A Multivariate Frequency-Domain Approach to Long-Lead Climatic Forecasting*

BALAJI RAJAGOPALAN

Lamont-Doherty Earth Observatory, Columbia University, Palisades, New York

MICHAEL E. MANN

Department of Geosciences, University of Massachusetts—Amherst, Amherst, Massachusetts

UPMANU LALL

Department of Civil and Environmental Engineering and Utah Water Research Laboratory, Utah State University, Logan, Utah

(Manuscript received 14 March 1997, in final form 13 November 1997)

ABSTRACT

Guided by the increasing awareness and detectability of spatiotemporally organized climatic variability at interannual and longer timescales, the authors motivate the paradigm of a climate system that exhibits excitations of quasi-oscillatory eigenmodes with characteristic timescales and large-scale spatial patterns of coherence. It is assumed that any such modes are superposed on a spatially and temporally autocorrelated stochastic noise background. Under such a paradigm, a previously described (Mann and Park) multivariate frequency-domain approach is promoted as a particularly effective means of spatiotemporal signal identification and reconstruction, and an associated forecasting methodology is introduced. This combined signal detection/forecasting scheme exhibits significantly greater skill than conventional forecasting approaches in the context of a synthetic example consistent with the adopted paradigm. The example application demonstrates statistically significant skill at 5–10-yr lead times. Applications to operational long-range climatic forecasting are motivated and discussed.

1. Introduction

Conventional means of climatic (e.g., Niño3.4 index) forecasting have provided little evidence of the ability to exceed a simple null forecast model (e.g., a low-order site-specific autoregressive forecast) at lead times longer than about 6 months, especially for Niño3.4 (see the review by Barnston et al. 1994). As interannual-to-decadal and even longer term climate variability becomes better understood, however, there is increasing recognition of the potential for more skillful climatic forecasting at interannual and longer timescales. While the El Niño–Southern Oscillation (ENSO) is perhaps the best understood low-frequency global-scale mode of climatic variability, with quasi-oscillatory character within the 3–7-yr timescale range, that shows evidence of longer-range predictability (e.g., Cane et al. 1986; Kepenne and Ghil 1992; Kepenne and Lall 1996), a number of observational and theoretical investigations have substantiated evidence for significant quasi-oscillatory climatic processes at decadal (e.g., Dettinger et al. 1995; Ghil and Vautard 1991; Mann and Park 1993, 1994, 1996) and interdecadal-to-century scale (Kushnir 1993; Schlesinger and Ramankutty 1994; Mann et al. 1995b; Mann and Park 1994, 1996), which may be associated with extratropical oceanic or coupled ocean–atmosphere climatic modes in the Pacific (e.g., Latif and Barnett 1994; Trenberth and Hurrell 1994) and Atlantic (Delworth et al. 1993) basins, respectively. Such modes of variability, with their broad inferred atmospheric teleconnections, are now recognized as a potential source of long-term predictability of large-scale climatic variations (see, e.g., the discussions in Ghil and Vautard 1991; Latif and Barnett 1994; Mann and Park 1996). Mann and Park (1994, 1996) show that a small number of distinct quasi-oscillatory modes of variability, combined with two statistically significant secular variations, describe nearly half of the interannual and longer-term climatic spatiotemporal variance during the past century. A forecasting framework that harnesses such potentially predictable sources of variability for improved long-lead (i.e., annual-to-decadal lead) forecasting is developed here. Madden (1981) explored

* Lamont-Doherty Earth Observatory Contribution Number 5757.

Corresponding author address: Dr. Balaji Rajagopalan, 103 Oceanography, Lamont-Doherty Earth Observatory, Columbia University, P.O. Box 1000, Rt. 9W, Palisades, NY 10964-8000.
E-mail: rbala@rosie.ldgo.columbia.edu

© 1998 American Meteorological Society

early on the possibility of long-range climatic prediction by seeking to isolate the potentially "predictable" climatic variance associated with peaks in the spectra of climate records that exceed the estimated noise variance. Our approach takes a philosophically similar attack, using recent, relatively sophisticated multivariate statistical methods.

The conceptual model of the climate we adopt, motivated by inferences from the observational studies cited above, envisions a continuously evolving large-scale state of the system that, at interannual and longer timescales, is well described as the superposition of a small number of potentially episodic quasi-oscillatory processes with preferred timescales and spatial expressions, and a temporally and spatially autocorrelated noise process. In the studies of Mann and Park (1994, 1996), a multivariate frequency domain statistical decomposition was performed with century duration spatially distributed climate datasets to detect and reconstruct "quasi-oscillatory" spatiotemporal climate signals that exhibit episodes of spatially correlated behavior. The episodic nature of the oscillations was treated by allowing for phase discontinuities and/or amplitude variation in modeling the modulating time envelope of a spatially correlated oscillatory signal through a multiple taper frequency domain, singular value decomposition (MTM-SVD) approach (Mann and Park 1994), and it was further refined through the application of a nonstationary generalization of the original procedure (Mann et al. 1995b; Mann and Park 1996). An analogous treatment with alternative frequency and time-domain localization tradeoffs has been developed based on an analogous "multiwavelet" statistical model (Park and Mann 1998). In each case, the motivation behind the methodology is the assumption that climatic signals should exhibit spatially correlated (but not phase synchronous) oscillatory character with a well-defined timescale and finite lifetime. Such behavior may arise from excitation of an independent set of linear or quasi-linear eigenmodes of the climate system by spatially correlated background stochastic variations. This background noise is assumed to exhibit a colored spectrum [the simplest case of which is first-order autoregressive or "AR(1)" red noise] and to provide an additive noise background from within which signals must be detected.

Qualitatively similar behavior to that described by such a linear model could arise from self-sustained nonlinear climate processes in the absence of noise, or from a combination of both linear and nonlinear sources of variability. In practice, it can be quite difficult to distinguish between which model best describes the true climate system (see, e.g., Saltzman et al. 1981). Particularly if the degree of nonlinearity is weak, a nonlinear system may often behave linearly in local regimes of the dynamical phase space. In that sense, the linear-based methodology is not without utility, even if the underlying system exhibits some degree of nonlinear behavior. When the nonlinear system is in a quasi-linear region of phase space, linear forecasting techniques will perform well. Conversely, when the nonlinear system is near a fundamentally nonlinear region of its phase space (i.e., an unstable fixed point), the linear forecast should, in a statistical sense, be no worse than a random forecast. Since the truly nonlinear system is intrinsically unpredictable at such a point, such a forecast is no worse than any other forecast. Generalizations to nonlinear methods (e.g., Lall et al. 1996) are, nonetheless, a worthwhile extension of the work outlined here.

The model of linear eigenmodes superposed on an additive colored noise background is readily amenable to linear forecasting strategies. Given such a model, it is meaningful to forecast the trajectory of the system based on a linear parametric (e.g., low-order autoregressive) model of the past behavior of the system and prediction based on application of that parametric model when initialized with a present state (see, e.g., Kepenne and Ghil 1992). We seek to employ such an approach to forecasting, using the multivariate frequency-domain spatiotemporal decompositions described earlier to provide the basis vectors for the application of linear prediction techniques. In section 2, we provide a background discussion of relevant issues in multivariate signal detection and statistical forecasting to provide a framework for the motivation of our particular forecasting approach. In section 3, we describe in detail our approach to signal detection, reconstruction, and forecasting. In section 4, we demonstrate the effectiveness of our methodology when applied to a synthetic dataset constructed to have attributes similar to those inferred for the true climate. We conclude in section 5 with a discussion of prospects for improved long-range climate forecasting.

2. Background

The goal of the methodology presented here is to use mutual information among a large spatially distributed set of climate indicators to both enhance signal detection and reconstruction, and to allow for efficient forecasting of a large number of independent interrelated times series. Of particular relevance in the realm of large-scale climatic forecasting are basic meteorological diagnostics such as sea level pressure and surface temperature. Of equal if not greater importance from a societal perspective is the forecasting of hydroclimatic variables (e.g., precipitation and streamflow). A useful feature of our methodology is that no restrictions are placed on the spatial organization of the sites or the nature of the variables under analysis; the intent is to simultaneously consider a variety of dynamically connected climatic and/or hydroclimatic state variables. The integrity of particular climatic processes is maintained, further, by the recognition of significant temporal patterns of variation through a fre-

quency domain approach. Variables that tend to temporally integrate climatic forcings (e.g., streamflow or lake levels) may intrinsically emphasize lower-frequency components of the climate system. Likewise, processes with varying characteristic timescales may exhibit a definite spatial dependence in their relative prominence [e.g., in a rough sense, interannual ENSO events may be emphasised in the Tropics, while extratropical feedbacks tend to emphasize decadal timescale processes—see Trenberth and Hurrell (1994) and Latif and Barnett (1994)]. For these reasons, the simultaneous consideration of a large spatially distributed set of diverse climatic variables may permit a more robust recovery of climatic signals. Mann et al. (1995a) used a frequency-domain multivariate analyses of gridded Northern Hemisphere sea level pressure and surface temperature data to explain local hydrological variations that lead to the dramatic rise and fall of the Great Salt Lake in the 1980s in terms of structured low-frequency climatic variability. Moon and Lall (1996) also identify significant low-frequency variability of the Great Salt Lake connected to large climate variability such as ENSO, Pacific–North America teleconnection, etc. The influence of this structured large-scale forcing was shown to provide significant long-term forecast skill for the Great Salt Lake (Lall et al. 1996; Moon 1995).

a. Spatiotemporal modeling and forecasting

There are a variety of statistical models that could be used to described the spatiotemporal structure of large-scale climatic variability. Among the most elementary and well-studied statistical approaches is the multivariate, autoregressive (AR) model. Such a model could be used to describe the vector time series, $\mathbf{x}_t$, $t = 1, \ldots, N$ and is described by

$$\mathbf{x}_{t+1} = \sum_{j=1}^{p} \mathbf{A}_j \mathbf{x}_{t-j+1} + \mathbf{e}_{t+1}, \qquad (1)$$

where $\mathbf{x}_t$ is a vector of length M, M is the number of sites, $\mathbf{A}_j$ is a $M \times M$ matrix of autoregressive coefficients for a lag j, p is the order or maximum number of lags considered, and $\mathbf{e}_t$ is a residual vector of length M.

