongs to the TENSORFLOW platform [25]. The VONMISESFISHER class draws samples using the nonrejection-based method detailed by Kurz and Hanebeck [11].

Fig. 5 Orthogonal modes of variation of an eigensystem. *Left*, in one variation mode, the *eigenvalues* of a new observation differ from those of the past, thereby stretching and/or compressing the point-cloud ellipse. *Right*, in another variation mode, it is the *eigenvectors* that vary in this way, thereby changing the orientation of the point-cloud ellipse with respect to the constituent basis

Fig. 6 Illustration of eigenvector wobble in $\mathbb{R}^3$ and its $\mathbb{R}^2$ subspace. *Left*, points illustrating eigenvector wobble for a stationary process. The points lie on the hypersphere $\mathbb{S}^2$, and there is one constellation for each eigenvector. Samples were drawn from a von Mises–Fisher distribution with concentration $\kappa = 100$. In this illustration, the point constella-

tions are identical, each having been rotated onto the mean direction of each vector. In general, the point constellations are mixtures having different component dispersions. *Right*, eigenvector v_2 may wobble in the $(\bar{v}_2, \bar{v}_3)$ plane independently from v_1 wobble. Points here illustrate wobble in this $\mathbb{R}^2$ subspace

informs us regarding how to define mean direction (the vector analogue of location) and dispersion in a consistent manner for point constellations such as these [9,13,15]. Directional statistics provides a guide on how to use the vector information available from $\mathcal{V}$ and embedded angle information recorded in θ , see (21). Recent work focuses on the application of these statistics to machine learning [21]. Nonetheless, application of directional statistics to eigenvector systems appears to be underrepresented in the literature.

The following discussion only applies to eigensystems that have been consistent orientated.

9.2.1 Estimation of mean direction

Both mean direction and directional dispersion are measurable statistics for eigensystems, and determining the mean direction is the simpler of the two. The reason this is so is

because the eigenvectors are orthogonal in every instance of an eigenbasis; thus, the mean directions must also be orthogonal. Consequently, the eigenvectors can be treated uniformly.

It is a tenet of directional statistics that the *mean direction* is calculated from unit vectors, not from their angles.⁴ For an ensemble of N unit vectors $x[i]$, the mean direction $\bar{x}$ is defined as

$$\bar{x} \equiv \|x_S\|^{-1} x_S \quad \text{where} \quad x_S = \sum_{i=1}^N x[i], \quad (30)$$

with the caveat that the mean direction is undefined for $\|x_S\| = 0$. The *resultant vector* x_S is the vector sum of component vectors $x[i]$ and has length $\|x_S\|$ under an L_2 norm

⁴ Simple examples exist in the cited literature that show that an angle-based measure is not invariant to the choice a “zero” angle reference.

Fig. 7 Resultant vector $x_S \equiv \sum_i x[i]$ for two levels of concentration, $\kappa = (1, 10)$. Lower concentration, which is akin to higher dispersion, leads to a smaller expected norm of the resultant vector $\|x_S\|$, whereas higher concentration leads to a larger expected norm. Sample vectors

in $\mathbb{R}^2$ were generated from the von Mises–Fisher distribution (Note that the difference in horizontal and vertical scales distorts the length of the unit vectors in the plots)

(see Fig. 7). The mean direction is thus a unit-vector version of the resultant vector.

Extending this construction for mean direction to an ensemble of oriented eigenvector matrices $\mathcal{V}$, the mean location of the eigenvectors is

$$\bar{\mathcal{V}} = \|\mathcal{V}\|^{-1} \mathcal{V}_S \quad \text{where} \quad \mathcal{V}_S = \sum_{i=1}^N \mathcal{V}[i], \quad (31)$$

and

$$\|\mathcal{V}\| \equiv \text{diag}(\|v_{S,1}\|, \|v_{S,2}\|, \dots, \|v_{S,N}\|) \quad (32)$$

for normalization.

Using the methodologies above, the average basis $\bar{\mathcal{V}}$ can be rotated onto I with a suitable rotation matrix $R^T \bar{\mathcal{V}} = I$, and in doing so, the Cartesian form of $\bar{\mathcal{V}}$ can be converted into polar form. Using the generator function $\mathcal{G}()$ from (22) to connect the two, we have

$$\bar{\mathcal{V}} = \mathcal{G}(\bar{\theta}). \quad (33)$$

Note that $\bar{\theta}$ is *not* the arithmetic average of angles $\theta[i]$ but is exclusively derived from $\bar{\mathcal{V}}$. In fact, an alternative way to express the basis sum in (31) is to write

$$\mathcal{V}_S = \sum_{i=1}^N \mathcal{G}(\theta[i]), \quad (34)$$

and from this, we see the path of analysis:

$$\theta[i] \rightarrow \mathcal{V}[i] \rightarrow \bar{\mathcal{V}} \rightarrow \bar{\theta}.$$

