A Caveat Concerning Singular Value Decomposition

MATTHEW NEWMAN AND PRASHANT D. SARDESHMUKH

Cooperative Institute for Research in the Environmental Sciences, University of Colorado, Boulder, Colorado

(Manuscript received 28 September 1993, in final form 25 May 1994)

ABSTRACT

An assessment is made of the ability of the singular value decomposition (SVD) technique to recover the relationship between two variables x and y from a time series of their observations. It is shown that SVD is rigorously successful only in the special cases when either (i) the transformation linking x and y is orthogonal or (ii) the covariance matrix of either x or y is the identity matrix. The behavior of the method when these conditions are not met is also studied in a simple two-dimensional case.

That this caveat can be relevant in a meteorological context is demonstrated by performing an SVD analysis of a time series of global upper-tropospheric streamfunction and vorticity fields. Although these fields are linked by the two-dimensional Laplacian operator on the sphere, it is shown that the pairs of singular patterns resulting from the SVD analysis are not so related. The problem is apparent even for the first SVD pair and generally becomes worse for succeeding pairs. These results suggest that any physical interpretation of SVD pairs may be unjustified.

1. Introduction

Singular Value Decomposition (SVD) is one of several methods (see Bretherton et al. 1992, hereafter BSW) for the simultaneous statistical analysis of two quantities x and y that vary in both space and time. It involves performing the singular value decomposition of the temporal cross-covariance matrix of x and y as $\mathbf{C_{xy}} = \mathbf{U}\sigma\mathbf{V}^T$, hence its name. (We will follow standard practice by using SVD to refer to both the statistical technique and the matrix decomposition.) Here $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices whose columns are the left and right singular vectors, and σ is a diagonal matrix of nonnegative singular values. It is easily shown that the left and right vectors are eigenvectors of the symmetric matrices $\mathbf{C_{xy}}\mathbf{C_{xy}^T} = \mathbf{U}\sigma^2\mathbf{U}^T$ and $\mathbf{C_{xy}^T}\mathbf{C_{xy}} = \mathbf{V}\sigma^2\mathbf{V}^T$, respectively (Smith 1989). Each singular value is associated with a pair of left and right singular vectors. When x and y are projected on to such pairs, one obtains coefficient time series whose covariance is equal to the associated singular values. The total squared covariance of x and y, defined as the sum of the squares of the elements of $\mathbf{C_{xy}}$, is equal to the trace of $\mathbf{C_{xy}^T}\mathbf{C_{xy}}$ and therefore to the sum of the squared singular values. Each singular vector pair can therefore be thought of as accounting for a certain fraction of the total squared covariance. The pair associated with the largest singular value accounts for the largest fraction, the pair associated with the second largest value for the next largest fraction, and so on. SVD thus finds pairs of spatial patterns of x and y that are best correlated in time. It is one way of answering the question: What are the most frequently occurring pairs of patterns of x and y?

The widely used technique of principal component analysis (PCA; see Kutzbach 1967) of a single variable x can be thought of as a special case of SVD in which x and y are the same. Thus, one writes $\mathbf{C_{xx}} = \mathbf{E}\lambda\mathbf{E}^T$. The singular vectors, or the columns of $\mathbf{E}$, are now called empirical orthogonal functions (EOFs). Projecting x on to the EOFs yields coefficient time series (the principal components, or PCs) whose variance is equal to the associated eigenvalues λ. As in the general case above, each pair of singular vectors (that is, each EOF) accounts for a certain fraction of the total squared covariance of x. Additionally, and more simply, each EOF accounts for a certain fraction of the total variance of x, defined as the sum of the temporal variance over all points or the trace of $\mathbf{C_{xx}}$, with the EOF associated with the largest eigenvalue accounting for the largest fraction, and so on. PCA thus finds orthogonal spatial patterns whose standing oscillations account for the variance of x in decreasing order of importance. It is one way of answering the question: What are the most frequently occurring patterns of x?

PCA is powerful as a data reduction device, but precisely what it reveals about the underlying dynamical system is unclear. It is tempting to identify the EOFs with the free modes of the dynamical equation for x; however, the latter are often not orthogonal. Also, the free modes of such an equation are not relevant in cases where x is part of a larger system, comprising x

Corresponding author address: Dr. Matthew Newman, Cooperative Institute for Research in Environmental Sciences, University of Colorado at Boulder, Boulder, CO 80309-0449.

and say y: the free modes of this larger system can be very different. In any event, if free modes are what one is after, then one also has to consider time lag-correlation statistics, in addition to the zero-lag or simultaneous correlation statistics embodied in $\mathbf{C_{xx}}$. As discussed for example by Penland (1989) and Penland and Magorian (1993), the free modes of a linear system are most directly identified with the eigenvectors of the matrix $\mathbf{C_{xx}}(\tau)\mathbf{C_{xx}}(0)^{-1}$. If the system is a truly linear system with stationary statistics, then performing this analysis at different lags τ yields the same result. A similar analysis can of course also be performed for the extended system comprising x and y.

Like PCA, singular value decomposition is a useful exploratory device, but precisely what it reveals about the relationship between x and y in the underlying dynamical system is also unclear. The assumption is always made that the relationship between the two data fields being analyzed is fundamentally a *linear* one. (It should be stressed that "linear" does not mean "uncomplicated.") Otherwise, the search for coupled patterns using a cross-covariance matrix would have little meaning. Thus, setting aside nonlinear dynamics for a moment, one can think of three possibilities: (I) y is diagnostically related to x as $\mathbf{y}(t) = \mathbf{Lx}(t)$, where $\mathbf{L}$ is a linear operator, (II) y and x are parts of a larger system, or (III) y represents an external forcing for x; that is, the dynamical equation for x involves y but not vice versa. Meteorological examples can readily be found of all three cases.

The question is can SVD distinguish between these three possibilities? For example, what does an SVD analysis of 500-mb geopotential height and SST (e.g., Wallace et al. 1992) tell one about the relationship between these quantities? Although the two fields may be related, other factors such as orographic waves and high-frequency baroclinic eddies could also be contributing to the variability of either one or both of these fields. In that case, the answer to our question would be "no," since there is no direct information about other variables in the cross-covariance matrix.

We consider case I in some detail in this paper. We show in section 2 that even in this apparently simple case, successful recovery of the relationship between x and y is possible only under some rather severe restrictions upon either the operator $\mathbf{L}$ and/or the data, the most interesting of which occur when either $\mathbf{L}$ is orthogonal or $\mathbf{C_{xx}}$ or $\mathbf{C_{yy}}$ are identity matrices. The interpretation of SVD products in cases II and III is presumably even more problematic.