To specify this model, $p \times M \times M$ coefficients need to be estimated from the data, and an appropriate order p has to be selected. Clearly, this can lead to a very large number of parameters to be estimated from a limited amount of data. Several sites may be only weakly correlated with each other, leading to either near-zero or spurious elements in the matrices $\mathbf{A}_j$. As the number of series increases, these matrices can also be ill conditioned, since there may be little marginal information added by new series. Determination of the appropriate order, p, of the AR model in the multivariate setting can also be tenuous. These factors translate into an inability to properly identify such models and a consequent loss of predictability. Empirical orthogonal function (EOF) analysis, principal component analysis (PCA) or canonical correlation analysis (CCA) are related techniques that can be applied to the vectors $\mathbf{x}_t$ to identify a reduced set of variables, $\mathbf{y}_t$, of dimension $d \ll M$, which carry most of the significant spatial information in the entire dataset (e.g., Preisendorfer 1988). Each of these variables is orthogonal or independent by construction and is obtained via an SVD of the multivariate time series. Univariate AR models may be used to forecast each such component. Forecasting schemes based on these general ideas have been used for seasonal forecasts of meteorologic fields (e.g., the review by Barnston et al. 1994 and references therein). There are two types of problems with such schemes. First, the number of time ordinates, nt, needs to be greater than the number of series, M, for PCA to provide a unique set of eigenvectors and, hence, unique coordinates $\mathbf{y}_t$. This may not be feasible as M increases, for a fixed N. Second, while these methods work with the spatial structure of the time series at lag 0 (or sometimes for a few lags), there is no explicit attempt to exploit the frequency structure in time or space and time. CCA-based methods (e.g., described in Barnston 1994) may implicitly capture some of the frequency structure by stacking consecutive predictor seasons to provide a decomposition of the space–time correlation structures. Given that the different series may have incommensurate frequency structure, or at least exhibit lagging or leading behavior (i.e., shifts in phase) relative to each other, and that no attempt is made to reconcile this information, the resulting model often has a predictability horizon of only a few months. These considerations have motivated forecasting approaches based on extended EOF or multivariate correlation-space eigenvector methods (e.g., multivariate singular spectrum analysis or "MSSA"—see Vautard et al. (1996) and Kepenne and Lall (1996)] as a basis for decomposing the original data matrix into a reduced set of empirical orthogonal data vectors that have dimension $d \ll M$ and are based on a decomposition of the space and time correlation structure across the series. In this case, a decomposition is achieved by forming an embedded data matrix, $\mathbf{S}$, of dimension $[(W + 1)M](N - W)$ whose columns are the original data vectors and their W lagged copies {e.g., the $W + 1$ columns from the first data series x_t, $t = 1, \ldots, N$, would be $x(W + 1{:}N)$, $x(W{:}N - 1)$, $x(W - 1{:}N - 2)$, $\ldots$, $x(1{:}N - W)$}. This amounts to choosing a time window of length W across all the time series over which temporal correlation structure is considered. An SVD of $\mathbf{S}$ provides a decomposition of the space–time correlation structure across the time series. Significant space–time patterns corresponding to the leading eigenvectors (the ones that explain the most variance of $\mathbf{S}$) are identified and each of the original data series is reconstructed in terms of these basis functions. The

subsequent forecast is obtained as the sum of linear or nonlinear autoregressive forecasts of the individual reconstructed series. With this design, one can potentially capture long-range predictability that is manifest as slowly evolving time structures across the set of data series, using space and time information simultaneously. However, such methods can safely be used only to analyze quasi-oscillatory structures such that $W = N/3$ and for periods in the range $(W/5, W)$ (see Vautard et al. 1992). Consequently, there are rather severe restrictions on the range of frequency bands over which temporal structure can be simultaneously reconstructed. For instance, to recover interdecadal patterns (approximate period 20 yr), one may want to choose $W = 30$ yr with a 100-yr record. This window width will not allow the reliable decomposition of time structures corresponding to ENSO with periods of less than 6 yr. Furthermore, in the multivariate case, given a fixed duration, nt, the number of rows of $\mathbf{S}$ will quickly be smaller than the number of columns. A unique set of eigenvectors for decomposing $\mathbf{S}$ cannot then be identified.

Mann and Park (1996), Lall and Mann (1995), and Thomson (1982) note that correlation domain analyses, such as MSSA, are less efficient for identifying narrowband low-frequency structure than their frequency domain counterparts. The methodology developed in this paper is similar in spirit to the approach of Vautard et al. (1996), Sarda et al. (1996), and Kepenne and Lall (1996) but uses the MTM-SVD approach of Mann and Park (1994, 1996) for the spatiotemporal decomposition of the ns climate times series. The problems indicated above are circumvented, although this approach is not without its own set of complications and limitations.

b. Frequency domain MTM-SVD approach

The simultaneous use of space and frequency-domain information from multiple climate state variables provides a novel framework for decomposing a high dimensional dataset into components that are potentially predictable, while preserving spatially coherent structure. The frequency domain MTM-SVD methodology is described in detail elsewhere (e.g., Mann and Park 1994, 1996). We briefly review the approach, which involves the decomposition of a set of time series into its orthogonal spectral and spatial components by means of a singular value decomposition. In this manner, spatially coherent variability is sought within all distinct frequency bands that are resolvable in the dataset. Under the assumption that climate signals are reasonably narrowband in character, a spectral domain approach enhances signal detectability. A critical feature that is assumed of the data is that the power spectrum of the noise varies smoothly, so that it can be considered as roughly constant over the narrow bandwidth of the spectral estimates [this is the so-called locally white assumption discussed by Thomson (1982)]. Such an assumption holds trivially for the moderately "red" spectra that typify climatic data series (see, e.g., Mann and Lees 1996 and references therein). The noise is allowed to be spatially correlated. However, it is almost universally true that the noise spatial decorrelation scale is small compared to that of any climatic signals present in the data (e.g., ENSO) that exhibit large-scale teleconnections. This is the primary reason (see Mann and Park 1994, 1996) that multivariate frequency-domain analyses can increase the effective signal-to-noise ratio involved in signal detection and reconstruction.

To identify broadly coherent structure across a multivariate dataset, each constituent series is first standardized by converting it to an "anomaly" series if it is not already (i.e., removing its climatological seasonal cycle), removing the resulting series mean, and dividing by the standard deviation. The time series are transformed to the spectral domain using multitaper spectral analysis (see, e.g., Thomson 1982; Park et al. 1987; Mann and Lees 1996). For a given time series y, a set of K orthogonal data tapers and K associated tapered Fourier transforms or "eigenspectra" is determined at each frequency, f, through

$$Y_k^{(m)}(f) = \sum_{t=1}^{N} w_t^{(k)} y_n e^{i2\pi f t \Delta t}, \tag{2}$$

where $\Delta t = 1$ month is the sampling interval; $\{w_t^{(k)}\}_{t=1}^{N}$ is the kth member in an orthogonal sequence of tapers (also called Slepian tapers), $k = 1, \ldots, K$; $m = 1, \ldots, M$ are the number of sites; and N is the length of the time series. As a side note, the power spectral density at this frequency f can be calculated through an appropriately weighted average of the K eigenspectra.

The Slepian tapers provide statistically independent spectra estimates $Y_k^{(m)}(f)$, $k = 1, \ldots, K - 1$, leading to multiple degrees of freedom within a given narrow frequency band. This is the primary motivation behind using the multitaper decomposition, as alternative frequency domain decompositions of multivariate data [e.g., complex harmonic PCA—see Preisendorfer (1988)] make use of spectral estimates (adjacent values of the discrete Fourier transform or "DFT") that are statistically correlated. Each eigentaper is orthogonal and represents independent information in a frequency band of half-bandwidth of pf_R about a given frequency, f, where $f_R = (N\Delta t)^{-1}$ is the Rayleigh frequency (the minimum resolvable frequency range for the time series). Since a larger p averages over a greater bandwidth, the choice between K and p represents the fundamental trade-off between spectral resolution and degrees of freedom/variance. The choices $p = 2$ and $K = 3$ seem to provide an excellent compromise between frequency resolution and multiple degrees of freedom in the frequency domain in climate studies (Mann and

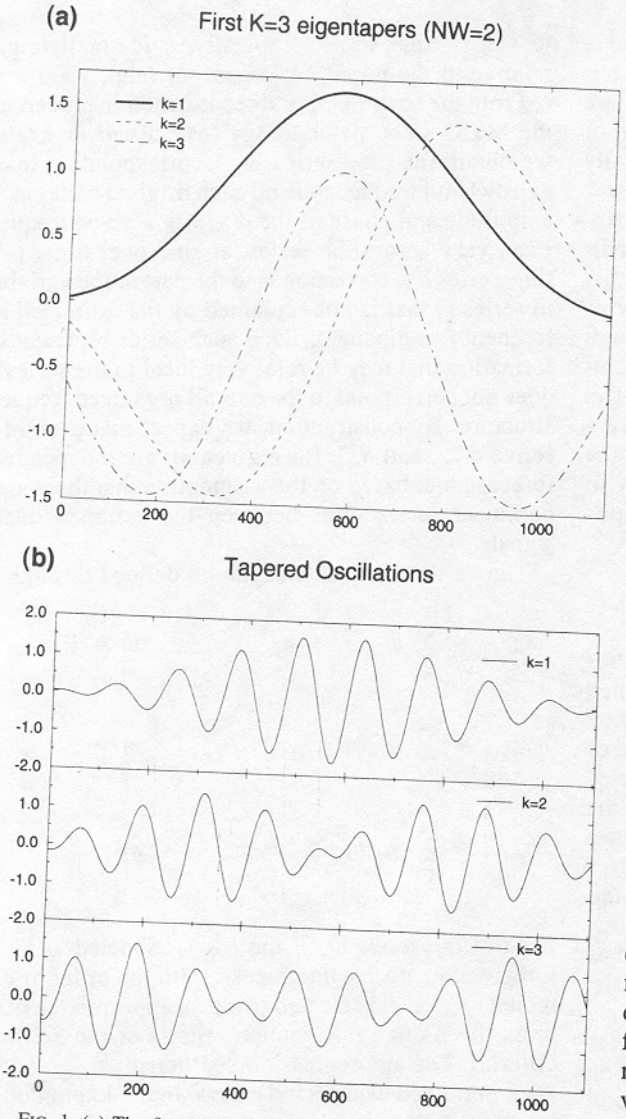

FIG. 1. (a) The first three orthogonal Slepian tapers for the case $K = 3$, $p = 2$ and (b) the components of a fixed-amplitude phase-coherent sinusoidal oscillation explained by each of the three data tapers.

Park 1994, 1996). Figure 1 shows the first $K = 3$ eigentapers for the case $p = 2$, and the component of a pure sinusoidal variation that is explained by each of the $K = 3$ eigentapers. It readily becomes evident that a variety of phase-/amplitude-varying oscillatory signals can be optimally modeled by a suitable combination of independently modulated components available in a multiple-taper decomposition. Rather than performing such an optimization in a univariate context (see e.g., Park and Maasch 1993), the MTM-SVD approach seeks to do so in a multivariate context.