9.2.2 Estimation of dispersion

Dispersion estimation is more involved for an ensemble of eigenbases than for an ensemble of single vectors because both common and differential modes of variation exist. Referring to Fig. 6 left, consider a case where the only driver

of directional variation for the eigenbasis is a change of the pointing direction of v_1 . As v_1 scatters about its mean direction, vectors $v_{2,3}$ will likewise scatter, together forming three constellations of points, as in the figure. Thus, wobble in v_1 imparts wobble in the other vectors. Since dispersion is a scalar independent of direction, the dispersion estimates along all three directions in this case are identical.

Yet, there is another possible driver for variation that is orthogonal to v_1 , and that is motion within the (v_2, v_3) plane. Such motion is equivalent to a pirouette of the (v_2, v_3) plane about v_1 . Taking $(\bar{v}_2, \bar{v}_3)$ as a reference, variation within this $\mathbb{R}^2$ subspace needs to be estimated, see Fig. 6 right. When viewed from $\mathbb{R}^n$, however, the point constellation distribution about $\bar{v}_2$ is a mixture of variation drivers.

Regardless of the estimation technique for dispersion, it is clear that the pattern developed in Sect. 6.3 for walking through subspaces of $\mathbb{R}^n$ to attain $R^T \mathcal{V} S = I$ occurs here as well. Thus, rather than calculating a dispersion estimate for each eigenvector, since they are mixtures, a dispersion estimate is made for each subspace.

Returning to the generator for subspace k in (22), given an ensemble $\theta[i]$ the ensemble of the k th subspace is

$$R_k[i] = \mathcal{G}(\theta[i], k), \quad (35)$$

where $R_k \in \mathbb{R}^{n \times n}$. From $R_k[i]$, the k th column is extracted and the lower k entries taken to produce column vector $x_k[i] \in \mathbb{R}^{(n-k+1) \times 1}$. Figuratively,

$$R_k = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \vdots & \vdots & \vdots \\ 0 & \vdots & \vdots & \vdots \\ 0 & \vdots & \vdots & \vdots \end{pmatrix} \rightarrow x_k = \begin{pmatrix} \vdots \\ \vdots \\ \vdots \\ \vdots \end{pmatrix}. \quad (36)$$

The reason behind of removing the top zero entries from the column vector is that parametric models of dispersion are isotropic in the subspace, and therefore keeping dimensions with zero variation would create an unintended distortion.

Now that the vectors of interest $x_k[i]$ have been identified, their dispersions can be considered. Unlike the mean direction, there is no one measure of dispersion. Two approaches are touched upon here, one being model free and the other based on a parametric density function. Both approaches rely on the resultant vector x_S .

Returning to Fig. 7, the two panels differ in that on the left there is a higher degree of randomness in the pointing directions of the unit vectors along the sequence, while on the right there is a lower degree of randomness. The consequence is that the expected *lengths* of the result vectors are different: The lower the randomness, the longer the expected length. Define the *mean resultant length* for N samples as

$$\bar{r} = \|x_S\|/N, \quad 0 \leq \bar{r} \leq 1. \quad (36)$$

For distributions on a circle, Mardia and Jupp define a model-free *circular variance* as $\bar{V} \equiv 1 - \bar{r}$ (see [15]). Among other things, this variance is higher for more dispersed unit vectors and lower for less dispersed vectors. Yet, after the generalization to higher dimensions, it remains unclear how to compare the variance from one dimension to another.

As an alternative, a model-based framework starts by positing an underlying parametric probability distribution and seeks to define its properties and validate that they meet the desired criteria. In the present case, eigenvectors are directional, as opposed to axial; their evolution-induced scatter is probably unimodal, at least across short sequences; and their dimensionality is essentially arbitrary. An appropriate choice for an underlying distribution therefore is the von Mises–Fisher (vMF) distribution. The vMF density function in $\mathbb{R}^n$ is [21]

$$p_{\text{vmf}}(x; n, \mu, \kappa) = c_n(\kappa) e^{\kappa \mu^T x} \quad \text{where} \quad c_n(\kappa) = \frac{\kappa^{n/2-1}}{(2\pi)^{n/2} I_{n/2-1}(\kappa)}. \quad (37)$$