The consideration of similar objections to two other commonly used methods is briefly discussed in section 3. In section 4, the behavior of the method when the above conditions are not met is studied in a two-dimensional setting. A meteorological example is considered in section 5. We perform an SVD analysis of an observed time series of streamfunction ψ and vorticity ζ fields of the horizontal flow at an upper-tropospheric level, which are linked through the two-dimensional spherical Laplacian operator as $\zeta = \nabla^2\psi$, and ask to what extent the singular vector pairs of streamfunction and vorticity obey this relationship. Some concluding remarks, particularly concerning the physical interpretation of patterns resulting from the general application of SVD in cases II and III, given the failure of SVD in case I, are made in section 6.

2. A caveat concerning SVD

Let $\mathbf{x}(j, n)$ denote the data matrix of x with $j = 1, 2, \cdots, J$ as the space index and $n = 1, 2, 3, \cdots, N$ as the time index. The symmetric covariance matrix $\mathbf{C_{xx}}$ is written as

$$\mathbf{C_{xx}} \equiv \frac{\mathbf{xx}^{\mathrm{T}}}{N} \equiv \langle \mathbf{xx}^{\mathrm{T}} \rangle = \mathbf{E}\lambda_{\mathbf{x}}\mathbf{E}^{-1}. \qquad (2.1)$$

Here $\mathbf{E}$ is an orthogonal ($\mathbf{E}^{-1} = \mathbf{E}^{\mathrm{T}}$) matrix of the eigenvectors of $\mathbf{C_{xx}}$ and $\lambda_{\mathbf{x}}$ is a diagonal matrix of eigenvalues. The columns of $\mathbf{E}$, the EOFs, form an orthonormal set $\mathbf{e}_i^{\mathrm{T}}\mathbf{e}_j = \delta_{ij}$. For the data matrix $\mathbf{y}(j, n)$ of y one can similarly write

$$\mathbf{C_{yy}} \equiv \frac{\mathbf{yy}^{\mathrm{T}}}{N} \equiv \langle \mathbf{yy}^{\mathrm{T}} \rangle = \mathbf{F}\lambda_{\mathbf{y}}\mathbf{F}^{-1}. \qquad (2.2)$$

A fact about PCA that is often unappreciated is that if y and x are linked by a linear transformation $\mathbf{y} = \mathbf{Lx}$, the EOFs of y and x are not necessarily so linked. Thus, in general $\mathbf{F} \neq \mathbf{LE}$. For if $\mathbf{y} = \mathbf{Lx}$, then

$$\mathbf{C_{yy}} = \langle \mathbf{yy}^{\mathrm{T}} \rangle = \mathbf{L}\langle \mathbf{xx}^{\mathrm{T}} \rangle \mathbf{L}^{\mathrm{T}} = \mathbf{LE}\lambda_{\mathbf{x}}\mathbf{E}^{-1}\mathbf{L}^{\mathrm{T}}. \qquad (2.3)$$

Comparing (2.2) and (2.3) gives $\mathbf{F} = \mathbf{LE}$ if $\mathbf{L}$ is an orthogonal matrix, that is, $\mathbf{L}^{-1} = \mathbf{L}^{\mathrm{T}}$, in which case $\lambda_{\mathbf{y}} = \lambda_{\mathbf{x}}$.

As described in section 1, SVD is concerned with the cross-covariance matrix $\mathbf{C_{xy}}$ of x and y. Thus, one writes

$$\mathbf{C_{xy}} \equiv \frac{\mathbf{xy}^{\mathrm{T}}}{N} \equiv \langle \mathbf{xy}^{\mathrm{T}} \rangle = \mathbf{U}\sigma\mathbf{V}^{\mathrm{T}}, \qquad (2.4)$$

where σ is a diagonal matrix of singular values, and $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices whose columns are the orthonormal left and right singular vectors. Thus,

$$\mathbf{u}_i^{\mathrm{T}}\mathbf{u}_j = \delta_{ij} \quad \text{and} \quad \mathbf{v}_i^{\mathrm{T}}\mathbf{v}_j = \delta_{ij}. \qquad (2.5)$$

Each singular vector pair $\mathbf{u}_i$, $\mathbf{v}_i$ is associated with a singular value σ_i.

Suppose again that $\mathbf{y} = \mathbf{Lx}$. We will now show that SVD is rigorously successful in giving $\mathbf{v}_i = \mathbf{Lu}_i$ only under certain limits upon the $\mathbf{L}$ operator and/or upon the data itself.

a. *Restriction upon* $\mathbf{L}$

To obtain $\mathbf{V} = \mathbf{LU}$, one must have

$$\mathbf{v}_i^{\mathrm{T}}\mathbf{v}_j = (\mathbf{Lu}_i)^{\mathrm{T}}(\mathbf{Lu}_j) = \mathbf{u}_i^{\mathrm{T}}\mathbf{L}^{\mathrm{T}}\mathbf{Lu}_j = \delta_{ij}, \qquad (2.6)$$

that is, $\mathbf{L}$ must be orthogonal. In this case, it turns out that SVD is redundant, since by comparing Eqs. (2.1), (2.2), and (2.4) we see that

$$\mathbf{U}\sigma\mathbf{U}^{\mathrm{T}} = \mathbf{E}\lambda_{\mathbf{x}}\mathbf{E}^{\mathrm{T}} \qquad (2.7a)$$

and

$$\mathbf{LU}\sigma(\mathbf{LU})^{\mathrm{T}} = \mathbf{F}\lambda_{\mathbf{y}}\mathbf{F}^{\mathrm{T}}. \qquad (2.7b)$$

In other words, the singular vectors of x are the same as the EOFs of x, and the singular vectors of y are the same as the EOFs of y; the eigenvalues of $\mathbf{C_{xx}}$ and $\mathbf{C_{yy}}$ are the same, and are equal to the singular values. Thus, the SVD process merely gives us the EOFs of the two fields involved.

The condition that $\mathbf{L}^{-1} = \mathbf{L}^{\mathrm{T}}$ is rather restrictive, and it seems unlikely that in general two meteorological fields of interest would be related in such a manner. For example, if x represents streamfunction and y represents relative vorticity, then in spectral space $\mathbf{L}$ is a diagonal matrix whose elements are as $n(n+1)$, where n is the spherical wavenumber. The example used by BSW represents a rather special case, in which SVD is performed upon two time series that each contain a different random component and a second term that is the same apart from the sign. In this case, $\mathbf{L}$ is essentially $-\mathbf{I}$, and as such Eq. (2.6) is satisfied.

