Error and Sensitivity Analysis of Geophysical Eigensystems

CÉCILE PENLAND AND PRASHANT D. SARDESHMUKH

CDC/CIRES, University of Colorado, Boulder, Colorado

(Manuscript received 20 May 1994, in final form 9 January 1995)

ABSTRACT

The first-order perturbation technique is reviewed as a tool for investigating the error and sensitivity of results obtained from the eigenanalysis of geophysical systems. Expressions are provided for the change in a system's eigenfunctions (e.g., normal modes) and their periods and growth rates associated with a small change $\delta\mathbf{L}$ in the system matrix $\mathbf{L}$. In the context of data analysis, these expressions can be used to estimate changes or uncertainties in the eigenstructure of matrices involving the system's covariance statistics. Their application is illustrated in the problems of 1) updating a subset of the empirical orthogonal functions and their eigenvalues when more data become available, 2) estimating uncertainties in the growth rate and spatial structure of the singular vectors of a linear dynamical system, and 3) estimating uncertainties in the period, growth rate, and spatial structure of the normal modes of a linear dynamical system. The linear system considered in examples 2 and 3 is an empirical stochastic-dynamic model of tropical sea surface temperature (SST) evolution derived from 35 years of SST observations in the tropical Indo-Pacific basin. Thus, the system matrix $\mathbf{L}$ is empirically derived. Estimates of the uncertainty in $\mathbf{L}$, required for estimating the uncertainties in the singular vectors and normal modes, are obtained from a long Monte Carlo simulation. The analysis suggests that the singular vectors, which represent optimal initial structures for SST anomaly growth, are more reliably determined from the 35 years of observed data than are the individual normal modes of the system.

1. Introduction

This article is concerned with estimating how the eigenvalues and eigenvectors of a real matrix $\mathbf{L}$ are affected by a small change $\delta\mathbf{L}$ in it. There exists a well-known first-order perturbation technique for doing this, which is widely used in some physical sciences but rarely in meteorology or oceanography. A complete treatment can be found in such classic works as that of Wilkinson (1965). In the meteorological literature, North et al. (1982) have used perturbation theory to estimate the uncertainties inherent in empirical orthogonal function (EOF) analysis. The technique is reviewed here, and some examples of its application are provided that have particular practical use.

The most direct application of the technique is in the sensitivity analysis of linear (or linearized) dynamical systems of the form $d\mathbf{x}/dt = \mathbf{L}\mathbf{x} + \mathbf{f}$, where $\mathbf{x}$ is the state vector, $\mathbf{L}$ is the system matrix, and $\mathbf{f}$ represents external forcing. The free solutions (normal modes) of such a system are $\mathbf{u}_\alpha \exp(\lambda_\alpha t)$, where $\mathbf{u}_\alpha$ and λ_α are the αth eigenfunction and eigenvalue of $\mathbf{L}$. The real and imaginary parts of λ_α represent the growth rate and angular frequency of the mode, respectively. Because the behavior of the system can be characterized completely in terms of such modes, it is useful to determine how they are affected by small changes in the system's parameters, incorporated in $\mathbf{L}$. The first-order perturbation technique enables one to do this. It also enables one to investigate the sensitivity of other quantities of interest such as the singular vectors of the system [i.e., the vectors obtained from a singular value decomposition of the Green function matrix $\mathbb{G}(\tau) \equiv \exp(\mathbf{L}\tau)$]. The right singular vectors, which are the eigenfunctions of $\mathbf{G}^T\mathbf{G}$, represent initial state vectors that maximize growth of the variance $\mathbf{x}^T\mathbf{x}$ over the time interval τ in the absence of forcing. For this reason they are also called optimal initial perturbations. The amplification factor for the variance is given by the associated singular values (i.e., the eigenvalues of $\mathbf{G}^T\mathbf{G}$).

The perturbation technique can also be applied usefully in data analysis. In recent years, eigenvalue techniques of data analysis have come to be appreciated for their ability to extract and compress the rich amount of information contained in multivariate data records, with relatively little effort expended on the part of the researcher. The key word here is "relatively." Some techniques, such as EOF analysis, require a nontrivial amount of computation when data volumes and the apparent number of degrees of freedom are large. EOFs are the eigenfunctions of a simultaneous, or zero-time lag, covariance matrix $\mathbf{C} = \langle \mathbf{x}(t)\mathbf{x}^T(t) \rangle$, where the angle brackets denote an ensemble average that is usually estimated as a time average. We consider the problem of updating a small subset of the EOFs as more

Corresponding author address: Dr. Cécile Penland, CIRES, University of Colorado, Campus Box 449, Boulder, CO 80309-0449.

data become available. Because EOFs are meaningful only to the extent that the statistics incorporated in $\mathbf{C}$ are stationary (i.e., that $\mathbf{C}$ does not depend upon time), one hopes that the additional data do not drastically alter $\mathbf{C}$. Still, some change can always be expected. It is then desirable to have a method of estimating the change in the leading EOFs without having to repeat the entire eigenanalysis with the modified $\mathbf{C}$. The perturbation technique enables one to do this and also gives, as a by-product, an assessment of the assumption of stationarity.

The above examples represent applications of the perturbation technique where the perturbation $\delta\mathbf{L}$ to $\mathbf{L}$ is either specified or known. We envisage another kind of application, particularly in data analysis and in the inverse modeling of $\mathbf{L}$ from data, where $\delta\mathbf{L}$ is not known but is estimated by a Monte Carlo technique described below. The perturbation technique then enables one to estimate the uncertainty or error in the eigenstructure of $\mathbf{L}$ using this estimate of $\delta\mathbf{L}$.

The inverse modeling example is particularly relevant. Linear inverse modeling consists of approximating any dynamical system $d\mathbf{x}/dt = \mathbf{L}\mathbf{x} + \mathbf{n}(\mathbf{x}) + \mathbf{f}$, where $\mathbf{n}(\mathbf{x})$ represents nonlinear terms, as a linear system driven by (spatially coherent) white noise $d\mathbf{x}/dt = \mathbf{B}\mathbf{x} + \xi$ and estimating the matrix $\mathbf{B}$ from the observed zero-lag and time-lag covariance matrices. This is done via multiple linear regression of the observations. If the statistics of the system are stationary, the eigenvalues of $\mathbf{B}$ have negative real parts. A fluctuation–dissipation relation $\mathbf{B}\mathbf{C} + \mathbf{C}\mathbf{B}^{\mathrm{T}} + \mathbf{Q} = 0$, where $\mathbf{C}$ is the zero-lag covariance matrix, is then used to estimate the covariance matrix $\mathbf{Q}$ of the noise forcing. The eigenfunctions of $\mathbf{Q}$, also called "noise EOFs," represent patterns of stochastic forcing in decreasing order of importance. The eigenfunctions of $\mathbf{B}$ are called empirical normal modes, and sometimes also principal oscillation patterns (e.g., Hasselmann 1988). As above, the empirical optimal perturbations are the right singular vectors of the Green function matrix $\mathbf{G}(\tau) \equiv \exp(\mathbf{B}\tau)$. Of course, the normal modes and singular vectors thus derived are physically meaningful only to the extent that the dynamical system can be reasonably approximated in this way. A more extensive discussion of linear inverse modeling, and methods of checking its validity, can be found in previous articles (Penland and Magorian 1993; Penland and Matrosova 1994; Penland and Sardeshmukh 1995).

In any event, it is desirable to estimate the uncertainty in the empirically derived normal modes and optimal initial perturbations of the system due solely to the length of the data record. To do this, one first needs to estimate the uncertainty $\delta\mathbf{B}$ in $\mathbf{B}$ resulting from considering an observational data record of finite length. To this end the stochastic-dynamic model $d\mathbf{x}/dt = \mathbf{B}\mathbf{x} + \xi$ of the system is integrated forward in time using the empirically derived system matrix $\mathbf{B}$ and a combination of the noise EOFs as forcing, whose amplitudes at each time step are specified using a random number generator. The details are given in Penland and Matrosova (1994) and will not be repeated here. Now the output from the model run can be processed in exactly the same way as the observations, to estimate a "model-output $\mathbf{B}$," which, if the run is sufficiently long, should be identical to the input $\mathbf{B}$ for self-consistency. One can then also estimate $\mathbf{B}$ from smaller segments of the model output, say of the same length as the length of the observational record used to generate the input $\mathbf{B}$. Because of sampling errors, these are not identical to the input $\mathbf{B}$, and their deviations from it can be used to generate an estimate of $\delta\mathbf{B}$. This estimate can then be used to generate estimates of the uncertainty in the empirical normal modes and optimal initial perturbations.

