Identifying Low-Dimensional Nonlinear Behavior in Atmospheric Data

D. A. S. PATIL AND BRIAN R. HUNT

Department of Mathematics, and Institute for Physical Science and Technology, University of Maryland at College Park, College Park, Maryland

JAMES A. CARTON

Department of Meteorology, University of Maryland at College Park, College Park, Maryland

(Manuscript received 31 July 2000, in final form 22 November 2000)

ABSTRACT

Computational modeling is playing an increasingly vital role in the study of atmospheric–oceanic systems. Given the complexity of the models a fundamental question to ask is, How well does the output of one model agree with the evolution of another model or with the true system that is represented by observational data? Since observational data contain measurement noise, the question is placed in the framework of time series analysis from a dynamical systems perspective. That is, it is desired to know if the two, possibly noisy, time series were produced by similar physical processes.

In this paper simple graphical representations of the time series and the errors made by a simple predictive model of the time series (known as residual delay maps) are used to extract information about the nature of the time evolution of the system (in this paper referred to as the dynamics). Two different uses for these graphical representations are presented in this paper. First, a test for the comparison of two competing models or of a model and observational data is proposed. The utility of this test is that it is based on comparing the underlying dynamical processes rather than looking directly at differences between two datasets. An example of this test is provided by comparing station data and NCEP–NCAR reanalysis data on the Australian continent.

Second, the technique is applied to the global NCEP–NCAR reanalysis data. From this a composite image is created that effectively identifies regions of the atmosphere where the dynamics are strongly dependent on low-dimensional nonlinear processes. It is also shown how the transition between such regions can be depicted using residual delay maps. This allows for the investigation of the conjecture of Sugihara et al.: sites in the midlatitudes are significantly more nonlinear than sites in the Tropics.

1. Introduction

The early work of Lorenz (1963) showed that a simple convective model could generate extremely complicated behavior such as sensitive dependence on initial conditions (''chaos''). Current models of the atmosphere tend to be much more complex, often involving coupled systems of partial differential equations with stochastic terms and reanalysis techniques. Still, they have the potential to exhibit behavior similar to that which Lorenz observed. Due to the level of complexity of these model systems, it is difficult to study their behavior analytically.

Thus we are left with the question of how one characterizes the behavior (which we will call the *dynamics*) of models in relation to one another and to the true system. For example, if we say that the underlying dynamics of a system are linear, then we mean that the system can be modeled by a set of equations for which the nonlinear terms are insignificant compared to the linear terms and external input. The question that we will be concerned with is the following: is it possible to compare the dynamics of different time series, especially when at least one of the time series includes observational noise?

This question is not unique to atmospheric sciences and has been studied in depth by the nonlinear dynamics community [see Kantz and Schreiber (1997) and references therein]. Some results of this research have been successfully applied in the atmospheric sciences (Jin et al. 1993, 1994; Tziperman et al. 1994). This has led to the conjecture that while the atmospheric system has a state space of high dimension, the dynamics takes place on an attractor of significantly lower dimension. That is, while all possible physical behaviors can be explained by a system with a large number of degrees of freedom, the behavior exhibited by the system can be accurately described by a small number of variables at

Corresponding author address: Dhanurjay A. Patil, Department of Mathematics, University of Maryland at College Park, College Park, MD 20742.
E-mail: dpatil@ipst.umd.edu

any given time (Mo and Ghil 1987; Selten 1993; Sugihara et al. 1999).

While this conjecture can be tested on models, it is fundamentally difficult to test using observational data (barometric pressure, temperature, etc.). The reason is that observational data is both noisy and usually of a limited sample size, while many of the techniques that would be required to investigate the conjecture (Lyapunov exponents, fractal dimension, and other statistical quantities known as dynamical invariants) require large amounts of data with little noise. Generally speaking, the amount of data required to estimate the dimension of an attractor grows exponentially with the embedding dimension.

Sugihara et al. (1999) proposed an alternative method for analyzing such data. Instead of trying to isolate dynamical invariants, they proposed using a qualitative approach to classify the underlying dynamics (e.g., linear vs nonlinear, etc.). The general principle is to look for a relationship between the data and the errors of some simple predictive model of the data. Graphical representations of this relationship are called residual delay maps (hereafter referred to as RDMs). While RDMs have been known to the financial community as a simple analytical tool, Sugihara et al. (1999) and Casdagli et al. (1996) demonstrated that residual delay maps can be used for the extraction of dynamical information from noisy data.

The motivation for this method is that if a short-term predictive model accurately describes the dynamics that produced the data, then errors between the data and the results from the predictive model should exhibit no relationship. In this paper we are primarily concerned with the identification of the low-dimensional, nonlinear components in the dynamics; thus we chose our predictive model to be linear combination of previous values in the data (typically known as an autoregressive model). The RDM is a lagged plot of the residuals (errors between the data and the predictive model) as a function of the previous data. If the data were the result of some linear process, then there should be no statistical relationship between the variables in the RDM. In the same manner, if the dimensionality is sufficiently large, the data can appear to be noise resulting in no observable relationship in the RDM. However, if there are low-dimensional nonlinear processes that are dominating the dynamics, then a nonlinear statistical relationship may be observed in the RDM, indicating that the linear predictor could be improved by adding nonlinear terms. This approach contrasts with the usual empirical orthogonal eigenvector approach based on covariance matrices that apply most naturally to linear systems. The purpose of this paper is to expand on the framework presented by Sugihara et al. (1999) and to provide concrete applications of RDMs for the analysis of atmospheric data.

Synoptic disturbances that characterize weather have timescales of a few days and approximately linear geostrophic dynamics with a surface streamfunction proportional to air pressure. At the same time, we anticipate that nonlinearity will play a role since, for example, low pressure systems tend to be more intense than high pressure systems. We begin examining these effects using the RDM approach with the same set of 25 Australian daily averaged surface air pressure time series as considered by Sugihara et al. (1999), but with refined statistics.