The single argument to the density function is the vector $x \in \mathbb{R}^{n \times 1}$, and the parameters are the dimension n , mean direction $\mu \in \mathbb{R}^{n \times 1}$, and concentration $\kappa \in [0, \infty)$. The mean direction parameter is the same as in (30): $\mu = \bar{x}$. The concentration κ parameter, a scalar, is like an inverse variance. For $\kappa = 0$, unit vectors x are uniformly distributed on the hypersphere $\mathbb{S}^{n-1}$, whereas with $\kappa \rightarrow \infty$, the density concentrates to the mean direction μ . Maximum likelihood estimation yields

$$\hat{\mu} = \|x_S\|^{-1} x_S \quad \text{and} \quad \hat{\kappa} = A_n^{-1}(\bar{r}),$$

where the nonlinear function $A_n(\cdot)$ is

$$A_n(\kappa) = \frac{I_{n/2}(\kappa)}{I_{n/2-1}(\kappa)} = \bar{r}$$

and $I_\alpha(x)$ is the modified Bessel function of the first kind. A useful approximation is

$$\hat{\kappa} \simeq \frac{\bar{r}(n - \bar{r}^2)}{1 - \bar{r}^2}.$$

The key aspect here is the connection of the concentration parameter κ to the mean resultant length $\bar{r}$ from (36). Notice as well that κ is related to the dimensionality n , indicating that concentration values across different dimensions cannot be directly compared.

Sampling from the vMF distribution is naturally desirable. Early approaches used rejection methods (see Introduction in [11]), these being simpler, yet the runtime is not deterministic. Kurz and Hanebeck have since reported a stochastic sampling method that is deterministic, see [11], and it is this method that is implemented in TENSORFLOW (see footnote 3 on page 11). The samples in Fig. 6 were calculated with TENSORFLOW.

To conclude this section, in order to find the underlying drivers of variation along a sequence of eigenbasis observations, concentration parameters κ , or at least the mean resultant lengths $\bar{r}$, should be computed for each descending subspace in $\mathbb{R}^n$. The n concentration parameters might then be stored as a vector,

$$\kappa_{\text{basis}} = (\kappa_1, \kappa_2, \dots, \kappa_n)^T. \quad (38)$$

9.3 Rank-order change of eigenvectors

Along an evolving system, the rank order of eigenvectors may change. This paper proposes no mathematically rigorous solution to dealing with such change—it is unclear that one exists, nor is its absence a defect of the current work—yet there are a few remarks that may assist the analyst to determine whether rank-order change is a property of the data itself or a spurious effect that results from an upstream issue.

Eigenvectors are “labeled” by their eigenvalue: In the absence of eigenvalue degeneracy, each vector is mapped one-to-one with a value. The concept of rank-order change of eigenvectors along evolving data is only meaningful if a second label can be attached to the vectors such that the second label does not change when the first does. Intuitively (for evolving data), a second label is the pointing direction of the vectors. This choice is fraught: Two time series are not comparable if they begin at different moments along a sequence because the initial labeling may not match. Additionally, validation of secondary labeling stability needs to be performed, and revalidation is necessary for each update because stability cannot be guaranteed.

With an assumption that a second labeling type is suitably stable, rank-order changes may still be observed. Before it can be concluded that the effect is a trait of the data, modeling

assumptions must be considered first. There is a significant assumption embedded in the estimation of the eigensystem itself: Using an `eig(...)` function call presupposes that the underlying statistical distribution is Gaussian. However, data are often heavy tailed, so copula or implicit methods are appropriate since Gaussian-based methods are not robust (see [16]). Another modeling concern is the number of independent samples per dimension used in each panel: Having too few samples will inherently skew the eigenvalue spectrum, and when coupled with sample noise, may induce spurious rank-order changes (again, see [16]). Lastly, PCA ought only be performed on homogeneous data categories because the mixing of categories may well lead to spurious rank-order changes. Better to apply PCA independently to each category and then combine.

If rank-order change is determined to be a true trait of the data, then axial statistics are applied in place of directional statistics. Referring to Fig. 6, two or more constellations will have some of their points mirror-imaged through the origin, creating a “barbell” shape along a common *axis*. The juxtaposition of the Watson density function, which is an axial distribution, to the von Mises–Fisher distribution,

$$p_{\text{watson}}(x; \mu, \kappa) = z_1 e^{\kappa(\mu^T x)^2} \quad \text{vs} \\ p_{\text{vmf}}(x; \mu, \kappa) = z_2 e^{\kappa(\mu^T x)},$$

shows that the axial direction is squared in order to treat the dual-signed nature of the pointing directions [22]. With this, directional statistics can be applied once again.