The first-order perturbation technique can thus be combined with linear inverse modeling to generate estimates of uncertainty in quantities derived from a dynamical system's observed statistics through eigenanalysis. It should be noted that such estimates are not always less expensive than a brute force calculation of ensembles of normal modes, and we discuss criteria for choosing which technique is desirable.

The article proceeds as follows. Section 2 is a review of the first-order perturbation theory of eigensystems. Both self-adjoint (symmetric) and non-self-adjoint (nonsymmetric) matrices are considered. Section 3 is an application of the theory to symmetric matrices, with particular emphasis on the updating of EOFs with the inclusion of additional data and on calculation of uncertainties in the empirically derived optimal initial perturbations of a linear system. Section 4 applies the theory to nonsymmetric matrices. Uncertainties in the structure and timescales of the empirically derived normal modes of one of the linear systems considered in section 3 are estimated. The note closes with a discussion of the technique and the physical insight it has provided.

2. Perturbation theory

Consider the eigenvalue problem for a real square matrix $\mathbf{L}$ of dimension S. Let $\mathbf{u}_\alpha$ be an eigenfunction

$$\mathbf{L}\mathbf{u}_\alpha = \lambda_\alpha \mathbf{u}_\alpha, \tag{1}$$

and $\mathbf{v}_\alpha$ its corresponding adjoint:

$$\mathbf{L}^{\mathrm{T}}\mathbf{v}_\alpha = \lambda_\alpha \mathbf{v}_\alpha. \tag{2}$$

Note that, in general, the matrix $\mathbf{U}$ having $\mathbf{u}_\alpha$ as its αth column and the matrix of corresponding adjoints $\mathbf{V}$ are not individually orthogonal but form a biorthogonal set:

$$\mathbf{U}^{\mathrm{T}}\mathbf{V} = \mathbf{I}. \tag{3}$$

and

$$\mathbf{U}\mathbf{V}^{\mathrm{T}} = \mathbf{I}. \tag{4}$$

Now let us assume that there is some change or error in the original operator $\mathbf{L}$ and that we wish to update the eigenvalues, eigenvectors, and adjoints accordingly. Denote the changes in $\{\mathbf{L}, \lambda_\alpha, \mathbf{u}_\alpha, \mathbf{v}_\alpha\}$ as $\{\delta\mathbf{L}, \delta\lambda_\alpha, \delta\mathbf{u}_\alpha, \delta\mathbf{v}_\alpha\}$. Then

$$(\mathbf{L} + \delta\mathbf{L})(\mathbf{u}_\alpha + \delta\mathbf{u}_\alpha) = (\lambda_\alpha + \delta\lambda_\alpha)(\mathbf{u}_\alpha + \delta\mathbf{u}_\alpha) \quad (5)$$

or, to first order,

$$(\mathbf{L} - \lambda_\alpha\mathbf{I})\delta\mathbf{u}_\alpha = (\delta\lambda_\alpha\mathbf{I} - \delta\mathbf{L})\mathbf{u}_\alpha. \quad (6)$$

Since $(\mathbf{L} - \lambda_\alpha\mathbf{I})$ has a zero eigenvalue, the adjoint of the eigenvector $\mathbf{u}_\alpha$ lying in its null space must be orthogonal to any inhomogeneity ("Fredholm's Alternative": (Boyce and DiPrima 1969). Thus, we obtain the change in the eigenvalue

$$\delta\lambda_\alpha = \mathbf{v}_\alpha^{\mathrm{T}}\delta\mathbf{L}\,\mathbf{u}_\alpha. \quad (7)$$

Now, if $\mathbf{U}$ has a unique inverse and $\{\mathbf{u}_\alpha\}$ is complete, then $\{\mathbf{v}_\alpha\}$ is also complete. We assume this to be the case and solve for the changes $\{\delta\mathbf{u}_\alpha, \delta\mathbf{v}_\alpha\}$ by first expressing them in terms of the original vectors as follows:

$$\delta\mathbf{u}_\alpha = \sum_{\gamma=1}^{S} \mathbf{u}_\gamma a_{\gamma\alpha} \quad (8a)$$

$$\delta\mathbf{v}_\alpha = \sum_{\gamma=1}^{S} \mathbf{v}_\gamma b_{\gamma\alpha}, \quad (8b)$$

where the sum over γ does not include α. This is justified since $(1 + a_{\alpha\alpha})\mathbf{u}_\alpha$ is an eigenvector of the unperturbed matrix $\mathbf{L}$, $(1 + b_{\alpha\alpha})\mathbf{v}_\alpha$ is an eigenvector of the unperturbed matrix $\mathbf{L}^{\mathrm{T}}$, and these quantities may be substituted for $\mathbf{u}_\alpha$ and $\mathbf{v}_\alpha$ anywhere in the perturbation theory applied up to this point. The factors $a_{\alpha\alpha}$ and $b_{\alpha\alpha}$ are then implicitly determined by choosing a normalization for the perturbed eigenvectors and multiplying the corresponding adjoints by a factor so that the sets are binormal. If $\mathbf{U}$ is self-adjoint, then $a_{\alpha\alpha} = 0$ to first order, since the perturbed eigenvector would differ from $\mathbf{u}_\alpha$ only by contributions orthogonal to it.

We use the biorthogonality property of the updated eigenvectors:

$$(\mathbf{u}_\alpha + \delta\mathbf{u}_\alpha)^{\mathrm{T}}(\mathbf{v}_\beta + \delta\mathbf{v}_\beta) = \delta_{\alpha\beta}$$

$$\mathbf{u}_\alpha^{\mathrm{T}}\mathbf{v}_\beta + \delta\mathbf{u}_\alpha^{\mathrm{T}}\mathbf{v}_\beta + \mathbf{u}_\alpha^{\mathrm{T}}\delta\mathbf{v}_\beta + \text{second-order term} = \delta_{\alpha\beta}$$

or

$$\delta\mathbf{u}_\alpha^{\mathrm{T}}\mathbf{v}_\beta + \mathbf{u}_\alpha^{\mathrm{T}}\delta\mathbf{v}_\beta = 0.$$

Substituting the expansions (8) gives

$$\sum_{\gamma=1}^{S} a_{\gamma\alpha}\mathbf{u}_\gamma^{\mathrm{T}}\mathbf{v}_\beta + \sum_{\gamma=1}^{S} b_{\gamma\beta}\mathbf{u}_\alpha^{\mathrm{T}}\mathbf{v}_\gamma = 0$$

or

$$a_{\beta\alpha} + b_{\alpha\beta} = 0. \quad (9)$$

Note that if $\mathbf{U}$ is self-adjoint, then $a_{\alpha\beta} = -a_{\beta\alpha}$, so that leaving the term $a_{\alpha\alpha}$ in the expansion (8a) would have no effect. It is easy to show that given (9), the corresponding perturbation completeness relation

$$(\mathbf{u}_\alpha + \delta\mathbf{u}_\alpha)(\mathbf{v}_\beta + \delta\mathbf{v}_\beta)^{\mathrm{T}} = \mathbf{I} \quad (10)$$

holds to first order identically. Now replacing $\delta\mathbf{u}_\alpha$ with its expansion (8a) in Eq. (6) gives

$$(\mathbf{L} - \lambda_\alpha\mathbf{I}) \sum_{\gamma=1}^{S} \mathbf{u}_\gamma a_{\gamma\alpha} = (\delta\lambda_\alpha\mathbf{I} - \delta\mathbf{L})\mathbf{u}_\alpha$$

or

$$\sum_{\gamma=1}^{S} (\lambda_\gamma - \lambda_\alpha)\mathbf{u}_\gamma a_{\gamma\alpha} = (\delta\lambda_\alpha\mathbf{I} - \delta\mathbf{L})\mathbf{u}_\alpha. \quad (11)$$

Taking the inner product of $\mathbf{v}_\mu$ and Eq. (11) and rearranging terms gives, finally,

$$a_{\mu\alpha} = -\mathbf{v}_\mu^{\mathrm{T}}\delta\mathbf{L}\,\mathbf{u}_\alpha/(\lambda_\mu - \lambda_\alpha). \quad (\mu \neq \alpha). \quad (12)$$

Equations (9) and (12) give the expansion perturbations for updating the eigenfunctions and their adjoints in (8). The eigenvalues are updated using (7). Note that these expressions reduce to the results of North et al. (1982) for $\mathbf{U} = \mathbf{V}$.

