The Abdus Salam
International Centre for Theoretical Physics

United Nations
Educational, Scientific
and Cultural Organization

International Atomic
Energy Agency

SMR.1676 - 27

8th Workshop on Non-Linear Dynamics and Earthquake Prediction

3 - 15 October, 2005

Weather Regime Prediction using Statistical Learning

Weather Regimes and Preferred Transition Paths in a Three-Level Quasigeostrophic Model

Michael Ghil
Ecole Normale Superieure
Laboratoire de Meteorologie Dynamique (LMD)
24, rue Lhomond
F-75231 Paris
France

Weather regime prediction using statistical learning

A. Deloncle[1], R. Berk[2], F. D'Andrea, M. Ghil[3]

Département Terre-Atmosphère-Océan and Laboratoire de Météorologie Dynamique du CNRS/IPSL, Ecole Normale Supérieure, Paris, France

J. Atmos. Sci., *submitted,*

September 26, 2005

[1] *Current affiliation and corresponding author address:* LadHyX, CNRS, Ecole Polytechnique, F-91128 Palaiseau Cedex, France. E-mail: axel.deloncle@ladhyx.polytechnique.fr

[2] *Permanent affiliation:* Department of Statistics, University of California, Los Angeles, CA 90095-1554

[3] *Additional affiliation:* Department of Atmospheric and Oceanic Sciences and Institute of Geophysics and Planetary Physics, University of California, Los Angeles, CA 90095-1565

Abstract

Two novel statistical methods are applied to the prediction of transitions between weather regimes. The methods are tested using a long, 6 000-day simulation of a three-layer, quasi-geostrophic (QG3) model on the sphere at T21 resolution.

The two methods are the k nearest-neighbor classifier and the random-forest method. Both methods are widely used in statistical classification and machine learning; they are applied here to forecast the break of a regime and subsequent onset of another one. The QG3 model has been previously shown to possess realistic weather regimes in its Northern Hemisphere and preferred transitions between these have been determined. The two methods are applied to the three more robust transitions; they both demonstrate a skill of 35–40% better than random and are thus encouraging for use on real data. Moreover, the random-forest method allows, while keeping the overall skill unchanged, to efficiently adjust the ratio of correctly predicted transitions to false alarms.

A long-standing conjecture has associated regime breaks and preferred transitions with distinct directions in the reduced model phase space spanned by a few leading empirical orthogonal functions of its variability. Sensitivity studies for several predictors confirm the crucial influence of the exit angle on a preferred transition path. The present results thus support the paradigm of multiple weather regimes and of their association with unstable fixed points of atmospheric dynamics.

1. Introduction and motivation

The low-frequency, intraseasonal variability of the extratropical atmosphere involves phenomena with time scales that are longer than the baroclinic-eddy life cycles and shorter than the change of seasons, that is 10 to 100 days. This variability is characterized by the existence of large-scale, persistent and recurrent flow patterns called weather regimes (Ghil and Robertson 2002; Molteni 2002). Several regimes have been identified in a consistent way by using diverse statistical and statistico-dynamical methods, which have been applied to observed atmospheric data, as well as to output from numerical models. A review of methods and results is included, for example, in Smyth et al. (1999) and Ghil and Robertson (2002).

The concept of weather regimes has been used successfully in different fields of the atmospheric sciences, from predictability through the downscaling of general circulation model (GCM) results to climate change impact assessment. In this paper, we examine the possibility that, because of their persistence, weather regimes provide a coarse-grained, predictable component of the atmosphere (Mo and Ghil 1988; Ghil et al. 1991) capable of circumventing the deterministic predictability barrier of 10 to 15 days (Lorenz 1969).

Markov chains of multiple regimes have been shown to provide extended predictability, at the cost of less detail in the predicted variables (Fraedrich and Klauss 1983; Ghil and Robertson 2002). Moreover, the most advanced numerical weather prediction models still have problems at forecasting regime transitions. This shortcoming has been investigated, for example, in the context of atmospheric blocking inception. Tibaldi and Molteni (1990) showed that much of the forecast error of the European Centre for Medium-Range Weather Forecasts (ECMWF) forecast model was due to its inability to enter a blocked state 3–4 days into the forecast. This reflected a general underestimation of blocking frequency in GCMs (D'Andrea and Coauthors 1998). Although much

progress has been made since, forecasts of blocking inception still have no skill starting from a lead time of 6 days (Pelly and Hoskins 2003).

The purpose of this article is to present a novel strategy, based on advanced statistical methods, to forecast regime breaks and subsequent onsets. The goal is to show the applicability and promise of such a strategy, rather than to establish an operational forecast system. For this reason, we will work with the output of an intermediate-complexity, quasi-geostrophic, three-layer (QG3) model introduced by Marshall and Molteni (1993). This model has been widely used to investigate the Northern Hemisphere atmosphere's low-frequency variability (D'Andrea and Vautard 2001; D'Andrea 2001; Ferreira and Frankignoul 2005).

More important, the QG3 model has been recently shown to have interesting regime dynamics. Kondrashov et al. (2004) carried out a long-time integration of this model and studied its properties in a phase space spanned by its three leading empirical orthogonal functions (EOFs). Using two distinct clustering procedures, these authors obtained four statistically significant weather regimes: the two phases of the North Atlantic Oscillation (NAO^+, NAO^-) and the two phases of a more hemispheric and zonally symmetric mode, which they identified with the Arctic Oscillation (AO^+, AO^-). They found that these four regimes were in good agreement with previous results (Kimoto and Ghil 1993a,b; Michelangeli et al. 1995; Corti et al. 1997). By studying the Markov chain of transitions between regimes, they identified five highly significant transitions that could be organized into two cycles: $NAO^- \rightarrow NAO^+ \rightarrow AO^+ \rightarrow NAO^-$ and $AO^+ \leftrightarrow NAO^+$.

They also showed that several specific transitions were characterized by preferential directions in phase space. To do so, they defined for every transition an exit point on the regime boundary; the exit vector, pointing from the regime centroid to the exit point, could then be described by two angles on the unit sphere around the centroid. The joint probability density function (PDF) of these two angles for the five identified transitions exhibited one or two maxima.

The existence of such preferential directions, along which the system's trajectory leaves a regime, has been conjectured by Legras and Ghil (1985), based on the nonlinear dynamics of their barotropic model on the sphere. In this model, certain regimes were associated with slowing down of the trajectories in the neighborhood of unstable fixed points. These trajectories were then ejected along the small number of unstable directions. Finding traces of similar behavior in the much more realistic, baroclinic QG3 model used here renders its investigation even more interesting in the present context.

In this article we make use of the same clustering methodology as Kondrashov et al. (2004) to define weather regimes and the preferred transition paths between them. Statistical learning techniques are then applied to exploit this knowledge for forecasting purposes.

The paper is organized as follows. In section 2, the atmospheric model and the pre-processing performed to obtain the weather regimes and the transition paths are briefly described; some details on the model appear in appendix A. In section 3, we present the two main statistical tools of this study: the k nearest-neighbor classifier and the "random-forest" technique. Further details about the latter are given in appendix B.

Section 4 is devoted to the main results of this study, in two cases of increasing complexity. In section 4*a*, we forecast the three specific regime breaks that constitute the first transition cycle identified by Kondrashov et al. (2004). In section 4*b*, we extend our study to any possible transitions starting from the NAO^- regime. In both situations, we show that our statistical methods have verifiable predictive skill. The performance of the random-forest algorithm can also be modulated according to the different weights one gives for different type of error: false alarms vs. failure to predict. A sensitivity study of the forecast skill to the predictors demonstrates the critical influence of preferred transition directions. A summary and discussion of the results follow in section 5.

2. The QG3 model and its weather regimes

a. The QG3 model

The model used in this study was first proposed and investigated by Marshall and Molteni (1993). It consists in the quasi-geostrophic (QG) potential vorticity (PV) equations, integrated on the sphere; the horizontal discretization is spectral, with a T21 truncation, and there are three levels in the vertical (200, 500 and 800 hPa); hence the QG3 abbreviation. At each vertical level, the prognostic equations for PV read:

$$\frac{\partial q}{\partial t} = -J(\psi, q) - D(\psi) + S, \tag{1}$$

where q is the potential vorticity, ψ the streamfunction and J the Jacobian operator on a pair of two-dimensional fields. The term $D(\psi)$ is a linear operator representing the effects of Newtonian relaxation of temperature, linear drag on the lower level (with drag coefficients depending on the nature of the underlying surface), and horizontal diffusion. The spatially varying, time-independent forcing S is designed to represent PV sources that result from processes not explicitly included in the model. This source term is constructed empirically, as in Marshall and Molteni (1993), to keep the model's mean state close to that of an observed wintertime climatology; see appendix A.

Despite its simplicity, the model has a remarkably good climatology and low-frequency variability, with a plausible stationary-wave pattern, Pacific and Atlantic storm tracks, and maxima in low-frequency activity at the end of the storm tracks. The model also produces wintertime weather regimes that are very similar to the observed ones (Corti et al. 1997; D'Andrea and Vautard 2000; Kondrashov et al. 2004).

b. The weather regimes

The main steps to calculate the weather regimes are only summarized here; further details are given by Kondrashov et al. (2004). A 54 000-day-long, perpetual-winter integration of the QG3 model is first carried out. In order to reduce the dimension of the phase space in which the coarse graining will be carried out, we perform an EOF analysis on the unfiltered, daily 500-hPa streamfunction anomalies over the model's Northern Hemisphere. We keep the first three EOFs, thus capturing 27% of the total day-to-day variance.

Weather regimes are then identified as areas of higher probability density in this three-dimensional phase space by applying the Gaussian mixture classification method of Smyth et al. (1999). To do so, we assume that every weather regime (or cluster) is described by a Gaussian density function. The total PDF is then modeled by a weighted linear combination of the individual weather regime density functions. With the QG3 output data, we obtain four regimes that we call, following Kondrashov et al. (2004): $\mathrm{NAO}^+, \mathrm{NAO}^-$, AO^+ and AO^-.

The next step is to determine the Markov chain of transitions between regimes. Each weather regime is defined in phase space as an ellipsoid whose centroid and semi-axes are given by the mean and the covariance matrix of the corresponding Gaussian density component. The exact volume of every cluster is fixed by a scaling factor $\sigma = 1.25$ along each axis of the ellipsoid; the axes are the principal directions of the covariance matrix, and $\sigma = 1$ corresponds to the associated standard deviations.

A data point is assigned to a weather regime if it lies within the corresponding ellipsoid. When a data point belongs to several ellipsoids, we assign it to a regime according to the maximum probability value found. With this classification, about 11% of the points are in the NAO^+ regime, 13% in NAO^-, 15% in AO^+ and 9% in AO^-; the remaining 52% of the points do not belong to

any cluster.

c. The preferred transition paths

Each transition is characterized by an exit point. The exit point is the mid-point between two consecutive trajectory points that lie on the opposite side of the cluster boundary, as defined in section 2*b*. The exit vector is then defined as the vector pointing from the cluster centroid to the exit point. In the three-dimensional phase space spanned by EOFs 1, 2 and 3, the coordinates of an exit point are (x, y, z) and the unit vector in its direction can be fully described by two angles θ and ϕ with:

$$\begin{aligned} \tan\theta &= \frac{z}{\sqrt{x^2+y^2}} \quad , \quad -\frac{\pi}{2} < \theta < \frac{\pi}{2}, \\ \tan\phi &= \frac{y}{x} \quad , \quad 0 < \phi < 2\pi, \end{aligned} \tag{2}$$

the positive pole being aligned with EOF–3. Computing the two-dimensional PDF of these two angles using a Gaussian kernel estimator (Silverman 1986), we obtain the preferred exit directions as the maxima of this PDF.

In Fig. 1 the PDFs of θ and ϕ are shown for the three transitions that will be analyzed in section 4*a*: $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+ \rightarrow \mathrm{AO}^+ \rightarrow \mathrm{NAO}^-$. For two of them, $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ and $\mathrm{NAO}^+ \rightarrow \mathrm{AO}^+$, the PDF has two sharp maxima close to each other; the regime break consequently occurs along either one of two paths. In the third case, $\mathrm{AO}^+ \rightarrow \mathrm{NAO}^-$, there is only one maximum, which is much less pronounced. Kondrashov et al. (2004) described these three transitions as "the first cycle of significant transitions"; they provide good examples of the two kinds of regime breaks that these authors observed on a larger set of highly significant transitions: on the one hand sharp and pronounced maxima, on the other less peaked angular PDFs. The first type of transition was found to be more frequent, on the whole. We chose to study this transition cycle because it

allowed us to compare the results here with those of Kondrashov et al. (2004) and also because they illustrate rather well the different situations in terms of exit-angle PDFs.

[Figure 1 about here.]

3. Methodology

a. Predictands and predictors

For each individual transition we are trying to forecast, we define a data point as an event or a non-event. Let us consider the transition $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$. For this transition, a point belonging to NAO^- is considered as an "event" if it is going to exit the NAO^- cluster the following day, and to enter the destination cluster NAO^+ at some moment in the future, after possibly having spent one or several days outside any regime boundary. Any other point of the NAO^- regime is considered as a "non-event." Non-events can be points not leaving the NAO^- regime the next day (staying longer in the regime) or leaving NAO^- to reach a different regime than NAO^+. Forecasting the $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ regime break means to classify NAO^- points into one of the two possible outcomes: "event" or "non-event."

Our predictors are based on the position and the velocity of a data point. In order to exploit the preferential paths of regime breaks identified by Kondrashov et al. (2004) and in section 2*c* here, we use the spherical coordinates (r, θ, ϕ) centered on the regime centroid and with the polar axis aligned with the preferred transition path, rather than with EOF–3. When the transition under consideration has two local maxima ($\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ and $\mathrm{NAO}^+ \rightarrow \mathrm{AO}^+$), we use the global maximum as the pole. Figure 2 illustrates this change of coordinates.

[Figure 2 about here.]

In these modified spherical coordinates, the "deviation" angle formed by the current state vector and the preferred transition direction is given by a single variable, θ. A value $\theta = \pi/2$ means the state vector is perfectly aligned with the preferred exit vector, while a value of $\theta = -\pi/2$ indicates that it is in the opposite direction. The coordinate r is the distance to the center of the regime centroid. The cartesian velocity components $dx/dt, dy/dt, dz/dt$, given by the QG3 model, are also called "tendencies" and will be expressed in the spherical coordinate system by (v_r, v_θ, v_ϕ). In summary, our predictors are daily data points in spherical coordinates and their tendencies $(r, \theta, \phi, v_r, v_\theta, v_\phi)$.

We have at our disposal a very long model simulation of 54 000 days, but wish to evaluate our method in a manner that is consistent with the amount of data one can obtain from a re-analysis data set. To do so, we will keep in the following only 6 000 days of the simulation and thus obtain a fair estimate of our method's forecast performance when using a realistic number of training data. Although the QG3 model was run in a perpetual winter mode, these 6 000 days can be thought to correspond to 50 winters of 120 days (mid-November to mid-March).

b. K nearest-neighbor classifier

We have used two forecast methods to classify events or non-events from the six predictors described in section 3*a*. The first is a classical analog procedure. We dispose of a library of 6 000 days that correspond to past observed data and that constitutes a training data set. From these, we can build a "look-up table" of predictors, classified into events and non-events.

We now consider a new point that is in NAO^- at initial forecast time and we want to determine if it is an event or not. We first search for its k nearest neighbors in the look-up table in terms of Euclidean distance in the space of the six predictors $(r, \theta, \phi, v_r, v_\theta, v_\phi)$. Once the k nearest neighbors

are identified, we count the number of "events" and "non-events" in these k table members. The forecast then assigns the new point to the category that is the best represented among its k nearest neighbors. It is easy to check if the forecast was correct by looking at the simulated days that follow in time the point that we just classified. The number k of analogs kept in the procedure is not fixed and several values, from 1 to 20, were tested to determine the one that gives the highest probability of correct forecasts.

c. Random forests

Random forests is a more advanced classification procedure, introduced in the past fifteen years; it is based on a generalization of classification and regression trees (CART). To the best of our knowledge, the present work is the first use of random forests to forecast meteorological phenomena. As in section 3*b*, the key idea is to assign a given point to a class based on information contained in a set of predictors. Random forests is largely based on recursive partitioning of a training data set by logical splits that permit accurate classifications.