We next explore the use of RDMs in comparing time series, using as an example the Australian station RDMs and corresponding RDMs of time series from the National Centers for Environmental Prediction–National Center for Atmospheric Research (NCEP–NCAR) reanalysis project (Kalnay et al. 1996). The reanalysis represents a middle ground between the station data, which is assimilated, and a pure numerical simulation. Differences between the station data and the reanalysis may result from several causes including spatial and temporal interpolation, data rejection, and poorly resolved or inaccurate model physics. Indeed, in their earlier study, Sugihara et al. (1999) found a number of differences between the station data and a shorter record of European Centre for Medium-Range Weather Forecasts (ECMWF) forecasts.

In the third part of our results we expand our view to consider RDMs over the entire globe. Here we identify regions of nonlinearity by qualitatively identifying low-dimensional dynamics in the time series of individual stations. This approach presents an alternative to evaluating the magnitude of advective terms that evaluates linearity. We use the approach detailed in section 3 to explore Sugihara et al.'s (1999) conjecture that the midlatitudes are significantly more nonlinear than the Tropics. Since the structure of the synoptic wind field changes dramatically with height, we explore this by repeating the analysis on tropopause-level air pressure.

2. Data

Our station (observational) dataset consists of observed daily averaged barometric pressure spanning roughly 37 years from 25 sites throughout the Australian continent. This is the same dataset that Sugihara et al. (1999) used in their preliminary residual delay map analysis.

We also examine daily averaged surface air pressure and pressure at tropopause heights from the NCEP–NCAR reanalysis project (Kalnay et al. 1996). The data are on a $2.5° \times 2.5°$ grid globally, four times daily from 1957 to 1996. For this paper we use the daily averages supplied by NCEP–NCAR. For an in-depth analysis of the methods behind the project, see Kalnay et al. (1996) and the references therein. As indicated in the introduction, surface pressure is used because it is related through the hydrostatic approximation to the mass and through the geostrophic approximation to the circulation (Gill 1982; Holton 1992; Sugihara et al. 1999). To ex-

amine the differences between the surface and the upper troposphere we repeat the analysis using pressure at the tropopause height.[1]

3. Methods

a. Residual delay map formulation

The underlying concept behind RDMs is to identify some qualitative relationship between a given time series and residuals (errors) from a short-term predictive model of the time series. The motivation for this procedure is that if the model accurately describes the dynamics of the data, then the residuals behave like random noise and thus have no meaningful relationship to the data. It is the combination of the identification of a relationship (or lack of relationship) and the careful choice of the short-term predictive model that will allow us to isolate different components of the underlying dynamics. For example, one conclusion we will make below is that if we use a short-term linear predictive model (whose parameters are fit to the model) and find that there is some relationship between the residuals and the data, then we know that the underlying dynamics must come from some low-dimensional, nonlinear process. This is because if the dynamics of the system are linear and the short-term model is chosen to be linear, then only noise would be plotted by the RDM (since the model is fitted, it rules out any linear correlations between the residuals and the data). The other possibility is that the dynamics are so high dimensional that the time series appears as noise, resulting in an RDM that exhibits no relationship between the residuals and the data.

RDMs can be considered in the following generalized format: Suppose $\{x_t\}$ is a time series of finite length. Given any point x_t (and its past history x_{t-1}, x_{t-2}, etc.), we can try to estimate $x_{t+\tau}$ by some predictive method (we will elaborate on the predictive model later in this section). Let $\hat{x}_{t+\tau}$ be the predicted value. We define the residual to be $r_{t+\tau} = \hat{x}_{t+\tau} - x_{t+\tau}$. The relationship between the residuals and the time series, if any, can be observed by plotting $r_{t+\tau}$ as a function of x_t. This can be considered a map between the residuals and the time series delayed by τ, hence the name residual delay map.

Since the time series can be lengthy and noisy (implying that the RDM will have a significant amount of variance), we bin points of the RDM in groups of size N according to their values of x_t. This is done by first sorting the pairs $(x_t, r_{t+\tau})$ according to the values of x_t and then partitioning the sorted values in bins (groups) of size N. For each bin, the average of the $r_{t+\tau}$ is plotted as a function of the average of the x_t. The RDM will now be a map between the binned values of the time series and the average predictive error over N points.

For the RDMs in this paper we use the same short-term predictive model as Sugihara et al. (1999): the next day's barometric pressure is predicted using a linear combination of the current barometric pressure, the previous day's barometric pressure, and the barometric pressure from two days prior. This corresponds to a prediction by an autoregressive model of order 3 and $\tau = 1$. The predictive model can be written in the following form:

$$\hat{x}_{t+1} = a_1 x_t + a_2 x_{t-1} + a_3 x_{t-2}, \tag{1}$$

where the coefficients a_1, a_2, and a_3 are determined by least squares (Press et al. 1992). The order of the autoregressive model was determined by the Akaike information criteria (Akaike 1974) and no significant changes in the structure occurred in the RDMs when the order was increased. We found empirically that a bin size of $N = 150$ gave good results when dealing with a time series containing more than 10 000 points. Varying N by 75 did not produce any dramatic changes to the results. We remark that the key aspect of choosing N is to visually extract the structure in the RDM that would be normally hidden due to a large degree of variability in time series.

We demonstrate an example of RDM analysis in Fig. 1 where we present the RDMs for the observed daily surface pressure at Sydney and Townsville, Australia. The RDM from Sydney shows a significant amount of structure in the form of a V shape, while the RDM from Townsville shows little, if any, structure. From the arguments presented in the previous sections and the visible structure in the RDM at Sydney, we conclude that there must be some low-dimensional, nonlinear relationship between the residuals and the time series. Since there is little visible structure in the Townsville RDM and little variance in the residuals, we can conclude that the short-term linear predictor was effective; thus, the RDM gives no evidence that the underlying dynamics is low dimensional and nonlinear.