Before going on to specific examples, it is worth noting the limitations of the technique. The theory is strictly valid only for small perturbations $\delta\mathbf{L}$ of $\mathbf{L}$, although what is "small" depends upon the problem being considered. If $\delta\mathbf{L} = c\mathbf{L}$, where c is a constant, then the expressions (7), (9), and (12) give the correct answer for arbitrarily large c. In general, however, one expects the theory to fail if the elements of $\delta\mathbf{L}$ have magnitudes comparable to the standard deviation of the eigenvalues λ_α, as in this case one expects the expansion perturbations $a_{\mu\alpha}$ in Eq. (12) to be O(1). In many problem of interest in which there exist small eigenvalues and/or eigenvalues that are close to one another, Eq. (12) suggests that the theory is likely to fail for the associated eigenvectors.

Another limitation involves the expense of calculation. For an S-dimensional matrix $\mathbf{L}$, the evaluation of (7) involves $\sim S^2$ multiplications for updating one eigenvalue and $\sim S^3$ multiplications for updating all S eigenvalues. The eigenvector updates (8) and (12) involve $\sim S^3$ multiplications for one eigenvector and $\sim 3S^3$ multiplications for all eigenvectors. The complete eigenanalysis of $\mathbf{L}$ also involves $O(S^3) = kS^3$ multiplications, where k is some constant. It would therefore appear that perturbation theory has no obvious computational advantage over repeating the complete eigenanalysis with $\mathbf{L} + \delta\mathbf{L}$. However, the value of k is important. Experiments performed with 100- and 300-dimensional real general matrices suggested that $k \sim 30$ for softward routines in the Numerical Algorithms Group eigenanalysis package. For these matrices, at least, the perturbation calculations were

faster by an order of magnitude. Of course, if one is interested in updating only a subset $M \ll S$ of the eigenvalues, the perturbation calculations involve only $\sim MS^2$ multiplications. If one is interested in updating only the corresponding M eigenvectors, then assuming that the sum in Eq. (8a) involves contributions mainly from this subset, the total number of calculations is again $\sim MS^2$. For $M \ll S$ this can result in a significant saving. Although for many systems the assumptions involving truncation to a subset M are valid, they may sometimes appear unjustified in other systems. However, if the purpose of the perturbation calculations is to assess the need to repeat the complete eigenanalysis, then if even a highly truncated calculation indicates a large change, that purpose is fulfilled.

3. Self-adjoint systems

a. *Example 1: Updating EOFs*

Our first example consists of estimating the small changes in EOF patterns and eigenvalues that result when more data are added to the set. We consider two cases: 250-mb Northern Hemisphere zonal-mean zonal winds and tropical sea surface temperatures (SSTs) in the Indo-Pacific basin. The first case was chosen because the nature of the EOFs allows small changes in them to be easily displayed. In practice, this would not be a situation where one would use the perturbation technique to update the EOFs since the dimensionality is so small; we include this case only for pedagogical reasons. The second case is more realistic. The multivariate time series represent data at 342 locations; however, most of the dynamics can be described by a much smaller number ($\sim$15) of degrees of freedom. Instead of retaining the entire multivariate record, therefore, we project the data onto a number of leading EOFs and investigate only the dynamics contained in their particular time series coefficients. As more data become available, we wish to know whether perturbations to our leading EOFs are large enough to necessitate recalculation of the entire field of EOFs.

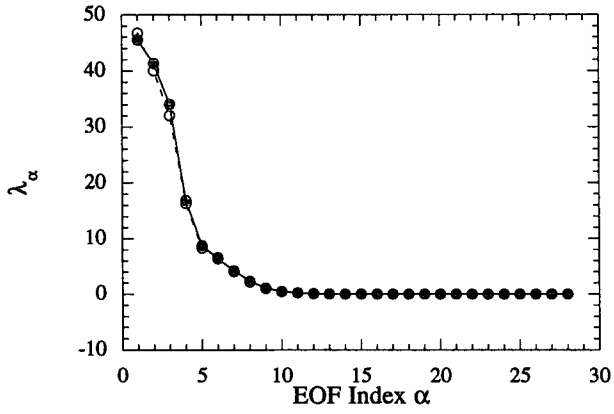

FIG. 1. Eigenvalues associated with the leading 28 EOFs of Northern Hemisphere 250-mb monthly zonal-mean zonal wind anomalies from the annual cycle, computed using 1965–83 data (open circles) and 1965–89 data (filled circles).

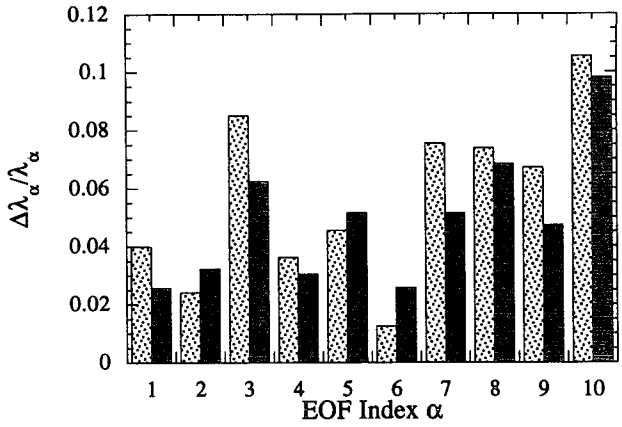

FIG. 2. Solid bars: fractional difference between the eigenvalues displayed in Fig. 1. Shaded bars: fractional difference estimated by perturbation theory. Results are shown only for the 10 leading EOFs.

In the first, pedagogical, case, zonal-mean zonal winds between latitudes 20.0° and 87.5°N, spaced at 2.5° intervals, were determined from zonal wind data provided by the National Meteorological Center. The twice-daily data were consolidated into monthly means for the period April 1965 to June 1989. The monthly means for the period April 1965 to June 1983 inclusive are called the "original" dataset in this exercise, while those for the period April 1965 to June 1989 are called the "full" dataset. The EOFs of the original data were determined first. Then, changes in the EOFs due to the additional 6 years of data were estimated using the perturbation technique outlined above. Finally, the EOFs updated using these estimated changes were checked against the EOFs of the full dataset.

In this example, $\mathbf{L}$ is the covariance matrix of monthly mean zonal-mean wind anomalies, defined as departures from a "climatological" (i.e., long-term mean) annual cycle. Estimation of the covariance matrix requires subtraction of the mean value at each location, so obtaining the perturbation $\delta\mathbf{L}$ requires knowledge of the means from the original dataset. In particular, the mean annual cycle must be correctly accounted for. The details of calculating $\delta\mathbf{L}$ are given in the appendix.

The EOFs and their corresponding eigenvalues were computed for both the original and full datasets. Happily for the assumption of stationarity, the eigenvalues for the two sets are very similar (Fig. 1). In both cases, more than 99% of the variance is contained in the 10 leading patterns. The solid bars in Fig. 2 show the fractional differences between the 10 leading eigenvalues of the full dataset and the corresponding eigenvalues of the original dataset. The shaded bars indicate analogous differences between updated (Eq. 7) and original

eigenvalues. If the perturbation theory were perfect, the shaded bars would be identical to the solid bars. The comparison is not very good for eigenvalues that are four orders of magnitude smaller than the leading one (i.e., λ_α, $\alpha \geq 14$) (not shown). However, it is satisfactory for eigenvalues that correspond to EOFs of any importance in the dataset.

The EOFs of the original set were updated using Eqs. (8a) and (12). The six leading EOFs, explaining 95% of the total variance, are shown in Fig. 3. Each EOF has been normalized to unity. As with the eigenvalues, the updated EOFs (fine solid line) are closer to the EOFs of the full dataset (heavy dashed line) than to those of the original dataset (heavy solid line). The updated patterns $\mathbf{u}_3$–$\mathbf{u}_6$ are virtually indistinguishible from the patterns they were meant to approximate. The solid bars in Fig. 4 show the rms difference between the EOFs of the full and original datasets, each weighted by $(\lambda_\alpha)^{1/2}$. The shaded bars show analogous quantities calculated using the updated EOFs. As in Fig. 2, the comparison is again reasonable.