Classical classification trees use successive if–then conditions to obtain a unique deterministic tree. A "random forest" is constructed from a set of K such deterministic trees, each based on a random sample of training data and on using at each split within a given tree, a random sample of predictors. Data points are then classified through a majority vote over all of the trees in the forest. Classification trees and their extension, random forests, are usually very effective statistical methods for classifying complex data structures when no simple relationship (linear for example) between predictands and predictors is apparent. Random forests is described in greater detail in appendix B here and in Breiman (2001).

4. Forecast results

For the sake of simplicity, we first concentrate on one specific transition. This will also allow us to introduce contingency tables and the forecast score used. The transition chosen is $NAO^- \rightarrow NAO^+$, as anticipated in section 3*a*. For given points belonging to the NAO^- cluster, we forecast their regime transition to NAO^+ with the two methods above. There are only two outcomes possible in this case: either there is a transition to NAO^+ or not; these two outcomes are classified as an event or a non-event. We then briefly compare the results obtained for two other single transitions, $NAO^+ \rightarrow AO^+$ and $AO^+ \rightarrow NAO^-$, with the ones we got in the $NAO^- \rightarrow NAO^+$ case.

In section 4*b*, we forecast all the possible transitions from cluster NAO^-. In this case, there are five possible outcomes: (i) no transition: the point does not leave NAO^- in the next 24 h; (ii,iii,iv) transition to one of the three other clusters; and (v) re-entry, with the trajectory exiting the NAO^- cluster and then returning to it.

a. Single-transition forecasts

1) K NEAREST-NEIGHBOR CLASSIFIER

We apply this classifier with our data library of 6 000 days and then test it on 1 000 independent points belonging to the NAO^- weather regime. The results are summed up in a 2×2 contingency table that gives the discrete joint sample distribution of forecasts and validating observations. Table 1 summarizes the definition of contingency tables, and of user and model errors. As their name indicates, the former errors provide mainly information to the user of the forecast model, the latter mainly to the modeler.

[Table 1 about here.]

The contingency table found with this data set for the $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ transition is presented in Table 2. A basic difficulty of any regime-based forecast method is that a transition from a given Regime A to a given Regime B is essentially a rare event. We immediately see in the table that the event points are much less numerous than the non-event points: the former represent only $11\% \approx 7.5\% + 3.3\%$ of the total. This is not surprising because we consider as events only the points that are about to leave their original weather regime in the next 24 h. As we will see later, this makes the forecast much more difficult.

[Table 2 about here.]

To estimate the skill of this statistical predictor compared to a random guess, we use the Heidke Skill Score (HSS) (Von Storch and Zwiers 1999) H:

$$H = \frac{S - S_r}{N - S_r}, \tag{3}$$

with S the number of correct forecasts, S_r the number of correct forecasts that a random predictor would give, and N the number of assessment points. A perfect predictor would get a score of 1, whereas a value of 0 means the evaluated predictor demonstrates no skill over a random guess.

Another convenient definition of H can be given in terms of the numbers a, b, c, d introduced in Table 1:

$$H = \frac{2(ad - bc)}{(a+b)(b+d) + (a+c)(c+d)} \tag{4}$$

In the case of our regime transition forecast in Table 2, we find $H = 0.40$, meaning that the k nearest-neighbor classifier is 40% better than a random guess. This result shows that the variables we used as predictors do contain useful information for the break of the NAO^- regime and subsequent transition to NAO^+.

To better understand how this score is obtained, we must study more closely the contingency table. The user error is especially useful in practical applications of a forecasting system. When the model forecasts a non-event, it is wrong in 7.8% of the cases; this percentage becomes 25% when a transition to NAO^+ is forecast to occur. Both of these scores are very encouraging and the overall user error rate is low, only $8.6\% = 7.5\% + 1.1\%$.

But these user errors must be taken with caution. The complementary point of view is to consider the model error, which indicates how well the statistical model performs: respectively 1.2% and 69% of the non-event and event points are forecast incorrectly. Thus, in spite of its very low rate of false alarms, the k nearest-neighbor predictor is handicapped by a relatively low detection rate: only about one third of the transitions are forecast.

How can we explain these apparently contradictory results? In the k nearest-neighbor classifier, we do not assign any particular cost to the two possible types of error, false negative vs. false positive. More precisely, we implicitly consider them to be equal when we choose to classify a point in the category best represented among its k nearest neighbors. The ratio of false negatives to false positives is actually imposed, in this algorithm, by the data, that is by the underlying dynamics and the variables used to forecast it. In the case of a rare event like an $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ transition, the overall error is dominated by the misses compared to the false alarms, with a ratio of about 7:1. One implication of this shortcoming is the relatively low detection rate of events, which may not be acceptable for a practical user. Random forests may be a good way to address this issue, as we shall see forthwith.

2) Random forests

In our first run of the random-forest algorithm, we let the data determine the default ratio of false negatives to false positives. As described in appendix B, a contingency table is built with data points not used to construct the classifier. The results, presented in the light grey cells of Table 3, are qualitatively similar to the previous ones and the HSS, $H = 0.36$, is also quite comparable. In this case, neither statistical classifier demonstrates a significant advantage over the other.

[Table 3 about here.]

An interesting property of random forests, though (see again appendix B), is the algorithm's ability to impose unequal cost weights on false negatives and false positives, and yield therewith different ratios between the two types of outcomes. One way of achieving this is by allowing the bootstrap samples used in generating each random tree to overrepresent transition events vs. the non-events.

In the previous experiment, the data gave a default ratio of about 7:1, with many more false negatives than false positives. The results so far suggest two additional experiments, in which we give a much greater weight to the misses than to the false alarms. The ratio of the two types of error is now inverted; more precisely, we tried to get them as close as possible to 1:4 and 1:8. The results of these two experiments are also shown in Table 3 in white cells (1:4 ratio) and dark grey cells (1:8 ratio).

The detection rate increases considerably as greater weight is given to the misses: it was initially only $28\% = 100\% - 72\%$ in the default case and it is now $72\% = 100\% - 28\%$ in the 1:4 ratio case and $82\% = 100\% - 18\%$ in the 1:8 case. The classifier is now much better at correctly predicting transitions, which was our initial goal.

But this improvement of detection rate comes at the detriment of the number of false alarms. It was only 1.7% in the default-ratio case and it rises to respectively 13% and 17% in the two new experiments. This modification of the detection and false-alarm rates have of course direct consequences for the errors that a user would expect. For a given forecast that indicates a transition, the probability to be wrong rises from 34% to 61% and 65%, respectively.

Note, finally, that the HSS remains of comparable size: it is now 0.43 and 0.40, in the two unequal-weight cases. It means the general skill of the classifier is not modified, what is modified is only the distribution of the error.

3) Optimizing predictor choice

In the experiments of Table 3, a subset of the predictors is sampled at random for each split within each tree (see section 3*c* and appendix B). This increases the flexibility of the fitting algorithm by allowing predictors that are important for small fractions of the data to enter the model. To evaluate the relative impact of each predictor on the forecasts, we present in Fig. 3 a plot of forecast sensitivity to the predictors.

[Figure 3 about here.]

This is an "importance plot" that shows the decrease of detection rate when using the random-forest algorithm, as each one of the six variables $(r, \theta, \phi, v_r, v_\theta, v_\phi)$ is rendered irrelevant to the forecasting process. More precisely, when forecasts are made, we keep the values of five predictors unchanged, while randomly shuffling all the values of the sixth variable, namely the one whose importance is being evaluated. The predictor is not removed but the shuffling randomizes its values, making them uncorrelated on the average with the class to which the point is supposed to belong, event or non-event. This process is repeated for each predictor. When each predictor is

shuffled in turn, we expect a decrease in the detection rate for each, because information is lost in the shuffling. The larger the drop in the detection rate, the more critical for the forecast is the shuffled variable.

Figure 3 was built with the 1:8 weight ratio between false positives and false negatives, but other choices of the weights (not shown) produce only very slight differences in the results and lead to the same conclusions. Namely, for the $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ transition (Fig. 3a), two variables, v_r and θ, are much more important that the four others. This result is consistent with and expands upon the conclusions of Kondrashov et al. (2004): it confirms the inhomogeneity of the transitions in phase space and the crucial influence of a preferred direction.

The importance of v_r may indicate that the points that are moving out of the cluster and thus away from the centroid are characterized by specific radial velocities which are presumably larger that the radial velocities of the other points. To assess this hypothesis, we built the PDFs of v_r for the two groups of interest, events and non-events, by using a Gaussian kernel estimator (Silverman 1986). We present the results in Fig. 4a. As expected, the transition points show, on average, larger values of v_r than the non-event points.

[Figure 4 about here.]

The sensitivity of the classifier to the high-impact variable θ is investigated by producing the "partial-dependence plot" in Fig. 5. This plot provides an estimate of the conditional probability of the forecast (in Log-Odds Units or "logits") with respect to the angular variable θ.

[Figure 5 about here.]

In general, the impact on classifier results of one particular variable depends on the values of the other predictor variables as well and cannot, therefore, be represented in a simple plot. The partial-dependence plot in Fig. 5 isolates the dependence of correctly forecasting an event on the

value of θ, by averaging over the values of the other predictors. In effect, the other predictors are held constant. The algorithm for computing the results in Fig. 5 is given in appendix C.

Of the two sensitivity plots, the importance plot (Fig. 3) indicates that θ is a critical predictor in the forecasting process, while the partial-dependence plot (Fig. 5) tells which values of θ are most likely to yield a transition forecast. The curve in Fig. 5a shows a fairly sharp peak for θ around $\pi/2$. It means that, as expected, a transition is more likely to be forecast for vectors that point in direction of the preferred transition path.

4) Other transitions

We carried out a similar study for the two other transitions of the Kondrashov et al. (2004) cycle (see section 2*c* above) : $\mathrm{NAO}^+ \rightarrow \mathrm{AO}^+$ and $\mathrm{AO}^+ \rightarrow \mathrm{NAO}^-$. We used only the random-forest algorithm, since the results when using the k nearest-neighbor classifier (not shown) were quite similar to those obtained when allowing the weights to be determined by the data in the random-forest case.

As in section $4a(2)$, we first let the data determine the ratio of false positives to false negatives and then we prescribe the relative weights of false outcomes so that this ratio equal about 1:4 and 1:8, respectively. Tables 4 and 5 are the contingency tables for these two transitions and they are both quite similar to the one already discussed. We can expect approximately the same performance in forecasting these two transitions as in Table 3 and the issue of detection rate is still critical.

[Table 4 about here.]

[Table 5 about here.]

The sensitivity plots in Figs. 3b,c differ more substantially from Fig. 3a than Tables 4 and 5

from Table 3. In the case of the $NAO^+ \rightarrow AO^+$ transition, the angle θ is clearly more important than all the other variables. The situation is very close to the one presented in the previous section. As seen in section 2*c*, this regime break is characterized by a sharp peak in the angular PDF of exits (Fig. 1b), which explains the importance of the angle θ, but the variable v_r is less important than in Fig. 3a. This state of affairs is confirmed by Fig. 4b, which shows that the values of v_r associated with the regime breaks are less well separated, in this case, from those of the non-events than in Fig. 4a.

The $AO^+ \rightarrow NAO^-$ transition has different properties still: a group of four variables have larger importance than the other two, with v_r still the first and θ being only the third in order of importance. As discussed in section 2*c*, the preferred exits are not confined in this case to a narrow solid angle but are much more widely spread out (Fig. 1c). The dynamics of this transition probably has a degree of complexity that requires several predictors, rather than just one or two.

We have also plotted in Figs. 5b,c the partial dependence plots for these two additional transitions. In spite of the differences noted between the three panels in Fig. 3 and those in Fig. 4, the results in these two panels resemble quite well those obtained for the first transition we studied, namely a large, albeit broader peak for large values of θ with a maximum close to $\theta = \pi/2$. Transitions are thus more likely to be forecast when the state vector is aligned with the preferred transition path, in all three cases.

b. Multiple-transition forecasts

We study here all the possible transitions of a point belonging to a given cluster. This leads to distinguishing five categories, or outcomes, for the forecast. On the one hand, when a transition does occur, the point leaves the cluster within the next 24 h to reach one of the four clusters,

including re-entry; this gives four possible outcomes, one per cluster. On the other hand, when the point remains in its cluster for at least 24 h more, we classify it into the fifth category called the "non-event."

The random-forest method is applied to the points that are in the NAO^- cluster at initial time. The same number of data, 6 000 days, is used as in the previous chapter. In the present situation, the state vector cannot be expressed in the same system of coordinates as above. Since there are four possible transitions for a given point, with four different preferred exit directions, it makes no sense to choose one or another of these directions as the pole of the coordinate system. Thus, the spherical coordinates were computed with the pole being aligned with EOF–3.

The results are shown in Table 6 which is a generalized contingency table that allows five possible outcomes. The rows still contain the observations and the columns the forecasts. The cells on the diagonal thus still correspond to forecasts that are correct. Although the different possible errors and their interpretation become more complex, we can define all the same two important types of errors: the false positives and the false negatives. The first type corresponds to the points that are actually non-events and that are forecast as transitions. They are located in the first row of the contingency table. The second type includes the points that are transitions and that are forecast as non-events. These correspond to the first column of the contingency table. In addition, we have now a new type of error that did not exist in the two-outcome case: a transition point whose destination cluster is not correctly forecast. A point that is going to cluster AO^+ and that has been classified in the AO^- transition group would fall into this category.

[Table 6 about here.]

We performed only two multiple-outcome experiments with different ratios of false positives to false negatives. One is the control experiment, which lets the data set the weights, and the second

is an experiment that assigns a higher cost to false negatives, so as to achieve a higher detection rate. The control experiment yields the same result as in the two-outcome case: the false negatives are much more numerous than the false positives and the detection rate is low.

In this more general case, the overall user error is the complement of the correct forecasts, that is the complement of the diagonal elements. This error equals $26\% \approx 100\% - (62\% + 4.7\% + 4.9\% + 2.2\% + 0.1\%)$ and it is much higher in Table 6 than in Tables 2–5, where it does not exceed 16%. Indeed, the forecast of multiple outcomes is much more difficult than for only two outcomes, especially when each type of transition is a relatively rare event.

In the other experiment, with a higher weight on false negatives, we get a better rate of detection, and thus succeed in forecasting about half the transitions. The accuracy of the forecasts differs from transition to transition: the best results are obtained for the AO^- destination cluster, with a model error of only 32%, while the worst results are for NAO^+, with a model error of 57%. Once again, the results are considerably worse than in the two-outcome case, where the model error at predicting a transition was about 20%. The practical interest of a multi-outcome statistical forecast is therefore more limited than for a simpler case.

5. Concluding remarks

In this article, we have studied the predictability of the Northern Hemisphere's low-frequency variability in an intermediate-complexity model: the quasi-geostrophic, three-layer (QG3) model of Marshall and Molteni (1993). This model (section 2*a*) exhibits four significant weather regimes in a low-dimensional subspace spanned by the three leading EOFs of its variability. Kondrashov et al. (2004) showed that certain regime transitions in the QG3 model are characterized by preferred-direction paths in this phase space (see Fig. 1).

Our goal here was to use these specific features in phase space to forecast the regime breaks in advance. To do so, we used two statistical tools: the classical k nearest-neighbor classifier (section 3b) and the novel random-forest method (section 3c). The application of both methods to medium-to-long-range prediction of large-scale flow patterns appears to be new.

The model's EOFs and weather regimes (section 2b) were computed using a 54 000 day, perpetual-winter simulation. To put the statistical forecast methods under study to a more severe test, we used only a 6 000 day segment of this simulation as a learning set; this corresponds to 50 winters, each 120 days long, which could be obtained from the existing reanalysis of atmospheric observations.