9.4 Regression revisited

Let us apply the averages developed in this section to the equations in Sect. 8. The SVD and orientation steps remain the same. The linear regression of (26) is modified to use the local average of the eigenvalues from (29),

$$y = US \bar{\Lambda}^{1/2} \hat{\beta} + \epsilon. \quad (39)$$

The prediction step is also modified from (27) to read

$$Ey = P_{\text{out}} \bar{V} \hat{\beta}. \quad (40)$$

Here, the mean direction of the eigenvectors in (31) replaces the single-observation rotation matrix R in the original.

Use of eigenvalue and eigenvector averages will reduce sample-based fluctuation and therefore may improve the predictions. However, from a time-series perspective, averages impart delay because an average is constructed based on a lookback interval. For stationary underlying processes, the delay may not be a problem, yet for nonstationary processes, the lag between the average and the current state can lead to lower quality predictions.

In any event, at the very least, this section has presented a methodology to decompose variation across a sequence of eigensystem observations into meaningful quantities.

10 Reference implementation

The Python package `thucyd`, written by this author, is freely available from PyPI⁵ and CONDA-FORGE⁶ and can be used directly once installed. The source code is available on at <https://gitlab.com/thucyd-dev/thucyd>.

There are two functions exposed on the interface of `thucyd.eigen`:

- `orient_eigenvectors` implements the algorithm in Sect. 6, and
- `generate_oriented_eigenvectors` implements

$$R_k = \mathcal{G}(\theta, k), \text{ Eq. (22).}$$

The pseudocode in listing 1 outlines the reference implementation for `orient_eigenvectors` that is available at the above-cited source-code repositories. Matrix and vector indexing follows the Python Numpy notation.

10.1 Eigenvector orientation

A simple orientation example in $\mathbb{R}^3$ reads as follows:

```
import numpy as np
from thucyd.eigen import orient_eigenvectors
from thucyd.eigen import generate_oriented_eigenvectors

# setup an example eigenvector mtx with qr
full_dim = 3; A = np.eye(full_dim)
A[:, 0] = 1. / np.sqrt(full_dim)
V, _ = np.linalg.qr(A)
# setup trivial eigenvalues, descending order
E = np.diag(np.arange(full_dim)[: -1])
# reflect the second eigenvector to build Vin
S = np.diag(np.array([1., -1., 1.]))
Vin = V.dot(S)
# orient the original matrix V
Vor, Eor, signs, theta_matrix, sort_indices =
    orient_eigenvectors(Vin, E)
```

Execution of this snippet⁷ converts the original basis

⁵ Available at <https://pypi.org/project/thucyd/>.

⁶ Available at <https://github.com/conda-forge/thucyd-feedstock>.

⁷ The snippet was run with Numpy version 1.16.4 on MacOS 10.14.5. OpenBlas and Intel MKL produced the same results, as expected.

```

def OrientEigenvectors(V, E):
    Vsort, Esort, SortIndices ← SortEigenvectors(V, E)
    Vwork ← copy(Vsort), N ← dim(V)
    for i = 0 to N - 1 do
        SignFlipVec[i] ← (Vwork[i, i] >= 0) ? 1 : -1
        Vwork[:, i] ← Vwork[:, i] · SignFlipVec[i]
        Vwork, AnglesMtx[i, :] ←
            ReduceDimensionByOne(N, i, Vwork)
    end
    Vor ← Vsort · diag(SignFlipVec), Eor ← Esort
    return Vor, Eor, AnglesMtx, SignFlipVec, SortIndices

def SortEigenvectors(V, E):
    SortIndices ← argsort(abs(diag(E)))
    return Vsort[:, SortIndices], diag(Esort[SortIndices]),
    SortIndices

def ReduceDimensionByOne(N, i, Vwork):
    Vcol ← Vwork[:, i]
    AnglesCol ← SolveRotationAnglesInSubDim(N, i,
    Vcol)
    R ← ConstructSubspaceRotationMtx(N, i,
    AnglesCol)
    Vwork ← RT Vwork
    return Vwork, AnglesColT

def SolveRotationAnglesInSubDim(N, i, Vcol):
    AnglesWork ← zeros(0 : N)
    r ← 1
    for j = N to i + 1 step -1 do
        y ← Vcol[j]
        r ← r * cos(AnglesWork[j + 1])
        AnglesWork[j] ← (r != 0) ? arcsin(y/r) : 0
    end
    return AnglesWork[: N]

def ConstructSubspaceRotationMtx(N, i,
AnglesCol):
    R ← I(N)
    for j = N to i + 1 step -1 do
        R ←
            MakeGivensRotation(N, i, j, AnglesCol[j]) · R
    end
    return R

def MakeGivensRotation(N, i, j, θ):
    R ← I(N), c ← cos(θ), s ← sin(θ)
    R[i, i] ← c, R[j, j] ← c, R[i, j] ← s, R[j, i] ← -s
    return R