A similar analysis was performed on tropical Indo-Pacific (30°N–30°S, 30°E–70°W) SSTs. The 2° × 2° gridded data were consolidated onto a 4° × 10° grid and subjected to a three-month running mean. Monthly mean SST anomalies were defined as the departure of these three-month running means from the long-term mean annual cycle. In this exercise, the "original" dataset consisted of the monthly anomalies for February 1950 to January 1985 inclusive (i.e., a total of 420 months). The "full" dataset consisted of monthly anomalies, relative to the same 420-month climatology used in the original set, for February 1950 to November 1991 inclusive, a total of 502 months.

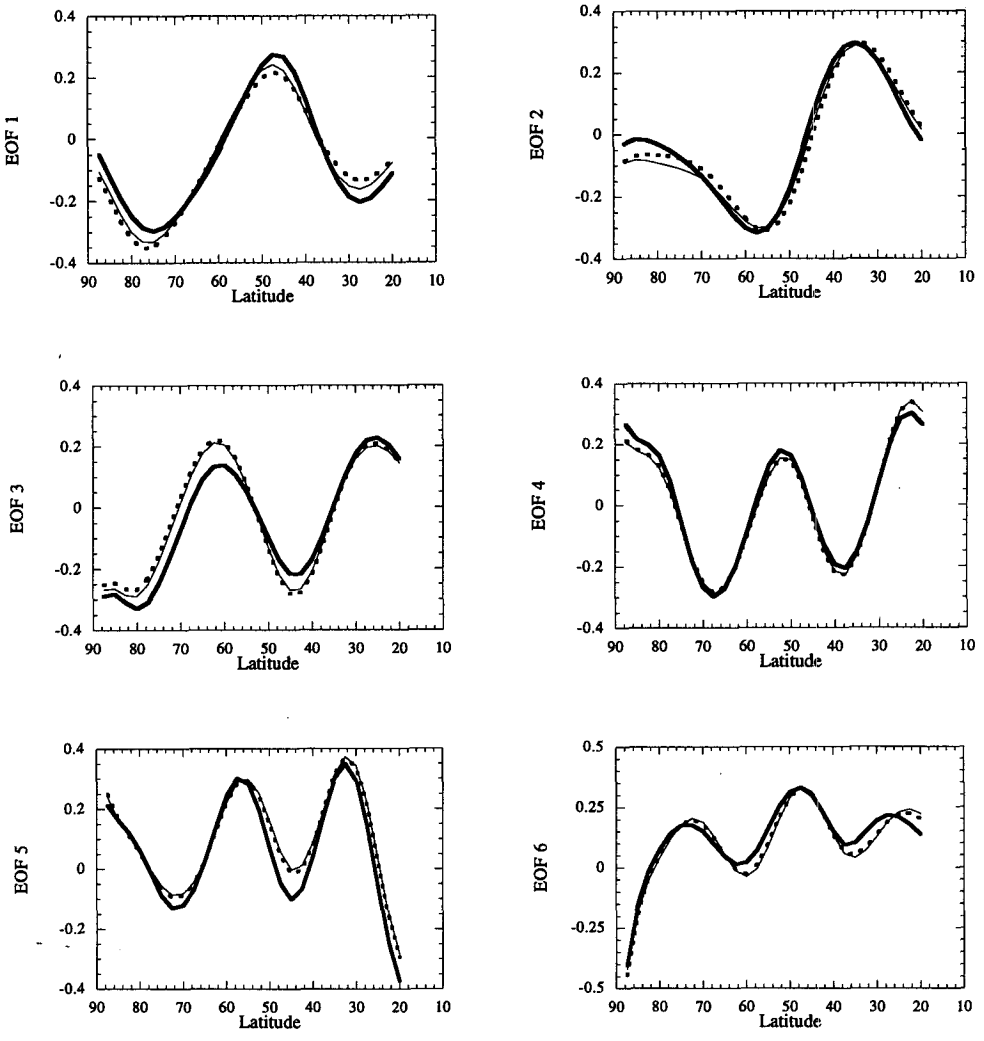

FIG. 3. The six leading EOFs of the zonal-mean zonal winds. Heavy solid lines: EOFs of 1965–83 data. Heavy dashed lines: EOFs of 1965–89 data. Thin solid lines: EOFs of 1965–89 data estimated by updating EOFs of 1965–83 data with perturbation theory.

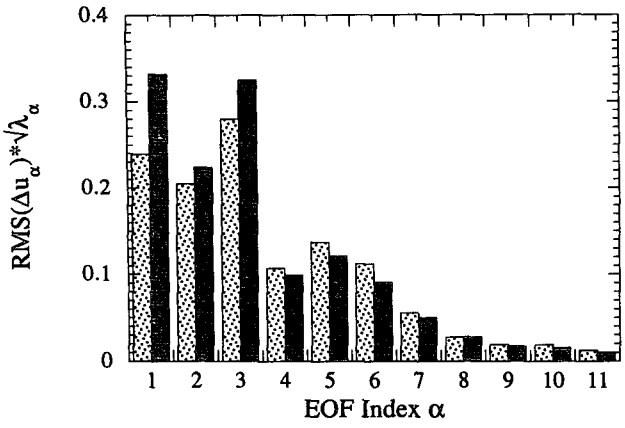

FIG. 4. Solid bars: rms difference of the zonal-mean zonal wind EOFs of 1965–89 data and 1965–83 data, weighted by $(\lambda_0)^{1/2}$. Shaded bars: the same quantity estimated using perturbation theory.

The 15 leading eigenvalues of the covariance matrices of the original and the full datasets, as well as the corresponding eigenvalues updated using perturbation theory, are shown in Fig. 5. There is practically no distinction among them. Differences between any two sets of eigenvalues are within their standard errors as calculated by the International Mathematical and Statistical Library (IMSL) routines. However, the standard errors $\delta\lambda_\alpha = \lambda_\alpha(2)^{1/2}/\text{ndf}$, where ndf = number of degrees of freedom, provided by IMSL are obtained under the assumption that the eigenvalues are distinct (Kendall and Stuart 1966; Girshick 1939), and this is true only for the leading eigenvalue in the SST data.

The EOFs for this dataset are normalized to unity and maps of them (not shown) typically use a contour interval of 0.03. Maximum loadings are typically 0.15. Figure 6 shows the rms difference between the original EOFs and the EOFs of the full dataset (solid bars) and also the rms difference between the original EOFs and the updated set (shaded bars). In this calculation, we have assumed that only the 15 leading EOFs of the original set (containing 63% of the variance) have been retained, and the summation over γ in Eq. (8a) includes contributions only from the retained EOFs. The solid and shaded bars would have equal height if the perturbation theory were perfect. We see that perturbation theory underestimated the change to the leading EOFs; however, for the four leading EOFs that contain nearly half (43%) the variance, this underestimation is small. EOF 5 contains only 3% of the variance and each of the higher-indexed EOFs contains even less. If the change in the original EOFs is required to remain smaller than one contour interval, then perturbation theory correctly indicates that the EOFs must be recalculated. If a change of two contour intervals is tolerated, perturbation theory allows one to forego recalculation for the time being.

The changes in the EOFs arising from the inclusion of additional data are small in the two cases considered here, so it is reasonable to ask why one would want to recalculate the EOFs at all. The problem, of course, is that one does not know this in advance. Perturbation theory, though inexact, does provide an indication of the magnitude of the change to expect, and if the change is large, the need to recalculate the eigenvectors. We have chosen rather low-dimensional systems to illustrate this. The calculation of the perturbations for *these* systems is not much less trouble than simply recalculating the EOFs. However, for dynamical systems with even moderately many degrees of freedom (e.g., a T42 barotropic vorticity model that has $S = 42 \times 44 = 1848$ degrees of freedom), the difference between MS^2 and S^3 can be substantial if most of the variance is contained in, say, the first $M = 20$ EOFs.

b. *Example 2: Uncertainty in optimal structure of a linear system*

Our second example consists of estimating the uncertainty in the empirically derived optimal initial structure for growth in a linear system. The system is a linear, stochastically driven model of monthly SST anomalies fitted to SST observations in the tropical Indo–Pacific basin for the period February 1950 to January 1985 inclusive. The SST anomalies are defined in the same manner as in the example above. The SST anomalies (SSTAs) were projected onto 15 leading EOFs. All calculations were performed in the 15-dimensional EOF space and then transformed to geographical space for interpretation.