We first focused on forecasting single transitions and obtained surprisingly good predictability, even with this short learning set. We considered the cycle of three transitions $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+ \rightarrow \mathrm{AO}^+ \rightarrow \mathrm{NAO}^-$ and, for each of the three, the statistical prediction is about 35 to 40% better than random (see Tables 3–5).

A major obstacle in correctly predicting regime transitions is the fact that these are fairly rare events. In practical situations, though, misses and false alarms may be given different weights, in particular when the two types of forecast outcomes are qualitatively different. The random-forest method allows one to easily assign distinct costs to false positives vs. false negatives. Of course, any improvement in the detection rate of transitions is inevitably associated with a larger number of false alarms and vice versa. Eventually it is the end-user's choice to define precisely what risk is acceptable according to the prospective application of the forecast. Since the transitions of interest are rare events, we were able to obtain higher detection rates by assigning higher weights to the misses than to the false alarms, while keeping the overall skill unchanged.

The preferred transition paths identified by Kondrashov et al. (2004) were found to carry predictive information on regime transitions. Sensitivity studies to different predictors, through

importance (Fig. 3) and partial-dependence (Fig. 5) plots showed the key role of the deviation angle θ formed by the exit vector with the preferred exit direction. These studies indicate that a transition is more likely to be forecast for points aligned with the preferred transition direction. We also found that the influence of θ is more crucial when the preferred transition path is confined within a fairly sharp solid angle: underlying exit dynamics seems to be largely dependent on θ in this case, although the velocity component v_r along the preferred exit direction also plays a role. The role of θ decreases when the exit-vector PDF is not limited to a narrow angle but is more spread out.

The results for the single-transition case are encouraging in view of a practical use of statistical methods in medium-to-long-range forecasting. These results provide further support for the Legras and Ghil (1985) conjecture that (i) certain atmospheric-flow regimes are associated with unstable fixed points in the flows' phase space; and, hence, (ii) exit from such regimes and subsequent transitions to other regimes originate along preferred directions of unstable growth of perturbations.

The natural development of the present work would be to study observed data, where preferred transition paths were also hypothesized by Kimoto and Ghil (1993a,b). This will make it possible to compare the skill of statistical and dynamical models on specific transitions like those between zonal and blocked states. Such transitions are of real meteorological interest and remain a problem for numerical weather prediction models.

Acknowledgements

It is a pleasure to thank D. Kondrashov for his providing the code to preprocess the data used in this article and for his advice in using it. This work grew up out of a DEA diploma project

carried out by A. Deloncle under the guidance of M. Ghil and O. Talagrand; both A. Deloncle and M. Ghil appreciate O. Talagrand's input. R. Berk would like to thank the Ecole Normale Supérieure and the Ecole Polytechnique for their hospitality and support during a sabbatical term in Paris. Support for this work was provided by NSF grants SES04-37169 ("Ensemble Methods for Data Analysis in the Behavioral, Social and Economic Sciences," R. Berk) and ATM00-82131 (M. Ghil), as well as by the European Commission's Project No. 12975 (NEST) "Extreme events: Causes and consequences (E2-C2)" (F. D'Andrea and M. Ghil).

APPENDIX A

Average source term

In the atmospheric model governed by Eq. (A1), the time-independent forcing S represents sources of potential vorticity that result from processes not explicitly included in the equations: radiative forcing, other diabatic heat fluxes (linked, for example, to precipitation), and the effect of divergent flow. In addition the forcing implicitly contains the effects of subgrid-scale processes. The forcing term has been estimated here empirically, following Marshall and Molteni (1993), as follows.

From a long series of wintertime analyzed states, one can substitute $\hat{q}$ and $\hat{\psi}$ into Eq. (A1), for every day of observed fields available; the hat indicates observed fields. Equation (A1) holds for observed fields and gives a value of S for that day. Taking then the time average, represented by the overbar, an equation for a mean field S is obtained:

$$S = \overline{J(\hat{\psi}, \hat{q})} + D(\overline{\hat{\psi}}). \tag{A1}$$

Daily streamfunction fields were obtained from the ECMWF operational analysis for the months of January and February of the years 1984–92.

APPENDIX B

Random forests

With categorical predictands such as those used in this paper, random forests provides a classification method. The intent is to assign classes to observations using information contained in a set of predictors. A random forest is constructed from a large number of classification trees, each tree based on a random sample (with replacement) of the data, and for each partitioning of the data for each tree, a random sample of predictors. Classification trees will be described briefly, before explaining random forests. For ease of exposition, and with no major loss of generality, we consider in the following only a binary response variable: only two outcomes are possible, for instance "event" and "non-event."

a. Classification trees

Each classification tree provides a recursive partitioning of a training data set. The goal is to construct contiguous subsets within the space defined by the predictors that are less heterogeneous than the data before the partitioning. All possible predictors are screened before a potential partitioning of the data is selected; the predictor eventually used at each step is the one that decreases heterogeneity the most. Two popular measures of heterogeneity are entropy E, defined in the binary outcome case as $E = -p \log p - (1-p)\log(1-p)$, and the Gini index G, defined as $G = p(1-p)$. In section $4a$ here, p is for instance the proportion of "event" points in a data partition, with $1-p$ the proportion of "non-events."

Figure 6 represents a simple example. There is a binary response coded "A" or "B," and two predictors x and y. The single vertical line at $x = 3$, say, produces the first partition. The double horizontal line at $y = 6$ produces the second partition. The triple horizontal line at $y = -4$

produces the third partition. Partition boundaries must be straight lines perpendicular to the predictor axes.

[Figure 6 about here.]

In this simple illustration, the upper-left set and the lower-right set are fully homogeneous. There remains considerable heterogeneity in the other two sets and, in principle, their partitioning could continue. When there are no longer any ways to further partition the data to make them more homogeneous, the algorithm stops. Each final set is then assigned a class, based on a majority vote of the observations in that set. Here either class A or class B would be assigned to a set according to which has a greater proportion of observations in that set. The classification of a new point not included in the training data set requires only to determine in which set the observation lies and the associated class.

b. Random forests

Random forests generalizes classification trees by considering a large set of trees generated by a process that introduces random factors. Let n be the number of training observations on hand. The random-forest method then operates with the following steps:

1. Take a random sample of size n *with replacement* from the total data set on hand.
2. Take a random sample *without replacement* of all the possible choices of predictors included in the data.
3. Construct the first data partition of a classification tree.
4. Repeat Step 2 and Step 3 for each subsequent split, until the classification tree is as deep as desired. Do not prune the tree.

5. Drop the data *not* included in the sample from Step 1 down the tree. Store the class assigned to each observation along with each observation's predictor values.

6. Repeat Steps 1–5 a large number of times (we used 500 trees in this paper), so that there is a large number of trees, which constitute a random forest.

7. Using only the class assigned to each observation when that observation is *not* used to build a tree, count the number of times over trees that the observation is classified in one outcome category and the number of times over trees it is classified in the other outcome category.

8. Assign each observation to one of the two outcome classes by a majority vote over the set of trees.

Random forests has five demonstrable assets. First, for the kinds of data analyzed in this paper, there are no classifiers to date that will consistently classify and forecast more accurately. Most will do worse, especially when the true relationships with the response are highly nonlinear and noisy. Second, one can prove (Breiman 2001) that random forests does not overfit. This is very important because it implies that the results will generalize well to new random samples from the same population (i.e., data with the same characteristics except for random sampling error). Third, because performance is determined by a contingency table computed from observations not used to construct a given tree (i.e., observations not selected in Step 1), performance rests on real forecasting skill. Fourth, random forests provides a means by which the relationships between inputs and outputs can be represented in an instructive way, using importance plots and partial-dependence plots.

Finally, there are several systematic ways in which the relative costs of false negatives and false positives can be taken into account. The approach used in this paper gives more weight to observations in which a transition occurs, so that if such observations are misclassified, the

consequences are greater. This is accomplished by oversampling transition events when bootstrap samples are drawn for each tree; in other words, transition events are made more common in the analysis than they are in the data. The presence of such a random element in determining the weights of events vs. non-events is the reason for achieving a targeted weight ratio only approximately in Tables 3–5.

This is an improvement over the more classical k nearest-neighbor algorithm. Indeed, we tried to force different cost weights with the latter method by modifying the classification process. A point was forecast as an event if $k_{event}/k > a$ with k_{event} the number of event points in its k nearest neighbors and a a parameter setting the relative cost of false positives to false negatives. Equal costs corresponds to $a = 0.5$, while $a < 0.5$ gives more weight to the false negatives compared to the false positives. We could not get the same results than random forests especially for the larger ratio 1:8; the overall skill dropped making the forecast of no practical interest. This is probably a consequence of the limited size of the data set that imposes to choose a small value of k and thus, does not allow to finely tune the value of a without being subject to critical sampling problems. This is not an issue for random forests because the classification is made through a majority vote over a large number of trees (we used 500 trees in this study) and not among a small number of k nearest neighbors.

Breiman (2001) give a formal exposition of classification and regression trees (CART), while Breiman et al. (1984) provides a formal exposition of random forests. An excellent reference to statistical learning in general is Hastie et al. (2001).

APPENDIX C

Partial-dependence plots

Partial-dependence plots display in Log-Odds Units ("logits") how the probability of a particular event (here, a transition) is related to a given predictor, the values of all other predictors being fixed. A partial-dependence plot is constructed in the following manner.

1. Grow a forest.

2. Suppose x has v distinct values in the training data set. Construct v data sets as follows. For each of the v values of x, make up a new data set where x only takes on that value, leaving all other variables untouched.

3. For each of the v data sets thus obtained, predict the response using random forests.

4. For each of the v data sets, average these predictions determining the proportions p and $1-p$ of trees that respectively forecast an event and a non-event. Compute in logits the ratio of these proportions, $R = 0.5 \log[p/(1-p)]$.

5. Finally, plot this ratio R (expressed in logits) for each of the v values of x.

Thus, partial-dependence plots show the relationship between a given predictor x and the response averaged over the joint values of the other predictors as they are represented in the tree structure. In this way, the other predictors are being "held constant" by matching, so that no assumptions are being made about how the predictors are related to one another or to the response variable. More details about partial-dependence plots can be found in Hastie et al. (2001).

References

Breiman, L., 2001: Random forests. *Machine Learning*, **45**, 5–32.

Breiman, L., J. Friedman, R. Olshen, and C. Stone, 1984: *Classification and Regression Trees*. Wadsworth Press, Monterey, CA., 368 pp.

Corti, S., A. Giannini, S. Tibaldi, and F. Molteni, 1997: Patterns of low-frequency variability in a three-level quasi-geostrophic model. *Climate Dyn.*, **13**, 883–904.

D'Andrea, F., 2001: Extratropical low-frequency variability as a low-dimensional problem. Part II: Stationarity ans stability of large-scale equilibria. *Quart. J. Roy. Meteor. Soc.*, **128**, 1059–1073.

D'Andrea, F., and R. Vautard, 2000: Reducing systematic errors by empirically correcting model errors. *Tellus*, **52A**, 21–41.

——, and ——, 2001: Extratropical low-frequency variability as a low-dimensional problem. Part I: A simplified model. *Quart. J. Roy. Meteor. Soc.*, **127**, 1357–1374.

D'Andrea, F., and Coauthors, 1998: Northern Hemisphere Atmospheric Blocking as simulated by 15 general circulation models in the period 1979–1988. *Climate Dyn.*, **14**, 385–407.

Ferreira, D., and C. Frankignoul, 2005: The transient atmospheric response to midlatitude SST anomalies. *J. Climate*, **18(7)**, 1049–1067.

Fraedrich, K., and M. Klauss, 1983: On single station forecasting: Sunshine and rainfall Markov chains. *Beitr. Phys. Atmosph.*, **56**, 108–134.

Ghil, M., and A. W. Robertson, 2002: "Waves" vs "Particles" in the atmosphere's phase space: A pathway to long-range forecasting? *Proc. Natl. Acad. Sci.*, **99**, 2493–2500.

Ghil, M., M. Kimoto, and J. D. Neelin, 1991: Nonlinear dynamics and predictability in the atmospheric sciences. *Rev. Geophys., Suppl., U.S. National Report to the International Union of Geodesy and Geophysics*, **29**, 46–55.

Hastie, T., R. Tibshirani, and J. Friedman, 2001: *The Elements of Statistical Learning*. Springer-Verlag, New York., 552 pp.

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

——, and ——, 1993b: Multiple flow regimes in the Northern Hemisphere winter. Part II: Sectorial regimes and preferred transitions. *J. Atmos. Sci.*, **50**, 2645–2673.

Kondrashov, D., K. Ide, and M. Ghil, 2004: Weather regimes and preferred transition paths in a three-level quasigeostrophic model. *J. Atmos. Sci.*, **61**, 568–587.

Legras, B., and M. Ghil, 1985: Persistent anomalies, blocking and variations in atmospheric predictability. *J. Atmos. Sci.*, **42**, 433–471.

Lorenz, E. N., 1969: Three approaches to atmospheric predictability. *Bull. Amer. Meteor. Soc.*, **50**, 345–349.

Marshall, J., and F. Molteni, 1993: Towards a dynamical understanding of planetary-scale flow regimes. *J. Atmos. Sci.*, **50**, 1792–1818.

Michelangeli, P. A., R. Vautard, and B. Legras, 1995: Weather regimes: Recurrence and quasi stationarity. *J. Atmos. Sci.*, **52**, 1237–1256.

Mo, K. C., and M. Ghil, 1988: Cluster analysis of multiple planetary flow regimes. *J. Geophys. Res.*, **93D**, 10927–10952.

Molteni, F., 2002: Weather regimes and multiple equilibria. *Encyclopedia of Atmospheric Sciences*, J. Holton, J. Cury, and J. Pyle, Eds., Academic Press, Volume 6, 2577–2586.

Pelly, J., and B. Hoskins, 2003: How well does the ECMWF ensemble prediction system predict blocking. *Quart. J. Roy. Meteor. Soc.*, **129**, 1683–1703.

Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 175 pp.

Smyth, P., K. Ide, and M. Ghil, 1999: Multiple regimes in Northern Hemisphere height fields via mixture model clustering. *J. Atmos. Sci.*, **56**, 729–752.

Tibaldi, S., and F. Molteni, 1990: On the operational predictability of blocking. *Tellus*, **42a**, 343–365.

Von Storch, H., and F. Zwiers, 1999: *Statistical Analysis in Climate Research*. Cambridge Univ. Press, 484 pp.