This multivariate optimization requires us to form the $M \times K$ matrix $\mathbf{A}(f)$, from the K eigenspectra estimates [from (2) above] at all the M sites as

$$\mathbf{A}(f) = \begin{bmatrix} w_1 Y_1^{(1)} & w_1 Y_2^{(1)} & \cdots & w_1 Y_1^{(1)} \\ w_2 Y_1^{(2)} & w_2 Y_2^{(2)} & \cdots & w_2 Y_K^{(2)} \\ \cdots & \cdots & \cdots & \cdots \\ w_M Y_1^{(M)} & w_M Y_2^{(M)} & \cdots & w_M Y_K^{(M)} \end{bmatrix}, \quad (3)$$

where the w represent gridpoint-specific weightings to accommodate varying latitude of gridpoint sampling, etc. We perform a complex SVD of the above matrix through

$$\mathbf{A}(f) = \sum_{k=1}^{K} \lambda_k(f)\mathbf{u}_k(f) \otimes \mathbf{v}_k^*(f) \quad (4)$$

into its dominant modes, where λ_k describes the relative fraction of total variance explained by the kth mode, its associated left eigenvector $\mathbf{u}_k$ describes its complex spatial loadings, and the complex right eigenvector $\mathbf{v}_k$ describes the complex combination of the K eigenspectra that contain the information about the modulating envelope of the signal. This information is used to reconstruct the full spatiotemporal signal $\phi_{f,m,t}$, at each grid point m, and at frequency band f, which describes a phase- and amplitude-modulated quasi oscillation with spatially dependent phase and amplitude. For explicit details on reconstruction see Mann and Park (1994). Details regarding the boundary conditions in temporal reconstruction (see Mann and Park 1996) are discussed later. The spatial variance structure is restored to the previously standardized data series, upon reconstruction.

Outside the "secular band," which includes all frequencies unresolvable from trend ($f < pf_R$), we typically expect at most one significant signal will be present within a given narrow frequency band. Thus, the fractional variance explained by the *principal* eigenmode as a function of frequency or "local fractional variance" (LFV) spectrum is a useful normalized multivariate spectrum (which varies between $1/K$ and unity) within which prominent peaks indicate the presence of a statistically significant narrowband spatiotemporal signal. Under the assumption of a smoothly varying noise spectrum, the confidence levels on the LFV spectrum are uniform with frequency outside the secular band [within which they are somewhat higher—see Mann and Park (1994)] and can be estimated by either parametric methods [e.g., based on the assumption of an ensemble of N independent locally white time series, where N is the estimated number of spatial degrees of freedom in the noise—see Mann and Park (1994)] or through a nonparametric bootstrap approach wherein no assumptions need be made regarding spatial degrees of freedom, but serial correlation of the noise must be appropriately accounted for (see Mann and Park 1996). When possible, significant testing based on both approaches provides a more robust criterion.

Finally, it is often useful to employ an "evolutive" version of the above procedure in which the analysis is repeated in a large sequence of moving overlapping

windows through the multivariate data. In particular, if the LFV spectrum as determined above yields multiple significant peaks within a narrow band, a more appropriate statistical model may involve a single or small number of distinct components with temporally drifting frequency characteristics. Mann and Park (1996) show that this provides a parsimonious description of ENSO in terms of two distinct but period-drifting bands of interannual variability. For a signal that is most efficiently described in these terms, the temporal reconstruction of the signal is performed through an averaging of all overlapping segments of the temporal reconstructions from adjacent windows. The width of the window in the evolutive analysis is typically chosen so that it includes multiple periods of the oscillatory signal of interest but is short enough to capture the slower evolution of frequency and amplitude features over the duration of the record.

3. Forecasting methodology

Once a small set of statistically significant spatiotemporal signals (including trend) have been identified using the above procedure, each site or gridpoint time series is reconstructed by summing the contribution of each of these spatiotemporal components projected onto that site. For each site, the difference between the original series and summed signal component provides the "residual" series. Each signal component series and the residual series is then forecast forward using the time series forecasting approach discussed below, and a composite forecast at each site is determined as the sum of the component and residual series forecasts. Thus the slowly evolving components of the climate system that are dynamically consistent across the series analyzed are identified and forecast site by site. The short-term predictability associated with the stochastic (e.g., red noise) component is accounted for by forecasting the residuals through an optimally determined autoregressive process.

Let k denote the number of distinct narrowband spatiotemporal signals isolated and reconstructed by application of the MTM-SVD or evolutive MTM-SVD procedure described above to the series $x_{m,t}$. Consider the following decomposition of the data series $x_{m,t}$:

$$x_{m,t} = \sum_{i=1}^{k} \phi_{i,m,t} + h_{m,t}$$

$$t = 1, \ldots, N; \qquad m = 1, \ldots, M. \tag{5}$$

Here, the $\phi_{i,m,t}$ are time series that correspond to each of k decompositions of the frequency structure of x at grid point m. These decompositions represent low-frequency components (interannual and longer) of the climate system that are spatially coherent across the m series. The assumption is that if there is a quasi-oscillatory phenomena with a characteristic large-scale spatial structure, it is unlikely to be restricted to a small locale or time series. Typically, k is small (e.g., 5) relative to the number of series, M (e.g., 1000).

From the spatiotemporal reconstruction procedure of the MTM-SVD methodology (Mann and Park 1994), we obtain the time series $\phi_{i,m,t}$ corresponding to each narrowband frequency i and each original series m. The amplitude and phase of the $\phi_{i,m,t}$ for a given frequency i can vary across the series, m, and over time, t. The time series $h_{m,t}$ corresponds to the part of the variability in series m that is not explained by the estimated low-frequency component. Each such series represents information that may be relatively local to the series and does not correspond to the overall organized frequency structure. By construction, we expect that each of the series $\phi_{i,m,t}$ and $h_{m,t}$, for a given m, are independently forecastable, based on the assumption that there are no nonlinear interactions between the assumed distinct signals.

Thus a forecast scheme can be defined through

$$x_{m,t+1} = \sum_{i=1}^{k} \phi_{i,m,t+1} + h_{m,t+1} \qquad m = 1, \ldots, M, \tag{6}$$

$$\phi_{i,m,t+1} = \sum_{j=1}^{p_{i,m}} \alpha_{i,m,j}\phi_{i,m,t-j+1} + \beta_{i,m,t} \qquad m = 1, \ldots, M, \tag{7}$$

$$h_{m,t+1} = \sum_{j=1}^{q_m} \gamma_{m,j}h_{m,t-j+1} + \eta_{m,t} \qquad m = 1, \ldots, M. \tag{8}$$

Each of the series $\phi_{i,m,t}$ and $h_{m,t}$ is modeled as an autoregressive stochastic process, with the order of each model (p_{im} or q_m) selected using an appropriate criteria (e.g., the Akaike information criteria or the Schwartz criteria). The autoregression coefficients $\alpha_{i,m,j}$ and $\gamma_{m,j}$ are determined using standard maximum likelihood estimation methods. The model resulting from such a determination of coefficients is termed the best or "optimal" AR model. The residual series $\beta_{i,m,t}$ and $\eta_{m,t}$ are presumed to be independent and identically distributed for a given m and/or i.

A spatial ensemble jackknife procedure is used to assess the uncertainty in the forecasts for individual grid points and provides an important test of the robustness in estimated forecasting skill. A test for statistical consistency with respect to multiple ensemble noise realizations would represent a worthwhile extension of the robustness tests we have performed. Such calculations have not been performed due to the intensive nature of the analysis required for multivariate signal reconstruction and forecasting given just a single spatiotemporal noise realization.

4. Synthetic example and results

To demonstrate the utility of the methodology in a comprehensible setting, we construct a synthetic ex-

ample that shares many of the salient features of empirically inferred climate signals and noise. A synthetic grid represented by $5 \times 5 = 25$ grid points uniformly covering a synthetic global domain was defined. Three 100-yr-long signals of monthly resolution were chosen to represent the variety in spatial and temporal characteristics one expects to find (and hopes to be able to recover) in actual climatic signals. These spatiotemporal signals are superposed on underlying spatially correlated red noise of equal variance. The temporal model for the noise is a first-order autoregressive AR(1) noise process with lag-one correlation coefficient chosen to provide the roughly 1-yr noise decorrelation timescale estimated for a variety of surface temperature data (see, e.g., Mann and Lees 1996). The noise, further, is specified in such a way that it is spatially correlated at distances of one grid spacing but not beyond that, qualitatively mimicking the synoptic-scale noise spatial decorrelation scale observed in instrumental surface temperature data (see Madden et al. 1992; Mann and Park 1993).

The characteristics of the signals and noise are summarized in Table 1. The secular trend describes rapid "warming" (or, in the case of a couple of grid points, cooling) midway through the 100-yr period, an interdecadal (roughly 15-yr period) amplitude-modulated quasi-oscillatory signal with phase that varies spatially in such a way that the global average of the signal is identically zero at all times, and a slowly amplitude- and frequency-modulated interannual (3–5-yr period) oscillatory signal, with spatially variable phase, but a substantial global-average projection. The spatial patterns and associated reference site (the "central" grid point of the domain) temporal reconstructions are shown in Fig. 2. The length of the vector centered at a grid point indicates the amplitude of the corresponding signal, while its direction indicates its relative phase (i.e., the lag relative to the reference time signal (e.g., +90° indicates that the series is delayed by a quarter period or 1 yr for a 4-yr oscillation). Figure 3 shows the total "signal" component (the sum of the three signals) for

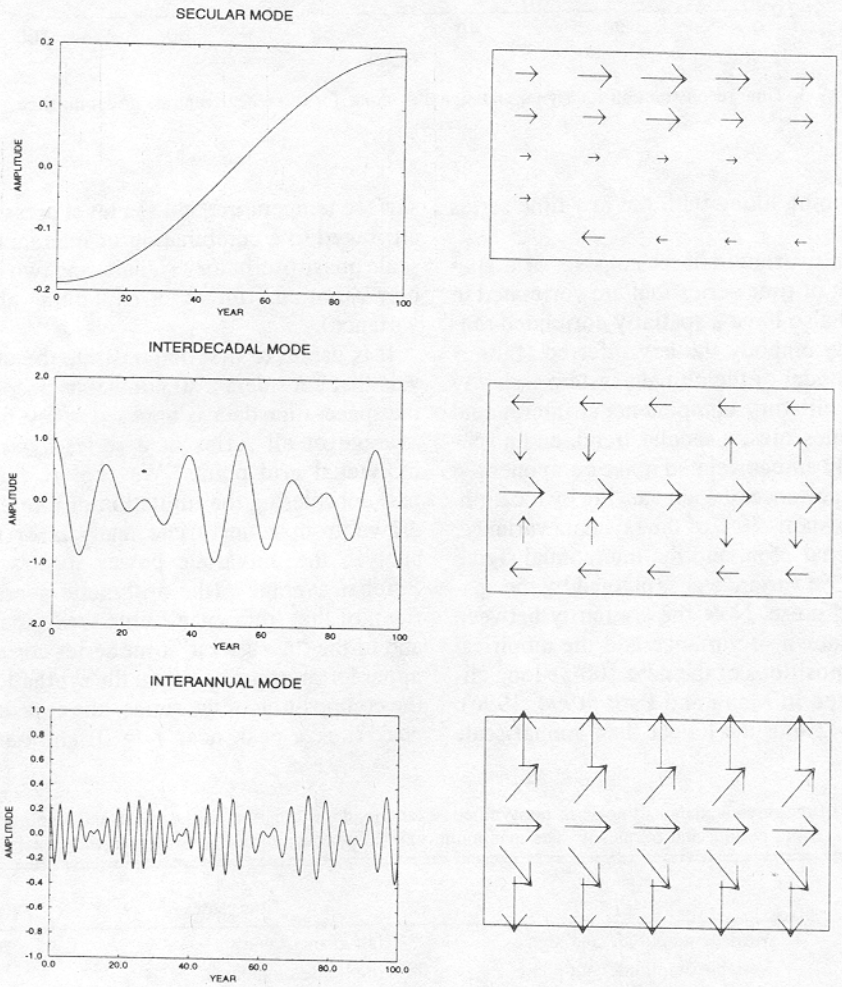

FIG. 2. Temporal (left) and spatial (right) patterns of synthetic signals showing secular, interdecadal, and interannual modes. Conventions are described in the text.