As outlined in section 1, the evolution of the ith component of a stable, centered, M-dimensional linear system $\mathbf{x}(t)$ driven by rapidly varying forcing can be written

$$\frac{dx_i}{dt} = \sum_{j=1}^{M} B_{ij}x_j + \xi_i, \qquad (13)$$

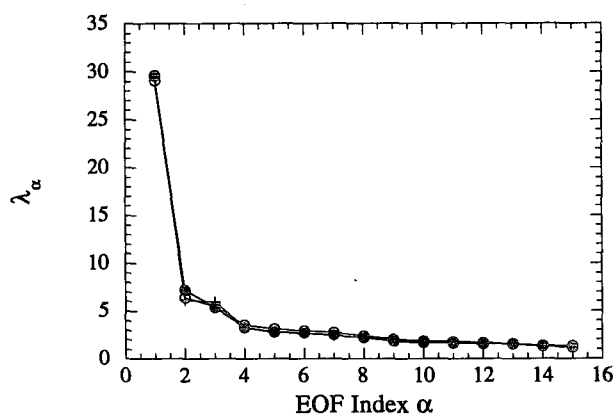

FIG. 5. Eigenvalues associated with the 15 leading EOFs of tropical Indo–Pacific monthly SST anomalies from the annual cycle, computed using 1950–85 data (open circles) and 1950–91 data (filled circles). The crosses indicate eigenvalues for 1950–91 data estimated by updating the eigenvalues for 1950–85 data with perturbation theory.

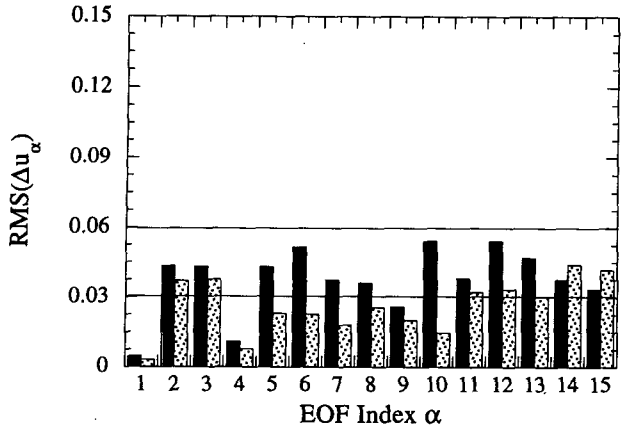

FIG. 6. Solid bars: rms difference of EOFs of 1950–91 SST data and 1950–85 data. Shaded bars: the same quantity estimated using perturbation theory.

where ξ_i is white noise with covariance matrix

$$\langle \xi_i(t+\tau)\xi_j(t)\rangle = Q_{ij}\delta(\tau)$$

and $\mathbf{B}$ is a constant matrix. In this example, $x_i(t)$ represents the ith principal component of the SST anomaly (i.e., the time coefficient of the ith SST EOF) at time t. The empirical normal modes (eigenvectors of $\mathbf{B}$) along with the corresponding eigenvalues are obtained from the covariance structure of the data as described below. The Green function matrix $\mathbf{G}(\tau) \equiv \exp(\mathbf{B}\tau)$ has the same eigenvectors and adjoints as $\mathbf{B}$, and the eigenvalues $\{g_\alpha(\tau)\}$ of $\mathbf{G}(\tau)$ are related to the eigenvalues $\{\beta_\alpha\}$ of $\mathbf{B}$ through a simple exponential relation $g_\alpha(\tau) = \exp(\beta_\alpha\tau)$. We shall consider uncertainty in the eigenvectors $\{\mathbf{u}_\alpha\}$ of the non-self-adjoint matrix $\mathbf{G}(\tau)$ in section 4.

It can be verified for Eq. (13) that

$$\mathbf{G}(\tau) = \langle \mathbf{x}(t+\tau)\mathbf{x}^{\mathrm{T}}(t)\rangle\langle \mathbf{x}(t)\mathbf{x}^{\mathrm{T}}(t)\rangle^{-1} \quad (14)$$

and that the most probable prediction $\hat{\mathbf{x}}(t+\tau)$ of the field $\mathbf{x}(t+\tau)$ given initial condition $\mathbf{x}(t)$ is $\hat{\mathbf{x}}(t+\tau) = \mathbf{G}(\tau)\mathbf{x}(t)$. The initial condition leading to the maximum growth (or minimum decay) of field variance $\mathbf{x}^{\mathrm{T}}\mathbf{x}$ between times t and $t+\tau$ is called the *optimal initial structure for growth* in the linear system at lead time τ. It is the leading eigenvector of the self-adjoint operator $\Gamma(\tau) \equiv \mathbf{G}^{\mathrm{T}}\mathbf{G}(\tau)$:

$$\Gamma(\tau)\phi_1 = \phi_1\gamma_1, \quad (15)$$

where the leading eigenvalue γ_1 of $\Gamma(\tau)$ is the factor by which the field variance grows over time τ after the occurrence of the initial condition ϕ_1. Given random uncertainties in $\Gamma(\tau)$, a measure of the uncertainty in γ_1 and ϕ_1 can be obtained using Eqs. (7), (8), and (12) with $\mathbf{L}$ set to $\Gamma(\tau)$:

$$\langle(\delta\gamma_1)^2\rangle = \sum_{mnpq} \phi_{m1}\phi_{n1}\langle\delta\Gamma_{mn}\delta\Gamma_{pq}\rangle\phi_{p1}\phi_{q1} \quad (16)$$

and

$$\langle\delta\phi_1\delta\phi_1^T\rangle = \sum_{\alpha\beta\neq 1} \phi_\alpha\phi_\beta^T\langle a_{\alpha 1}a_{\beta 1}\rangle. \quad (17)$$

In Eq. (17), ϕ_α is the eigenvector of $\Gamma(\tau)$ corresponding to γ_α and

$$\langle a_{\alpha 1}a_{\beta 1}\rangle = \sum_{mnpq} \phi_{m\alpha}\phi_{n1}\langle\delta\Gamma_{mn}\delta\Gamma_{pq}\rangle\phi_{p\beta}\phi_{q1}/$$
$$[(\gamma_\alpha-\gamma_1)(\gamma_\beta-\gamma_1)], \quad \alpha,\beta\neq 1. \quad (18)$$

Penland and Sardeshmukh (1995) estimated $\Gamma(\tau)$ from the 420 months of observed SST anomaly data. A graph of the leading eigenvalue of $\Gamma(\tau)$ as a function of τ showed that maximum growth of about 500% could be attained at a lead time of 7 months given the optimal structure as an initial condition. Note that the normal modes of this system (the eigenfunctions of $\mathbf{B}$) are all individually damped; however, growth can occur for as long as 18 months from their constructive interference. The optimal structure represents the best possible combination of the normal modes at the initial time to maximize growth over 7 months. The importance of this structure in the tropical Indo-Pacific Ocean is discussed in detail in Penland and Sardeshmukh's paper. Our primary concern here is to estimate the uncertainty in the optimal structure and its associated growth factor, given that it is derived from a relatively short observed SST time series of 420 months.

To estimate the uncertainty in the optimal structure and its growth, a Monte Carlo calculation was used to estimate $\langle\delta\Gamma_{mn}\delta\Gamma_{pq}\rangle$. Following the procedures described in section 1, a numerical model of the stochastic–dynamic system (13) was constructed (Penland and Matrosova 1994) and integrated forward in time. After exercising the model for 1000 months to eliminate memory of the initial condition, which was set to the peak of the multivariate stationary distribution at $\mathbf{x} = \mathbf{0}$, data were recorded once per model month for a total of 42 000 months. The long model time series was divided into 100 segments of 420 months, each equal in length to the original data time series. The matrices $\mathbf{C}$, $\mathbf{G}(\tau)$, and $\Gamma(\tau)$ were estimated for each segment, using $\tau = 7$ months. Ensemble averages of $\mathbf{C}_{\mathrm{av}}$, $\mathbf{G}_{\mathrm{av}}(\tau)$, $\Gamma_{\mathrm{av}}(\tau)$, $\{\langle\delta C_{mn}\delta C_{pq}\rangle\}$, $\{\langle\delta G_{mn}\delta G_{pq}\rangle\}$, and $\{\langle\delta\Gamma_{mn}\delta\Gamma_{pq}\rangle\}$ were then constructed from the 100-member ensemble. Fourier spectra of individual normal modes calculated from the time series in these segments approached the theoretically expected spectra when averaged over the 100 samples, giving confidence that 100 samples were adequate to estimate the desired covariance structures. Note that $\langle\delta\Gamma_{mn}\delta\Gamma_{pq}\rangle$ is a fourth moment of the Green function. The self-consistency of the linear model was verified through satisfactory comparisons of $\mathbf{B}$ and $\mathbf{Q}$ obtained from the data with $\mathbf{B}_{\mathrm{av}}$ and $\mathbf{Q}_{\mathrm{av}}$ obtained from the model output (not shown).