List of Figures

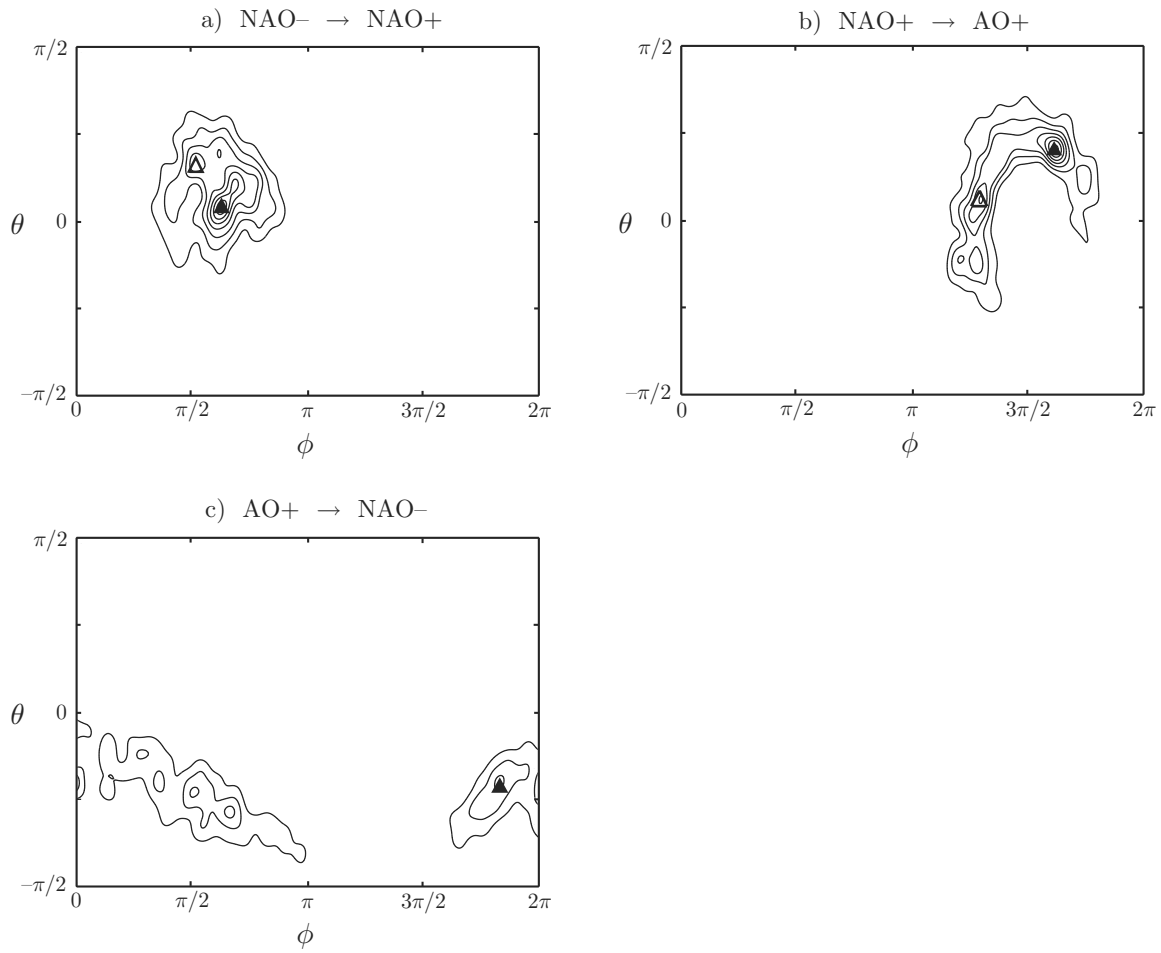

Figure 1: Probability density functions (PDFs) of regime exit angles θ and ϕ. Filled triangles correspond to global PDF maxima, while open triangles mark strong secondary maxima. The contour interval for all panels is equal to 0.2 in non-dimensional units. This figure was produced following the procedure described in Kondrashov et al. (2004).

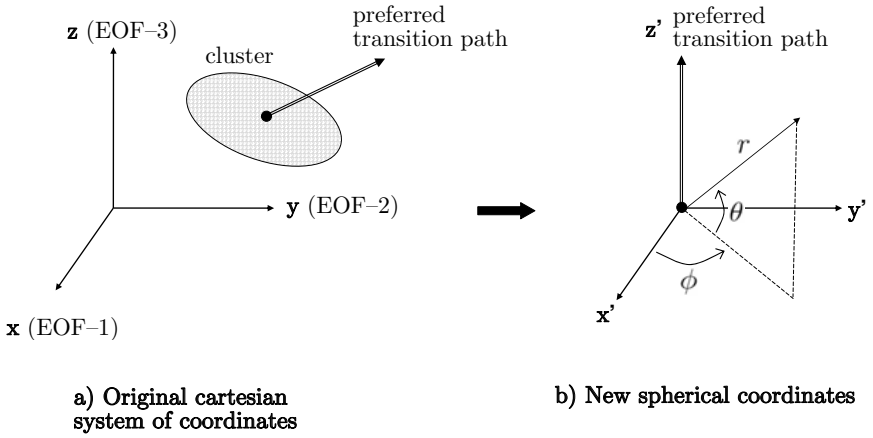

Figure 2: Change of coordinate system to take into account the existence of a preferred direction of transition. In the new coordinate system, θ is related to the angle formed by the state vector with the preferred direction of transition.

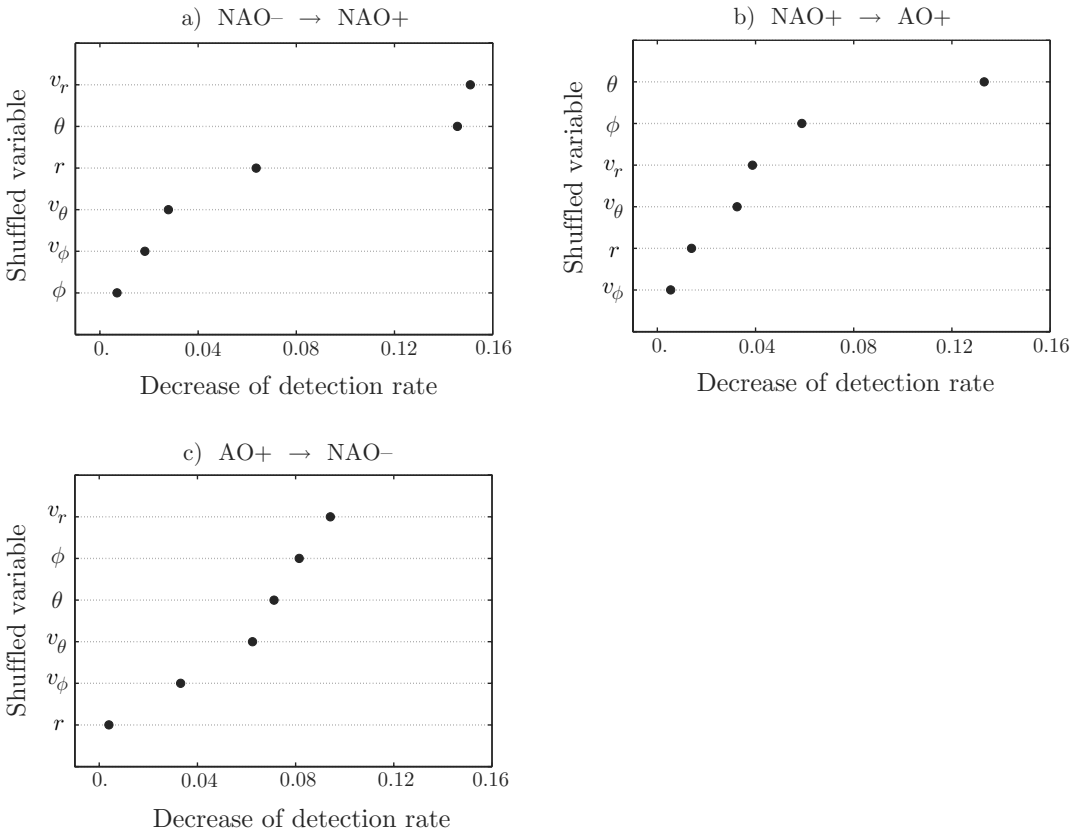

Figure 3: Relative importance of the predictors. The plot shows the decrease in detection rate when a variable is shuffled and measures the importance of each variable in the forecasting process.

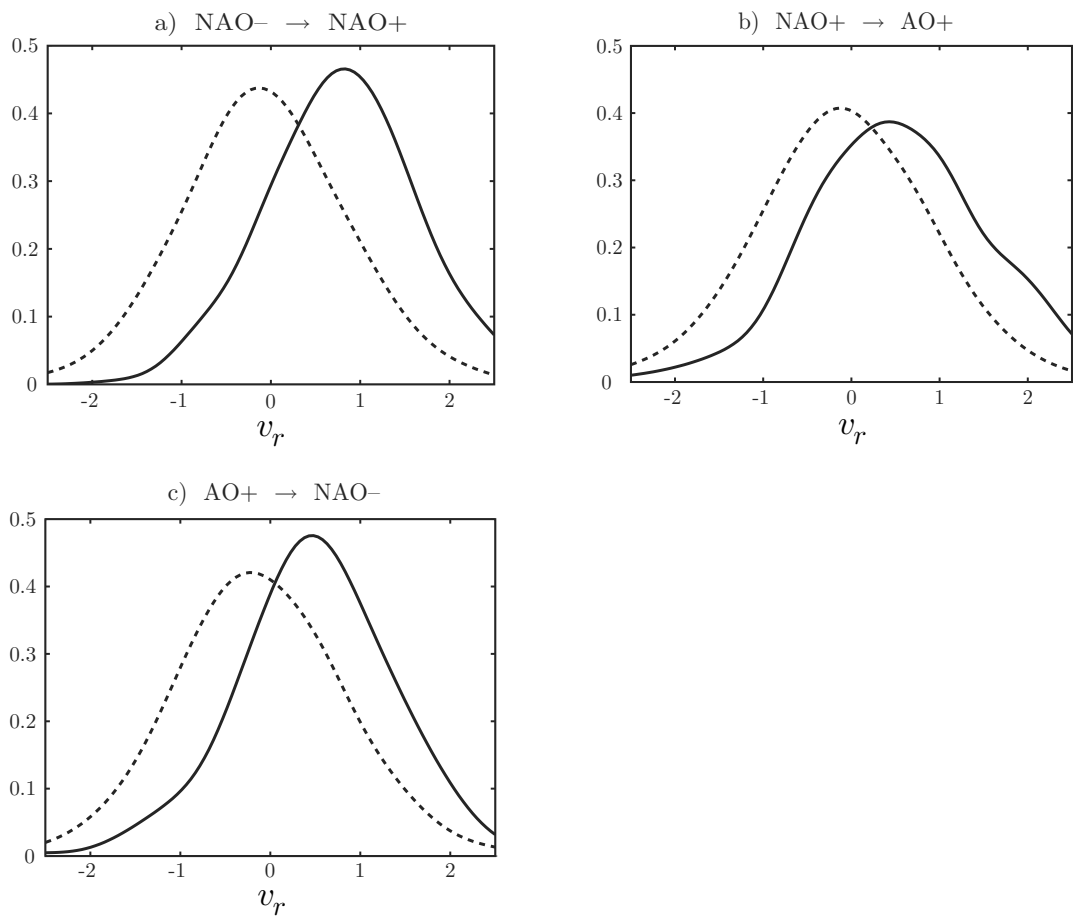

Figure 4: Probability density functions (PDFs) of v_r for event points (solid line) and non-event points (dashed line). The event points, which are associated with regime breaks have on average larger values of v_r.

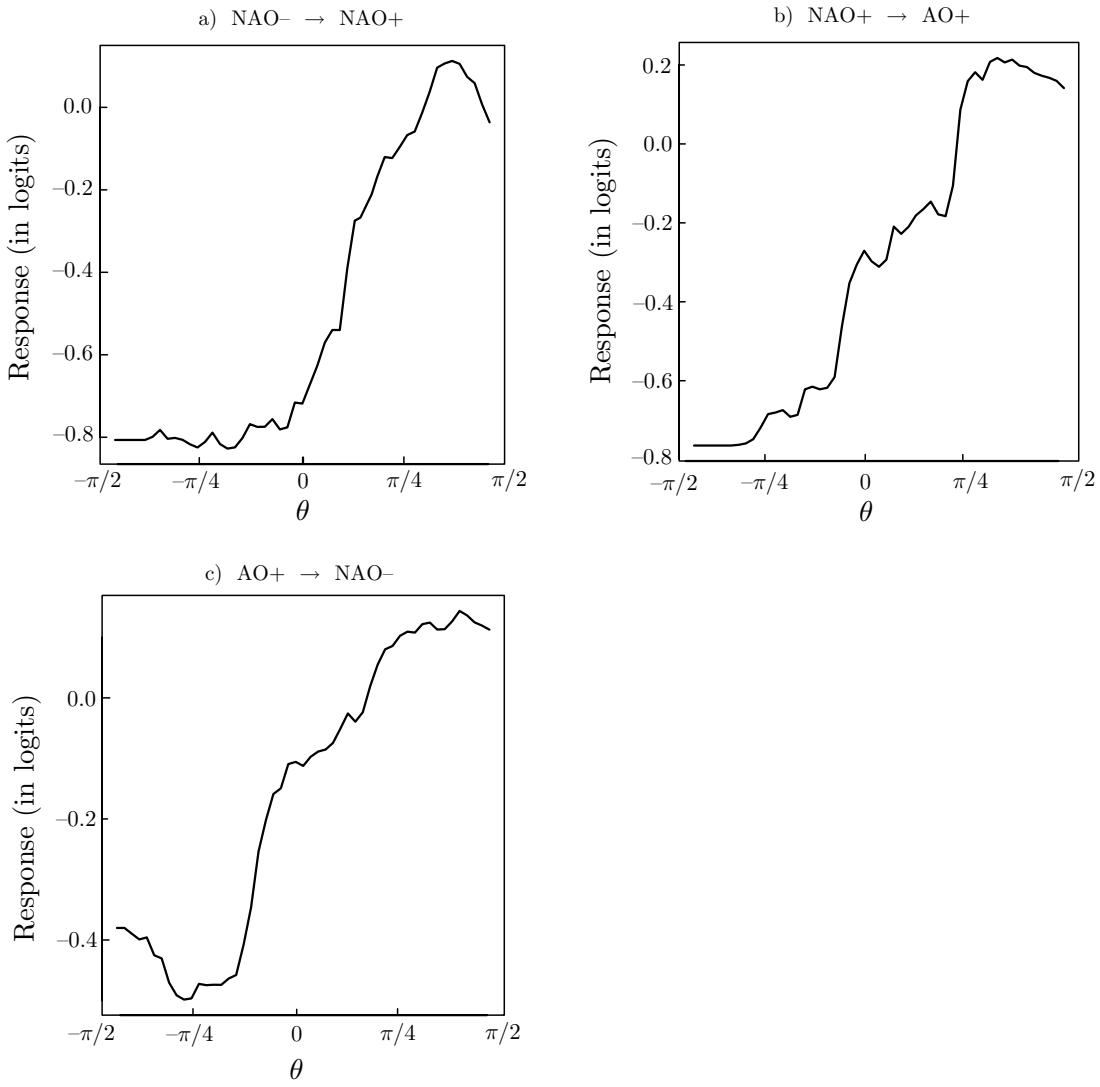

Figure 5: Partial-dependence plots for θ. The ordinate is the logit of the conditional probability of event detection with respect to the high-impact variable θ.

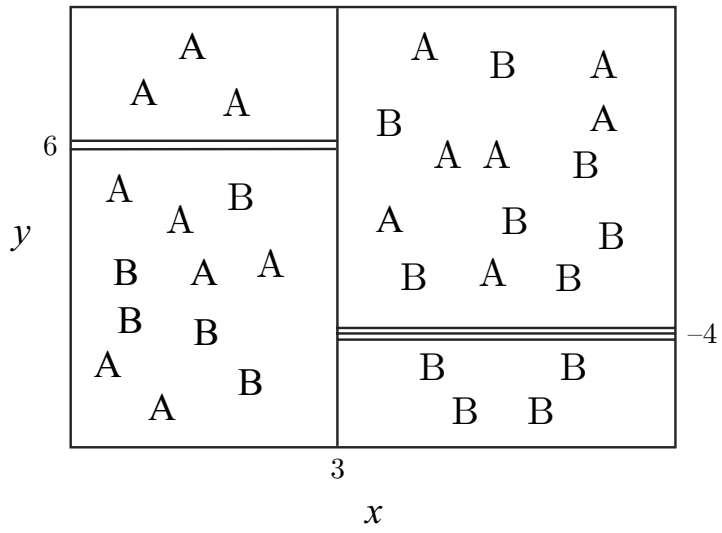

Figure 6: Recursive partitioning used in classification trees. There is a binary response coded "A" or "B" and two predictors x and y.

List of Tables

		Forecast		Model error
		non-event	event	
Observed	non-event	*a (true negatives)*	*b (false alarms)*	*b/(a+b)*
	event	*c (misses)*	*d (true positives)*	*c/(c+d)*
User error		*c/(a+c)*	*b/(b+d)*	

Table 1: Definition of a 2×2 contingency table. The observations (actual category) of the points are in the rows and the forecasts in the columns. The numbers a, b, c, d are the percentages of each case obtained on the assessment data set ($a+b+c+d = 100$). Thus, true forecasts are on the diagonal and correspond to true negatives a and true positives d. The misclassified points are off the diagonal and consist in the false positives (false alarms) b and the false negatives (misses) c; the overall user error is the sum of the off-diagonal elements $b + c$.