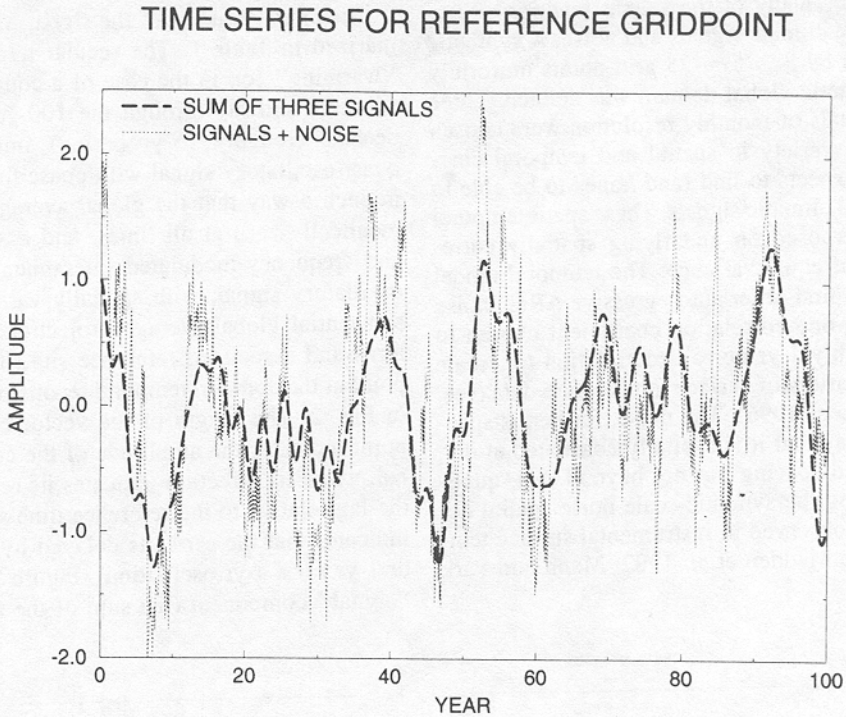

FIG. 3. Time reconstruction for signal and signal + noise for reference (central) gridpoint time series.

the reference grid point, along with the raw time series (signal + noise).

The resulting multivariate dataset consists of a spatially distributed set of time series that are correlated in time and space and also have a spatially correlated random character. They embody the key inferred features of our conceptual model of the climate system—slowly modulated quasi-oscillatory components at interannual and interdecadal timescales, a secular trend, and a spatially and temporally autocorrelated noise component—with the relative importance of each varying by location. The secular trend explains 56% of the raw data variance, the interdecadal signal 25%, and the interannual signal 8%. The residual 43% variance is explained by the spatially correlated red noise. Note the similarity between this imposed breakdown of variance and the empirical signal/noise decompositions of the near 100-yr-long climate records reported in Mann and Park (1994, 1996). In the latter study, 42% of the longer-than-annual-scale surface temperature and sea level pressure variance was attributed to a combination of interannual and decadal-scale quasi-oscillatory signals, and two independent secular variations (the latter explaining about 20% of the variance).

It is useful to first demonstrate the utility of an analysis that considers a simultaneous decomposition of the the space–time data as opposed to one based on a spatial average of all series or a series-by-series analysis of individual grid points. We explore this issue through first considering the limitation of inferences that can be drawn from a univariate analysis of the dataset. We analyze the univariate power spectra (Fig. 4) of the "global average" (the arithmetic average over all series), of the "reference" time series (center grid point), and of the "northwest" time series corresponding to the upper leftmost grid point in the synthetic domain. Given the composition of the series, one expects a secular band peak (i.e., a peak near $f = 0$) and narrowband peaks

TABLE 1. Description of the three signals and noise in the synthetic example, describing the spatial phase and amplitude pattern, temporal modulation, period, and frequency. The maximum regional peak amplitude is 1.0 for all the signals.

Signal	Signal pattern	Time pattern	T (yr)	f(cycles yr^{-1})
Secular trend	Variable amplitude and sign	Half-cosine trend	200	0.005
Interdecadal oscillation	Variable amplitude and phase	Amplitude modulated	15	0.065
Interannual oscillation	Uniform amplitude variable phase	Amplitude and frequency modulated	3–5	—
Red noise	Near-neighbor spatial correlation	AR(1) red noise	—	—

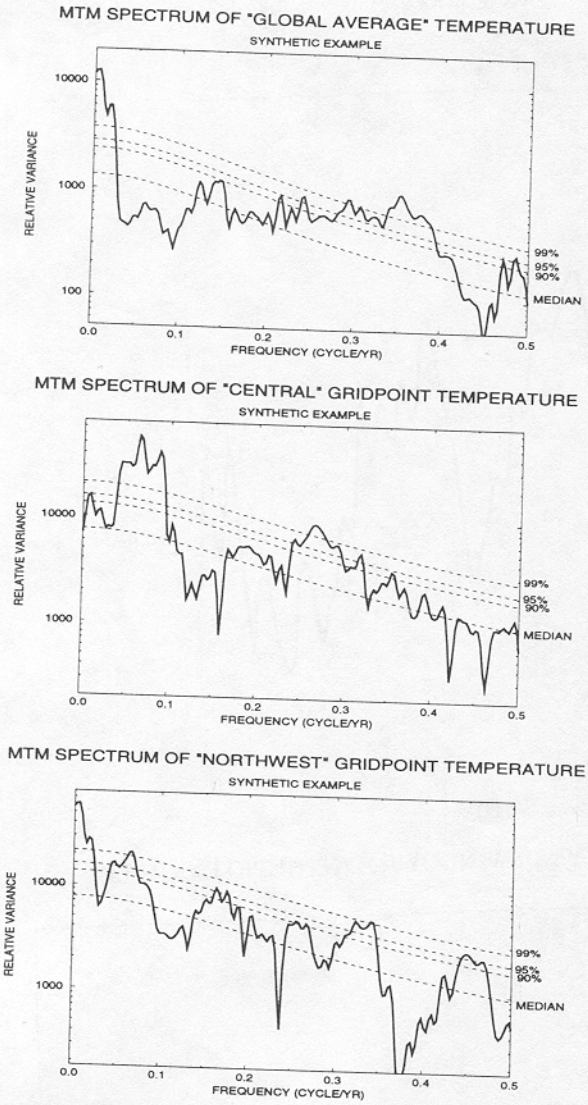

FIG. 4. Multitaper spectra of timeseries from (top) global average over domain, (middle) central grid point, and (bottom) extreme northwest grid point, along with robustly estimated median red noise level and 90%, 95%, and 99% confidence limits for significance relative to red noise.

centered near $f = 0.065$ cycles yr^{-1} (roughly 15-yr "interdecadal" period), and within the band $f = 0.33$–0.2 cycles yr^{-1} (3–5-yr drifting "interannual" period). The power spectra and confidence limits relative to a red noise null hypothesis are estimated using the robust method of Mann and Lees (1996). Since the phase of each oscillation varies over the spatial grid, there is some degree of cancellation of the oscillatory components in the process of forming an average across series. Note that the MTM spectra of the "global average" does not detect the interdecadal signal, constructed to vanish in a global average, as significant. Although interdecadal and interannual peaks are detected for the reference grid point, the secular trend, small at that grid point, is not recognized as significant at the >95% level. The secular trend and interdecadal peak are clearly detected in the spectrum for the northwest grid point, but it is difficult to identify any consistent interannual peak(s) based on the spectrum of this and the reference grid point, and there are several spurious peaks (sampling fluctuation from the noise background) that rival the true signals in their prominence. In short, spatially consistent signals cannot be properly identified (and clearly, thus, cannot be effectively reconstructed) based on the application of conventional univariate methods.

Now, we consider the inferences from the application of the multivariate MTM-SVD methodology to this dataset. In Fig. 5a we show the LFV multivariate spectrum for the synthetic dataset. The signals (secular trend: near-zero frequency, interdecadal peak centered near 0.065 cycle yr^{-1}, multiple peaks within the $f = 0.5$–0.33 cycle yr^{-1} band) are significant well above the 99.5% level, while there is no evidence of spurious peaks detected at this confidence level. Confidence bounds in this case can be precisely determined because the noise model has exactly nine spatial degrees of freedom and an underlying AR(1) red noise temporal correlation structure. An estimate from a bootstrap resample of the confidence limits yields estimates consistent with the estimates based on the assumption of nine independent locally white noise series. A key feature of the synthetic dataset is that signal and noise are predetermined and are thus exactly separable. This allows us to perform a parallel analysis of the noise-only component of the multivariate dataset (not shown). This analysis yields precisely the expectations from chance coincidence: roughly 1% of the spectrum exceeds the 99% confidence level for significance, roughly 5% exceeds the 95% confidence limit for significance, etc. No peaks are detected above the 99.5% significance level, suggesting that peaks that cross this threshold of significance in the full (signal + noise) time series are indeed unlikely to be spurious. In essence, because there are a maximum of 100 independent discrete frequencies in the range 0–0.5 under investigation, chance fluctuations at much greater than $p = 0.01$ (i.e., 99% level) are relatively unlikely to be encountered in this sample

FIG. 5. (a) Local fractional variance (LFV) spectrum of SVD based on full 100 yr of monthly data and (b) based on evolutive analysis with a 40-yr moving window. In the former case, 90%, 95%, and 99% confidence levels are shown with dashed line. For latter case, the LFV spectrum amplitude is indicated in gray scale, filtered at the 90% level of significance. Darker contrast indicates greater amplitude and significance.