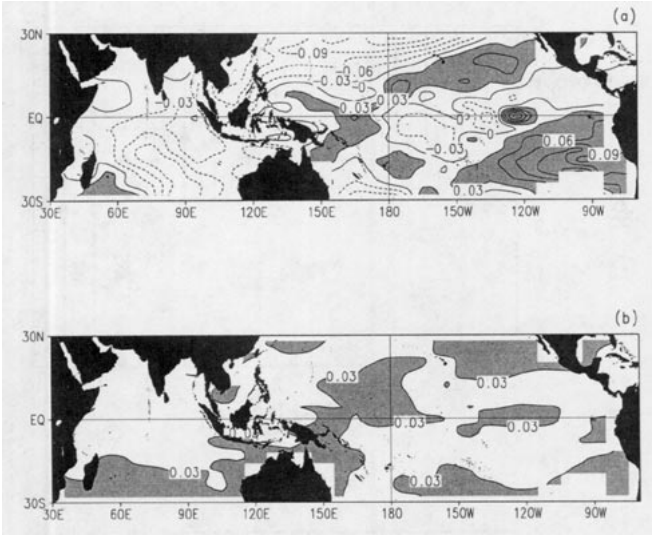

FIG. 7. (a) The optimal initial structure for growth of SST anomalies, normalized to unity, determined from a 42 000 month run of the stochastic-dynamic model. The contour interval is 0.03. Positive values are indicated by solid and negative by dashed contours. (b) The expected local error of the field in (a) due to inadequate sampling if determined from a 420-month series. The error is estimated using perturbation theory and Eq. (17). The contour interval is the same as in (a). Note that only the 0.03 contour is visible.

The leading eigenvector of $\Gamma_{av}(\tau = 7 \text{ months})$, normalized to unity, is shown in Fig. 7a. The rms difference between it and the one found by Penland and Sardeshmukh (1995) from the 420-month observed time series is 4.1%. Using Eqs. (17) and (18), an uncertainty map $\langle(\delta\phi_i)^2\rangle^{1/2}$ was generated (Fig. 7b). Only one contour (0.03) is visible on the uncertainty map, although this value is a large fraction of the loading of the optimal structure in the vicinity of the equatorial date line. Still, considering the short length, 420 months, of the observed time series from which the optimal structure was obtained by Penland and Sardeshmukh (1995), this accuracy is surprisingly good. The growth factor γ_1 was found to be 4.67 ± 1.23, where the standard error was obtained from Eq. (16). This compares with a growth factor of 5.08 found by Penland and Sardeshmukh.

The eigenvalues of a real symmetric matrix are more easily found than those of a real general matrix. It may therefore be more efficient to simply consider the statistics of ensembles of optimal structures, performing the eigenanalysis of Γ, say, 100 times and calculating $\langle(\delta\phi_i)^2\rangle^{1/2}$ and the growth rate uncertainty by brute force rather than using perturbation theory. The number of calculations of the brute force method is $O(S^3)$ as opposed to the more expensive calculation of $\langle\delta\phi_1\delta\phi_1^T\rangle$ (Eq. 17), which requires $O(S^6)$ calculations. The actual computation times for the two approaches were virtually identical for our system, each taking about three minutes *after* the ensemble time series (required by both approaches) was calculated. One disadvantage to the brute force method is that the arbitrary sign of the eigenvector in each member of the ensemble makes automation more complicated, and if the signs are not regularized before the statistics are taken, estimates of both $\langle\phi_1\rangle$ and $\langle\delta\phi_1\delta\phi_1^T\rangle$ are severely compromised. We performed the brute force calculation for the problem discussed above and, when the signs of the eigenvectors were regularized, the correlation between the brute force map $\langle(\delta\phi_i)^2\rangle^{1/2}$ and that obtained using perturbation theory exceeded 98%. Regularizing the signs can be automated if the average of the ensemble patterns $\{\phi_1\}$ is not too close to the null vector and if the members of the ensemble themselves are not too different. However, we found that the brute force method usually does require some human time and effort (i.e., checking by hand). Whether or not this outweighs the computational cost should be left to the individual scientist.

4. Non-self-adjoint systems

The eigenvalues $\{\beta_\alpha\}$ and eigenvectors $\{\mathbf{u}_\alpha\}$ of the nonsymmetric real matrix $\mathbf{B}$ occur in complex-conjugate pairs. Each eigenvalue β_α is associated with a period $T_\alpha = 2\pi/\text{Im}(\beta_\alpha)$ and a decay time $\mu_\alpha = -1/\text{Re}(\beta_\alpha)$. The uncertainty in the eigenvalue β_α is related to that in the eigenvalue g_α of the Green function $\mathbf{G}(\tau)$ as $\delta g_\alpha/(g_\alpha\tau)$, giving

$$\delta\beta_\alpha = \frac{\exp(-\beta_\alpha\tau)}{\tau}\,\mathbf{v}_\alpha^T\delta\mathbf{G}\,\mathbf{u}_\alpha. \tag{19}$$

The uncertainties in the real and imaginary parts of $\delta\beta_\alpha$ are needed separately. For notational facility, we define $\beta_\alpha^r = \text{Re}(\beta_\alpha)$, $\beta_\alpha^i = \text{Im}(\beta_\alpha)$, $\mathbf{v}_\alpha^r \equiv \text{Re}(\mathbf{v}_\alpha)$, $\mathbf{v}_\alpha^i \equiv \text{Im}(\mathbf{v}_\alpha)$, $\mathbf{u}_\alpha^i \equiv \text{Im}(\mathbf{u}_\alpha)$, $\mathbf{u}_\alpha^r \equiv \text{Re}(\mathbf{u}_\alpha)$. Also, let

$$\text{Term 1} = v_{j\alpha}^r v_{m\alpha}^r u_{n\alpha}^r u_{k\alpha}^r + v_{j\alpha}^i v_{m\alpha}^i u_{n\alpha}^i u_{k\alpha}^i - 2v_{j\alpha}^r v_{m\alpha}^i u_{n\alpha}^i u_{k\alpha}^r$$

$$\text{Term 2} = v_{j\alpha}^r v_{m\alpha}^r u_{n\alpha}^i u_{k\alpha}^i + v_{j\alpha}^i v_{m\alpha}^i u_{n\alpha}^r u_{k\alpha}^r + 2v_{j\alpha}^r v_{m\alpha}^i u_{n\alpha}^r u_{k\alpha}^i$$

$$\text{Term 3} = v_{j\alpha}^r v_{m\alpha}^r u_{n\alpha}^i u_{k\alpha}^r - v_{j\alpha}^i v_{m\alpha}^r u_{n\alpha}^i u_{k\alpha}^i - v_{j\alpha}^i v_{m\alpha}^i u_{n\alpha}^r u_{k\alpha}^i + v_{j\alpha}^r v_{m\alpha}^i u_{n\alpha}^r u_{k\alpha}^r.$$

Then, as usual, it is straightforward but tedious to show that

$$\langle(\text{Re}\delta\beta_\alpha)^2\rangle = \frac{\exp[-2\beta_\alpha^r\tau]}{\tau^2}\sum_{jkmn}\langle\delta G_{jk}\delta G_{mn}\rangle \times \{\cos^2(\beta_\alpha^i\tau)\times\text{Term 1} + \sin^2(\beta_\alpha^i\tau)\times\text{Term 2} + 2\sin(\beta_\alpha^i\tau)\cos(\beta_\alpha^i\tau)\times\text{Term 3}\} \tag{20}$$

and

$$\langle(\mathrm{Im}\delta\beta_\alpha)^2\rangle = \frac{\exp[-2\beta_\alpha^r\tau]}{\tau^2}\sum_{jkmn}\langle\delta G_{jk}\delta G_{mn}\rangle$$
$$\times\{\sin^2(\beta_\alpha^i\tau)\times\text{Term 1} + \cos^2(\beta_\alpha^i\tau)\times\text{Term 2}$$
$$-2\sin(\beta_\alpha^i\tau)\cos(\beta_\alpha^i\tau)\times\text{Term 3}\}. \quad (21)$$

Note that $\exp(-\beta_\alpha^r\tau)/\tau$ is minimized at $\tau = -1/\beta_\alpha^r$ (see also Penland and Ghil 1993).