b. Restrictions upon the data

If $\mathbf{y} = \mathbf{Lx}$, then Eq. (2.4) can be written

$$\mathbf{C_{xy}} = \langle \mathbf{xy}^{\mathrm{T}} \rangle = \langle \mathbf{xx}^{\mathrm{T}} \rangle \mathbf{L}^{\mathrm{T}} = \mathbf{C_{xx}}\mathbf{L}^{\mathrm{T}} \qquad (2.8)$$

and also

$$\mathbf{C_{xy}} = \mathbf{U}\sigma\mathbf{V}^{\mathrm{T}} = \mathbf{U}(\mathbf{V}\sigma)^{\mathrm{T}}, \qquad (2.9)$$

giving

$$\mathbf{V}\sigma = \mathbf{LC_{xx}U}. \qquad (2.10)$$

(We could similarly have obtained $\mathbf{U}\sigma = \mathbf{L}^{-1}\mathbf{C_{yy}V}$; the following discussion applies to this equation as well.) It is clear that if $\mathbf{C_{xx}} = \mathbf{I}$, then $\mathbf{V}\sigma = \mathbf{LU}$. Note that in this case both $\mathbf{C_{yy}}$ and $\mathbf{C}_{\mathbf{xy}}^{\mathrm{T}}\mathbf{C_{xy}}$ are equal to $\mathbf{LL}^{\mathrm{T}}$, so they have the same eigenvectors; that is, the right singular vectors are also the EOFs of y. Interestingly, this does not require the SVs of x to be the same as the EOFs of x (which are $e_{ij} = \delta_{ij}$ for the ith eigenvector), since we are not assuming here, as in Eq. (2.7a), that $\mathbf{LL}^{\mathrm{T}} = \mathbf{I}$. Also, since $\mathbf{C_{xx}} = \mathbf{I}$ implies that the value of x at any location is uncorrelated with the value of x at any other location, this condition is just as restrictive as that upon $\mathbf{L}$ obtained above. On the other hand, this could explain a curious result in Wallace et al. (1992), where it was noted that the first 2 SVs of SST were highly correlated with the first 2 EOFs of SST, although the same was not true of the 500-mb height EOFs and SVs. Given the limited domain used for SST, and the relatively large scale of those patterns as opposed to that of the height patterns, it is possible that the height covariance matrix is more diagonally dominant than the SST covariance matrix.

Finally, recall that we can always express any singular vector $\mathbf{u}$ in the x EOF space; that is, $\mathbf{U} = \mathbf{Es}$, where $\mathbf{s}$ is the projection coefficient matrix. Then, from Eqs. (2.1) and (2.10), we have

$$\mathbf{V}\sigma = \mathbf{LE}\lambda_{\mathbf{x}}\mathbf{s}. \qquad (2.11)$$

Thus, the effect of the interposition of $\mathbf{C_{xx}}$ in Eq. (2.10) is to amplify the contribution to $\mathbf{u}_i$ of each x EOF in proportion to its eigenvalue. For example, suppose the first left SV of an SVD analysis is $\mathbf{u}_1 = \mathbf{E}_1 + \mathbf{E}_2$, where $\mathbf{E}_1$ and $\mathbf{E}_2$ are the leading x EOFs, and $\lambda_1 = 10$ and $\lambda_2 = 5$. Then the corresponding right SV $\mathbf{v}_1 \propto \mathbf{L}(2\mathbf{E}_1 + \mathbf{E}_2)$. Clearly, the degree to which an SV pair maintains the correct linear relationship will depend upon the way in which either SV projects upon its corresponding EOF space, and the distribution of variance within that space.

Inspection of Eq. (2.10) shows that if $\mathbf{U} = \mathbf{E}$, Eq. (2.10) reduces to $\mathbf{V}\sigma = \mathbf{LU}\lambda_{\mathbf{x}}$, so that again $\mathbf{v}_i$ is proportional to $\mathbf{Lu}_i$. Thus, if the singular vectors of one of the datasets are equivalent to the EOFs of that dataset, the linear relationship will hold. It is hard to imagine this happening, however, unless either $\mathbf{C_{xx}}$ is the identity matrix or else $\mathbf{L}$ is orthogonal.

3. Combined PCA and canonical correlation analysis

There are two other commonly used methods of analysis that do not necessarily suffer the above restrictions. For example, the two data fields can be combined into the same PCA (Kutzbach 1967); thus

$$\mathbf{X} = \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix};$$

PCA of the $\mathbf{C_{XX}}$ covariance matrix (also known as combined PCA, or CPCA) is

$$\langle \mathbf{X}\mathbf{X}^{\mathrm{T}} \rangle \equiv \mathbf{C}\lambda_c\mathbf{C}^{\mathrm{T}}, \qquad (3.1)$$

where

$$\mathbf{C} = \begin{pmatrix} \mathbf{C_x} \\ \mathbf{C_y} \end{pmatrix}. \qquad (3.2)$$

Here $\mathbf{C_x}$ is the submatrix containing the portion of the EOF matrix $\mathbf{C}$ that represents x, and $\mathbf{C_y}$ is the submatrix containing the portion that represents y. Note that although the eigenvectors within $\mathbf{C}$ must be orthogonal between each other, there is no requirement for either submatrix to contain an orthogonal set. Although this does not prohibit the correct linear relationship between x and y, it does not mandate it either. Other problems with this method are discussed by BSW. Interestingly, the SVD of the combined data matrix $\mathbf{X}$ itself will produce the same results as CPCA and can in some instances be computationally more efficient

(Chambers 1977), having the added advantage of producing both the spatial patterns and associated time series in a single step. It is important to note that it is not, however, the same as the SVD of the cross-covariance matrix $\mathbf{C_{xy}}$.

A better choice may be to transform $\mathbf{x}$ and $\mathbf{y}$ into their respective EOF spaces: $\mathbf{x} = \mathbf{EA}$ and $\mathbf{y} = \mathbf{FB}$. Since the EOF time series $\mathbf{A}_i$ and $\mathbf{B}_i$ are orthogonal but not orthonormal, one then scales the matrices containing the EOFs and their respective time series so that

$$\mathbf{x} = (\mathbf{E}_1\sqrt{\lambda_{x_1}}\mathbf{E}_2\sqrt{\lambda_{x_2}}\cdots)\begin{pmatrix}\mathbf{A}_1/\sqrt{\lambda_{x_1}}\\ \mathbf{A}_2/\sqrt{\lambda_{x_2}}\\ \vdots\end{pmatrix} \equiv \mathbf{e}\boldsymbol{\alpha} \quad (3.3)$$

(and similarly for $\mathbf{y}$ so that $\mathbf{y} \equiv \mathbf{f}\boldsymbol{\beta}$). Both $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are now orthogonal matrices, and we can rewrite the relationship between x and y as $\boldsymbol{\beta} = \mathbf{M}\boldsymbol{\alpha}$, where $\mathbf{M} = \mathbf{f}^{-1}\mathbf{Le}$. Then if we compute the covariance matrix of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$,

$$\mathbf{C}_{\alpha\beta} = \langle \boldsymbol{\alpha}\boldsymbol{\beta}^{\mathrm{T}} \rangle = \mathbf{W}\sigma\mathbf{Z}^{\mathrm{T}},$$

we have, in analogy with (2.10), $\mathbf{Z}\sigma = \mathbf{MW}$ (since $\langle \boldsymbol{\alpha}\boldsymbol{\alpha}^{\mathrm{T}} \rangle = \mathbf{I}$). If we define the transform back into real space as $\mathbf{U} = \mathbf{eW}$ and $\mathbf{V} = \mathbf{fZ}$, then $\mathbf{V}\sigma = \mathbf{LU}$.