By design of the linear short-term predictive model, any structure that is observed in the RDM gives evidence that the underlying dynamical process is inherently low dimensional and nonlinear. This can be verified by using surrogates of the time series that are consistent with a linear process. That is, we can create a time series similar to the original time series, but consistent with with a linear process with Gaussian inputs. This type of surrogate is known as an amplitude adjusted surrogate (Theiler et al. 1992). The surrogate is created by taking the Fourier transform of the time series, randomizing the phases, and then taking the inverse of the Fourier transform. This will preserve the autocorrelation function of the time series (it is in this regard that we consider the time series and its surrogate to be similar). This process will preserve the linear components and will make the surrogate consistent with null hypothesis

[1] The tropopause field is output as a postprocess from the NCEP–NCAR reanalysis. It can be considered to be the lowest pressure level above 450 mb where the lapse rate is under 2 K per km but not above 85 mb.

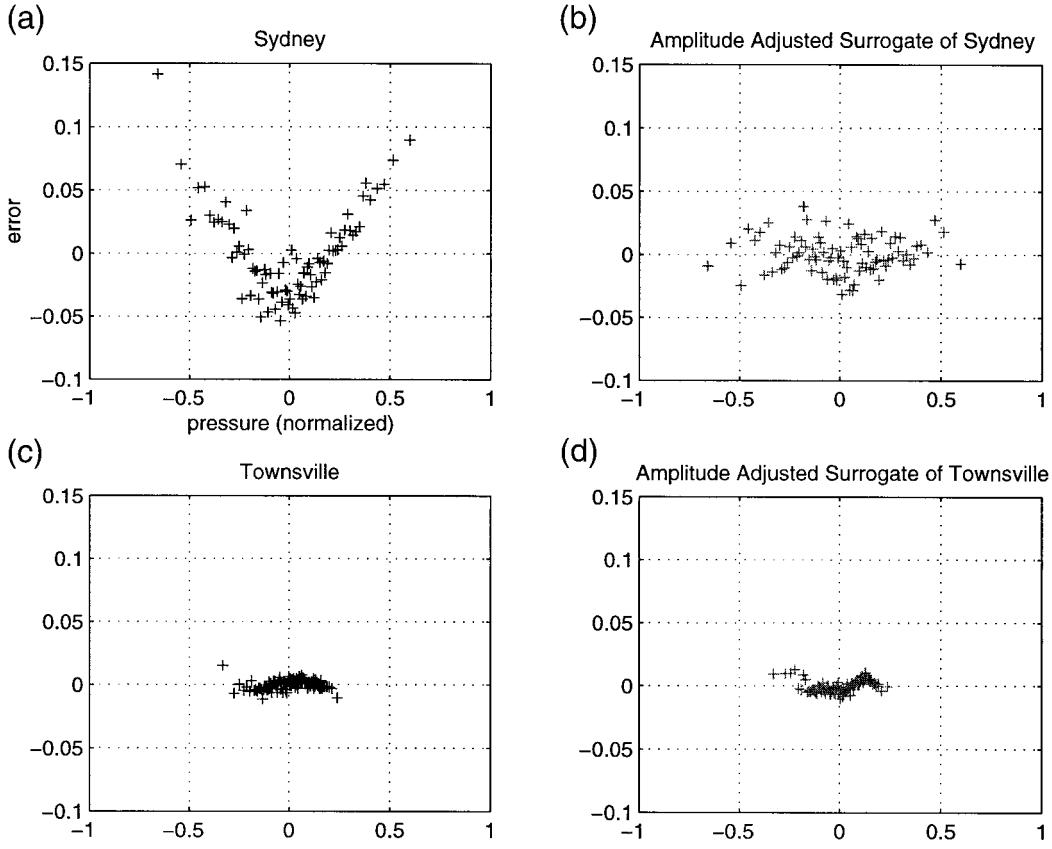

FIG. 1. (a), (c) RDMs using a short-term linear predictive model [Eq. (1)] of daily averaged observed surface pressure (spanning 37 yr) for Sydney (33.52°S, 151.13°E) and Townsville (19.16°S, 146.48°E), Australia (see section 3 for details of the RDMs are computed). (b), (d) RDMs of the corresponding RDMs from amplitude adjusted surrogates (similar, but consistent with a linear signal hypothesis; detailed in section 3a) of the respective time series. Note the destruction of the structure between (a) and (b).

that the data came from a linear system with Gaussian inputs [see Theiler et al. (1992) for an excellent in-depth review]. The results are illustrated in Fig. 1 where we also show the RDMs for the amplitude adjusted surrogates of the Sydney and Townsville time series. As expected, the structure that was present in the RDM of the original Sydney time series has been destroyed (since there is no longer a nonlinear component in the dynamics). The RDM of the surrogate for the Townsville time series exhibits little difference from the original RDM. Hence, it is likely that the dynamics involved in the Townsville time series was not dominated by low-dimensional, nonlinear dynamics. This can be interpreted several ways, the first is that the dynamics are dominated by linear terms. The other possibility is that the dynamics comes from such a high-dimensional system for which considering only the projection onto the barometric pressure variable yields a time series that may appear to be noisy with a small linear component. To separate this issue, further analysis is typically required.

b. Statistics for RDMs

When dealing with numerous time series from large datasets, it becomes difficult to visually compare the RDMs for the amount of structure present. Instead, a statistic can be defined that quantifies the amount of structure present in the RDM. Sugihara et al. (1999) proposed a statistic based on the ratio of the variance of the unbinned RDM to the binned RDM. The difficulty with that statistic is that an RDM with qualitative structure (i.e., the V shape vs no discernible shape) cannot be distinguished from an RDM with little coherent structure and a small difference in the variance between the binned and unbinned residuals. As Sugihara et al. (1999) stated, they by no means considered their statistic to be powerful enough to detect these differences in the structures that are observed in the RDMs.