As stated in the previous section, $\mathbf{G}_{\mathrm{av}}(\tau)$ and the covariance structure $\{\langle\delta G_{mn}\delta G_{pq}\rangle\}$ of $\mathbf{G}(\tau)$ were evaluated using the stochastic numerical model. Figure 8 shows the real and imaginary parts of $\{\beta_\alpha\}$ obtained from $\mathbf{G}_{\mathrm{av}}(\tau)$, along with error bars estimated using the expressions (20) and (21). The corresponding timescales and their ranges, where meaningful, are provided in Table 1, along with the timescale obtained from the observed SST anomaly data. The error bars are very large, and only a few [$\alpha = (6, 7), (8, 9), (12, 13)$] are small enough to satisfy a posteriori the assumptions of first-order perturbation theory. Particularly disturbing is the fact that few decay times are significantly longer than that of the most highly damped mode. Also, except for the few modal pairs already mentioned, the periods are so uncertain as to make the modes indistinguishable from purely exponentially decaying structures. While we do not know what the true uncertainty of these periods is, because the first-order perturbation theory breaks down, we *do* know that this breakdown occurs because the uncertainty in the eigenvalues is large. This itself is valuable information.

The uncertainty in the normal modes themselves is estimated from the covariance structure of Eq. (8a):

$$\langle\delta\mathbf{u}_\alpha^\dagger\delta\mathbf{u}_\alpha\rangle = \sum_{\gamma=1}^{S}\sum_{\mu=1}^{S}\mathbf{u}_\gamma^\dagger\mathbf{u}_\mu\langle a_{\gamma\alpha}^* a_{\mu\alpha}\rangle, \quad (22)$$

where (*) denotes complex conjugate and (†) denotes complex conjugate transpose. The quantity $\langle a_{\gamma\alpha}^* a_{\mu\alpha}\rangle$ is related to $\langle\delta G_{mn}\delta G_{pq}\rangle$ through Eq. (12) with $\mathbf{L}$ set to $\mathbf{G}(\tau)$. Each normal mode is normalized to unity (i.e., $\langle\mathbf{u}_\alpha^\dagger\mathbf{u}_\alpha\rangle = 1$). Unfortunately, the value $\langle\delta\mathbf{u}_\alpha^\dagger\delta\mathbf{u}_\alpha\rangle$ is greater than unity for nearly every mode $\mathbf{u}_\alpha$ of the 15-dimensional SSTA system (Table 1). Again, the first-order perturbation theory has broken down, and, again, the perturbation theory need not be valid to be valuable. The message is clear that the structure and timescales of the individual normal modes of SSTA evolution, as determined from 420 months of data, are highly uncertain. However, their collective behavior as manifested in the optimal initial structure for growth is more reliable. This is consistent with previous work and is discussed further below.

In the example presented here, the brute force approach is unfeasible. Many of the different normal modes within an ensemble are highly correlated spatially, and, given the large variations in the timescales obtained from eigenanalysis of the ensembles, we were unable to find a way of automating the analysis so that a particular mode from one ensemble could be unambiguously identified with a mode in another ensemble. This is not surprising since the uncertainties in the modes are so large that perturbation theory itself breaks down. However, perturbation theory does provide some quantitative measure of how the similarity between different normal modes affects their distinguishability.

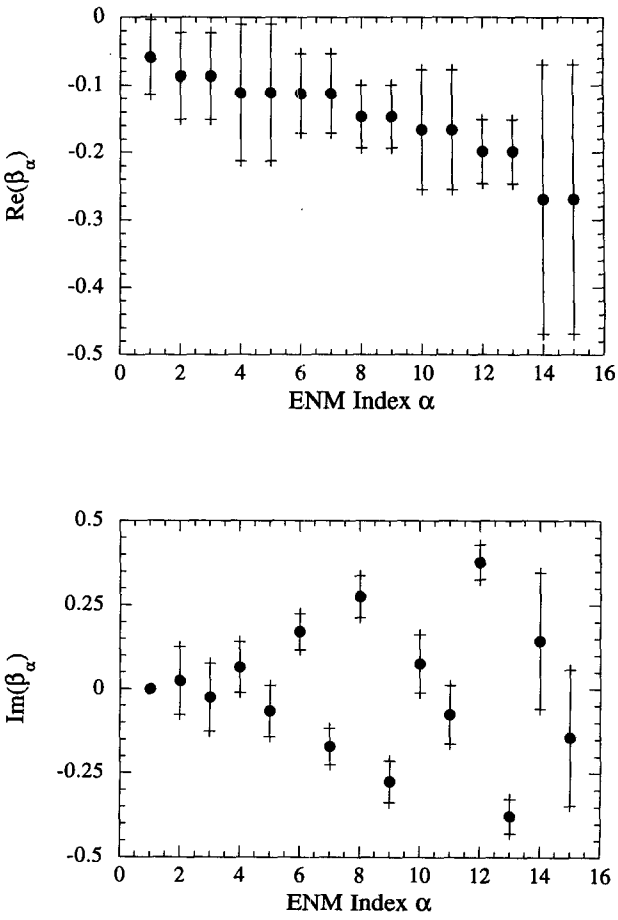

FIG. 8. The real and imaginary parts of the eigenvalues of the 15-dimensional nonsymmetric matrix $\mathbf{B}_{\mathrm{av}}$. Note that except for $\alpha = 1$, the other 14 eigenvalues occur in complex-conjugate pairs. The vertical line segments indicate the expected uncertainty in the plotted values due to inadequate sampling if determined using a 420-month series. The errors are estimated using perturbation theory and Eqs. (20) and (21).

5. Discussion

The first-order perturbation technique is a versatile tool for investigating the error and sensitivity of eigenvalue systems. When the modification $\delta\mathbf{L}$ to a matrix $\mathbf{L}$ is known, the operations are particularly simple. Alternatively, perturbation theory can be combined with Monte Carlo modeling, and often provides the only available estimate of the expected error in an eigenanalysis product derived from observed data.

TABLE 1. Ranges of decay times and periods of empirically derived normal modes. The first entry of each ordered triple is the lower bound, the second entry (boldfaced) is the value obtained from $\mathbf{G}_{av}$ (see text) and the third entry is the upper bound. The value in parentheses was obtained from the Coupled Ocean–Atmosphere Data Set data. Mode 1 is purely real. Other entries marked "N/A" indicate that the uncertainties are too large to be estimated by perturbation theory. Also shown is the rms uncertainty in each normal mode pair; the uncertainty is valid if $\langle \delta \mathbf{u}_\alpha^\dagger \mathbf{u}_\alpha \rangle \ll 1$. These values should be compared with $\langle \delta\phi_1^\dagger \delta\phi_1 \rangle = 0.56$ for the uncertainty in the first optimal eigenvector.

Mode index α	Decay time (month)	Period (month)	$\langle \delta \mathbf{u}_\alpha^\dagger \delta \mathbf{u}_\alpha \rangle$
1	8.8, **17.0**, 269.0 (23.8)	N/A	3.5
2/3	6.6, **11.6**, 44.5 (17.4)	N/A, **252.7**, N/A (364.4)	3.1
4/5	4.7, **9.0**, 97.3 (10.2)	N/A, **96.7**, N/A (159.2)	4.6
6/7	5.9, **8.9**, 18.9 (8.5)	27.8, **36.7**, 53.7 (45.9)	1.1
8/9	5.2, **6.8**, 10.0 (8.0)	18.7, **22.9**, 29.5 (24.6)	1.3
10/11	3.9, **6.0**, 13.0 (7.3)	N/A, **83.0**, N/A (84.5)	3.7
12/13	4.1, **5.0**, 6.6 (4.5)	14.7, **16.7**, 19.3 (18.0)	0.9
14/15	2.1, **3.7**, 14.4 (3.7)	N/A, **107.7**, N/A (72.2)	2.0

We have illustrated application of the technique in 1) updating EOFs and their eigenvalues with the inclusion of additional data, 2) estimating uncertainties in the eigenvalues and eigenfunctions of real symmetric matrices, and 3) estimating uncertainties in the eigenvalues and eigenfunctions of general real matrices. The results of the particular examples used in the illustrations are worth examining for the insight they provide.