→

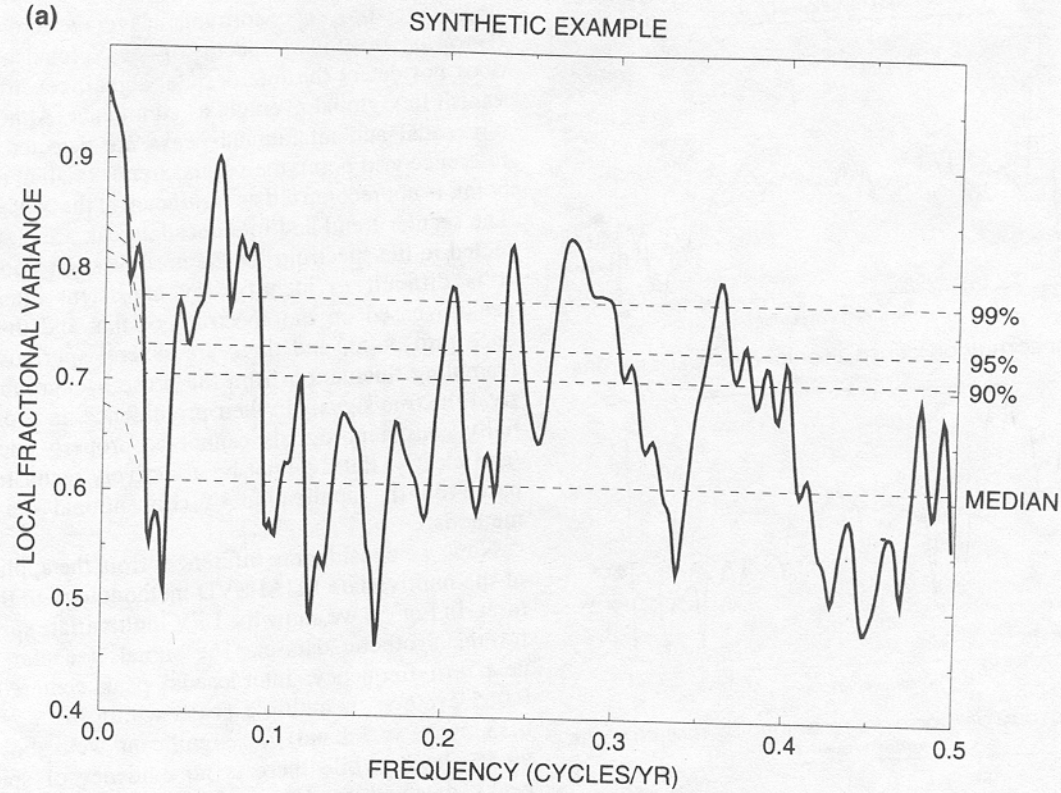
(a)
SYNTHETIC EXAMPLE
LOCAL FRACTIONAL VARIANCE
0.9
0.8
0.7
0.6
0.5
0.4
99%
95%
90%
MEDIAN
0.0
0.1
0.2
0.3
0.4
0.5
FREQUENCY (CYCLES/YR)

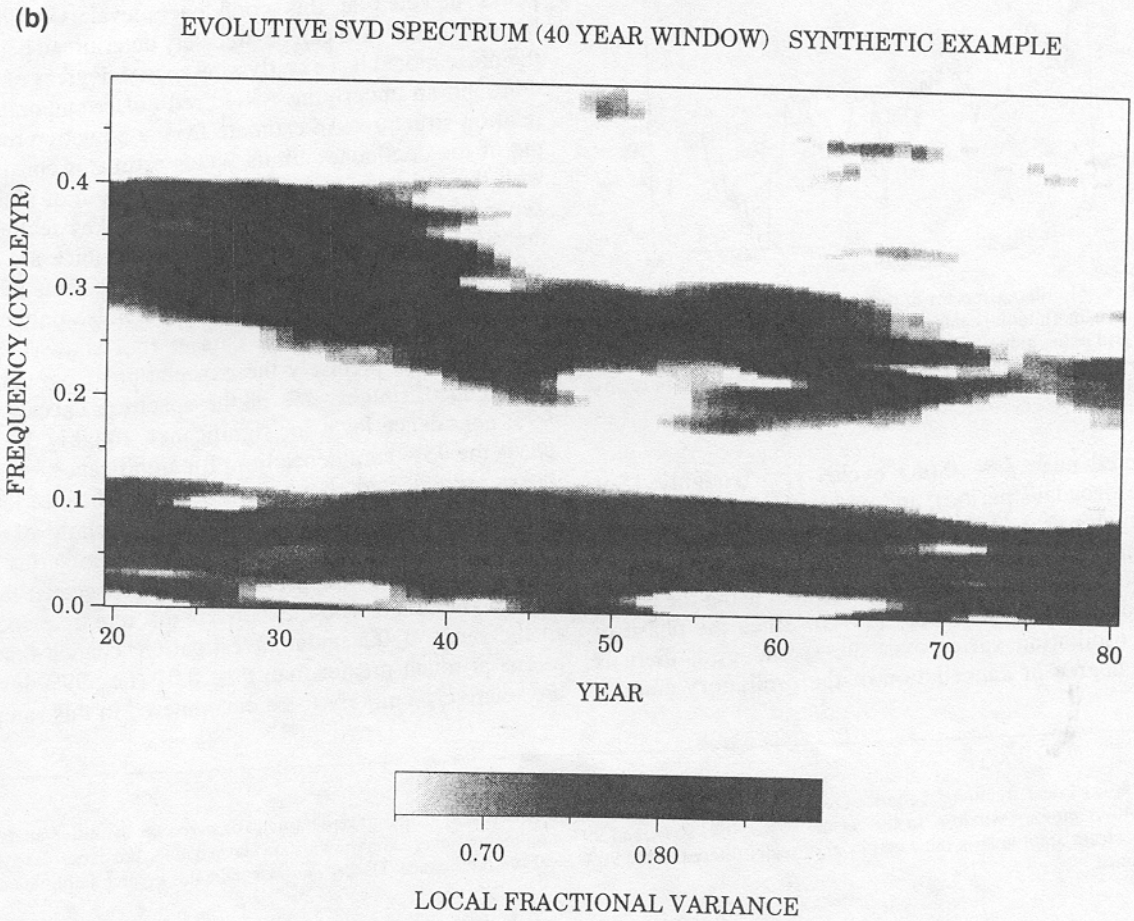
(b)
EVOLUTIVE SVD SPECTRUM (40 YEAR WINDOW) SYNTHETIC EXAMPLE
FREQUENCY (CYCLE/YR)
0.4
0.3
0.2
0.1
0.0
20
30
40
50
60
70
80
YEAR
0.70
0.80
LOCAL FRACTIONAL VARIANCE

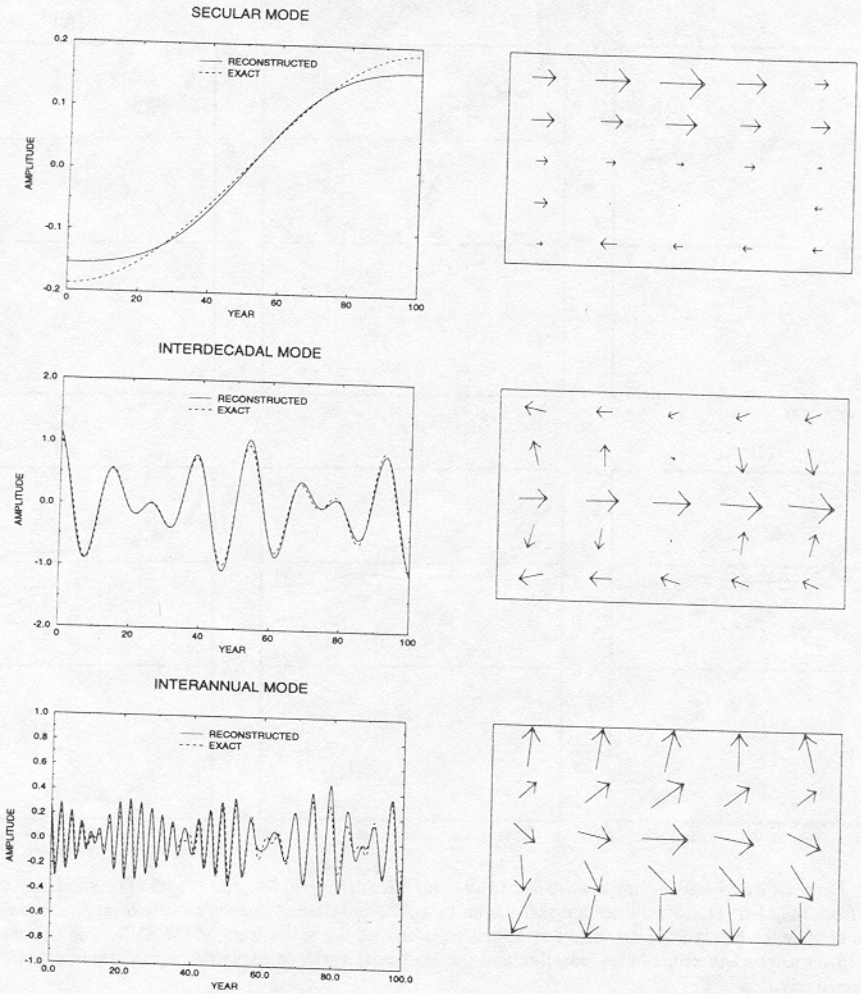

FIG. 6. Temporal (left) and spatial (right) patterns of reconstructed synthetic signals showing secular, interdecadal, and interannual modes. Actual reference time-domain signal (see Fig. 3) is shown by dashed curve for comparison to reconstructed reference time signal.

[see, e.g., the discussion of spurious peak detection of Thomson (1982)].

The multiple, closely spaced set of highly significant peaks in the $f = 0.2$–0.33 cycle yr^{-1} (3–5-yr period) range that were detected in the LFV spectrum are suggestive of a single frequency-evolving signal, as discussed earlier. Thus, it is useful to use an evolutive version of the analysis to see if a single frequency-modulated signal is indeed a more parsimonious model of the 3–5-yr period band variability. The evolutive SVD spectrum using a 40-yr window is shown in Fig. 5b. Note the drift over the record in the dominant band of interannual variance from being centered near $f = 0.35$ to near $f = 0.5$ cycle yr^{-1} over time. Amplitude modulation is also evident in the evolutive LFV spectrum itself, although a reliable estimate of the amplitude modulation is only possible through signal reconstruction.

The spatiotemporal structure of each of the three signals is shown in Fig. 6. The interannual signal is reconstructed based on the evolutive time-reconstruction approach described earlier. The left side of the figure shows the time series of the reconstructed and exact signal at the reference grid point, while the right side shows the spatial pattern in amplitude and phase of the reconstructed signal. Note that both the spatial and temporal patterns are quite close to the exact counterparts shown in Fig. 2. Note that small errors in both amplitude and phase are evident in the reconstructions. Such errors are associated with the inevitable contamination of detection and reconstruction by the underlying noise. Further, because the noise is spatially correlated, such errors (i.e., the phase errors evident in the spatial patterns) are also spatially correlated. In this sense, spatial correlation of noise can lead to short-range spatially coherent structure that is nonetheless spurious. However, such coherent noise structures are local and not global like the signal correlation structures.