```

Algorithm 1: Pseudocode implementation of orient_eigenvectors.

$$\text{Vin} = \begin{pmatrix} -0.577 & -0.408 & -0.707 \\ -0.577 & +0.816 & 0 \\ -0.577 & -0.408 & +0.707 \end{pmatrix} \quad \text{to} \quad
 \text{Vor} = \begin{pmatrix} 0.577 & -0.408 & -0.707 \\ 0.577 & +0.816 & 0 \\ 0.577 & -0.408 & +0.707 \end{pmatrix},$$

with signs = (-1, +1, +1). The cross product exists in $\mathbb{R}^3$ so the column vectors in Vin can be inspected to determine their handedness. In this case, Vin is a left-handed basis, so a pure rotation cannot align the basis of Vin to I. However, once the first eigenvector is reflected, a pure rotation is all that is required for alignment. The matrix of rotation angles theta_mtx, which follows (21) in form and is expressed here in degrees, is

$$\theta = \begin{pmatrix} 0 & 45 & 35.264 \\ 0 & 0 & -30 \\ 0 & 0 & 0 \end{pmatrix}.$$

The generator function consumes this array of angles to reconstruct the oriented eigenbasis, as in

```

# reconstruct Vor from rotation angle matrix
Vor_recon = generate_oriented_eigenvectors(theta_mtx)

```

Running this snippet will show that Vor_recon = Vor.

10.2 Eigenvector reconstruction from rotations

The reconstruction of an oriented eigenvector from rotations is better illustrated in $\mathbb{R}^4$. Here is a snippet that calls the generate_oriented_eigenvectors api function.

```

# setup example eigenvector and eigenvalue mtx
full_dim = 4; A = np.eye(full_dim)
A[:, 0] = 1. / np.sqrt(full_dim)
Vin, _ = np.linalg.qr(A)
E = np.diag(np.arange(full_dim)[::-1])
# cast Vin into an oriented basis
Vor, _, theta_matrix, _ = orient_eigenvectors(Vin, E)
# build W_k
W_k = np.eye(full_dim)
# iterate across columns in (Vor, W_k)
for cursor in np.arange(full_dim):
    # call focus fcn (the 2nd argument is in base(1))
    W_k = W_k.dot(
        generate_oriented_eigenvectors(
            theta_matrix, cursor + 1))
# reconstructed eigenbasis
Vor_recon = W_k

```

The output theta_matrix from the function orient_eigenvectors is the input to the focus function here. While we could simply call

```
Vor_recon = generate_oriented_eigenvectors(theta_matrix)
```

it is revealing to build up Vor_recon one subspace at a time.

The oriented eigenvector and angles matrices are

$$V_{\text{or}} = \begin{pmatrix} 0.5 & -0.289 & -0.408 & -0.707 \\ 0.5 & +0.866 & 0.000 & 0.000 \\ 0.5 & -0.289 & +0.816 & 0.000 \\ 0.5 & -0.289 & -0.408 & 0.707 \end{pmatrix} \quad \text{and}$$

$$\theta \simeq \begin{pmatrix} 0 & 45 & 35 & 30 \\ 0 & 0 & -30 & -19 \\ 0 & 0 & 0 & -30 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

We can now build up V_{or} through W_k from these embedded angles:

$$R_1 = \begin{pmatrix} 0.5 & -0.707 & -0.408 & -0.289 \\ 0.5 & 0.707 & -0.408 & -0.289 \\ 0.5 & 0.000 & 0.816 & -0.289 \\ 0.5 & 0.000 & 0.000 & 0.866 \end{pmatrix},$$

$$R_1 R_2 = \begin{pmatrix} 0.5 & -0.289 & -0.707 & -0.408 \\ 0.5 & +0.866 & 0.000 & 0.000 \\ 0.5 & -0.289 & 0.707 & -0.408 \\ 0.5 & -0.289 & 0.000 & 0.816 \end{pmatrix},$$

$$R_1 R_2 R_3 = \begin{pmatrix} 0.5 & -0.289 & -0.408 & -0.707 \\ 0.5 & +0.866 & 0.000 & 0.000 \\ 0.5 & -0.289 & +0.816 & 0.000 \\ 0.5 & -0.289 & -0.408 & 0.707 \end{pmatrix}, \quad \text{and}$$

$$R_1 R_2 R_3 R_4 = \begin{pmatrix} 0.5 & -0.289 & -0.408 & -0.707 \\ 0.5 & +0.866 & 0.000 & 0.000 \\ 0.5 & -0.289 & +0.816 & 0.000 \\ 0.5 & -0.289 & -0.408 & 0.707 \end{pmatrix}.$$

It is apparent that as the subspace rotations are concatenated, the rightward columns become aligned to the reference eigenbasis V_{or} . That the reconstructed matrix matches the reference matrix before the last rotation is simply because the last rotation, R_4 , is the identity matrix and is only included for symmetry.