The stationarity assumption inherent the calculation and interpretation of EOFs was found to be well grounded, both for the zonal-mean zonal wind and SST datasets considered. The SST result is particularly interesting since it indicates that even if the mean SST for 1985–91 is significantly different from that for 1950–85, this change may not be large enough to significantly alter results based on the original EOF analysis.

Considering empirically derived dynamical quantities, it was found that the collective behavior of the normal modes as manifested in the optimal initial structure for nonmodal growth was much more reliable than the structure or timescales of the individual modes themselves. Even with the relatively small, 15, degrees of freedom, 35 years of data were insufficient to discern unambiguously the normal modes of the sea surface temperature anomalies. This result may somewhat compromise studies that emphasize individual normal mode pairs. Nevertheless, there is reason to believe that the empirically derived modes do contain some physical information. The large estimated uncertainties in the modes are almost certainly affected by the fact that many of them are highly correlated and, therefore, very long time series are necessary to distinguish between separate mode pairs. This interpretation is consistent with Penland and Matrosova's (1994) conclusion that the similar geographical structure of several modes and their corresponding adjoints results in an imperfect projection of the multivariate time series onto the normal modes so that the modes contaminate each others' spectra.

We are aware that the uncertainties as we have calculated them require a numerical model to provide the covariance structure of the matrix operators, and that this is not always available. Detailed recipes for constructing such a model can be found in Penland and Matrosova (1994). The computational cost of running such a model is another consideration. For an S-dimensional system (13), the number of mathematical operations per model time step are $O(S^2)$. However, to ensure that the model is self-consistent and to generate reliable estimates of a quantity such as $\langle \delta L_{mn} \delta L_{pq} \rangle$ required in (16) and (18), the model has to be run for many time steps J. Note that $\langle \delta L_{mn} \delta L_{pq} \rangle$ has S^4 elements, so the number of operations required to evaluate (16) and (18) are $O(S^4)$. Estimation of uncertainties *by the Monte Carlo method* thus require $O(JS^2 + S^4)$ operations. For the 15-dimensional system analyzed here, running the model for $J \sim 10^7$ time steps was the most expensive part of the analysis, taking only a few hours on our Sun workstations.

For large $S \sim 10^3$, some compromise may be necessary. For example, instead of using $\langle \delta L_{mn} \delta L_{pq} \rangle$ in (16) and (18), one could use (7) and (12) directly, specifying some estimate of $\delta \mathbf{L}$. At the very worst, one could specify $\delta \mathbf{L}$ as $\mathbf{L}_2$-$\mathbf{L}_1$, with $\mathbf{L}_1$ and $\mathbf{L}_2$ determined from two halves of the data record or, if (as in the SST case) the data record is too short for this, from two randomly chosen members of the Monte Carlo ensemble. One might then use the corresponding $\{\delta \mathbf{u}_\alpha\}$ [Eqs. (8) and (12)] as a measure of the uncertainty. Note also that the model run need not be very long if one is only interested in generating a few samples of $\mathbf{L}$. The total number of mathematical operations required are then $O(jS^2)$, with j much smaller than the J above. Of course, this would represent a compromise. Compromise or not, we believe that such an analysis performed as a matter of course would give a realistic impression of the value of an eigensystem analysis.

Acknowledgments. We are pleased to acknowledge Dr. Mark Borges for his experiments concerning the timing of perturbation theory applied to real general matrices, and Ms. Ludmila Matrosova for graphical assistance. We are also grateful for the useful comments of an anonymous reviewer.

APPENDIX

Updating the Covariance Matrix of a System with Cyclically Varying Mean

Let N_s be the number of years for which there are data for month s in the original dataset. For example, the zonal-mean zonal wind data used in example 1 has $N_s = 19$ for s = April and $N_s = 18$ for s = July. Let $\mathbf{x}_s(t_y)$ be the data vector corresponding to month s in the year t_y, $\mathbf{m}_s^{\mathrm{O}}$ be the sample mean of the data vector, N be the total number of months in the original dataset, and ΔN_s be the number of times month s occurs in the additional data ($\Delta N_s = 6$ for all months s in this example). That is,

$$\mathbf{m}_s^{\mathrm{O}} = \frac{1}{N_s} \sum_{y=1}^{N_s} \mathbf{x}_s(t_y) \tag{A1}$$

and

$$N = \sum_{s=1}^{12} N_s. \tag{A2}$$

We additionally define $\mathbf{m}_s^u$ as the sample mean for the full dataset:

$$\mathbf{m}_s^u = \frac{1}{N_s + \Delta N_s} \sum_{y=1}^{N_s+\Delta N_s} \mathbf{x}_s(t_y) \tag{A3}$$

and the total number of additional months

$$\Delta = \sum_{s=1}^{12} \Delta N_s. \tag{A4}$$

Of course, the purpose of this exercise is to avoid having to use the entire dataset for updating the EOFs, so it is desirable to express $\mathbf{m}_s^u$ only in terms of $\mathbf{m}_s^{\mathrm{O}}$ and the new data. This is easily done:

$$\mathbf{m}_s^u = \frac{N_s}{N_s + \Delta N_s} \mathbf{m}_s^{\mathrm{O}} + \frac{1}{N_s + \Delta N_s} \sum_{y=N_s+1}^{N_s+\Delta N_s} \mathbf{x}_s(t_y). \tag{A5}$$

The sample covariance matrix calculated for the original dataset is

$$\mathbf{C}^{\mathrm{O}} = \frac{1}{N-1} \sum_{s=1}^{12} \sum_{y=1}^{N_s} [\mathbf{x}_s(t_y) - \mathbf{m}_s^{\mathrm{O}}][\mathbf{x}_s(t_y) - \mathbf{m}_s^{\mathrm{O}}]^{\mathrm{T}} \tag{A6}$$

and similarly for the sample covariance matrix $\mathbf{C}^u$ for the full dataset. It is straightforward but tedious to show that the perturbation matrix $\delta\mathbf{L} = \mathbf{C}^u - \mathbf{C}^{\mathrm{O}}$, in terms of the new data and previously calculated quantities, is

$$\delta\mathbf{L} = -\frac{\mathbf{C}^{\mathrm{O}}\Delta}{N+\Delta-1} + \frac{1}{N+\Delta-1} \times \sum_{s=1}^{12} \{N_s \mathbf{m}_s^{\mathrm{O}} \mathbf{m}_s^{\mathrm{O}T} - (N_s + \Delta N_s)\mathbf{m}_s^u \mathbf{m}_s^{uT}\} + \frac{1}{N+\Delta-1} \sum_{s=1}^{12} \sum_{y=N_s+1}^{N_s+\Delta N_s} \mathbf{x}_s(t_y)\mathbf{x}_s^{\mathrm{T}}(t_y), \tag{A7}$$

where $\mathbf{m}_s^u$ is obtained from Eq. (A5).

REFERENCES

Boyce, W. E., and R. C. DiPrima, 1969: *Elementary Differential Equations and Boundary Value Problems.* 2d ed., John Wiley and Sons, 533 pp.

Girshick, M. A., 1939: On the sampling theory of roots of determinantal equations. *Ann. Math. Stat.*, **10,** 203–224.

Hasselmann, K., 1988: PIPs and POPs—A general formalism for the reduction of dynamical systems in terms of principal oscillation patterns and principal interaction patterns. *J. Geophys. Res.*, **93,** 11 015–11 021.

Kendall, M. G., and A. Stuart, 1976: *The Advanced Theory of Statistics.* Vol. 3, *Design and Analysis, and Time Series.* 3d ed., Charles Griffin & Co.

North, G. R., T. L. Bell, R. F. Cahalan, and F. J. Moeng, 1982: Sampling errors in the estimation of empirical orthogonal functions. *Mon. Wea. Rev.*, **110,** 699–706.

Penland, C., and M. Ghil, 1993: Forecasting Northern Hemisphere 700-mb geopotential heights using principal oscillation patterns. *Mon. Wea. Rev.*, **121,** 2355–2372.

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

——, and L. Matrosova, 1994: A balance condition for stochastic numerical models with application to the El Niño–Southern Oscillation. *J. Climate,* **7,** 1352–1372.

——, and P. D. Sardeshmukh, 1995: The optimal growth of tropical sea surface temperature anomalies. *J. Climate,* **8,** 1999–2024.

Wilkinson, J. H., 1965: *The Algebraic Eigenvalue Problem.* Oxford University Press, 662 pp.