This is the well-known method of canonical correlation analysis, or CCA (BSW; Barnett and Preisendorfer 1987). Unlike SVD, it will always maintain the correct linear relationship. However, it has its own well-known drawback: since all EOF time series have equal weight (in α and β), the leading pairs of patterns will not necessarily explain significant amounts of the squared covariance. Barnett and Preisendorfer ameliorate this problem by truncating $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ to the time series of a specified group of leading EOFs; however, this introduces an element of subjectivity into the analysis process. Again, see BSW for a more detailed discussion of this method.

It should be noted that this truncation can also interfere with the ability of CCA to recover the correct linear relationship. If we call $\hat{\boldsymbol{\alpha}}$ and $\hat{\boldsymbol{\beta}}$ the truncated time series matrices of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, respectively, then the relationship $\mathbf{V}\sigma = \mathbf{LU}$ can only be exactly recovered if $\hat{\boldsymbol{\beta}} = \mathbf{M}\hat{\boldsymbol{\alpha}}$. In other words, $\boldsymbol{\beta}$ must be truncated with respect to $\hat{\boldsymbol{\alpha}}$ and not with respect to itself (or vice versa). (Of course, if $\mathbf{L}$ is orthogonal so that $\mathbf{f} = \mathbf{Le}$, then $\hat{\boldsymbol{\beta}} = \mathbf{M}\hat{\boldsymbol{\alpha}}$ will be satisfied identically.) Presumably, the sensitivity of the "apparent" $\mathbf{L}$ to truncation could be tested.

4. A simple example of SVD

As a simple demonstration of the potential pitfalls of the SVD method, consider the following example. Let $\mathbf{y} = \mathbf{Lx}$, where

$$\mathbf{L} = \begin{pmatrix} a & 0 \\ 0 & 1 \end{pmatrix} \quad (4.1)$$

and specify the covariance matrix $\mathbf{C_{xx}}$ of x as

$$\mathbf{C_{xx}} = \langle \mathbf{xx}^{\mathrm{T}} \rangle \equiv \begin{pmatrix} 1 & \epsilon \\ \epsilon & 1 \end{pmatrix} \quad (4.2)$$

so that

$$\mathbf{C_{xy}} = \mathbf{C_{xx}}\mathbf{L}^{\mathrm{T}} = \mathbf{U}\sigma\mathbf{V}^{\mathrm{T}}. \quad (4.3)$$

If we then define

$$\tilde{\mathbf{V}} = \left(\frac{\mathbf{Lu}_1}{|\mathbf{Lu}_1|}\ \frac{\mathbf{Lu}_2}{|\mathbf{Lu}_2|} \right),$$

we can correlate $\mathbf{V}$ and $\tilde{\mathbf{V}}$ with the measure

$$R(a, \epsilon) \equiv \det(\tilde{\mathbf{V}}^{\mathrm{T}}\mathbf{V}). \quad (4.4)$$

If $\mathbf{V}_i^{\mathrm{T}}\tilde{\mathbf{V}}_j = \delta_{ij}$ for both vectors, then $R = 1$. Results are displayed in Fig. 1a. Note the maxima located along $a = 1$, for which $\mathbf{L}$ is the identity matrix, and $a = -1$, for which $\mathbf{L}$ is orthogonal. Also, $R = 1$ on the $\epsilon = 0$ axis. In this case $\mathbf{L}$, $\mathbf{C_{xx}}$, and $\mathbf{C_{xy}}$ are all diagonal; hence

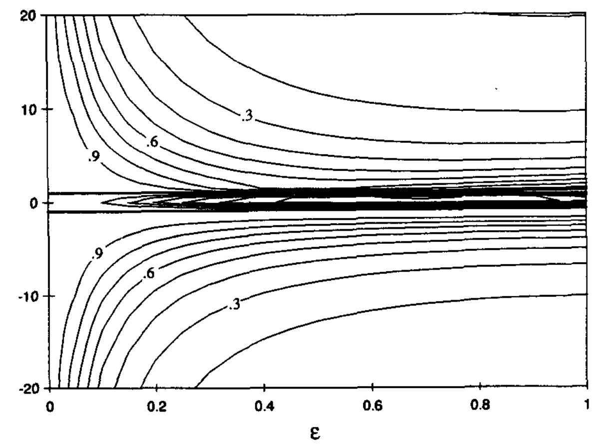

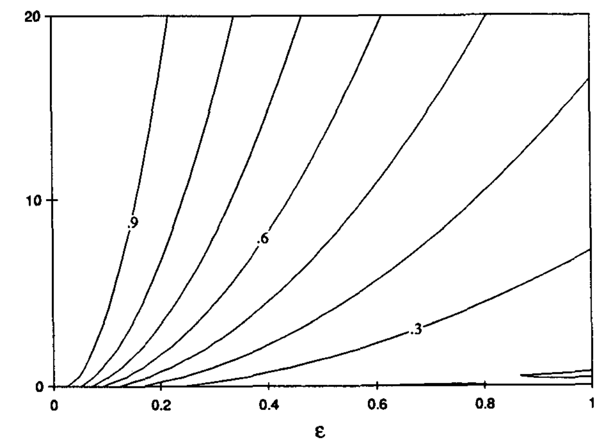

FIG. 1. (a) Contours of R [Eq. (4.4)] as a function of a and ϵ for the first example described in the text. The two solid lines represent $a = +1$ and $a = -1$. (b) Contours of R [Eq. (4.4)] as a function of b and ϵ for the second example described in the text.

$\mathbf{U}$ and $\mathbf{V}$ are also diagonal and equal to each other. In general, as $|\epsilon|$ increases, the correlation between $\mathbf{LU}$ and $\mathbf{V}$ decreases; this effect is more dramatic for large $|a|$ as $\mathbf{L}$ becomes less orthogonal. These results are not dependent upon the dimension of the matrices used here: we get the same results by taking the two EOFs and associated time series of $\mathbf{C_{xx}}$, evaluating them at eight points instead of two, and then repeating the above procedures with 8×8 matrices.