Since the series were generated by adding a random red noise component, the dataset is a multivariate re-

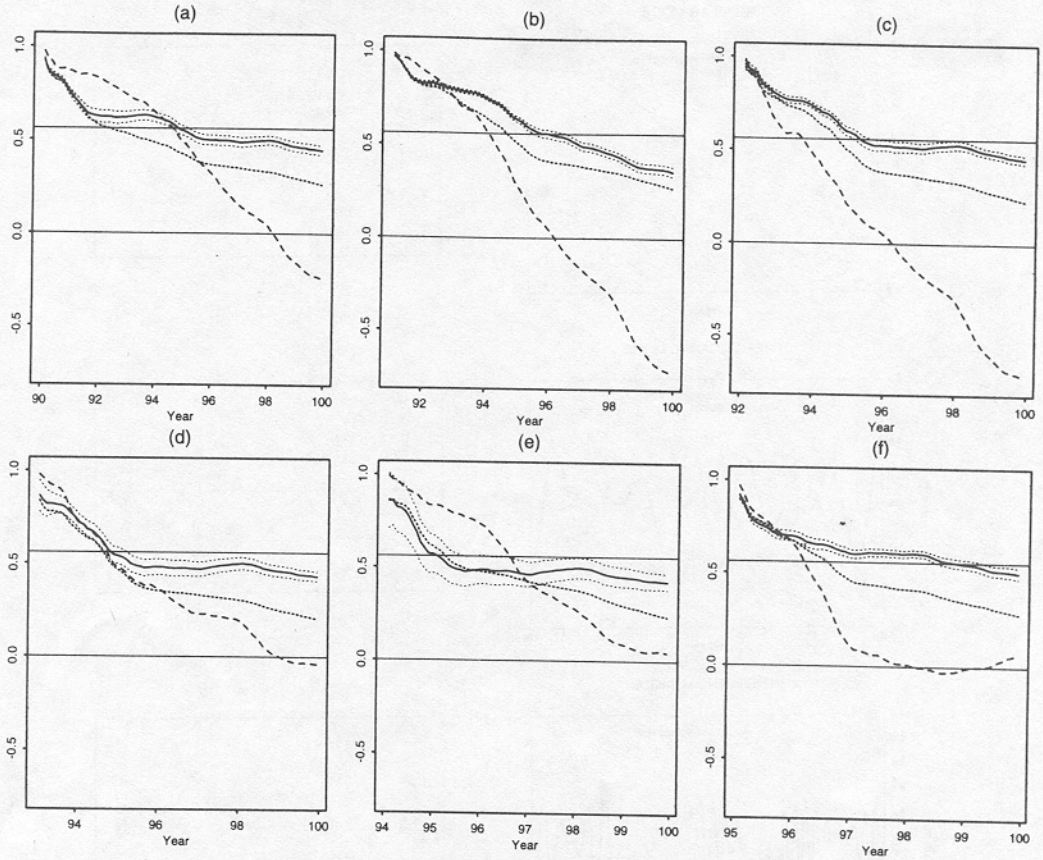

FIG. 7. Forecast skills starting in year (a) 90, (b) 91, (c) 92, (d) 93, (e) 94, and (f) 95. The solid line corresponds to skills from MTM-SVD, dotted lines are skills from best AR, and dashed lines are skills from persistence. The 95% jackknife confidence bands are also shown as dotted lines around the skills from MTM-SVD. The horizontal lines at 0 and 0.56 correspond to climatology baseline and the fractional variance explained by the three signals in the raw data, respectively.

alization of a stochastic process. The estimated spatial amplitudes and phases of the reconstructed signals are random variables whose bias and variance will determine the effectiveness of any subsequent forecasts. Note that significant biases can occur near the ends of the time series where temporal reconstruction is intrinsically less well constrained. Furthermore, when a signal is significantly amplitude modulated, its effective signal-to-noise ratio becomes prohibitively low when the modulating envelope diminishes or vanishes (e.g., near a phase discontinuity). Significant biases in the estimates of the gridpoint signal phase and amplitude will also generally be observed, due to the nature of the particular stochastic realization. Each of these problems, inevitable in any method of signal detection and reconstruction when noise levels are significant, will naturally impact the quality of any subsequent forecasts. An error in phase, for example, will lead to a fixed offset in the forecast for any lead time. An error in the amplitude has a similar connotation, while an error in the estimated frequency will lead to an increasing drift between the true and the estimated components and hence an increasing loss of forecast skill. An attempt is made to deal optimally with each of these types of problems in the context of signal reconstruction and forecasting, as described below.

First, consider the bias arising from phase discontinuities. Observe from Fig. 6 in the case of the temporal reconstruction of the interannual mode that the bias may be quite small subsequent to the phase discontinuity as the amplitude of the oscillation increases. Thus, there may be a temporary loss of predictability upon passage through such features. If a forecast begins near such a phase discontinuity the forecast will be poor. In the context of our conceptual model for the climate signal, this corresponds to the point where a particular excitation of a quasi oscillation has died out. The boundary bias discussed above may be ameliorated through the use of localized taper weighting procedures and constraints. We use (see Mann and Park 1996) a linear combination of zeroth-, first-, and second-order derivative constraints on the time envelope that minimizes the mean-square "error" with the actual data. Frequency drift bias is diminished by extending the temporal reconstruction optimization described above to a search throughout a significant narrowband peak detected in the multivariate

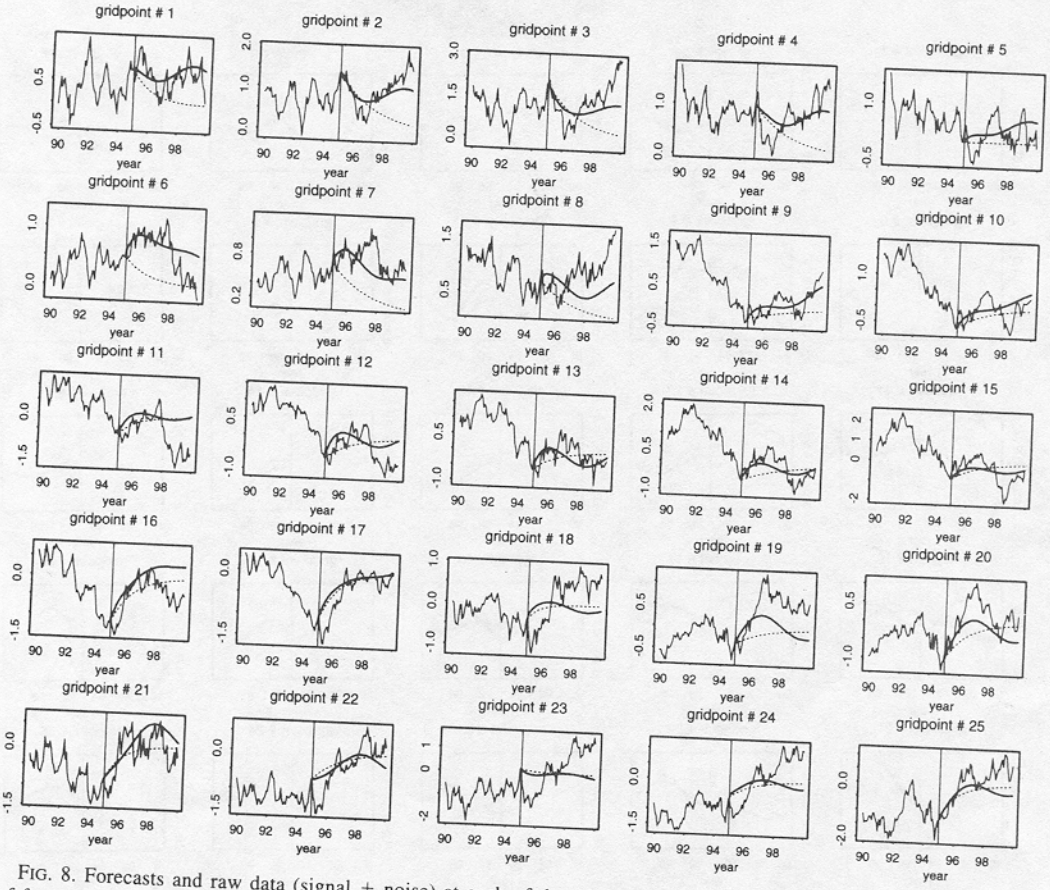

FIG. 8. Forecasts and raw data (signal + noise) at each of the 25 grid points. The vertical line indicates the start of forecast (i.e., in year 95), dotted lines are forecasts from best AR, heavy solid lines are those from MTM-SVD, and solid lines are the raw data.

spectrum, which may span several points of the DFT. In this way, frequency drift associated with a potential error in selecting the discrete frequency point of the DFT optimally associated with the narrowband signal is diminished. Note that some loss of information near boundaries is inevitable, and bias correction methods typically lead to increased variance in this context. This bias, along with errors in the estimated phase and amplitude pattern of the reconstructed signal, provides the primary barrier to perfect forecasting, as we shall see below.

Thus far, the multivariate dataset has been decomposed into a small number of statistically significant oscillatory components that offer potential long-lead predictability and a stochastic component that can provide some near-term (e.g., the roughly 1-yr lead skill expected from the imposed noise decorrelation timescale) predictability. The multivariate dataset has been forecast forward for between 10 and 5 yr (120–60 months), starting at years 90 to 95, respectively, using the optimal autoregressive prediction given by (6)–(8). We term this the "MTM-SVD" forecast. In each case, only data prior to the forecast start date is used to fit the MTM-SVD spectra and the forecast model. The subsequent data is used only to compare forecast performance for continuously increasing lead times. It is not used either for model fitting or for forecast initialization, and in this sense our forecasts are truly "blind" and indicative of real predictive skill. The MTM-SVD forecast is compared to two null forecasts: persistence (the forecast is the last value recorded for the series) and best or optimal AR [a univariate forecast based on (1), using only the time series at a given site; the order for the AR model is picked using Akaike's Information Criteria). The best AR in the synthetic example considered here generally resulted in a small order. Thus, long-range dependence indicated by the quasiperiodic components is not represented. In a classic time series modeling framework one would indentify and remove the nonstationarities and model the residual as an ARMA process. This is essentially what the MTM-SVD-based procedure tries to do. The best or optimal AR corresponds to a "naive" approach that is commonly used. Recall that the annual cycle was not included in any of the synthetic series since our purpose was to focus on interannual predictions. Consequently, seasonality considerations are not embodied in any of the forecasts. These can be readily accommodated. The performance