In order to alleviate this issue, we introduce an alternative statistic that has a maximum of 1, with larger values indicating the presence of coherent structure in the RDM, and a minimum of 0. We define our statistic, χ, to be the normalized root-mean-square of the first two points of the autocorrelation function of the RDM.

Suppose $\{z_i\}$ are points created by the RDM algorithm (expected residuals indexed from left to right by bin number), then the autocorrelation function of the RDM can be defined as

$$R(k) = \frac{\sum_{i=1}^{N-k} z_{i+k} z_i}{\sum_{i=1}^{N} z_i^2}. \tag{2}$$

Then the statistic is defined as

$$\chi = \sqrt{\frac{R(1)^2 + R(2)^2}{2}}. \tag{3}$$

Since $R(k)$ measures the linear correlation between z_i and z_{i+k}, considering $R(1)$ and $R(2)$ corresponds to considering correlations between an arbitrary point of $\{z_i\}$ and its four nearest neighbors (i.e., the correlations between $z_{i-2}, \ldots, z_{i+2}$). If there is structure in the RDM, both $R(1)$ and $R(2)$ will be large since nearby points will be forced to conform to the structure. If little structure is present, then nearby points need not obey any relationship to each other, and both $R(1)$ and $R(2)$ should be small. We found empirically that using the first two points of the autocorrelation function yielded the best results for the statistic. The $R(k)$ for $k \geq 3$ contributed little to the statistic; however, if the time series were longer it might have been beneficial to include more terms of the autocorrelation function.

Since our RDMs are constructed from a best-fit autoregressive model, there can be no linear structure present in the RDMs. Thus, while χ does not directly measure nonlinearity, any coherent structure detected by χ must be nonlinear. As with most statistics, it is important to emphasize that this statistic cannot be applied blindly. The purpose of developing this statistic was to aid in identifying which RDMs, of the numerous ones created, to visually investigate in detail.

4. Analysis and Results

a. Properties observed in RDMs

The shape of the structure in the RDM can be valuable for understanding the physical processes behind the dynamics. For example, Sugihara et al. (1999) demonstrated that the structure in the RDMs of the station data is independent of seasonal variability. They also conjectured that the V shape in the RDM at Sydney was due to the difference in timescales on which low and high pressures persist. If the high pressures persisted significantly longer than low pressures and the transitions from lows to highs was abrupt, then the linear predictor would accurately predict high pressures, giving the bottom of the V shape. However, because the transitions between low and high pressures are on a much faster timescale then the high pressure portion of the time series, the linear predictor will do poorly during the transitions. This will give the tails of the V shape. This can be verified by creating the RDM from a model time series with these properties. We agree that this is a plausible explanation for the structure; however, while the V shape appears in the midlatitude sites of the Australian station and model datasets, we do find different shapes in the RDMs at other regions from the model dataset at both the surface and at the tropopause (further detailed in section 4c and in Fig. 5). Thus, further investigation into the mechanisms behind the structures observed in the RDMs would be useful.

An important issue to address when trying to isolate the nonlinear components of a dynamical system is the timescale on which the nonlinearity occurs. In relation to the RDM, this question becomes a matter of determining the timescale for which the structure in the RDM persists relative to the short-term predictive model. For example, suppose a weather station could record the barometric pressure at different sampling rates such as every four hours, every day, or once a week. How would the structure in the corresponding RDMs differ? Sugihara et al. (1999) showed that the RDMs could change significantly by sampling at rates of 6, 24, or 36 h. While their emphasis was in how specific portions of the structure in the RDM change (specifically, that the V shape becomes flatter as the sampling rate increases), we found that the amount of total coherent structure in the RDMs varied at different samplings rates. We show RDMs in Fig. 2 from the Sydney time series using the full time series, every second day, every third day, and every fourth day. As the sampling rate is decreased it is evident that the structure is being destroyed. When using every fourth day of the time series, the structure is nearly gone. Thus it is important to emphasize that the timescale of the nonlinearity that we are trying to identify is dependent on the sampling rate and is relative to the short-term predictive model that was used to create the RDM.

b. Comparison of model data to observational data

We now apply this approach to examine the similarity of station and reanalysis time series. If the systems producing the time series contain noise, a traditional examination of difference statistics (e.g., root-mean-square statistics or correlations) may be unsatisfying. Under these conditions it is also difficult to isolate traditional dynamical invariants (statistical properties that remain constant as the system evolves in time) such as dimension estimates and the full frequency domain. Here, instead, we compare the low-dimensional nonlinear components of the dynamics that reflect the dynamical invariants. This is the basis for the test we propose for model verification. Indeed, Sugihara et al. (1999) used a version of this test to show that the operational ECMWF forecasts did not contain all of the underlying dynamics that were present in the observational data over the Australian continent. Here we begin with a similar comparison of the NCEP–