Looking ahead, an optimized implementation can avoid trigonometry by carrying the sine functions $\sin \theta$ rather than the angle values θ . On the range $\theta \in [-\pi/2, \pi/2]$, the function is invertible, and thus either representation will do. Elimination of trigonometry leaves only arithmetic and matrix multiplication, thus making the `orient` algorithm a candidate for LAPACK implementation.

11 An example from weather data

Figure 1 illustrates the ambiguity of the eigenvector sign for a single point cloud in $\mathbb{R}^3$, and Fig. 2 shows how the ambiguity imprints itself on a time series of weights β generated from

a sequence of linear regressions. These figures illustrate—in a simplified, pedagogical way—this author’s experience with proprietary datasets that are typical of the capital markets. The markets offer a natural source for data that evolves temporally.

This section turns to a publicly available dataset to illustrate the central point of this work, being that the ambiguity of eigenvector signs along an evolving dataset hinders interpretability, whereas post-processing each `eig(...)` or `SVD(...)` call with an orientation step can create a stable, interpretable temporal path.

Weather, like the markets, evolves temporally, and the United States National Oceanographic and Atmospheric Agency (NOAA) collects weather data worldwide and publishes summaries publicly. Daily maximum-temperature data reported by the Global Historical Climatology Network [17] from three NOAA sites along the east coast of the USA and one at Bermuda were assembled, and Fig. 8 *left* pinpoints the site locations. The east-coast sites are used to build an a data panel P in what follows. An example of the temperature readings from Virginia appears in Fig. 8 *center*. Daily readings from 2010–2019, inclusive, were used in this study, from which 26 records were removed because of missing data.⁸

It is apparent from the temperature data that the strong seasonality component should be removed, and in doing so, we are left with temperature *deviations* from the expected temperature. Rather than build a sophisticated model of seasonality for this otherwise simple example, a local average spanning six weeks was taken over the raw data.⁹ The differences between a temperature reading and its local average value are then used going forward. A data panel $P \in \mathbb{R}^{m \times 3}$, as illustrated in Fig. 8 *right*, is then built, where the entries are temperature deviations and each column represents an east-coast site. A sliding window three years in length is then translated to “evolve” the data. Denote by P_k a slice of P that starts on the k th record. For each P_k slice, SVD is performed,

$$P_k = U_k \Lambda_k^{1/2} V_k^T, \quad (41)$$

(see (2)) and the sequence of (V_k, Λ_k) is analyzed. To give a sense of the embedded correlation structure, the correlation matrix from a typical slice is

⁸ For data and code, visit <https://gitlab.com/thucyd-dev/thucyd-eigen-working-examples>.

⁹ The averaging window was noncausal so that the resultant local average has no delay with respect to the input.

Fig. 8 Temperature data overview: *Left*, locations of the four US weather stations used in this example: Maine, Virginia, Florida, and Bermuda. (This map was created with Google Maps using NOAA-station coordinates.) The site codes are USC00080228, USC00053005, USC00458773, and BDM00078016. *Center*, example of the daily

maximum-temperature reading from Virginia, and a contemporaneous local average used as a crude seasonality adjustment. *Right*, illustration of data panel P that contains approximately 10 years of daily data for the east-coast three sites. A rolling 3-year slice is translated down the panel in order to “evolve” the system

Fig. 9 Time series of eigenvalues shows relative stability and, in particular, no crossings

$$\text{Corr} \simeq \begin{pmatrix} \text{FL} & \text{VA} & \text{ME} \\ 1 & 0.20 & 0.04 \\ & 1 & 0.16 \\ & & 1 \end{pmatrix} \begin{matrix} \text{FL} \\ \text{VA} \\ \text{ME} \end{matrix}.$$

Given the low FL/ME correlation, we do not expect a pronounced third mode in the eigensystem. Indeed, a typical (oriented) eigenvector matrix sorted by decreasing eigenvalue,

$$V_k \simeq \begin{pmatrix} \text{mode-1} & \text{mode-2} & \text{mode-3} \\ 0.18 & -0.23 & 0.96 \\ 0.69 & -0.66 & -0.29 \\ 0.70 & 0.71 & 0.04 \end{pmatrix} \begin{matrix} \text{FL} \\ \text{VA} \\ \text{ME} \end{matrix},$$

shows that mode 1 is a common mode, mode 2 is a differential between (FL, VA) and ME, and mode three—typically a butterfly—is essentially a differential between FL and VA.