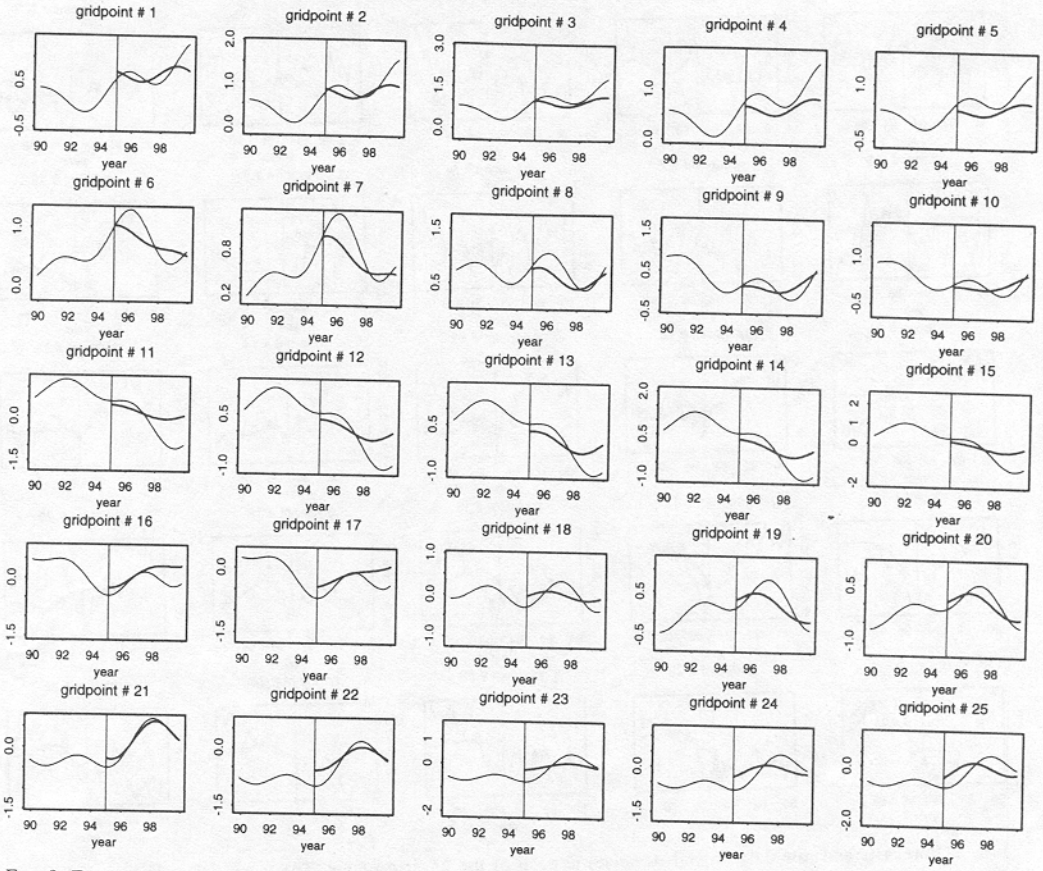

FIG. 9. Forecasts of the signals and the true signal at each of the 25 grid points. The vertical line indicates the start of forecast (i.e., in year 95), heavy solid lines are the forecasts from MTM-SVD, and solid lines are the true signal.

of each of the three forecast methods is evaluated using a skill measure that computes the fraction of the cumulative variance explained by the forecast. The skill $S_{f,\tau}$ is defined as

$$S_{f,\tau} = \frac{1}{M}\sum_{m=1}^{M}\left[1 - \frac{\sum_{n=1}^{\tau}(x_{m,t+n} - x^{f}_{m,t+n})^2}{\sum_{n=1}^{\tau} x^2_{m,t+n}}\right], \qquad (9)$$

where τ is the lead time of the forecast and $x^f_{m,t+n}$ is the forecast using method f (e.g., MTM-SVD, persistence, or best AR). Skills are evaluated each month for 10 yr for forecasts starting in year 90, for 9 yr for those starting in year 91, and for 5 yr for those starting in year 95. A jackknife procedure to provide a sense of spatial variation in the skill scores was also implemented. The skill is recomputed $M - 1$ times using $(M - 1)$ of the sites, omitting a different site each time. The standard deviation $S_{f,\tau}$ of these skill scores is computed for each τ and the approximate spatial variability of forecast skill at each lead time is characterized through the range $S_{f,\tau} \pm 2S_{f,\tau}$. The skills for the six sets of forecasts with different starting dates for each of the three methods are shown together with the jacknife error bounds for the MTM-SVD forecast in Figs. 7a–f. The following observations are apparent. The performance of the persistence forecast is highly variable as is to be expected. Typically, it has very poor skill as the forecast lead time increases. The best AR forecast is competitive with respect to the MTM-SVD forecast for short lead times. However, as expected, its relative skill drops off continuously as the lead time increases. The skill of the MTM-SVD forecast also decreases as a function of forecast lead time, albeit relatively slowly. The skill of the best AR forecast is below the jackknife skill variability limits for the MTM-SVD forecast as the forecast lead time increases beyond about 2 yr. It is rather interesting that the MTM-SVD forecast exhibits very nearly the same skill from 3 to 10 yr into the future. Recall that the skill scores used represent the average fractional cumulative variance explained by a forecast. So if the skill score does not change much, the forecast skill at the longer lead times is comparable to the skill early on. On the other hand, if the skill score continues to decrease (e.g., persistence or best AR), the actual marginal skill at long leads is substantially worse than the skill early on or relative to the forecast where skill is preserved. Also note that the skills from MTM-SVD in all the six cases are very close to the variance explained

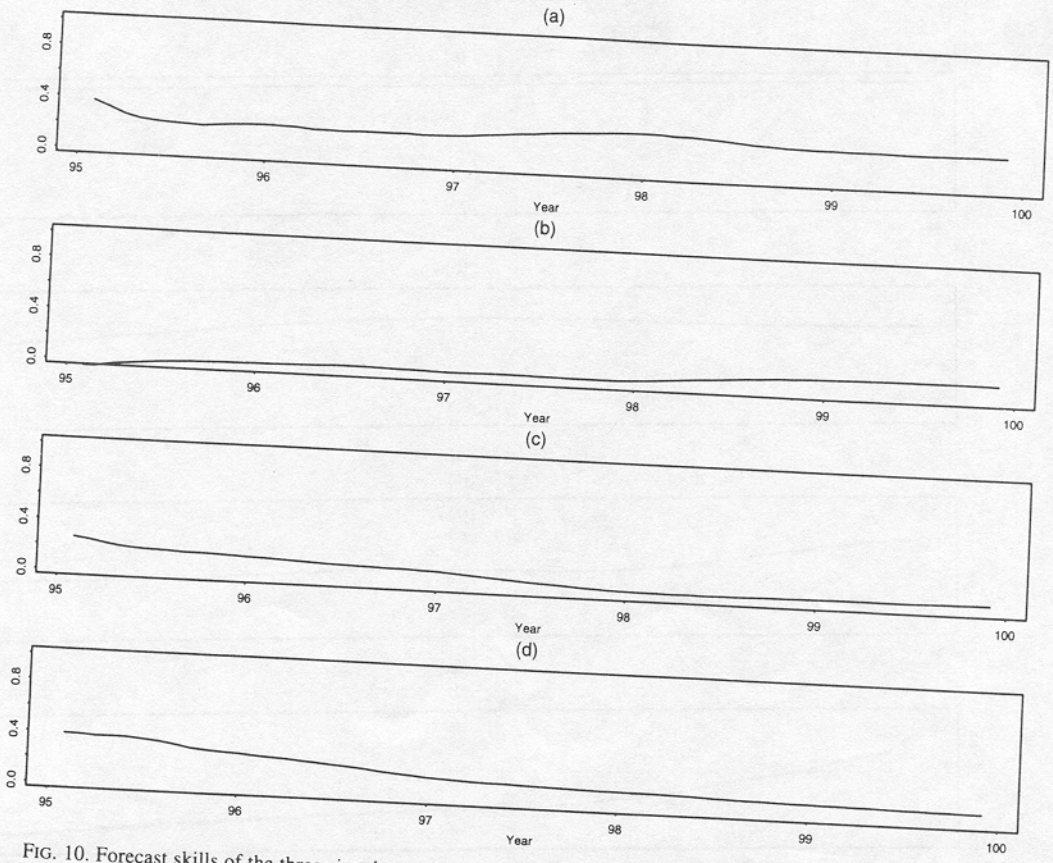

FIG. 10. Forecast skills of the three signals and the residual (noise), with respect to their respective true components, starting from year 95: (a) secular, (b) interdecadal, (c) interannual, and (d) residual (noise). Dotted line is the theoretical skill.

by the signals in the raw data (solid horizontal line at 0.56 in all the plots of Fig. 7) This suggests that the MTM-SVD performs well in terms of capturing the "signal variance" even at longer lead times.

It is instructive to view the actual forecasts to gauge more directly the quality of forecasts of individual series. The actual forecasts starting in year 95 for each of the 25 grid points are shown in Fig. 8. The following observations from these figures are apparent. The best AR and MTM-SVD forecasts for some series are quite similar. For others, the MTM-SVD forecast is invariably superior with a better representation of the oscillatory or quasi-oscillatory components in the data; it turns the corner where the AR forecast monotonically increases or decreases. However, there are significant biases in some cases in both the amplitude and phase of the MTM-SVD forecasts. The forecasts of the individual components $\phi_{i,m,t}$ were examined to understand why this was happening. Invariably, the explanation for a poor MTM-SVD forecast was a poor reconstruction of a particular narrowband signal. This is due to the instrinsic random error of signal phase and amplitude reconstruction, some unavoidable boundary constraint bias, and in part due to low envelope amplitudes and correspondingly poor signal–noise separation at certain forecast initialization times. In short, for a skillful forecast, it is essential to correctly estimate the curvature and magnitude of each reconstructed component near the initial forecast point. Some of this reconstruction error, however, can be picked up by the autoregressive model fit to the residual series $h_{m,t}$. However, the residual narrowband "signal" contribution is in this case inextricably mixed in with the red noise component, and the autoregressive model will not be able to effectively decompose the long- and short-term predictable components mixed into this residual; both contributions will be combined into a short-term "noise" component, albeit possibly with more persistence and skill than the pure red noise component alone. This is evident in the comparisons shown below. Further, the forecast of the residual component will be biased since the mean value of $h_{m,t}$ is biased relative to the mean appropriate for the boundary region. Formal techniques for automatically detecting and correcting for such errors still need to be developed. However, the average long-lead skill associated with the MTM-SVD forecast procedure is quite remarkable and is the source of some optimism in terms of operational long-lead climatic forecasting. Similar