		Forecast		*Model error*
		non-event	event	
Observed	non-event	**88.1**	**1.1**	*1.2*
	event	**7.5**	**3.3**	*69.4*
User error		*7.8*	*25.0*	

Table 2: Contingency table with k nearest-neighbor classifier for the transition $\mathrm{NAO}^- \to \mathrm{NAO}^+$. A value of $k = 9$ nearest neighbors was found to give the best results.

		Forecast						Model error		
		non-event			event					
Observed	non-event	**88.3**	**78.5**	**74.4**	**1.5**	**11.3**	**15.4**	*1.7*	*12.6*	*17.2*
	event	**7.3**	**2.9**	**1.9**	**2.9**	**7.3**	**8.3**	*72*	*28.0*	*18.3*
User error		*7.7*	*3.5*	*2.4*	*34.3*	*60.7*	*64.9*			

Table 3: Contingency table with random-forest algorithm for the transition $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$; 500 trees were used and two variables were tried at each split. Results are shown for three different ratios of false negatives to false positives: the default ratio imposed by the data (light grey cells) and two other ratios, approximately 1:4 (white cells) and approximately 1:8 (dark grey cells).

		Forecast						Model error		
		non-event			event					
Observed	non-event	**83.3**	**67.9**	**63.7**	**2.0**	**13.3**	**21.6**	*2.3*	*20.3*	*25.4*
	event	**11.5**	**4.0**	**2.8**	**3.2**	**10.7**	**12.0**	*78.1*	*27.1*	*18.8*
User error		*12.1*	*5.5*	*4.2*	*38.2*	*61.7*	*64.4*			

Table 4: Contingency table with random forests for the transition $\mathrm{NAO}^+ \rightarrow \mathrm{AO}^+$. Same algorithm and presentation as in Table 3.

		Forecast						Model error		
		non-event			event					
Observed	non-event	**77.0**	**65.9**	**59.9**	**4.1**	**15.2**	**21.2**	*5.0*	*18.7*	*26.2*
	event	**12.0**	**3.7**	**2.6**	**6.9**	**15.2**	**16.3**	*63.5*	*19.8*	*13.8*
User error		*13.5*	*5.4*	*4.2*	*37.1*	*50.0*	*56.5*			

Table 5: Contingency table with random forests for the transition $\mathrm{AO}^{+} \rightarrow \mathrm{NAO}^{-}$. Same algorithm and presentation as in Table 3.

		Forecast										Model error	
		non-event		NAO^+		AO^+		AO^-		NAO^-			
Observed	non-event	**61.7**	**39.7**	**0.6**	**2.5**	**1.0**	**8.6**	**1.0**	**6.5**	**0.2**	**7.5**	*4.6*	*38.7*
	NAO^+	**4.1**	**1.1**	**4.7**	**4.4**	**1.0**	**1.1**	**0.2**	**0.4**	**0.1**	**3.2**	*53.7*	*57.3*
	AO^+	**6.1**	**2.4**	**0.6**	**0.6**	**4.9**	**7.1**	**0.0**	**0.5**	**0.6**	**1.6**	*60.2*	*41.8*
	AO^-	**4.4**	**0.4**	**0.7**	**0.4**	**0.2**	**0.2**	**2.2**	**5.3**	**0.1**	**1.5**	*71.4*	*31.7*
	NAO^-	**2.9**	**0.7**	**1.2**	**1.2**	**0.4**	**0.0**	**0.5**	**0.7**	**0.1**	**2.4**	*97.6*	*53.7*
User error		*22.1*	*10.4*	*40.6*	*52.1*	*35.0*	*58.4*	*45.5*	*60.2*	*90.0*	*85.4*		

Table 6: Contingency table with random forests for every possible transition starting from the NAO^- cluster. Two different experiments are presented: in the first one (light grey), we let the data impose the detection rate; in the second one (white), we tried to get a higher detection rate.

Weather Regimes and Preferred Transition Paths in a Three-Level Quasigeostrophic Model

D. KONDRASHOV, K. IDE, AND M. GHIL*

Department of Atmospheric Sciences, and Institute of Geophysics and Planetary Physics, University of California, Los Angeles, Los Angeles, California

(Manuscript received 28 February 2003, in final form 29 August 2003)

ABSTRACT

Multiple flow regimes are reexamined in a global, three-level, quasigeostrophic (QG3) model with realistic topography in spherical geometry. This QG3 model, using a T21 triangular truncation in the horizontal, has a fairly realistic climatology for Northern Hemisphere winter and exhibits multiple regimes that resemble those found in atmospheric observations. Four regimes are robust to changes in the classification method, *k*-means versus mixture modeling, and its parameters. These regimes correspond roughly to opposite phases of the Arctic Oscillation (AO) and the North Atlantic Oscillation (NAO), respectively.

The Markov chain representation of regime transitions is refined here by finding the preferred transition paths in a three-dimensional (3D) subspace of the model's phase space. Preferred transitions occur from the positive phase of the NAO (NAO^+) to that of the AO (AO^+), from AO^+ to NAO^-, and from NAO^- to NAO^+, but not directly between opposite phases of the AO. The angular probability density function (PDF) of the regime exits that correspond to these preferred transitions have one or, sometimes, two fairly sharp maxima. These angular PDF maxima are, in most cases, not aligned with the line segments between regime centroids in phase space and might point to heteroclinic or homoclinic connections between unstable equilibria in the model's phase space. Preferred transitions paths are also determined for a stochastically forced Lorenz system to help explain this striking feature of the QG3 model.

The episodic description of the model's low-frequency variability via the Markov chain of multiple regimes is complemented by an oscillatory description. Multichannel singular-spectrum analysis is applied to the trajectory in the same 3D subspace. Two statistically significant oscillations are found and have periods of 19 and 37 days, respectively. Both oscillations have four composites that include NAO^+, AO^+, and NAO^-, in this order. The fourth composite occurs between AO^+ and NAO^-; it resembles the Pacific–North American pattern, which is not captured by the model's episodic description. The two oscillations have similar spatial patterns, and are weakly phased locked. They have certain features in common with the westward-propagating Branstator–Kushnir wave, as well as with the standing oscillation that arises from the oscillatory topographic instability of Ghil and associates.

1. Introduction and motivation

Numerous studies of atmospheric observations have shown that low-frequency variability (LFV) is characterized by the existence of large-scale persistent and recurrent flow patterns, also called weather regimes. Weather regimes can be objectively identified in model results, as well as in observations, using various classification or clustering methods (see Table 1 in Ghil and Robertson 2002). The dominant regimes of Northern Hemisphere (NH) wintertime circulation are most often identified as the Pacific–North American (PNA) pattern, the reverse PNA (RNA), the North Atlantic Oscillation (NAO), and the Arctic Oscillation (AO). The positive and negative phases of the NAO correspond to zonal and blocked flow in the Atlantic sector, in the same way the PNA and RNA do so over the Pacific sector of the NH. The AO (Thompson and Wallace 1998; Wallace 2000) is a hemispheric, annular mode that has been strongly associated with the sectorial NAO.

Marshall and Molteni (1993) introduced a spectral, three-level, quasigeostrophic (QG3) model in spherical geometry and showed that it has a fairly realistic NH winter climatology; it also exhibits multiple weather regimes that resemble those found in observations. These authors computed so-called neutral vectors that have the smallest time derivative, by linearizing the model's equations around its mean state. These neutral vectors were shown to be associated with the subspace of the

* Additional affiliation: Département Terre-Atmosphère-Océan and Laboratoire de Météorologie Dynamique, Ecole Normale Supérieure, Paris, France.

Corresponding author address: Dr. Dmitri Kondrashov, Department of Atmospheric Sciences, Institute of Geophysics and Planetary Physics, University of California, Los Angeles, 405 Hilgard Ave., Los Angeles, CA 90095-1565.
E-mail: dkondras@atmos.ucla.edu

system's phase space in which two quasi-stationary states are located. The long-time integration of the full QG3 model with a specially chosen forcing generated weather regimes similar to those defined by the neutral vectors.

The QG3 model has been widely used for studies of NH atmospheric LFV. Corti et al. (1997) analyzed leading NH teleconnections patterns and blocking from a long run of the QG3 model. The forcing sources in this run were computed using the European Centre for Medium-Range Weather Forecasts (ECMWF) objective analyses of the streamfunction field in wintertime and the condition of zero tendency for the potential vorticity. The model reproduced the observed wintertime mean state, as well as low-frequency and high-frequency variations of the atmospheric streamfunction field. Both the PNA and the NAO patterns were realistically simulated; the statistics of blocking frequency and duration in the Euro-Atlantic and Pacific sectors compared also reasonably well with the observations. Molteni and Corti (1998) investigated the dynamical origin of long-term variations in the statistical properties of the QG3 model's LFV.

D'Andrea and Vautard (2001) and D'Andrea (2002) analyzed the correspondence between global quasi-stationary states found in a low-order model and the QG3 model's weather regimes, as identified by cluster analysis. Their low-order model was constructed by projecting the equations on a few leading EOFs of the QG3 model. These authors found the flow-dependent parameterization of a closure term to be essential for the good performance of their low-order model. Given such a closure, approximate correspondence between the QG3 weather regimes and quasi-stationary states of the low-order model was obtained.

Extended-range weather prediction depends in a crucial way on skill at forecasting the duration of a regime event or other persistent anomaly, which is under way at initial forecast time, and the subsequent onset of another persistent anomaly, after the break of the current one (Ghil 1987). We use the QG3 model to address extended-range predictability in terms of the multiple-regime paradigm (Ghil and Robertson 2002). Our main purpose is to describe more precisely the preferred transition paths between regimes.

In section 2, the model is described briefly. In section 3, we identify robust weather regimes by two independent clustering methods, *k*-means (Michelangeli et al. 1995) and Gaussian mixture modeling (Smyth et al. 1999). In section 4, we show that highly significant transitions take place along preferred paths in a low-dimensional phase space. For better understanding, we also examine in the appendix the preferred transitions paths in a stochastically forced Lorenz (1963a) system; in this case, the paths can be fully explained using known dynamical properties of the underlying nonlinear system.

In section 5 we complement the episodic description of the model's LFV, as given in sections 3 and 4, by an oscillatory one: multichannel singular-spectrum analysis is used to identify intraseasonal oscillations and connect them with the Markov chain of transitions between regimes. In section 6, we argue that the preferred transitions paths demonstrated in the QG3 model suggest that a similar analysis of numerical weather prediction models could help forecast regime breaks and subsequent onsets. To do so would require an efficiently designed observational system that is able to track spatial signatures of regime transitions.

2. Atmospheric model

The global model of Marshall and Molteni (1993) is governed by the equations for conservation of potential vorticity at the 200-, 500-, and 800-hPa pressure levels, written in shorthand notation as

$$\frac{\partial q}{\partial t} = -J(\psi, q) - D(\psi) + S, \tag{1}$$

where t is time; ψ is the streamfunction; q the potential vorticity; J the quadratic Jacobian operator; and D represents linear dissipation, in particular Newtonian cooling, Ekman dissipation on the 800-hPa wind, and hyperviscosity. The drag coefficients depend on the nature of the underlying surface. The model uses an expansion in spherical harmonics with a triangular truncation of T21 and realistic topography. The gridpoint values of the topography and sea–land mask represent averages over areas of 1000 km^2.

Equation (1) is forced by time-independent, but spatially varying sources S of potential vorticity. These sources represent the average effects of diabatic heating and advection by divergent flow, and are determined from a condition requiring that, for a given climatological dataset of observed fields, the time derivatives of vorticity vanish. The dataset of observed fields is given by ECMWF analyses of the streamfunction field at the three levels of the QG3 model, for the two winter months January and February, from 1984 to 1992. With this forcing, the QG3 model reproduces the observed climatology of the underlying dataset, as well as providing good simulations of wintertime midlatitude variability. Marshall and Molteni (1993) and Corti et al. (1997) provide details of the model and of its performance.

3. Cluster analysis

The dataset for this analysis was obtained from a 54 000-day perpetual-winter simulation of our QG3 model. In order to examine robust features of the model's phase-space structure, it is necessary to reduce the dataset's dimensionality. For this purpose, we apply empirical orthogonal function (EOF) analysis to the unfiltered 500-hPa-level streamfunction anomalies in the model's NH. The anomaly field is sampled once a

TABLE 1. Classifiability index of the *k*-means algorithm: *k* is the prescribed number of clusters, while *d* is the number of the EOFs retained for the analysis.

	$k = 3$	$k = 4$	$k = 5$
$d = 2$	1.000	0.973	1.000
$d = 3$	1.000	1.000	1.000
$d = 4$	1.000	1.000	0.819
$d = 5$	1.000	1.000	0.807

TABLE 2. Cross-validated log-likelihood for 18 000 data points, using only the first third of the full dataset: the lower the absolute value of this score, the more likely the corresponding *k* value is for a given *d*.

	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$
$d = 2$	−22 770	−22 296	−22 255	−22 248	−22 245
$d = 3$	−31 839	−31 314	−31 274	−31 233	−31 233
$d = 4$	−40 445	−39 889	−39 836	−39 815	−39 821
$d = 5$	−48 342	−47 737	−47 682	−47 681	−47 687

day and the data points are weighted by the cosine of their latitude.

The spatial patterns of the leading EOFs (not shown) are similar to those obtained by Corti et al. (1997) and by D'Andrea and Vautard (2001). The 10 leading EOFs are responsible for 51% of the variance of the dataset: the first mode captures 12.5%, the second one 6.5%, the third one 5.2%, and the tenth one 3.1%. There is a slight break of slope in the variance curve after EOF-4, with the following eigenvalues forming a rather flat spectrum; the cumulative variance captured by the leading four EOFs equals 29%. Analyzing the streamfunction fields at the 250- and 800-hPa levels (not shown) reveals very similar leading EOF patterns; this similarity indicates a predominantly barotropic structure of the model's LFV, captured by the leading EOFs of the unfiltered data, with a finer structure present at the lower pressure level.

In order to objectively identify weather regimes in the QG3 model simulation, we apply two independent clustering methods (see Table 1 of Ghil and Robertson 2002) and compare the results. One method uses the *k*-means algorithm of Michelangeli et al. (1995) and the other uses the Gaussian mixture model of Smyth et al. (1999). Both these studies applied their respective methods to the classification of NH weather regimes in observed geopotential height fields.

For a given number *d* of leading EOFs, both methods provide a number of clusters *k* and the cluster centroids in a *d*-dimensional subspace of the model's phase space. We want each cluster to correspond to a weather regime of the QG3. Therefore, it is critical for our study to optimize the classification into clusters over various subspaces. The number of EOFs *d* and clusters *k* can be used as parameters to measure the robustness of the classification, along with diagnosing the similarity of the patterns themselves.

The *k*-means algorithm is based on the dynamic cluster method and formulated as follows. Given a prescribed number *k* of clusters in a *d*-dimensional space, it attempts to find an optimal partition of the data into the *k* clusters that minimizes the sum of the variances within each cluster. A data point belongs to a cluster if its distance to the cluster centroid is less than one standard deviation of all distances within a cluster. In order to determine the optimal *k*, Michelangeli et al. (1995) proposed the use of a classifiability index. This index measures the stability of the cluster solutions as a function of *k*, across different initial (random) seeds of the algorithm, based on the correlation between the cluster centroids. Table 1 gives the classifiability index of our QG3 model simulation for $2 \leq d \leq 5$. Clearly, it is very high for both $k = 3$ and $k = 4$; it follows that, based on this classifiability index alone, we cannot identify the optimal set of clusters in the QG3 model.

The Gaussian mixture model uses a linear combination of *k* Gaussian density functions, and differs from the *k*-means algorithm in the following two important aspects. First, each data point in the *d*-dimensional space can have a degree of membership in several clusters, depending on its position with respect to the centroid and the weight of a cluster (Smyth et al. 1999; Hand et al. 2001). Second, the mixture model has a built-in criterion for determining the optimal number of clusters supported by the data. This criterion is based on the cross-validated log-likelihood, shown in Table 2: the higher its algebraic value for a given dimension *d*, the more likely it is that *k* is the correct number of clusters for that *d*.