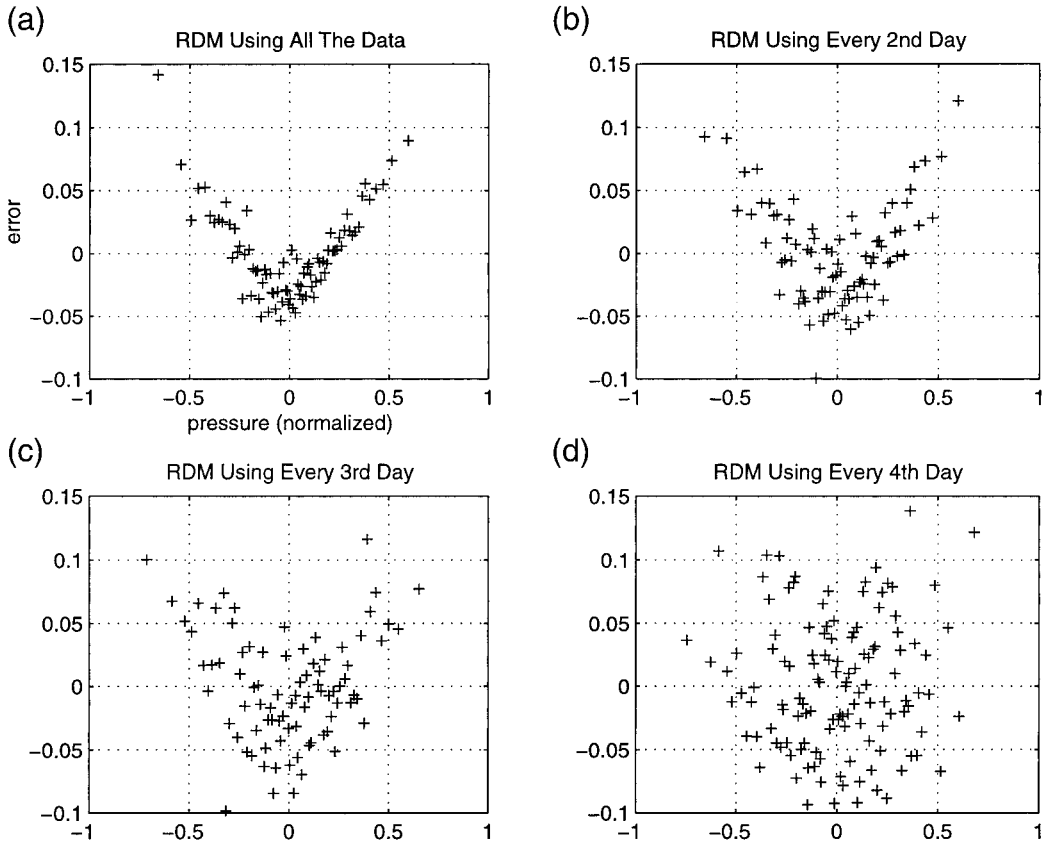

FIG. 2. RDMs are calculated on daily averaged observed surface pressure (spanning 37 yr) for Sydney (33.52°S, 151.13°E), Australia, at different sampling rates from every day to every fourth day. Note the break down in structure as the sampling rate decreases.

NCAR reanalysis carried out over a substantially longer period of time (39 yr).

In Fig. 3 the results of the Australian RDM comparisons are presented. In contrast to the earlier Sugihara et al. (1999) analysis of the ECMWF operational data, we find good agreement of the RDMs (although a few locations such as Wagga Wagga, Mount Tisa, and Albany have systematic differences). From these results we are able to conclude that the NCEP–NCAR reanalysis describes the underlying low-dimensional nonlinear component of the dynamics that is present in the station dataset.

As a final note we remind the reader that the form of the short-term predictive model is crucial to the RDM algorithm as it effectively determines the type of dynamics we are trying to isolate. It is most likely that there remain other, more specific characteristics of the underlying dynamics that are not captured by the model. In order to isolate such characteristics by the method of RDMs, the prediction scheme would need to be tailored to the specific hypothesis in question.

c. Global residual delay map analysis of the NCEP–NCAR reanalysis

In the previous section, the NCEP–NCAR reanalysis was found to accurately (with the exception of a few sites) describe the low-dimensional nonlinear component of the dynamics over the Australian continent. Here we apply the same techniques to global air pressure to isolate and identify regions of low-dimensional nonlinearity in the atmosphere. Our goal is to relate the spatial structure of these regions to the underlying dynamics.

We begin by computing the RDMs and the statistic, χ [defined in section 3b, Eq. (3)], on the surface and tropopause reanalysis pressure time series. The meridional structure of χ is illustrated in Fig. 4 along 155°E, a longitude just to the east of the Australian continent, extending southward from the equator. In the reanalysis data at the surface we find little coherent structure in the RDMs in the Tropics (low values of χ), increased structure in the midlatitudes, and less structure again at high latitudes. This behavior was found by Sugihara et al. (1999) in the station time series of surface pressure and led to their conjecture that sites in the midlatitudes should be significantly more nonlinear than in the Tropics. We investigate the validity of this conjecture by analyzing the RDMs calculated globally. Figure 5a presents a global map indicating which RDMs of the surface pressure time series from the reanalysis have coherent structure (a high value of the statistic is colored red) indicating the presence of low-dimensional nonlin-