Most cases of interest will have a more complicated $\mathbf{L}$ operator, of course, so that in general we do not know the potential for success (as measured by the closeness of R to 1). Also important is that this result shows that even for a "nice" $\mathbf{L}$ operator, the success of the results will depend upon the covariance structure of the data because of the interposition of $\mathbf{C_{xx}}$ between $\mathbf{L}$ and $\mathbf{U}$ in Eq. (2.10). As a second illustration of this latter point, we repeat this example, but this time fix $a = 10$ in Eq. (4.1) and rewrite $\mathbf{C_{xx}}$ as

$$\mathbf{C_{xx}} = \langle \mathbf{xx}^{\mathrm{T}} \rangle = \begin{pmatrix} b & \epsilon\sqrt{b} \\ \epsilon\sqrt{b} & 1 \end{pmatrix}. \quad (4.5)$$

Here $R(b, \epsilon)$ is defined the same as Eq. (4.4) and is contoured in Fig. 1b. As in Fig. 1a, R decreases as ϵ increases from zero.

One might think that the percentage of variance explained by the first EOF (say, ν_1) would have a bearing upon the results shown in Fig. 1b, but this does not appear to be the case. For example, traveling along any R contour from bottom left to top right is concurrent with a monotonic increase in both b and ϵ, and therefore with a monotonic increase in ν_1. And yet, the fact that one is positioned on a constant R contour implies that the results are just as unsatisfactory. Also, merely increasing the number of degrees of freedom in this example does not qualitatively change these results.

5. SVD performed upon two linearly related datasets

We next demonstrate these results by performing a series of analyses upon data fields derived from 21 December–February (DJF) monthly mean fields of 300-mb wind for the period between December 1984 and February 1991, and a second set derived from 21 June–August (JJA) monthly mean fields of 300-mb wind for the period between June 1985 and August 1991. Relative vorticity, ζ, is calculated from the winds and converted to a T21 spectral space. Streamfunction, ψ, is determined as $\nabla^{-2}\zeta$; both fields are then converted back to a Gaussian grid, which is repeated at 5.625° longitude intervals. Consistent with standard practice, the ψ and ζ fields are first weighted by $\sqrt{\cos\varphi}$, where φ is latitude, before performing the EOF yand SVD analyses. Given that $\zeta = \nabla^2\psi$, the $\mathbf{L}$ operator of section 2 essentially corresponds to the ∇^2 operator in the sphere; however, the linear relationship between the weighted variables is more precisely written as

$$\zeta_g = [\mathbf{S}^{-1}\mathbf{Y}\nabla_s^2\mathbf{Y}^{-1}\mathbf{S}]\psi_g,$$

where the subscript g means that the fields are in grid space, the subscript s refers to the fact that the ∇^2 operator is in spectral space, $\mathbf{Y}$ is the matrix containing the spherical harmonics vectors, and $\mathbf{S}$ represents the $\sqrt{\cos\varphi}$ weighting. The term within the brackets is the linear operator we are concerned with; henceforth, we shall just refer to it as the ∇^2 operator.

Four sets of dominant streamfunction patterns are produced:

1) ψ_{EOF}: the eigenfunctions of the ψ covariance matrix $\mathbf{C}_{\psi\psi}$
2) $\psi_{\zeta\mathrm{EOF}}$: the inverse Laplacian (∇^{-2}) of the eigenfunctions of the ζ covariance matrix $\mathbf{C}_{\zeta\zeta}$
3) ψ_{sv}: the left singular vectors obtained from the SVD analysis of the cross-covariance matrix $\mathbf{C}_{\psi\zeta}$
4) $\psi_{\zeta\mathrm{sv}}$: the inverse Laplacian (∇^{-2}) of the right singular vectors of $\mathbf{C}_{\psi\zeta}$.

Figure 2 shows the global pattern correlations of ψ_{sv} and $\psi_{\zeta\mathrm{sv}}$ for the first 16 singular vector pairs. Results are shown for both DJF and JJA cases. Clearly, the correlations are not 1; thus, SVD analysis is unable to recover the relationship between the ψ and ζ fields. The fact that the results differ for DJF and JJA further confirms that this failure is associated not only with the nonorthogonality of the ∇^2 operator, but also with the different structure of $\mathbf{C}_{\psi\psi}$ in the DJF and JJA seasons [see Eq. (2.10)].

Figure 3 contains a set of panels derived from the first ψ and ζ EOF and SV pairs. Shown are ψ_{EOF}, $\psi_{\zeta\mathrm{EOF}}$,

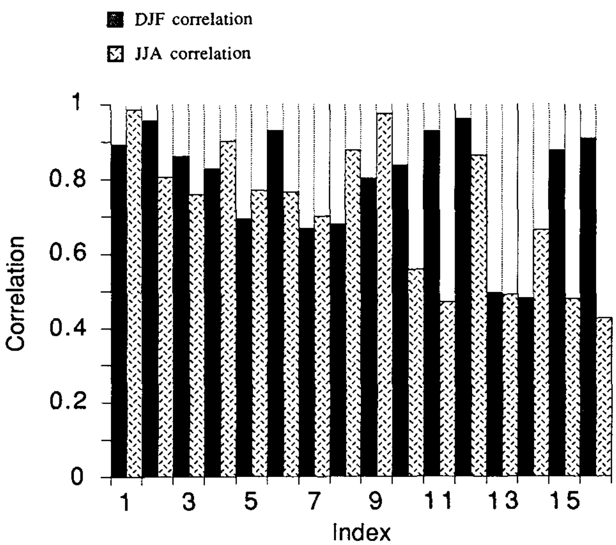

FIG. 2. Histogram of the correlation between the streamfunction SV (ψ_{SV}) and the inverted vorticity SV ($\psi_{\zeta\mathrm{SV}}$) for each SV pair. Correlations for SVD analysis of wintertime (DJF) data are heavy crosshatch; correlations for SVD analysis of summertime (JJA) data are light crosshatch. The DJF and JJA data are based upon NMC and ECMWF analyses, respectively.