The mixture model consistently gives log-likelihood curves that saturate at $k \geq 3$ and have a very weak maximum. Thus, we can choose either $k = 4$ or $k = 5$ as the optimal number of clusters, according to this method. Hannachi and O'Neill (2001) found that when the data are not Gaussian, the mixture model tends to overfit them, so that the apparently optimal number of clusters increases with the length of the dataset. This seems to be the case here, too. In fact, when using 54 000 data points (see Table 3), rather than the 18 000 points of Table 2, the cross-validated log-likelihood curves have a tendency to saturate at higher values of *k*.

TABLE 3. Cross-validated log-likelihood for the entire set of 54 000 data points.

	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$
$d = 2$	−67 749	−66 396	−66 286	−66 226	−66 214
$d = 3$	−94 685	−93 114	−92 926	−92 823	−92 760
$d = 4$	−120 233	−118 494	−118 311	−118 158	−118 113

TABLE 4. Correlation coefficients between pairs of clusters given by the mixture model, on the one hand, and *k*-means clusters, on the other, when retaining k clusters: (a) $k = 3$, (b) $k = 4$, and (c) $k = 5$. In each case, correlations are given for the pairs identified by best visual match, and $1 \leq n \leq k$.

(a)	$n = 1$	$n = 2$	$n = 3$		
$d = 2$	0.999	0.995	0.993		
$d = 3$	0.996	0.995	0.993		
$d = 4$	0.996	0.994	0.994		
$d = 5$	0.992	0.990	0.986		
(b)	$n = 1$	$n = 2$	$n = 3$	$n = 4$	
$d = 2$	0.991	0.959	0.947	0.913	
$d = 3$	0.997	0.986	0.982	0.970	
$d = 4$	0.997	0.996	0.995	0.989	
$d = 5$	0.996	0.994	0.991	0.985	
(c)	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$d = 2$	0.993	0.993	0.980	0.969	0.954
$d = 3$	0.994	0.990	0.925	0.919	0.628
$d = 4$	0.999	0.998	0.996	0.979	0.949
$d = 5$	0.999	0.997	0.993	0.990	0.882

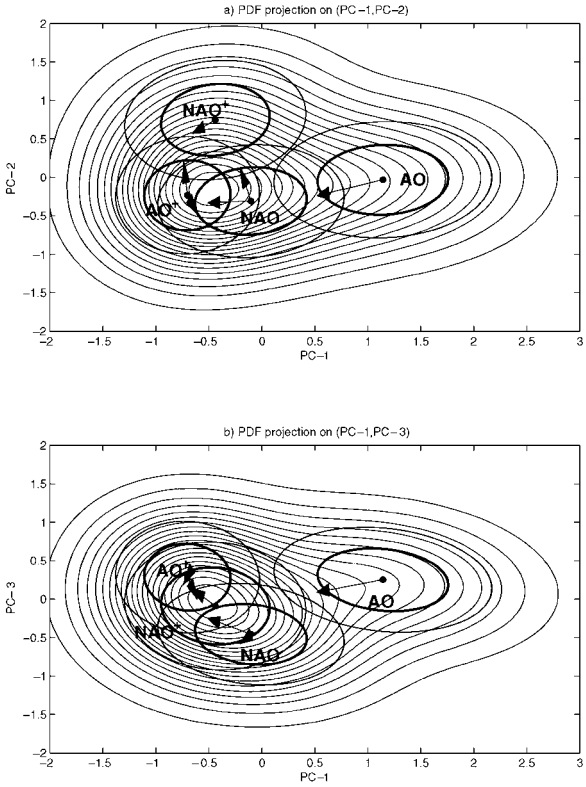

FIG. 1. PDF of the QG3 model's 500-hPa streamfunction field, as estimated by the mixture model for $d = 3$ and $k = 4$. The PDF is projected onto the plane spanned by (a) PC 1 and PC 2 and (b) PC 1 and PC 3; 20 contour levels are used. The heavy ellipses correspond to semiaxes equal to 0.75σ in each principal direction, while the light ones correspond to 1.25σ. Arrows indicate the preferred transition paths obtained in section 4b. The mixture model centroids have the following coordinates: $NAO^+ = (-0.44, 0.75, -0.09)$, $NAO^- = (-0.10, -0.31, -0.46)$, $AO^+ = (-0.71, -0.23, 0.29)$, and $AO^- = (1.14, -0.03, 0.25)$.

Since the cross-validated log-likelihood of Smyth et al. (1999) supports both $k = 4$ and $k = 5$, while the classifiability index of Michelangeli et al. (1995) supports either $k = 3$ and $k = 4$, we compare the anomaly maps of the centroids produced by both methods (see, e.g., Table 2 in Robertson and Ghil 1999) in order to choose the optimal number of clusters. To do so, we compute the pattern correlation coefficients of the cluster centroids in physical space for pairs of visually similar streamfunction anomaly maps produced by the two clustering methods and compare the results for different values of k. We obtain the maps that correspond to the cluster centroids in the d-dimensional subspace by computing the EOF expansion of the 500-hPa streamfunction field, that is the QG3 model's second level, truncated at a particular value of d.

Table 4 shows the correlation coefficients between corresponding pairs of clusters obtained from either method. Agreement between the two methods is very good for all values of d when using $k = 4$ (Table 4b), while for $k = 5$ (Table 4c) we find that only four of the clusters correlate well. The numerical correlation values are as good or better when using $k = 3$ (Table 4a); in this case, however, the actual patterns (not shown) do not match as well as those that were identified in previous work by other methods and from actual observational datasets as dominating NH LFV [see also the discussion in Smyth et al. (1999) and Ghil and Robertson (2002)].

Therefore, we choose $k = 4$ as the optimal set of clusters for our QG3 model. The probability density function (PDF) based on the mixture model for $d = 3$ and $k = 4$ is shown in Fig. 1, in the phase plane spanned by principal components (PCs) 1 and 2 (Fig. 1a), and by PCs 1 and 3 (Fig. 1b). All PCs are normalized by the standard deviation of PC 1. The coordinates of the cluster centroids for $d = 3$ are given in the figure caption; in higher dimensions ($d \geq 4$) they are very close to zero, which leads us to choose $d = 3$ for analyzing cluster properties. The size of the clusters is set by choosing a covariance ellipsoid around each cluster centroid.

Figure 1 shows the clusters' complex three-dimensional (3D) structure given semiaxes of the ellipsoid that correspond in length to 0.75 (heavy) and 1.25 (light) times the standard deviation in each direction. The names of the clusters are chosen based on the maps of the centroids, which will be discussed next. The arrows in Fig. 1 are projections of vectors that lie along preferred transitions paths between regimes and will be discussed in section 4.

The anomaly maps of the 500-hPa streamfunction centroids shown in Fig. 1 are plotted in Fig. 2. Cluster AO^- (Fig. 2d) occupies a distinct region on the PDF ridge that stretches along PC 1, while clusters NAO^+, NAO^-, and AO^+ are located around the global PDF

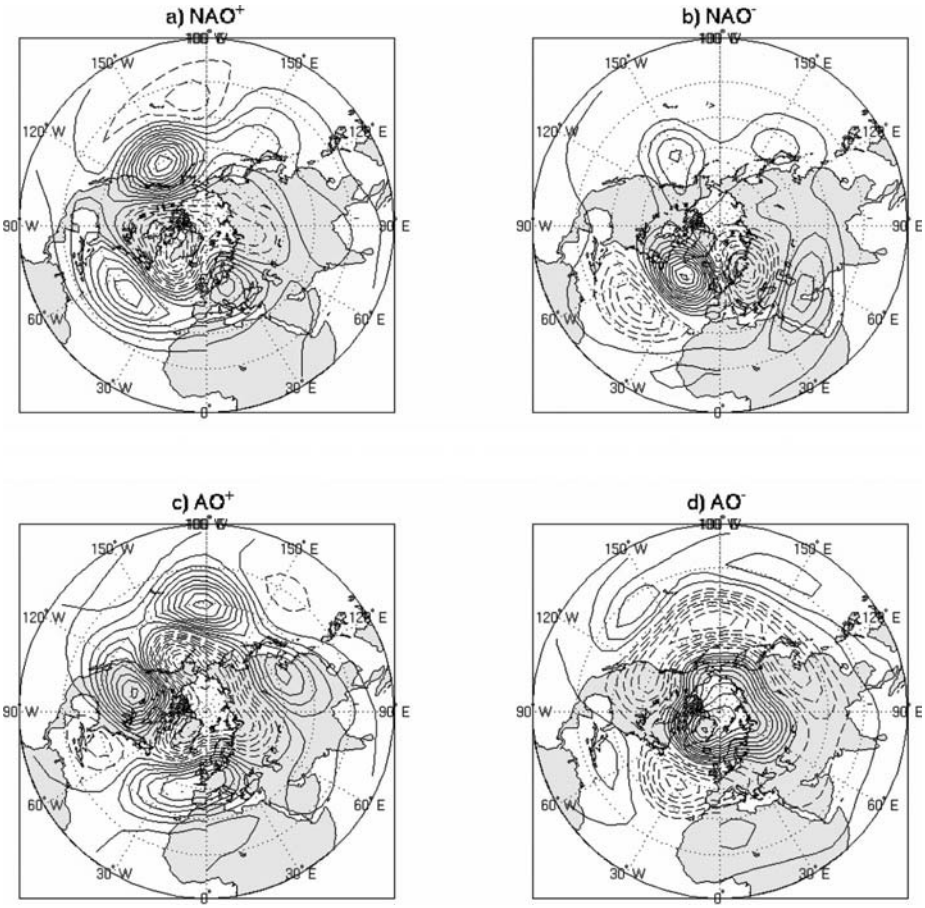

FIG. 2. Mixture model centroids for $d = 3$ and $k = 4$, showing streamfunction anomaly maps at 500 hPa: (a) NAO^+, (b) NAO^-, (c) AO^+, and (d) AO^-. Negative contours are dashed and land masses are shaded; twenty contour levels between maximum and minimum values are used, with the following intervals (in 10^6 m^2 s^{-1}): (a) 1.1, (b) 0.8, (c) 0.8, and (d) 1.1.

maximum. Cluster NAO^+ (Fig. 2a) is shifted away from the origin along PC 2, while clusters AO^+ (Fig. 2c) and NAO^- (Fig. 2b) are shifted away from the PC-1 axis, in opposite directions with respect to PC 3.

Each of the regimes in Fig. 2 represents, roughly, one of the opposite phases of two spatial patterns. The maps in Figs. 2a and 2b capture the two extreme phases of the NAO, with the pattern in Fig. 2a completing a wavenumber-3 pattern outside the Atlantic sector. The maps in Figs. 2c and 2d have a large central anomaly that extends over the whole Arctic and a pronounced zonally symmetric component. The two maps thus have important features in common with the opposite phases of the AO (Thompson and Wallace 1998) and with Mo and Ghil's (1988) north–south seesaw, while Fig. 2d also exhibits a substantial wavenumber-4 component.

We denoted these four regimes, therefore, by NAO^+ (Fig. 2a) and NAO^- (Fig. 2b), AO^+ (Fig. 2c), and AO^- (Fig. 2d). The spatial patterns in Fig. 2 resemble well four of the five clusters found by D'Andrea and Vautard (2001) in the QG3 model, by using the *k*-means method. We recall that NAO^+ and NAO^- correspond to the zonal and blocked phases of the jet in the western North Atlantic, while AO^+ and AO^- are composites of hemispherically high-index and low-index flow (in the sense used by J. Namias, C.-G. Rossby, and H. Willet in the 1950s). The lack of symmetry between NAO^+ and NAO^-, as well as between AO^+ and AO^-, points to the potentially nonlinear origin of these regimes.

4. Markov chain of transitions

a. Transition probability matrix

Given the clustering results of section 3, we proceed now to the main part of this study, namely, investigating in detail the Markov chain of transitions between the four regimes. In a *d*-dimensional subspace of the model's phase space, each weather regime is defined by the ellipsoid of covariance around the centroid, whose semiaxes equal the corresponding singular values, that is, the square roots of the eigenvalues of the covariance

TABLE 5. Regime statistics for different cluster sizes. The cluster size is determined by the scaling factor of the standard deviation σ along each semimajor axis.

Size		NAO^+	NAO^-	AO^+	AO^-	Total
0.75σ	Events	908	1137	1026	673	3744
	Days	1602	2030	1930	1322	6884
1.25σ	Events	1970	2375	2319	1665	8329
	Days	5248	6956	6740	5603	24 547

matrix of each Gaussian component (or a fixed multiple thereof; see Fig. 1). A data point is assigned to a weather regime if it lies within the corresponding ellipsoid. If a data point belongs to several ellipsoids, we assign it to a cluster according to the maximum probability value found from the mixture model. We can vary the size of all four ellipsoids, and therefore the number of data points that belongs to each of the four clusters, by using the same scaling factor along the semiaxes of all four.

A regime event is defined as the number of consecutive points (days) along the model trajectory that fall within a given cluster (Mo and Ghil 1987, 1988; Kimoto and Ghil 1993b). Table 5 shows the number of distinct events in each regime and the total number of days spent in the regime for scaling factors of 0.75 and 1.25. A total number of about 25 000 days out of the entire run of 54 000 days belong to one of the large-size clusters (1.25σ), while for the smaller clusters (0.75σ) the total regime population is of about 7000 days.

The distribution of residence times for each of the regimes is shown in Fig. 3 (solid), together with associated significance curves (dashed); the latter correspond to one standard deviation of the residence time distribution within a population of 100 red-noise surrogate data samples (following Dole and Gordon 1983). Each surrogate sample consists of three AR(1) time series that have the same length, variance and lag-1 autocorrelation coefficient as the QG3 model's three leading PCs; the surrogate's regime residence time is determined by the time spent by the sample trajectory in the cluster ellipsoides of our mixture model.

The persistence distributions are exponential in nature, and are similar in length and frequency of occurrence to those obtained by Corti et al. (1997) for the blocking episodes. The negative slopes of the residence curves become larger (in absolute value) and more equal to each other as the cluster size decreases (compare the curves for 0.75σ, labeled by squares, with those for 1.25σ, labeled by diamonds). For the large-size clusters (diamonds), the low-index cluster AO^- (Fig. 3d), albeit of relatively small total size (see Table 5), is visibly more persistent than the rest, while the sectorially zonal cluster NAO^+ (Fig. 3a) is visibly less persistent. The zonal NAO^+ regime is clearly less persistent than the

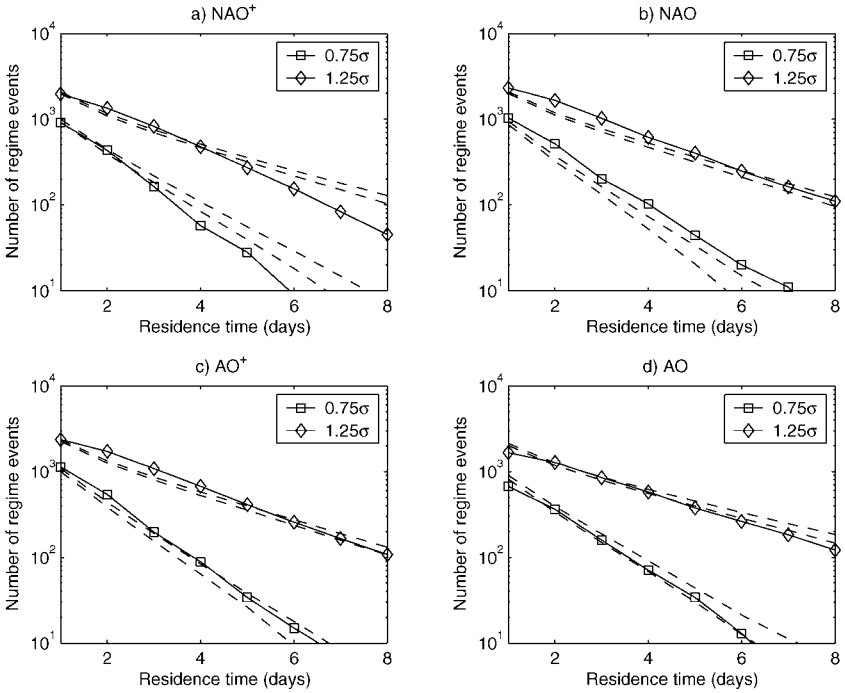

FIG. 3. Distribution of cumulative residence time (solid lines) in the model's regimes for cluster sizes of 0.75σ and 1.25σ, where σ refers to the standard deviation of each regime, along each semimajor axis (omitting subscripts). Dashed lines correspond to one standard deviation, above and below the mean of the residence time distribution, within a population of 100 red-noise surrogate data samples (see text for details). (a)–(d) The four regimes shown in Figs. 1 and 2.