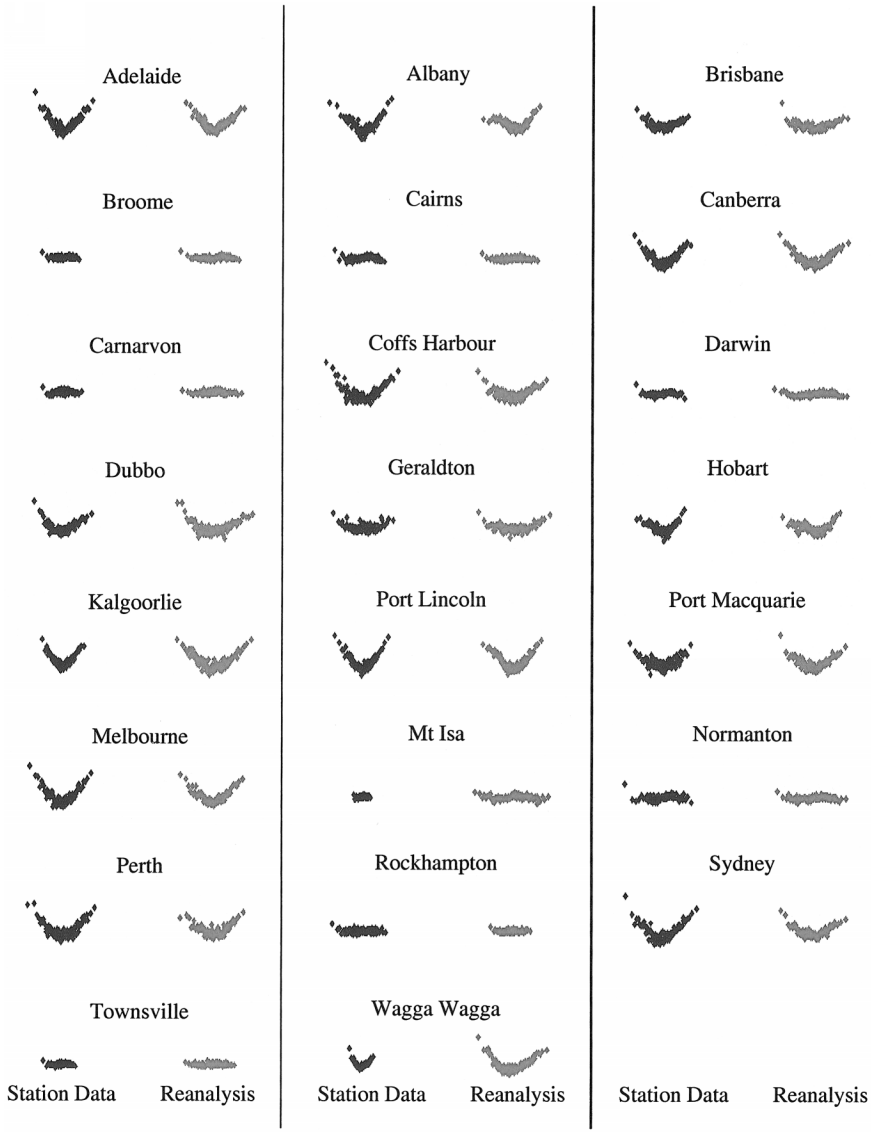

FIG. 3. RDMs for 23 sites on the Australian continent using daily averaged observed surface pressure spanning 37 yr are shown on the left. On the right are the corresponding RDMs from reanalysis data. Note the good agreement at most of the stations, which is in contrast to the results of Sugihara et al. (1996).

ear dynamics versus those RDMs that have little structure (a low value of the statistic is colored blue).

High values of χ in the Southern Hemisphere, at the surface, occur in a zonal band between 30° and 40°S. This band of latitudes span the most southernly edge of the Hadley circulation, a region with large temperature gradients and strong storm activity, particularly downwind (eastward) of continents (Lau 1988). These instabilities are highly nonlinear and would be poorly approximated by a linear predictor. They are consistent with the existence of intense low pressure events with rapid onset. Low values in the Tropics are consistent with the mild weather generally found there.

In the Northern Hemisphere at the surface, however, the connection between high values of χ and the position of the storm tracks seems much weaker. The highest values of χ occur in a region extending northeastward across the Eurasian landmass from the Ukraine to the Siberian Arctic, while smaller regions of strongly low-dimensional nonlinearity occur in the Middle East and western China. Strikingly, there are low values across the stormy North Pacific and Atlantic Oceans as well as across North America.

The vertical dependence of χ is examined by repeating the analysis at tropopause heights, well above most direct topographic influences (Kimoto and Ghil 1993; Holton 1992; Selten 1993; Trevisan 1995). The results (Figs. 4 and 5) show that the RDMs at tropopause