Turning to the evolution of the system, let us begin with the eigenvalues since they are scalar. The time series of eigenvalues generated by translating the fixed-length slice down the panel shows that there are no crossings (see Fig. 9), which in turn indicates that the rank-order change of eigenvectors, discussed in Sect. 9.3, is not a concern for this dataset.

Fig. 10 Unoriented and oriented eigenvectors from evolving data. *Left*, the arbitrary eigenvector signs create six point clouds on S^2 . Moreover, some V_k are in a right-hand basis while others are left-handed. *Right*, point clouds of oriented eigenvectors have a nearly fixed orientation and are always right-handed

Next, the eigenvectors have “unoriented” and “oriented” forms. Eigenvectors V_k from (41) are in unoriented form, and eigenvectors $\mathcal{V}_k$ generated by

$$\mathcal{V}_k, \dots, S_k, \dots = \text{orient_eigenvectors}(V_k, \dots)$$

are oriented. As seen in Fig. 10 *left*, the arbitrariness of the eigenvector signs creates three pairs of point clouds, where each pair has a point cloud on either end of its eigenaxis. The cloud pairs are not mirror images: Roughly one half of the points align in one orientation and the other half in the opposite. The time-series snippet of signs S_k plotted in Fig. 11 illustrates the variable nature of the signs. Also, while hard to show, the handedness of eigenvectors V_k is split nearly evenly between left and right. In contrast, Fig. 10 *right* shows that all $\mathcal{V}_k$ have a consistent pointing direction and handedness. By construction, the signs of the corrected eigenvectors are always +1.

Fig. 11 A time-series snippet of eigenvector signs S_k that corresponds to the S^2 point clouds in Fig. 10 left

The oriented eigenvectors in Fig. 10 right represent the most significant result of this section: The model-independent methodology in this paper orients raw eigenvectors V into $\mathcal{V}$ such that the latter is only a rotation away from I .

It is nonetheless apropos to conclude with an illustration of sequential regression in order to verify that the time series of regression weights evolves in an interpretable way akin to Fig. 2 bottom right. To do so, a dependent variable y is required, and a suitable target is the temperature deviation of Bermuda (NOAA station #4 in Fig. 8). For each step k in the evolution, and by using (41), two systems are solved¹⁰

$$y_k = (P_k V_k) \beta_k + \epsilon_k \quad \text{and} \quad y_k = (P_k \mathcal{V}_k) \beta'_k + \epsilon'_k.$$

Observe that a sign flip of an eigenvector in V_k flips the sign of the corresponding entry in β_k , whereas the same sign flip is absorbed by S_k in the oriented system, leaving $\mathcal{V}_k$ in with a consistently oriented basis.

Figure 12 shows the time series of β_k and β'_k entries. In the top plot, it is evident that the unoriented time series is replete with spurious sign flips, as originally indicated by Fig. 2 top right. In contrast, the bottom plot shows that the evolution of β'_k is clean.

12 Conclusions

Although eigenanalysis is an old and well-studied topic, linking eigenanalysis to an evolving dataset leads to unexpected results and new opportunities for analysis. The optimization of eigenanalysis codes for one-time solutions has admitted inconsistent eigenvector orientation because such inconsistency is irrelevant to the one-time solution. Yet for an evolving system, it is precisely the inconsistency that disrupts interpretability of an eigen-based model. This article reports a model-free method to correct for the inconsistency

¹⁰ No dimension reduction is used here. See the remarks associated with (28).

Fig. 12 Time series of unoriented β_k entries and oriented β'_k entries. *Top*, for any index k , the signs of the β_k entries assume arbitrary values, producing a time series that is roughly mirror imaged. *Bottom*, a time series of β'_k that is clean and ready for interpretation

and does so as a postprocessing step to existing eigenanalysis implementations. Once corrected, directional statistics brings a new avenue of inquiry to the behavior of the underlying factors in the dataset.

Acknowledgements Thank you to Professor A. Brace (National Australia Bank, University of Technology Sydney), Professor P. N. Kolm (New York University), Dr. A. Meucci (Advanced Risk and Portfolio Management), Professor H. R. Miller (Massachusetts Institute of Technology), and Dr. J. P. B. Müller (Harvard) for their helpful conversations about, and encouragement of, this publication. Thank you also to Professor L. Cao, Editor of this Journal, and my anonymous reviewers for their helpful suggestions.

Compliance with ethical standards

Conflict of interest I declare that I have no conflict of interest.

Code availability The Python package `thucyd` is available on PyPI and CONDA- FORGE at

<https://pypi.org/project/thucyd/>

<https://github.com/conda-forge/thucyd-feedstock>.