TABLE 6. Transition probabilities estimated using the mixture model. Transitions that are significantly higher at the 95% level with respect to 10 000 random shuffles of the sequence of regime events are in bold, while entries that are italicized are significantly lower at the 95% level for the same test. The entries in the table are for clusters of size 0.75σ, 1.0σ, and 1.25σ, in this order. Not all rows sum exactly to 1.00 because of round-off effects.

$0.75\sigma/1.00\sigma/1.25\sigma$	NAO^+	NAO^-	AO^+	AO^-
NAO^+	**0.33/0.29/0.25**	*0.27/0.23/0.21*	**0.34/0.41/0.44**	*0.06/0.07/0.10*
NAO^-	**0.26/0.27/0.26**	**0.32**/0.28/*0.25*	0.27/0.28/**0.30**	*0.15/0.17*/0.19
AO^+	0.24/**0.29/0.29**	**0.38/0.39/0.40**	**0.31**/0.27/*0.25*	*0.07/0.05/0.06*
AO^-	*0.11/0.09/0.10*	*0.20/0.22/0.26*	*0.13/0.10/0.10*	**0.56/0.59/0.53**

corresponding AR(1) process and the blocked NAO^- is slightly more persistent (although the difference is significant at larger durations only for 0.75σ), while AO^+ and AO^- cannot be said to be significantly different from the AR(1) process.

Table 6 shows transition probabilities between the clusters using the sequence of regime events along the trajectory. Note that when a trajectory exits a particular cluster, defined by the mixture modeling approach (as in Fig. 1), it may cross a ''fuzzy'' region of phase space, populated by data points that are not classified, before it ends up inside another cluster or the same one. Self-transitions are also counted, in addition to those that lead to another cluster, for computing transition probabilities between clusters.

Monte Carlo simulation was applied to provide a statistical significance test for the elements of the transition matrix, following Vautard et al. (1990). The test is designed to take into account the difference in size between regimes and uses random shuffling of the sequence of regime events in the model simulation, subject to the constraint of the number of events in each regime being fixed and equal to the one in Table 5. The transition probabilities between regimes that are higher or lower than the Monte Carlo result at the 95% level appear in Table 6 in bold or italics, respectively. The results are similar (not shown) when only regime events that last for 3 days or longer have been considered. Note that each row in the table sums to unity, for all three cluster sizes being reported (0.75σ, 1.00σ, and 1.25σ).

The transition probabilities and their significance are fairly similar for the small and large clusters. The fact that there are more reentries for smaller-size clusters is entirely expected, since events of a given length are more likely to be erroneously ''chopped up'' into two or more events that belong to the same cluster, when its size is smaller; this finding merely suggests, therefore, that 1.25σ is a somewhat better choice of scaling factor than 0.75σ or 1.00σ. All but one of the transitions were found to be significantly different from a random shuffle for the largest clusters. The AO^- regime has a higher reinjection rate than the other three regimes, independently of scaling factor. By the previous reasoning about reentries and cluster size, this scale independence of its higher reinjection rate agrees with AO^- having the highest persistence (see Fig. 3d).

Taking the highest transition probability, excluding reentries, in every row, we obtain the preferential cycle of transitions: $NAO^+ \rightarrow AO^+ \rightarrow NAO^- \rightarrow NAO^+$; the only disagreement between the estimates obtained with different scaling factors for this cycle is for the leg $NAO^- \rightarrow NAO^+$ and so we have selected the transition for which all three estimates agree. A second cycle that is apparent only for 1.0σ and 1.25σ is $AO^+ \leftrightarrow NAO^+$. Transitions from each of the three regimes that have a small or negative PC-1 component to AO^- are unlikely by themselves but, taken together, they obviously add up to the low-index hemispheric regime AO^- being reachable from the higher-index part of the phase space. Direct transitions between AO^- and AO^+ are very unlikely for all cluster sizes considered.

It is clear from this examination of the transition probability matrix that the model's LFV is quite nonlinear, with direct transition from the zonal to the blocked phase of the NAO, NAO^+ to NAO^-, as well as between the high- and low-index phases of the AO, AO^+ to AO^- or vice-versa, being quite unlikely. This agrees with the conclusions drawn from such transition matrices for NH observations by Mo and Ghil (1988) and Kimoto and Ghil (1993b). Given the model's degree of realism, on the one hand, and much greater availability of data points, on the other, we are thus motivated to explore certain details of its trajectory in phase space.

b. Preferred transition paths

To do so, we compute the exit angle from its regime for each events. When the model trajectory exits a cluster, we determine the phase space coordinates (x, y, z) of the exit point, which is defined as the midpoint between two consecutive trajectory points that lie on opposite sides of the cluster boundary. The exit vector is then defined to point from the cluster centroid to the exit point. In the 3D subspace spanned by EOFs 1, 2, and 3, with coordinates (x, y, z), each exit vector is uniquely defined by two angles: $\tan\phi = y/x$, $0 < \phi < 2\pi$, and $\tan\theta = z/(x^2 + y^2)$, $-\pi/2 < \theta < \pi/2$.

We then compute two-dimensional (2D) PDF distributions of the exit angles in the (ϕ, θ) coordinates for all of the regime events using a Gaussian kernel estimator (Silverman 1986). This estimator approximates a PDF by assigning a localized kernel density function of a given shape (Gaussian in our case) to each data point. The PDF estimate at any point in phase space is then

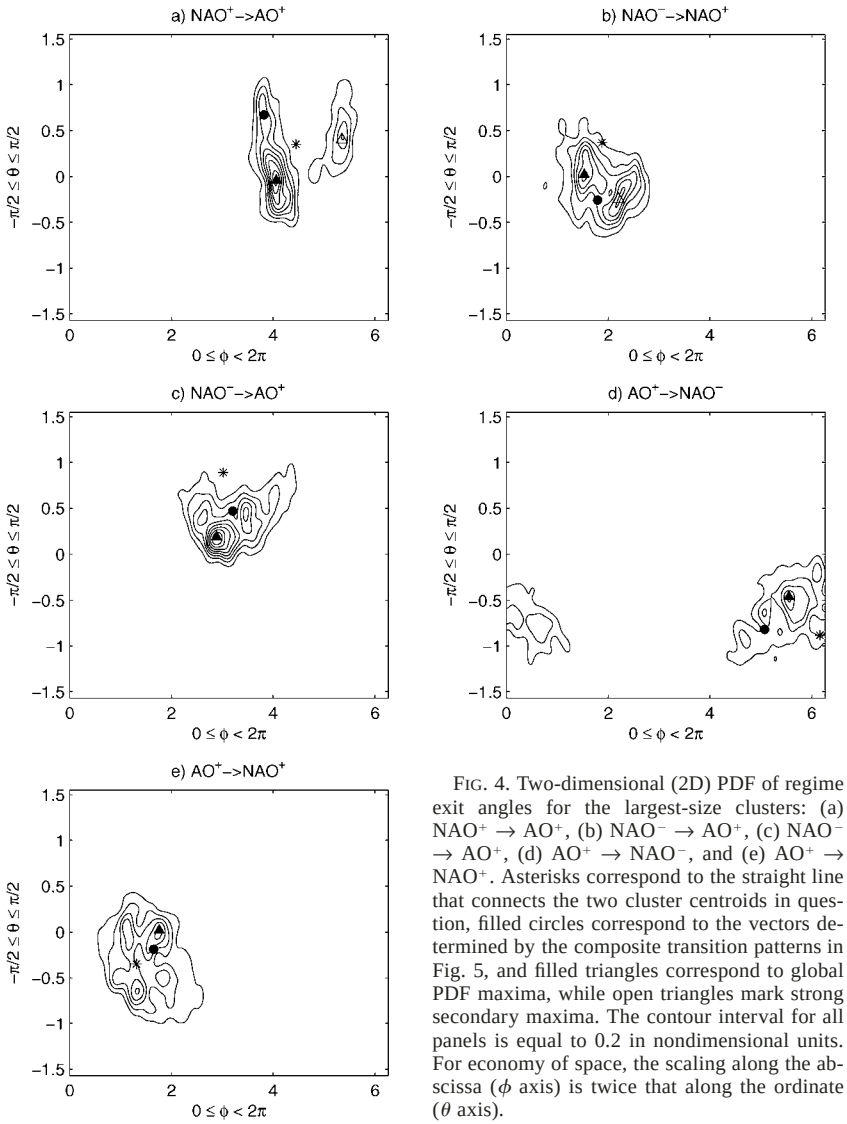

FIG. 4. Two-dimensional (2D) PDF of regime exit angles for the largest-size clusters: (a) $NAO^+ \rightarrow AO^+$, (b) $NAO^- \rightarrow AO^+$, (c) $NAO^- \rightarrow AO^+$, (d) $AO^+ \rightarrow NAO^-$, and (e) $AO^+ \rightarrow NAO^+$. Asterisks correspond to the straight line that connects the two cluster centroids in question, filled circles correspond to the vectors determined by the composite transition patterns in Fig. 5, and filled triangles correspond to global PDF maxima, while open triangles mark strong secondary maxima. The contour interval for all panels is equal to 0.2 in nondimensional units. For economy of space, the scaling along the abscissa (ϕ axis) is twice that along the ordinate (θ axis).

given by the sum of these kernel density functions. We use the data-adaptive version of the kernel estimator (Kimoto and Ghil 1993a) and modify it to account for the periodic nature of the PDF in the ϕ direction. The PDFs so obtained are too noisy for smaller-size clusters (0.75σ and 1.00σ), due to the small number of regime events. We show therefore in Fig. 4 only the exit-angle PDFs for the largest-size clusters (1.25σ), a choice that is also consistent with the earlier argument in section 4a on this size being less subject to reentries. The global PDF maxima in Fig. 4 are indicated by filled triangles, while open triangles mark strong secondary maxima, asterisks correspond to the straight line that connects the two cluster centroids in question, and filled circles correspond to the transition vectors determined by the transition patterns (see below).

Inspection of Fig. 4 demonstrates that most of the highly significant transitions (see again Table 6) occur along one or two narrow solid angles in phase space. For example, the NAO^+ to AO^+ and NAO^- to NAO^+ transitions (Figs. 4a,b) have a narrow double peak in the PDF distribution, while the NAO^- to AO^+ transition (Fig. 4c) has a single narrow maximum. The transitions from AO^+ to NAO^- and AO^+ to NAO^+ (Figs. 4d,e), while still highly significant, are not quite as sharply confined in this angular PDF. As expected from the conclusions of section 4a, the preferred directions of transition lean away from the straight line that connects the cluster centroids. To examine more closely the possible connections between a model's nonlinear dynamics and

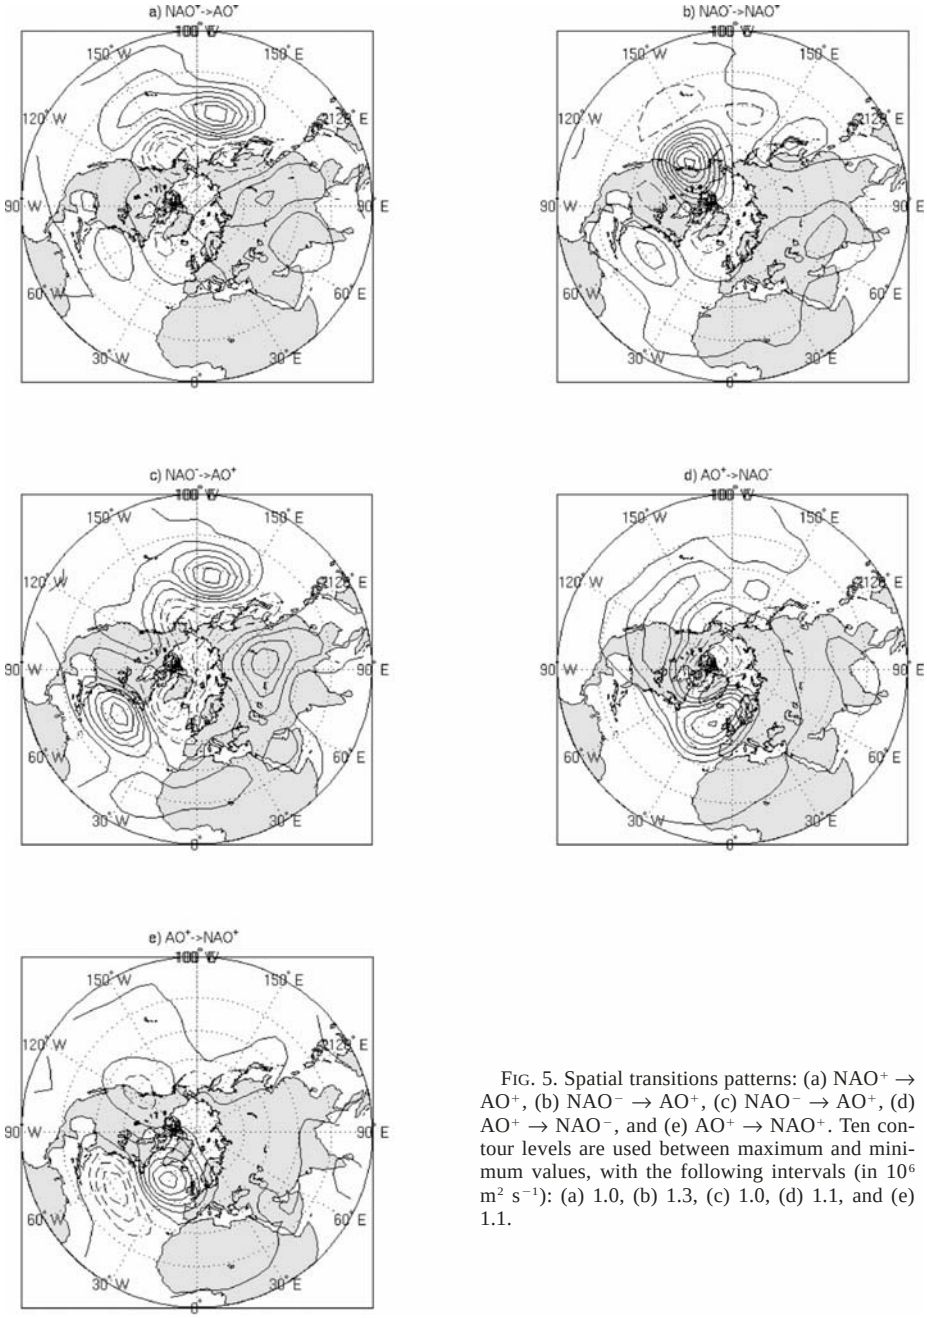

FIG. 5. Spatial transitions patterns: (a) $NAO^+ \rightarrow AO^+$, (b) $NAO^- \rightarrow AO^+$, (c) $NAO^- \rightarrow AO^+$, (d) $AO^+ \rightarrow NAO^-$, and (e) $AO^+ \rightarrow NAO^+$. Ten contour levels are used between maximum and minimum values, with the following intervals (in 10^6 m^2 s^{-1}): (a) 1.0, (b) 1.3, (c) 1.0, (d) 1.1, and (e) 1.1.

its preferred regime transition angles, we study these connections for the Lorenz (1963a) system in the appendix.