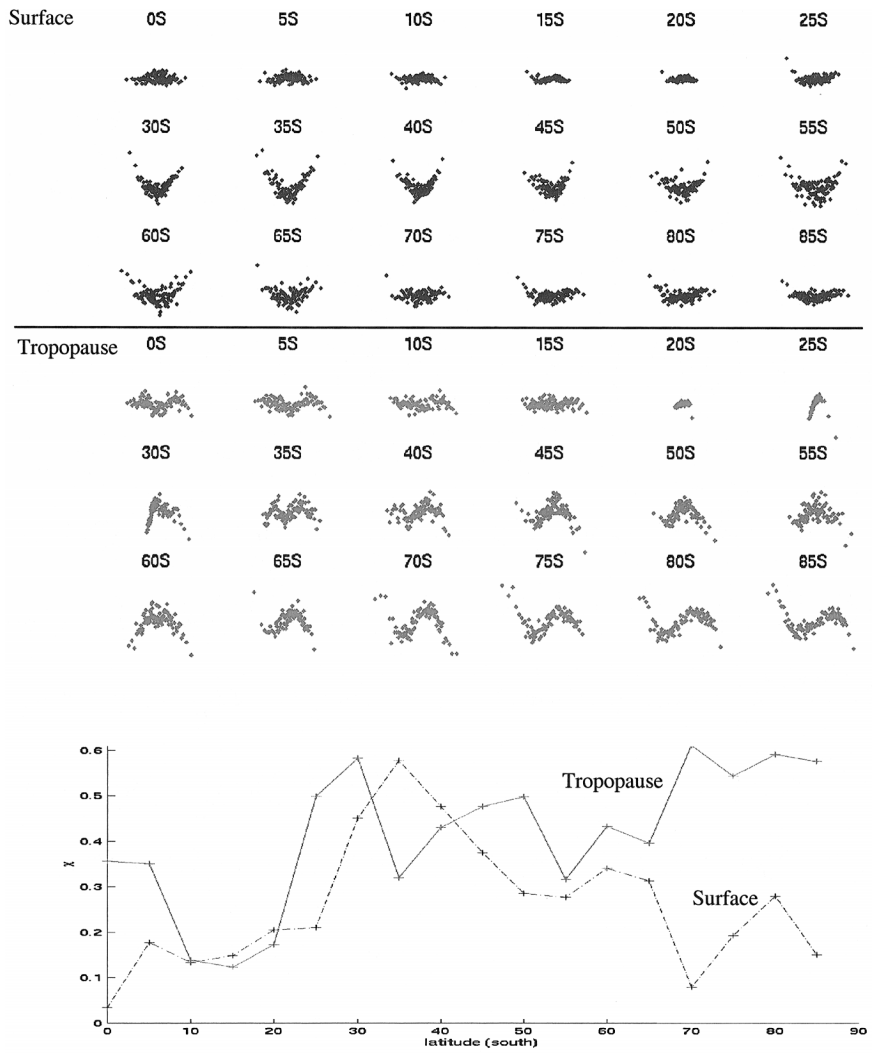

FIG. 4. RDMs are shown for the Southern Hemisphere along 155°E using NCEP–NCAR reanalysis data. Results using barometric pressure (a) at the surface and (b) at the height of the tropopause. (c) Amount of low-dimensional nonlinearity identified by the RDM (using a statistic defined in the paper) as a function of latitude.

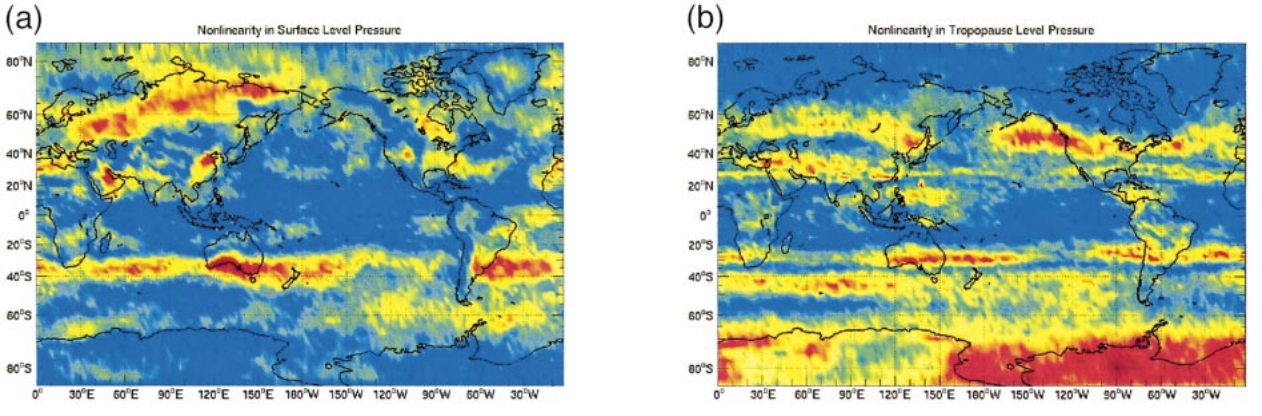

FIG. 5. RDMs were calculated using 39 yr (1957–96) of daily averaged surface pressure from NCEP–NCAR reanalysis data (a) at the surface and (b) at tropopause height. A statistic (detailed in section 3b) was then applied to each RDM to measure the amount of low-dimensional nonlinearity (structure). Red regions indicate RDMs with coherent structure.

levels can be quite different. A comparison between the two panels in Fig. 5 indicates large differences in the spatial distribution of the low-dimensional, nonlinear dynamics at the respective heights. The number of zonal bands that were present in the RDMs of the surface pressure have increased. The strong band that was between 30° and 40°S at the surface has been replaced with two bands: one sharp and narrow band just above 30°N and a thicker, less defined band at 40°S. Several thin bands between 25° and 40°N have emerged and the strong region of low-dimensional nonlinearity over the far northern portion of Europe and Asia has diminished in size. The strong zonal region of low-dimensional, nonlinearity in the Northern Hemisphere coincides with average position of the jet stream and agrees with the results from Kimoto and Ghil (1993). Over Antarctica, the region indicating nonlinear dynamics isolates the continent from the rest of the regions. Indeed, the RDMs have characteristically different shapes between the time series at the surface and the tropopause heights, such as in Fig. 4 at 30°S, 155°E. For example, in many of the RDMs at the tropopause heights, there is a pronounced downward V shape suggesting the presence of intense high pressure fronts.

While the driving mechanisms for the tropical atmosphere are inherently nonlinear (convection), the large-scale response is linear (Webster 1972; Gill 1980, 1982; Sugihara et al. 1999). This indicates that stations in the Tropics are likely to be dominated by linear dynamics (in contrast to high-dimensional nonlinear dynamics). From this and the global images in Fig. 5, Sugihara et al.'s (1999) conjecture is likely to be true at the surface in the Southern Hemisphere, but not in much of the Northern Hemisphere.

Due to the prior belief that the large-scale dynamics in the Tropics should be linear, we are able to make the distinction between linear and low-dimensional, nonlinear dynamics in the RDMs over the Tropics. However, outside the Tropics no such assumptions can be made due to baroclinic instabilities. Thus outside the Tropics only the identification of regions with low-dimensional dynamics can be made. It is important to emphasize that our analysis to test their conjecture relies on reanalysis data rather than actual station data. Also due to the nature of our technique, only qualitative assessments can be made without further rigorous statistical arguments and testing.

5. Concluding remarks

In this paper we have presented two primary uses for residual delay maps. The first is a method to isolate the low-dimensional nonlinear component of the dynamics from a given time series. This technique allows us to move away from trying to compute dynamical invariants and focus on qualitative aspects of the time series. We were able to show how this method can provide valuable insight into atmospheric dynamics. These include pattern formation of the global atmospheric system and the investigation of the transition of the dynamics between two spatial regions with characteristically different dynamics. The second use for RDMs that we presented dealt with the issue of model verification. We showed how the qualitative view that we adopt when using RDMs can be used to compare different models from dynamical rather than statistical perspectives.