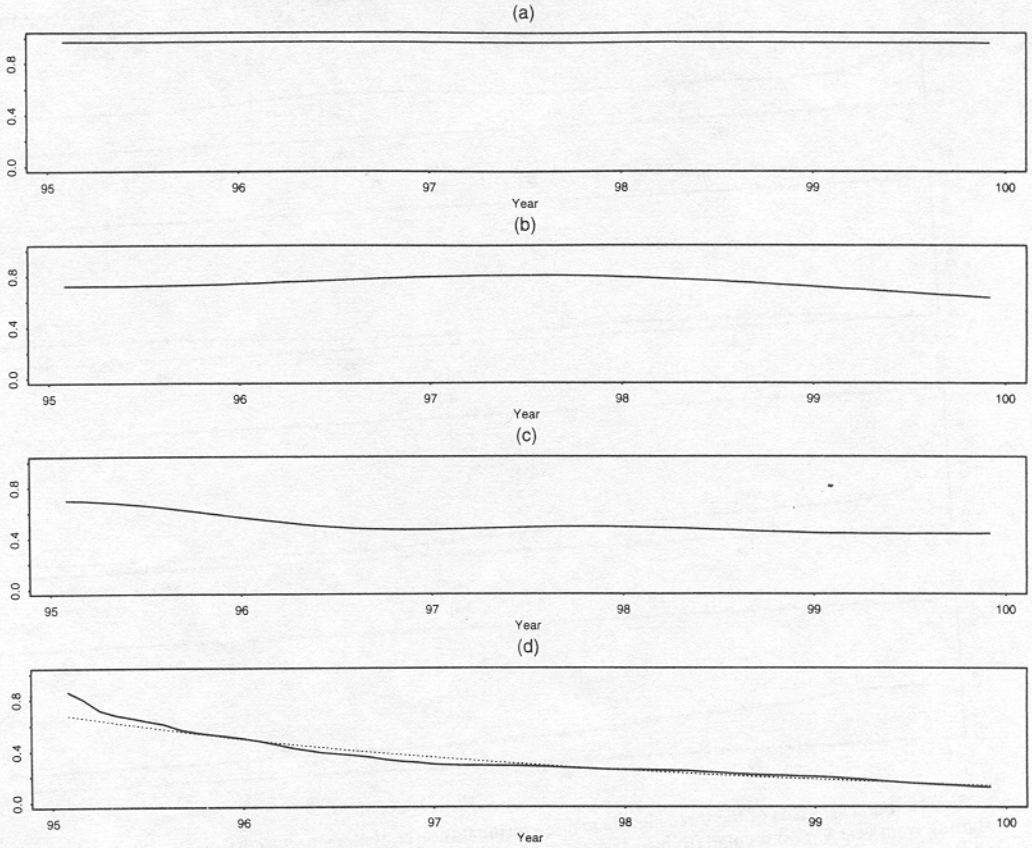

FIG. 11. Forecast skills with respect to the raw data (signal + noise) of the three signal components and the noise component, starting from year 95: (a) secular, (b) interdecadal, (c) interannual, and (d) residual (noise).

observations can be made from Fig. 9, which shows the MTM-SVD forecasts of the signal alone.

One important, and thus far unanswered question, is how much of the skills reported above are due to long-lead "signal" contribution, and how much due to the short-term "noise" residual component forecast. To address this question, we consider the skills of the forecasts initiated at year 95 for each of the three reconstructed signal components (secular, interdecadal, and interannual) relative to the "noiseless" component of the data (i.e., the residual between the raw synthetic data and the true specified noise process) and the residual component (Fig. 10), as well as the corresponding skills relative to the raw synthetic data series itself (Fig. 11). This comparison provides a measure of the skill of the forecasts as well as of the level to which the apparent forecast skill is degraded when measured against combined signal + noise background variance (Fig. 11). Several points become immediately evident in this comparison. The reconstructed signals are forecast quite effectively as compared against the known total signal contribution (Fig. 11), with skills that degrade, not suprisingly, with a timescale that is roughly proportional to the signal's own intrinsic timescale. This is a known argument that has been shown to hold true for 10-day weather forecasts as well by Van den Dool and Saha (1990). Each of the signals are forecast with a skill score in excess of 0.45 for the entire 5-yr forecasting interval (note that a score of 0.5 represents the midpoint between the skills of climatological and perfect forecasts). The forecast skill of the residual component between the raw data and reconstructed signals drops off precisely as we expect for red noise with a 1-yr autocorrelation timescale. In this case, the low-frequency spatially correlated noise ensemble can considerably obscure the contributions of the predictable components. The interannual and noise or "residual" component forecasts drop off rapidly relative to the lower-frequency signal-based forecasts. The interannual signal explains a relatively small fraction of skill, but the skill drops less quickly with lead time than the residual-component skill, indicating that something is gained by identifying the interannual signal as a quasi-oscillatory component distinct from the noise residual. It is noteworthy that the skills of the secular and interdecadal components (the low-frequency signals) either remain consistent or increase as a function of increasing lead time, even measured in the presence of the noise background.

5. Conclusions

We have introduced a synthetic example that encompasses one paradigm for the spatiotemporal signal and error structure of observational climate data at interannual and longer timescales. This paradigm assumes a climatic system with spatially correlated and temporally autocorrelated "red noise" upon which oscillatory, perhaps stochastically, excited quasi-oscillatory low-frequency modes are superposed. In the context of this paradigm for climate variability, the methods described in this article appear to provide impressive long-term (i.e., interannual to decadal lead) forecasting skill, far exceeding the skill available by conventional means (i.e., optimal autoregressive forecasting). For these reasons, we are optimistic about the usefulness of this methodology as applied to operational long-lead climatic forecasting. To the extent that nonlinear variability is is important on the timescales of interest, the linear methods invoked can provide an at best imperfect model of observed climate variability, and an accordingly imperfect basis for forecasting. Nonlinear generalizations of the forecasting methodology presented are worthy of investigation. Applications to actual climate data are the logical follow-up to the study described here and will be forthcoming.

Acknowledgments. This work was supported by NOAA Grants UCSIO CU0155 6601D and NA36GP0074-02 (BR), the National Science Foundation (MEM), and the Department of Energy (MEM). MEM acknowledges support through the Alexander Hollaender Distinguished Postdoctoral Research Fellowship program of the Department of Energy.

REFERENCES

Barnston, A. G., 1994: Linear statistical short-term climate predictive skill in the Northern Hemisphere. *J. Climate,* **7,** 1513–1564.

——, and Coauthors, 1994: Long-lead seasonal forecasts—Where do we stand? *Bull. Amer. Meteor. Soc.,* **75,** 2097–2109.

Cane, M., S. E. Zebiak, and S. C. Dolan, 1986: Experimental forecasts of El Niño. *Nature,* **321,** 827–832.

Delworth, T., S. Manabe, and R. J. Stouffer, 1993: Interdecadal variations of the thermohaline circulation in a coupled ocean–atmosphere model. *J. Climate,* **6,** 1993–2011.

Dettinger, M. D., M. Ghil, and C. K. Kepenne, 1995: Interannual and interdecadal variability in United States surface-air temperatures, 1910–1987. *Climate Change,* **31,** 33–66.

Ghil, M., and R. Vautard, 1991: Interdecadal oscillations and the warming trend in global temperature time series. *Nature,* **350,** 324–327.

Keppenne, C. L., and M. Ghil, 1992: Adaptive filtering and prediction of the Southern Oscillation index. *J. Geophys. Res.,* **97,** 20 449–20 454.

——, and U. Lall, 1996: Complex singular spectrum analysis and multivariate adaptive regression splines applied to forecasting the Southern Oscillation. *Experimental Long Lead Forecast Bulletin,* A. G. Barnston, Ed., NOAA/CPC, 54–56.

Kushnir, Y., 1994: Interdecadal variations in the North Atlantic sea surface temperature and associated atmospheric conditions. *J. Climate,* **7,** 141–157.

Lall, U., and M. E. Mann, 1995: The Great Salt Lake: A barometer of low-frequency climatic variability. *Water Resour. Res.,* **31,** 2503–2515.

——, T. Sangoyomi, and H. D. I. Abarabanel, 1996: Nonlinear dynamics of the Great Salt Lake: Nonparametric short term forecasting. *Water Resour. Res.,* **32,** 975–985.

Latif, M., and T. P. Barnett, 1994: Causes of decadal variability over the North Pacific and North America. *Science,* **266,** 634–637.

Madden, R. A., 1981: A quantitative approach to long-range prediction. *J. Geophys. Res.,* **86**(C10), 9817–9825.

——, D. J. Shea, G. W. Branstator, J. J. Tribbia, and R. O. Weber, 1993: The effects of imperfect spatial and temporal sampling on estimates of the global mean temperature: Experiments with model data. *J. Climate,* **6,** 1057–1066.

Mann, M. E., and J. Park, 1993: Spatial correlations of interdecadal variation in global surface temperatures. *Geophys. Res. Lett.,* **20,** 1055–1058.

——, and ——, 1994: Global scale modes of surface temperature variability on interannual to century time scales. *J. Geophys. Res.,* **99,** 25 819–25 833.

——, and J. Lees, 1996: Robust estimation of background noise and signal detection in climatic time series. *Climate Change,* **33,** 409–445.

——, and J. Park, 1996: Joint spatiotemporal modes of surface temperature and sea level pressure variability in the Northern Hemisphere during the last century. *J. Climate,* **9,** 2137–2162.

——, U. Lall, and B. Saltzman, 1995a: Decadal-to-century scale climate variability: Insights into the rise and fall of the Great Salt Lake. *Geophys. Res. Lett.,* **22,** 937–940.

——, J. Park, and R. S. Bradley, 1995b: Global interdecadal and century-scale oscillations during the past five centuries. *Nature,* **378,** 266–270.

Moon, Y., 1995: Large scale atmospheric variability and the Great Salt Lake. Ph.D. dissertation, Utah State University, 140 pp.

——, and U. Lall, 1996: Interannual and interdecadal variability in the Great Salt Lake and selected atmospheric indices. *ASCE J. Hydrologic Eng.,* **1**(2), 55–62.

Park, J., and K. A. Maasch, 1993: Plio-Pleistocene time evolution of the 100-kyr cycle in marine paleoclimate records. *J. Geophys. Res.,* **92,** 12 675–12 684.

——, and M. E. Mann, 1998: Interannual temperature events and shifts in global temperature: A multiple wavelet correlation approach. *Earth Interactions,* in press.

——, C. R. Lindberg, and F. L. Vernon III, 1987: Multitaper spectral analysis of high-frequency seismograms. *J. Geophys. Res.,* **92,** 12 675–12 684.

Preisendorfer, R. W., 1988: *Principal Component Analysis in Meteorology and Oceanography.* Elsevier, 425 pp.

Saltzman, B., A. Sutera, and A. Evenson, 1981: Structural stochastic stability of a simple autooscillatory climatic feedback system. *J. Atmos. Sci.,* **38,** 494–503.

Sarda, H., G. Plaut, C. Pires, and R. Vautard, 1996: Statistical and dynamical long-range atmospheric forecasts: Experimental comparison and hybridization. *Tellus,* **48A,** 518–537.

Schlesinger, M. E., and N. Ramankutty, 1994: An oscillation in the global climate system of period 65–70 years. *Nature,* **367,** 723–726.

Thomson, D. J., 1982: Spectrum estimation and harmonic analysis. *Proc. IEEE,* **70,** 1055–1096.

Trenberth, K. E., and J. W. Hurrell, 1994: Decadal atmosphere–ocean variations in the Pacific. *Climate Dyn.,* **9,** 303–319.

Van den Dool H., and S. Saha, 1990: Frequency dependence in forecast skill. *Mon. Wea. Rev.,* **118,** 128–137.

Vautard, R., P. Yiou, and M. Ghil, 1992: Singular spectrum analysis: A toolkit for short, noisy and chaotic series. *Physica D,* **58,** 95–126.

——, G. Plaut, and C. Pires, 1996: Long-range atmospheric predictability using space–time principal components. *Mon. Wea. Rev.,* **124,** 288–307.