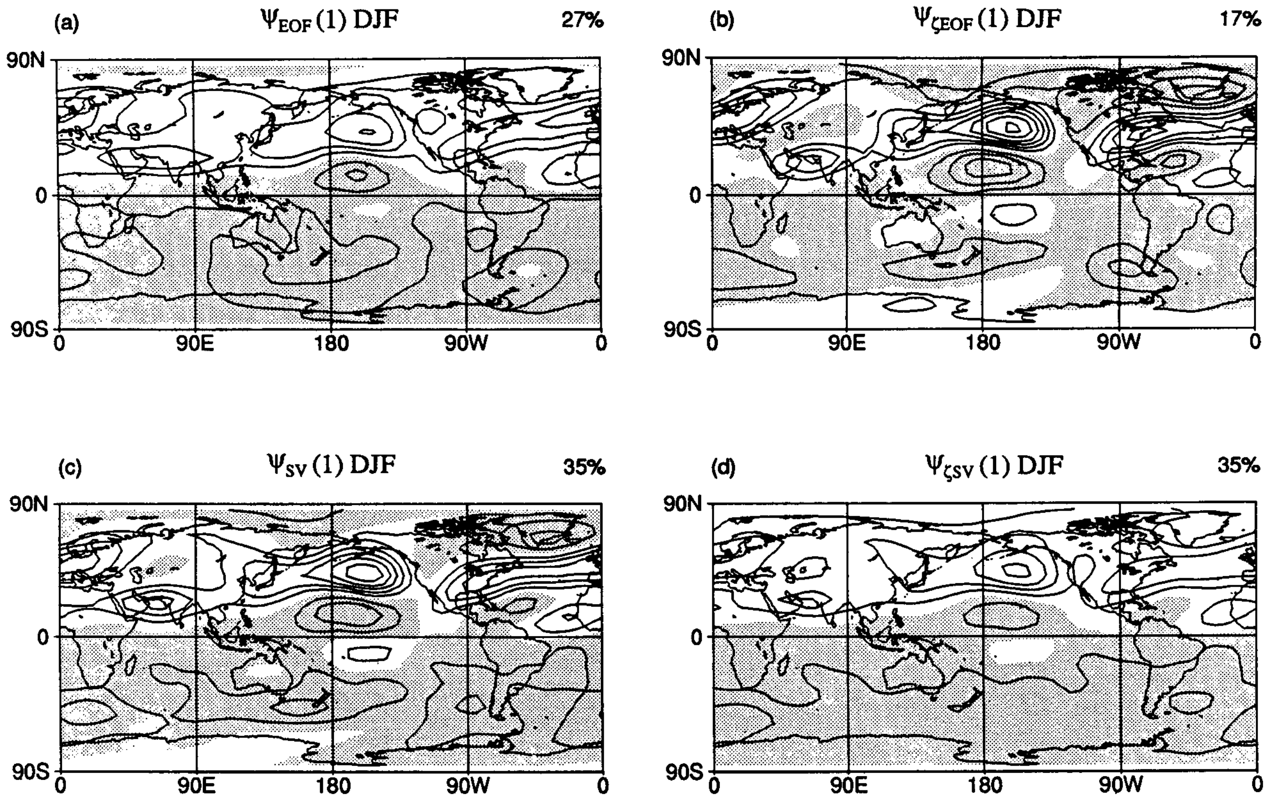

FIG. 3. Streamfunction patterns of EOF and SVD analysis of the DJF data. Negative values are shaded; the zero contour is removed for clarity. (a) Leading streamfunction EOF. Streamfunction variance explained by this pattern is in the upper right-hand corner. (b) Streamfunction calculated from the leading vorticity EOF. Vorticity variance explained by this pattern is in the upper right-hand corner. (c) First streamfunction SV. Squared covariance explained by this SV pair is in the upper right-hand corner. (d) Streamfunction calculated from the first vorticity SV. Squared covariance explained by this SV pair is in the upper right-hand corner.

ψ_{SV}, and $\psi_{\zeta SV}$. The strong zonally symmetric component evident in the Northern Hemisphere in ψ_{EOF} (Fig. 3a) is largely absent in $\psi_{\zeta EOF}$ (Fig. 3b), and the patterns in the Southern Hemisphere are dramatically different. These EOFS also explain differing portions of the variance in their respective fields. Clearly, the dominant ψ EOF is not the same as the ψ of the dominant ζ EOF. Likewise, one cannot conclude that the dominant wind EOFs will blow parallel to the contours of the dominant ψ EOF. As discussed in the beginning of section 2, one strictly cannot examine an EOF pattern of one field to draw conclusions about the EOF patterns of a linearly related field.

The ψ_{SV} and $\psi_{\zeta SV}$ fields are slightly better correlated than the ψ_{EOF} and $\psi_{\zeta EOF}$ fields. The difference between the two patterns is largely due to the maxima having different amplitudes, such as the wave train across the central Atlantic, although in a few regions such as Greenland there is a discrepancy of sign as well.

Less successful matches between the ψ_{SV} and $\psi_{\zeta SV}$ pairs are shown in Fig. 4, derived from the third and fifth SV pairs. Here the differences between the relative strength of centers are much more pronounced than for the first SV pair, and there are again discrepancies of sign. Also, a few of the features present in ψ_{SV} are absent in $\psi_{\zeta SV}$, such as the broad negative center over southern Asia in Fig. 4c.

The first two ψ_{SV} s produced by SVD for the summer case, and the corresponding $\psi_{\zeta SV}$ s, are displayed in Fig. 5. The first (ψ_{SV}, $\psi_{\zeta SV}$) pair is better correlated than any other pair resulting from either dataset, and differences between these patterns are generally insignificant (although not over Japan). The recovery of the correct linear relationship between the second pair, however, is much poorer. In particular, the PNA-like wavetrain in Fig. 5c is barely evident in Fig. 5d.

In all of these figures, the pattern of the inverted vorticity SV, $\psi_{\zeta SV}$, is equivalent to $\mathbf{C}_{\psi\psi}\psi_{SV}$ [see Eq. (2.10)]. Recall that the $\mathbf{C}_{\psi\psi}$ operator acts to amplify the contribution of the leading ψ EOFs to the $\psi_{\zeta SV}$ pattern relative to ψ_{SV}. In particular, it is clear that the primary difference between all the ψ_{SV} and $\psi_{\zeta SV}$ pairs is due to the strengthening of the leading streamfunction EOF (Fig. 3a) component of ψ_{SV}. Thus, the zonally symmetric component is relatively enhanced in all $\psi_{\zeta SV}$ panels of Figs. 3 and 4.