A key area of further work is to gain a better understanding of the vertical profile of regions of nonlinearity. In this paper we only considered how the low-dimensional nonlinearity varied at two heights: the surface and the tropopause. It can be seen in Fig. 4 that there are large differences in the RDMs between the two different heights at a given station. Since most atmospheric pressure data are stored as geopotential height at a given pressure level, analogs of Fig. 5 could be created for each pressure level using RDMs of the geopotential height time series. This would give a more complete description of regions of nonlinearity in the atmospheric system and how potentially different the dynamics can be at different heights [the importance of this has been stressed by Kimoto and Ghil (1993) in their discussion of planetary boundary layers as well as by Selten (1993) and Trevisan (1995)].

We reiterate the importance of the short-term predictive model used to calculate the RDM. As we discussed in section 3a, by choosing the short-term model from the class of autoregressive models we restricted our analysis to the identification to low-dimensional, nonlinear dynamics. In this paper we focused on the qualitative structure of the RDMs to motivate further work on the classification of dynamics from time series. It should be possible to establish, in a statistical language, a hierarchy of null hypotheses and statistical criteria based on the choice of the predictive model for the identification of dynamics isolated by the RDM technique. Such predictors might include multivariate predictors that would give a better insight into the dynamics both spatially and temporally. This would be useful in addition to the applications seen in this paper since the dynamics involving barometric pressure occur in waves and fronts that are best analyzed by considering groups of stations rather than each station individually.

As is the case in all time series applications, the methods cannot be applied blindly; rather, some care must be taken before applying a specific method (in our case, choosing the short-term predictive model). From our results we believe that RDMs can be a powerful new tool for the analysis of atmospheric data.

Acknowledgments. The authors would like to thank Dr. Reggie Brown, Dr. Jim Yorke, Dr. Eugenia Kalnay, Dr. V. Krishnamurthy, and members of the 1998 Grand Combin Summer School for insightful discussions. Thanks also goes to Dr. Paul Dixon and Dr. George Sugihara for numerous discussions and making their data available to us. Also thanks to NCEP and NCAR

for making their data available for public use. This project was made possible by the gracious support of the W. M. Keck Foundation. The first author would like to thank the Cosmos Club of Washington, D.C., for financial support. The second author was partially supported by the National Science Foundation (Divisions of Mathematical and Physical Science). The third author also gratefully acknowledges support from the National Science Foundation (OCE9812404).

REFERENCES

Akaike, H., 1974: A new look at the statistical identification model. *IEEE Trans. Automatic Control,* **19,** 716–723.

Casdagli, M. C., L. Iasemidis, J. Sackellares, S. Roper, R. Gilmore, and R. Savit, 1996: Characterizing nonlinearity in invasive EEG recordings from temporal lobe epilepsy. *Physica D,* **99,** 381–399.

Gill, A. E., 1980: Some simple solutions for heat-induced tropical circulation. *Quart. J. Roy. Meteor. Soc.,* **106,** 447–462.

——, 1982: *Atmosphere–Ocean Dynamics*. Academic Press.

Holton, J. R., 1992: *An Introduction to Dynamic Meteorology*. 3d ed. Academic Press, 511 pp.

Jin, F.-F., J. Neelin, and M. Ghil, 1993: El Niño/Southern Oscillation and the annual cycle: Subharmonic frequency-locking and aperiodicity. *Physica D,* **98,** 442–465.

——, ——, and ——, 1994: El Niño on the devil's staircase: Annual subharmonic steps to chaos. *Science,* **264,** 70–72.

Kalnay, E., and Coauthors,1996: The NCEP/NCAR 40-Year Reanalysis Project. *Bull. Amer. Meteor. Soc.,* **77,** 437–471.

Kantz, H., and T. Schreiber, 1997: *Nonlinear Time Series Analysis*. Cambridge University Press, 304 pp.

Kimoto, M., and M. Ghil, 1993: Multiple flow regimes in the Northern Hemisphere winter. Part 1: Methodology and hemispheric regimes. *J. Atmos. Sci.,* **50,** 2625–2643.

Lau, N., 1988: Variability of the observed midlatitude storm tracks in relation to low-frequency changes in the circulation pattern. *J. Atmos. Sci.,* **45,** 2718–2743.

Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.,* **20,** 130.

Mo, M., and M. Ghil, 1987: Statistics and dynamics of persistent anomalies. *J. Atmos. Sci.,* **44,** 877–901.

Press, W. H., S. A. Teukolsky, and B. P. Flannery, 1992: *Numerical Recipes*. 2d ed. Cambridge University Press, 994 pp.

Selten, F., 1993: Toward an optimal description of atmospheric flow. *J. Atmos. Sci.,* **50,** 861–877.

Sugihara, G., M. Casdagli, E. Habjan, D. Hess, G. Holland, and P. Dixon, 1999: Residual delay maps unveil global patterns of atmospheric nonlinearity and produce improved local forecasts. *Proc. Natl. Acad. Sci.,* **96,** 14 210–14 215.

Theiler, J., S. Eubank, A. Longtin, B. Galdrikian, and J. Farmer, 1992: Testing for nonlinearity in time series: The method of surrogate data. *Physica D,* **58,** 77–94.

Trevisan, A., 1995: Statistical properties of predictability from atmospheric analogs and the existence of multiple flow regimes. *J. Atmos. Sci.,* **52,** 3577–3592.

Tziperman, E., L. Stone, M. Cane, and H. Jarosh, 1994: El Niño chaos: Overlapping of resonances between the seasonal cycle and the Pacific ocean–atmosphere oscillator. *Science,* **264,** 72–74.

Webster, P., 1972: Response of the tropical atmosphere to local, steady forcing. *Mon. Wea. Rev.,* **100,** 518–541.