The code is covered by the APACHE 2.0 license and is freely available. In addition, the source code for the reference implementation is freely available at

<https://gitlab.com/thucyd-dev/thucyd>.

Data availability Section 11 was developed using Global Historical Climatology Network data collected by the United States National Oceanographic and Atmospheric Agency (see [17]). A subset of this data was consumed and processed. The code and processed data are freely available at

<https://gitlab.com/thucyd-dev/thucyd-eigen-working-examples>.

The project's `readme.md` files give links into the NOAA data.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material

is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alpaydin, E.: Introduction to Machine Learning, 2nd edn. MIT Press, Cambridge, MA (2010)
- Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.: LAPACK Users' Guide, 3rd edn. Society for Industrial and Applied Mathematics, Philadelphia (1999)
- BLAS Contributors: BLAS (Basic Linear Algebra Subprograms). Netlib.org. <http://www.netlib.org/blas/>
- Dash, J.W.: Quantitative Finance and Risk Management. World Scientific, River Edge, NJ (2004). See chapter 22, Correlation Matrix Formalism; the $\mathcal{N}$ -Sphere
- Golub, G.H., Loan, C.F.V.: Matrix Computations, 4th edn. Johns Hopkins University Press, Baltimore (2013)
- Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistics Learning. Springer, New York (2001)
- Hou, B.J., Zhang, L., Zhou, Z.H.: Learning with feature evolvable streams. In: Advances in Neural Information Processing Systems, pp. 1417–1427 (2017)
- IEEE 60559 Standards Contributors: atan2, atan2f, atan2l — arc tangent functions. In: The Open Group Base Specifications Issue, 7, 2018 edn. IEEE and The Open Group (2018). <http://pubs.opengroup.org/onlinepubs/9699919799/>. Accessed online, searched under atan2
- Jammalamadaka, S.R., SenGupta, A.: Topics in Circular Statistics. World Scientific, River Edge, NJ (2001)
- Kuhn, M., Johnson, K.: Applied Predictive Modeling. Springer, New York (2013)
- Kurz, G., Hanebeck, U.D.: Stochastic sampling of the hyperspherical von Mises–Fisher distribution without rejection methods. In: 2015 Sensor Data Fusion: Trends, Solutions, Applications (SDF), pp. 1–6. IEEE (2015). <https://doi.org/10.1109/SDF.2015.7347705>. <https://www.researchgate.net/publication/290937895>
- Leite, D.F., Costa, P., Gomide, F.: Evolving granular classification neural networks. In: 2009 International Joint Conference on Neural Networks, pp. 1736–1743 (2009)
- Ley, C., Verdebout, T.: Modern Directional Statistics. Chapman and Hall/CRC Press, New York (2017)
- Li, S.Z., Jain, A.K. (eds.): Handbook of Face Recognition, 2nd edn. Springer, New York (2011)
- Mardia, K.V., Jupp, P.E.: Directional Statistics. Wiley, New York (2000)
- Meucci, A.: Risk and Asset Allocation. Springer, New York (2007)
- National Centers for Environmental Information: Global Historical Climatology Network (GHCN). National Oceanic and Atmospheric Administration (NOAA). <https://www.ncdc.noaa.gov/data-access/land-based-station-data/land-based-datasets/global-historical-climatology-network-ghcn>
- Numpy Contributors: numpy.arctan2. Scipy.org (2019). <https://docs.scipy.org/doc/numpy/reference/generated/numpy.arctan2.html>
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes in C. The Art of Scientific Computing, 2nd edn. Cambridge University Press, New York (1992)
- Scholkopf, B., Smola, A.J.: Learning with Kernels. MIT Press, Cambridge, MA (2002)
- Sra, S.: Directional statistics in machine learning: A brief review. Applied Directional Statistics: Modern Methods and Case Studies p. 225 (2018)
- Sra, S., Karp, D.: The multivariate Watson distribution: maximum-likelihood estimation and other aspects. J. Multivar. Anal. **114**, 256–269 (2013)
- Strang, G.: Introduction to Applied Mathematics. Wellesley-Cambridge Press, Wellesley, MA (1986)
- Strang, G.: Linear Algebra and Its Applications, 3rd edn. Saunders College Publishing, New York (1988)
- TensorFlow Contributors: TensorFlow: Large-scale machine learning on heterogeneous systems. tensorflow.org (2015). <https://www.tensorflow.org/>
- TensorFlow Probability Contributors: Class VonMisesFisher. tensorflow.org. https://www.tensorflow.org/probability/api_docs/python/tfp/distributions/VonMisesFisher

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.