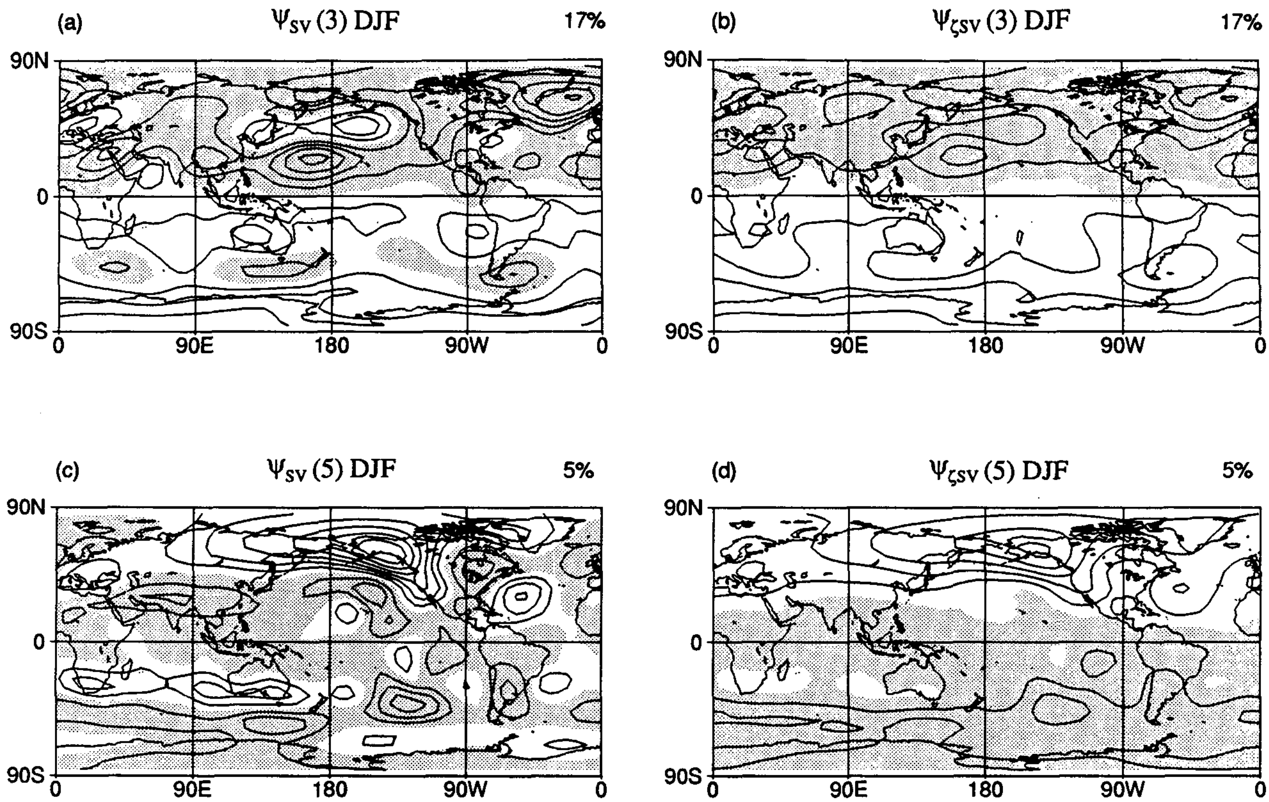

FIG. 4. (a) and (b) Same as Figs. 3c and 3d but for the third SV pair. (c) and (d) Same as Figs. 3c and 3d but for the fifth SV pair.

6. Discussion and concluding remarks

It has been shown that SVD cannot recover the correct linear relationship between two data fields except in some rather special cases. In our example with data, SVD analysis of streamfunction and vorticity produced pairs of patterns that were not related by the correct ∇^2 operator. To a casual user, the discrepancy between Figs. 3c and 3d may perhaps not appear too disturbing. It should be recognized, however, that even apparently minor differences in the relative strength of extrema on such plots of SV pairs could make a crucial difference in interpretation. If such a degree of uncertainty is acceptable, then it is unclear as to what one learns from an SVD analysis that one does not from simply correlating the PCs of one variable with the time series of the other. It should be emphasized too that the success of SVD in recovering the relationship between two variables depends not only on the relationship but also on the covariance structure of the data. Thus, two separate analyses with a given set of variables (say, precipitation and SST) could be relatively successful one time but unsuccessful the next.

Given that SVD is unable to recover a simple diagnostic relationship between y and x, one may ask to what extent it can diagnose the more complicated relationships described by cases II and III in section 1. Suppose that not all of the variability of y is related to that of x. One may write, quite generally, $\mathbf{y}(t) = \mathbf{L}\mathbf{x}(t) + \mathbf{z}(t)$, where $\mathbf{z}(t)$ represents contributions from all other system variables and external forcing. Note that the SVD analysis is carried out only upon $\mathbf{C}_{\mathbf{xy}}$; however, one now has $\mathbf{C}_{\mathbf{xy}} = \langle \mathbf{x}\mathbf{y}^{\mathrm{T}} \rangle = \langle \mathbf{x}\mathbf{x}^{\mathrm{T}} \rangle \mathbf{L}^{\mathrm{T}} + \langle \mathbf{x}\mathbf{z}^{\mathrm{T}} \rangle = \mathbf{C}_{\mathbf{xx}}\mathbf{L}^{\mathrm{T}} + \mathbf{C}_{\mathbf{xz}}$. If $\mathbf{C}_{\mathbf{xz}}$ is not null, then $\mathbf{L}$ cannot be determined from a knowledge of $\mathbf{C}_{xy}$ alone. If it is null, the problem reduces to that analyzed in section 2.

Case III of section 1 deserves further comment. To what extent can members of SV pairs, in say, Wallace et al. 1992 or Hsu 1994, be interpreted as representing dominant patterns of forcing and response? Let us say that y represents a forcing of x. One writes $d\mathbf{x}/dt = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y}$. An analysis that treats $d\mathbf{x}/dt$ as a separate variable will suffer the same drawbacks discussed in the previous paragraph. On timescales long enough that $d\mathbf{x}/dt \sim 0$, this equation can be cast into the generic form $\mathbf{y} = \mathbf{L}\mathbf{x}$ considered in this paper. Then, to the extent that the SVD analysis of x and y does not recover the correct operator $\mathbf{L}$, the singular vectors of x cannot be interpreted as a steady response to the singular vectors of y. In fact, the SVD analysis alone cannot determine whether y forces x or vice versa.

If SVD of the cross-covariance matrix cannot reveal the true nature of the coupling between two variables

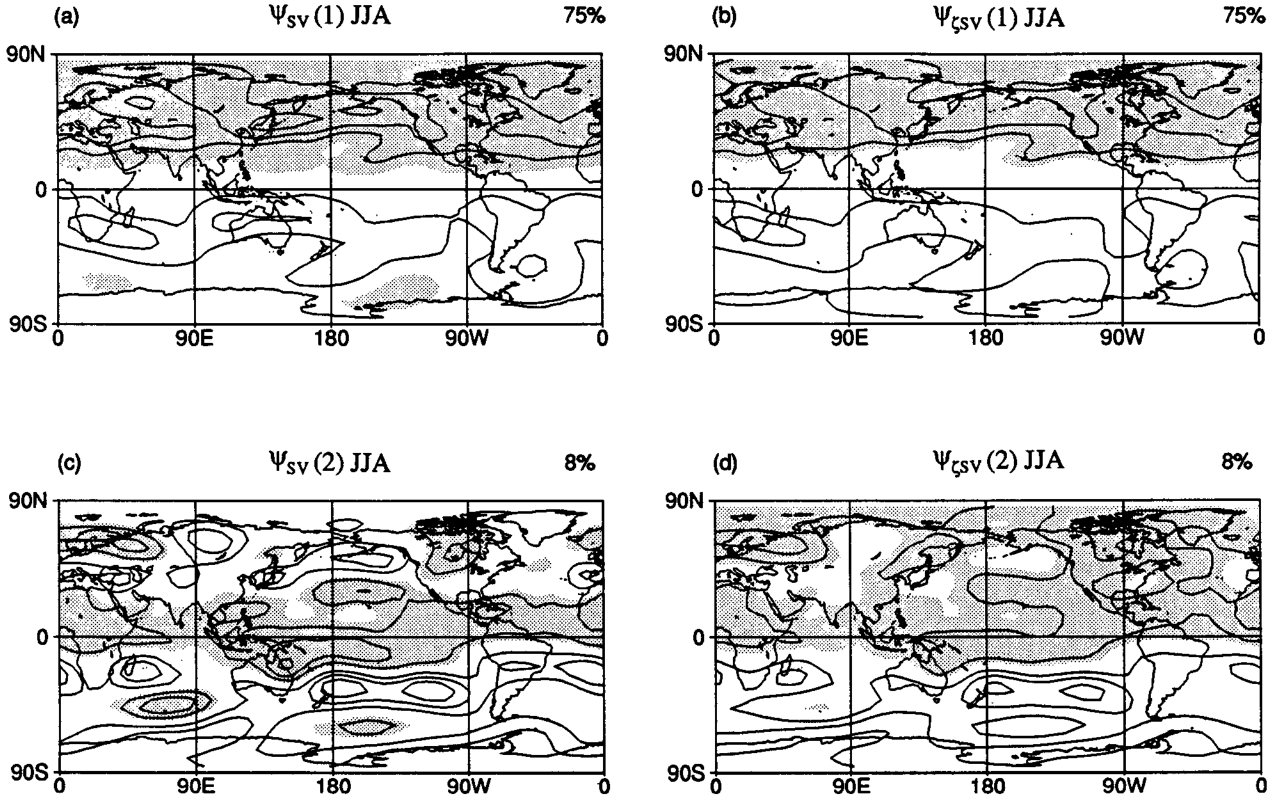

FIG. 5. (a) and (b) Same as Figs. 3c and 3d but for the JJA data. (c) and (d) Same as (a) and (b) but for the second SV pair.

x and y, one may ask precisely what it does reveal. It was stated in section 1 that it provides one means of isolating the most frequently occurring pairs of patterns of x and y. One could perhaps claim that members of an SV pair, say $\mathbf{x}_1$ and $\mathbf{y}_1$, occur "simultaneously" or "go together" in some sense. Even this interpretation needs to be treated with caution. One can think of three possibilities in this context. (i) The coefficient time series of $\mathbf{x}_1$ and $\mathbf{y}_1$ are not highly correlated. The notion of the simultaneous occurrence of $\mathbf{x}_1$ and $\mathbf{y}_1$ is then clearly dubious. (ii) The coefficient time series are highly correlated but $\mathbf{x}_1$ and $\mathbf{y}_1$ do not project strongly on the dominant EOFs of x and y. In this case $\mathbf{x}_1$ and $\mathbf{y}_1$ might occur simultaneously, but they are not important because the total covariability of x and y is small. (iii) The coefficient time series are highly correlated and $\mathbf{x}_1$ and $\mathbf{y}_1$ project strongly on the EOFs of x and y. SVD might appear useful in this case. The point is moot, however, if x and y are each dominated by a single EOF, because then patterns of x and y always look like those EOFs, and it is not necessary to do SVD to isolate simultaneous pairs of patterns since there is only one pair anyway. The notion of simultaneous occurrence is problematic even if x and y are not dominated by single EOFs. As our example of streamfunction and vorticity demonstrates, it is, in fact, *impossible* for the streamfunction and vorticity SVs to occur simultaneously. What really occurs simultaneously (or goes together) with the first SV of streamfunction is its Laplacian, not the first SV of vorticity.

SVD is often also employed to isolate "coupled modes of oscillation" by considering the time derivatives of the relevant variables as well as the variables themselves. If one variable (say $\mathbf{x}$) has two dominant EOFs and their PCs show a simple oscillation, and if the SVD analysis of $\mathbf{x}$ and $\mathbf{y}$ yields $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_2, \mathbf{y}_2)$ as SV pairs, then the SVD analysis of $d\mathbf{x}/dt$ and $\mathbf{y}$ will yield $(\mathbf{x}_2, \mathbf{y}_1)$ and $(-\mathbf{x}_1, \mathbf{y}_2)$ as SV pairs, *regardless of whether or not* $\mathbf{y}$ *has a simple oscillation.* Such a result, therefore, does not by itself prove the existence of a coupled mode. To do that one must also demonstrate that an SVD analysis of $\mathbf{x}$ and $d\mathbf{y}/dt$ yields $(\mathbf{x}_1, \mathbf{y}_2)$ and $(\mathbf{x}_2, -\mathbf{y}_1)$ as SV pairs.

In conclusion, SVD can be useful in compressing two data fields in much the same way that PCA can be used to compress a single data field. If the leading SV pairs explain significant portions of the variance of their respective data fields, then SVD may also be used to assess the strength of coupled variability present (although, again, correlating PCs with data may do this as well). SVD may not necessarily aid in improving one's understanding of the *nature* of this coupling, however, when it is unable to recover the correct relationship between the two data fields. If the resulting

SV pairs are not correctly related, then it is impossible for them to exist simultaneously. This was the case for the ψ_{SV} and ζ_{SV} patterns found in section 5.

Our caveat is clear. A method that fails a simple test, where the answer is known in advance, must be viewed with suspicion when applied to more complicated cases. Clearly, the scientist who chooses to use SVD to analyze two data fields must exercise caution in the physical interpretation of the results.

Acknowledgments. The authors wish to thank Catherine Smith and Cécile Penland for discussions that were very helpful in the development of this work, and Mark Borges who kindly allowed us the use of his graphics program.

REFERENCES

Barnett, T. P., and R. W. Preisendorfer, 1987: Origins and levels of monthly and seasonal forecast skill for United States surface air temperatures determined by canonical correlation analysis. *Mon. Wea. Rev.,* **115,** 1825–1850.

Bretherton, C. S., C. Smith, and J. M. Wallace, 1992: An intercomparison of methods for finding coupled patterns in climate data. *J. Climate,* **5,** 541–560.

Chambers, J. M., 1977: *Computational Methods for Data Analysis.* Wiley, 282 pp.

Hsu, H.-H., 1994: Relationship between tropical heating and global circulation: Interannual variability. *J. Geophys. Res.,* **99,** 10 473–10 489.

Kutzbach, J., 1967: Emperical eigenvectors of sea-level pressure, surface temperature, and precipitation complexes over North America. *J. Appl. Meteor.,* **6,** 791–802.

Penland, C., 1989: Random forcing and forecasting using principal oscillation pattern analysis. *Mon. Wea. Rev.,* **117,** 2165–2185.

——, and T. Magorian, 1993: Prediction of Niño 3 sea surface temperatures using linear inverse modeling. *J. Climate,* **6,** 1067–1076.

Wallace, J. M., C. Smith, and C. S. Bretherton, 1992: Singular value decomposition of wintertime sea surface temperature and 500-mb height anomalies. *J. Climate,* **5,** 561–576.