From the example of the Lorenz model we conclude that the observed maxima in the PDF distributions of the regime exit angles in our QG3 model are likely to be determined by the dynamical properties of the underlying nonlinear system of equations. These properties include the linearly stable and unstable directions of the steady states, as well as genuinely nonlinear, finite-amplitude behavior away from these steady states. Unlike in the Lorenz model, a distinct feature of both the $NAO^+ \rightarrow AO^+$ and $NAO^- \rightarrow NAO^+$ transitions is the appearance of a double peak in the exit-angle PDFs (see Figs. 4a,b). Such fine details of the regime transitions are being further examined in a nonlinear, stochastically forced system of 15 ordinary differential equations that represents a low-order version of the QG3 model (D. Kondrashov et al. 2003, unpublished manuscript).

The spatial patterns associated with the preferred tran-

sition paths are shown in Fig. 5. These patterns correspond to ''second-order anomalies,'' that is, to deviations from the cluster centroids in a preferred direction of transition. For any transition in Fig. 4, we compute a composite anomaly map of the exit points from the regime, and then subtract from it the anomaly map of the cluster centroid. The patterns so obtained are robust with respect to cluster size and the length of the simulated dataset (not shown). These patterns are significantly different from those obtained by just taking the differences between the cluster centroids: compare the position of the filled circles with that of the asterisks in Figs. 4a–d. Also, the patterns for transitions back and forth between a pair of clusters are strongly anticorrelated but the correlation is significantly distinct from -1.0; thus, for $\mathrm{NAO}^- \leftrightarrow \mathrm{AO}^+$, it is -0.53, while for $\mathrm{AO}^+ \leftrightarrow \mathrm{NAO}^+$, it is -0.62.

The transitions between NAO^+ and AO^+ and back (Figs. 5a,e) exhibit both a strong SW–NE-oriented dipole, extending over much of the central North Pacific and Gulf of Alaska. Those between NAO^- and AO^+ share a strong, N–S-oriented dipole (with opposite polarities, obviously, in Figs. 5c and 5d) but the transition in Fig. 5c also has a strong feature in the central North Pacific, which Fig. 5d lacks. Finally, the $\mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ transition (Fig. 5b) is characterized by a wave train across North America that suggests a downstream influence of the Rockies (Kalnay-Rivas and Merkine 1981).

The projection of these transition patterns on the three leading EOFs determines the preferred transition vectors indicated by arrows in Fig. 1; the magnitude of these vectors, though, is not related to velocity in phase space and thus gives no indication on the rapidity of the corresponding regime break. The relative direction of the arrows is consistent with the correlation coefficients between the patterns: for example the arrows that correspond to the back-and-forth transitions between the pairs $\mathrm{NAO}^- \leftrightarrow \mathrm{AO}^+$ and $\mathrm{AO}^+ \leftrightarrow \mathrm{NAO}^+$ are almost at right angles to each other in Fig. 1a, in both cases. Note that this agreement is not self-evident, since the pattern correlations are computed in the full physical space, while the arrows correspond to projections onto a low-dimensional subspace.

5. Intraseasonal oscillations

In sections 3 and 4, we have provided an episodic description of the model's LFV, based on the Markov chain of transitions between its four regimes. Given the overall realism of the model's LFV, it seems appropriate to complement this description, if possible, with an oscillatory one that uses unstable limit cycles to describe the preferred trajectories in the model's phase space.

Advanced spectral methods (Ghil et al. 2002) permit the statistically reliable identification and description of the least unstable limit cycles (see, in particular, Fig. 1 there). We apply singular spectrum analysis (SSA) to study the model's intraseasonal oscillations and attempt to connect them with the statistically significant closed cycle $\mathrm{NAO}^+ \rightarrow \mathrm{AO}^+ \rightarrow \mathrm{NAO}^- \rightarrow \mathrm{NAO}^+$ in regime transitions. Multichannel SSA (M-SSA) generalizes single-channel SSA to extract oscillatory modes of variability from a multivariate time series and is especially useful in the study of amplitude- and phase-modulated signals. Keppenne and Ghil (1993) and Plaut and Vautard (1994), among others, have applied M-SSA to isolate oscillatory components in observations of atmospheric LFV.

We apply M-SSA to the three leading PCs of the 500-hPa streamfunction anomalies of our model integration, $\{X_l(t)\colon t = 1, \ldots, N\}$, with $N = 54\,000$. The choice of three channels, $1 \leq l \leq L = 3$, is based on the conclusions of section 3, where we found $d = 3$ to be the optimal dimension of the subspace in which to compare the results of our two clustering methods. The notation here follows Ghil et al. (2002; see, in particular, Table 1 there) and we use the trajectory matrix approach of Broomhead and King (1986a,b), to compute the lag-covariance matrix.

The multichannel trajectory matrix $\mathbf{X}$ is formed by first augmenting each channel $X_l(t)$ of the vector time series $\mathbf{X}(t)$ with M lagged copies of itself, to yield $\mathbf{X}_l(t)$. The full augmented trajectory matrix is then given by

$$\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_L). \tag{2}$$

The window width M determines the range of oscillation periods to be detected; to encompass the full range of intraseasonal oscillations, we have used values of M up to and including 100 days. The grand lag-covariance matrix $\mathbf{C}_\mathbf{x}$ is given by

$$\mathbf{C}_\mathbf{x} = \frac{1}{N'}\mathbf{X}^\mathrm{T}\mathbf{X}, \tag{3}$$

where $N' = N - M + 1$.

The eigenvectors of the matrix $\mathbf{C}_\mathbf{x}$ are known as space–time EOFs (ST-EOFs). The associated space–time PCs (ST-PCs) are single-channel time series that are computed by projecting $\mathbf{X}$ onto the ST-EOFs. M-SSA identifies an oscillation in the multivariate time series when two consecutive eigenvalues of $\mathbf{C}_\mathbf{x}$ (ordered by size) are nearly equal, the corresponding ST-EOFs are periodic, with the same period and in quadrature, and the associated ST-PCs are in quadrature as well. Following Allen and Robertson (1996) we apply, in addition to the criteria above, a Monte Carlo test to ascertain the statistical significance of the oscillations detected by M-SSA.

Figure 6 shows the eigenvalues obtained by M-SSA with a window width of $M = 100$ days; these equal the variances associated with the corresponding eigenmode. The two arrows in Fig. 6 point to the M-SSA modes 8–9 and 15–16, which form two significant oscillatory pairs, with a period of 37 and 19 days, respectively. In addition, these pairs pass the lag-correlation test (not shown) of Plaut and Vautard (1994): on both sides of

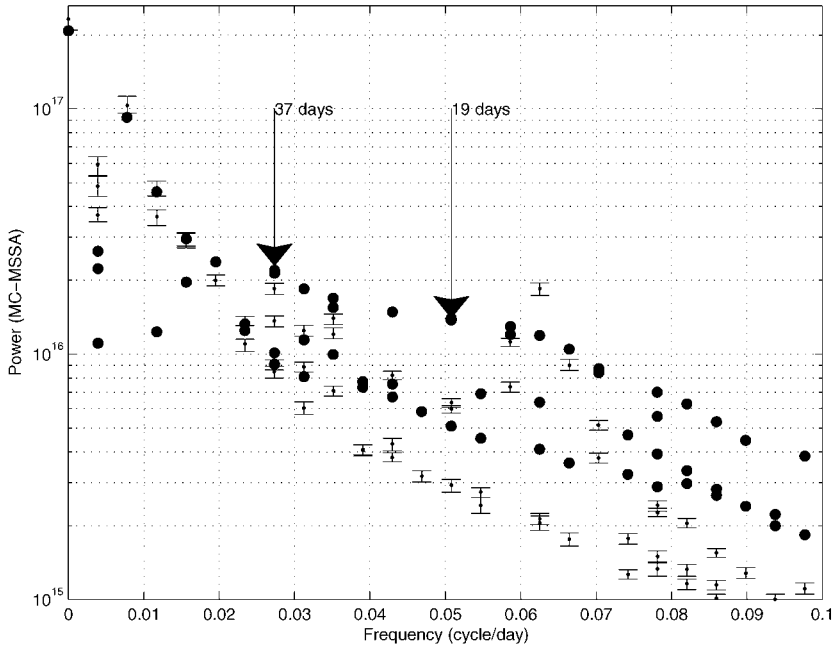

FIG. 6. Singular spectrum obtained by M-SSA of our model's 54 000-day model integration, in a 3D subspace. Filled circles show M-SSA eigenvalues, plotted against the dominant frequency associated with the corresponding space–time PCs. Arrows indicate two significant pairs with periods of 19 and 37 days, respectively; note that the two circles in either pair are almost indistinguishable. The confidence intervals correspond to the 5% and 95% percentiles of 100 Monte Carlo simulations of a red-noise process with the same length, variance and lag-one autocorrelation as the model time series being tested; see text and section 4.2 of Ghil et al. (2002) for details.

lag 0, the two successive extrema of the lag-autocorrelation of the corresponding ST-PCs exceed 0.5 in absolute value. The reconstructed components (RCs) of the 8–9 and 15–16 pairs describe the unstable limit cycle associated with the corresponding pair in the reduced phase space.

We analyze the behavior of the oscillation by using composite maps keyed to the phases of the oscillation, as in Plaut and Vautard (1994). Figures 7a–d show four successive, equally populated composites that correspond to one half-cycle of the 37-day oscillation. The cycle of regime transitions $NAO^+ \rightarrow AO^+ \rightarrow NAO^-$ is easily identified. After passing through the NAO^+ (Fig. 7a) and AO^+ (Fig. 7b) phases, however, the next phase of the oscillation resembles the PNA pattern (Fig. 7c), before reaching the NAO^- phase (Fig. 7d).

There is a certain similarity between some phases of this oscillation and certain phases of the westward-propagating Branstator (1987)–Kushnir (1987) wave, as documented also by Dickey et al. (1991), Ghil and Mo (1991), and Plaut and Vautard (1994); see especially Fig. 5 of Branstator (1987). There are also marked differences, however, such as the 37-day period here versus that wave's period of 23–25 days, and the very pronounced features of Fig. 7 in the North Atlantic–European sector. We constructed therefore a movie (not shown, but available upon request from the corresponding author; also see our Web site http://www.atmos.ucla.edu/tcd), based on 100 composite maps of the cycle, rather than the eight composites for the full cycle used to construct Fig. 7.

The movie makes it clear that the 37-day cycle here is mainly a standing oscillation, with occasional northeastward propagation of certain features. In particular it shows strong interaction between the Rockies, on the one hand, and a low and a high that switch places in a counterclockwise fashion between lying upstream and downstream of this mountain range, on the other. The period and standing features of this oscillation, as well as its interaction with the Rockies, resemble certain characteristics of the oscillatory topographic instability of Ghil and associates (Legras and Ghil 1985; Ghil 1987; Ghil and Mo 1991; Ghil et al. 1991; Ghil and Robertson 2002).

Phase composites of the 19-day oscillation (not shown) are almost identical to Figs. 7a–d. This similarity in spatial patterns, as well as the ratio between the periods of the oscillations being so close to 2, make one suspect that they are related as (sub-) harmonics of each other. Following Plaut and Vautard (1994) and Koo et al. (2002), we calculated the histogram of simulta-

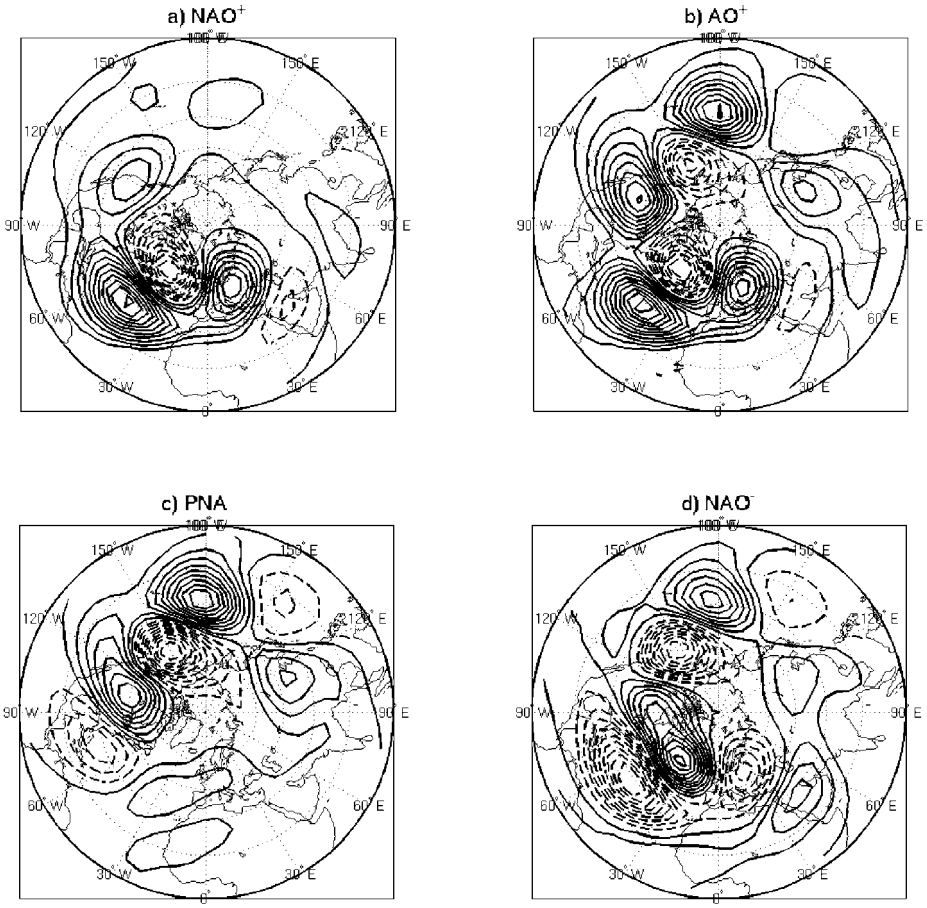

FIG. 7. Phase composites of the 37-day oscillatory mode, as given by the RCs 8 and 9 of the M-SSA. The phase angle φ in the plane spanned by the sum of RCs 8 and 9, on the one hand, and by the derivative of this time series, on the other, is obtained for each day t of the simulation; the circle $(0, 2\pi)$ in this plane is then subdivided into eight equally populated arcs of circle, each containing 6750 points [RC-8 + RC-9, Δ(RC-8 + RC-9)] with $\varphi(t)$ inside that arc. (a)–(d) Four composite maps of all the days t that fall within four such successive arcs; (a)–(d) shown thus cover one-half of the oscillation cycle, while the four additional panels (not shown) have very similar patterns, with the sign reversed. The composite maps are those of streamfunction anomalies at 500 hPa, as in Fig. 2, with the same plotting conventions and contour intervals.

neous phase occurrences of the two oscillations. This calculation shows that they are weakly phase-locked to each other; that is, they appear to be so when the calculation is done for certain window widths.

Another possible connection between the episodic and oscillatory description of preferred trajectories in phase space is to determine the average time taken by the model to complete the Markov-chain cycle $NAO^+ \rightarrow AO^+ \rightarrow NAO^- \rightarrow NAO^+$. As suggested by Fig. 1 in Ghil et al. (2002), we computed this average only for those segments of trajectory that actually complete such a cycle by starting in one of these three regimes, passing through the other two in succession, and returning to the starting regime. The average cycle period computed in this way equals approximately 19 days, which is consistent with the shorter period of oscillation detected by M-SSA.

Legras and Ghil (1985) highlighted the slowing down of trajectories near unstable equilibria. Following this suggestion, Vautard (1990) used trajectory speeds as clustering criteria in observational data, while Mukougawa (1988) and Vautard and Legras (1988) used them in intermediate models of lesser complexity than the present ones; see also Table 1 in Ghil and Robertson (2002). We want to know, therefore, whether there is a systematic connection between any of the regimes and the slow phases of the oscillations described by M-SSA.

To investigate this connection, if any, we compute the velocity in the reduced phase space spanned by the three leading EOFs, along the trajectory of RCs 8–9, during a particular 100-day-long oscillatory spell. The velocity components are obtained by central differencing the time series for each of the three channels in the RCs. Figures 8a–c display the three velocity compo-

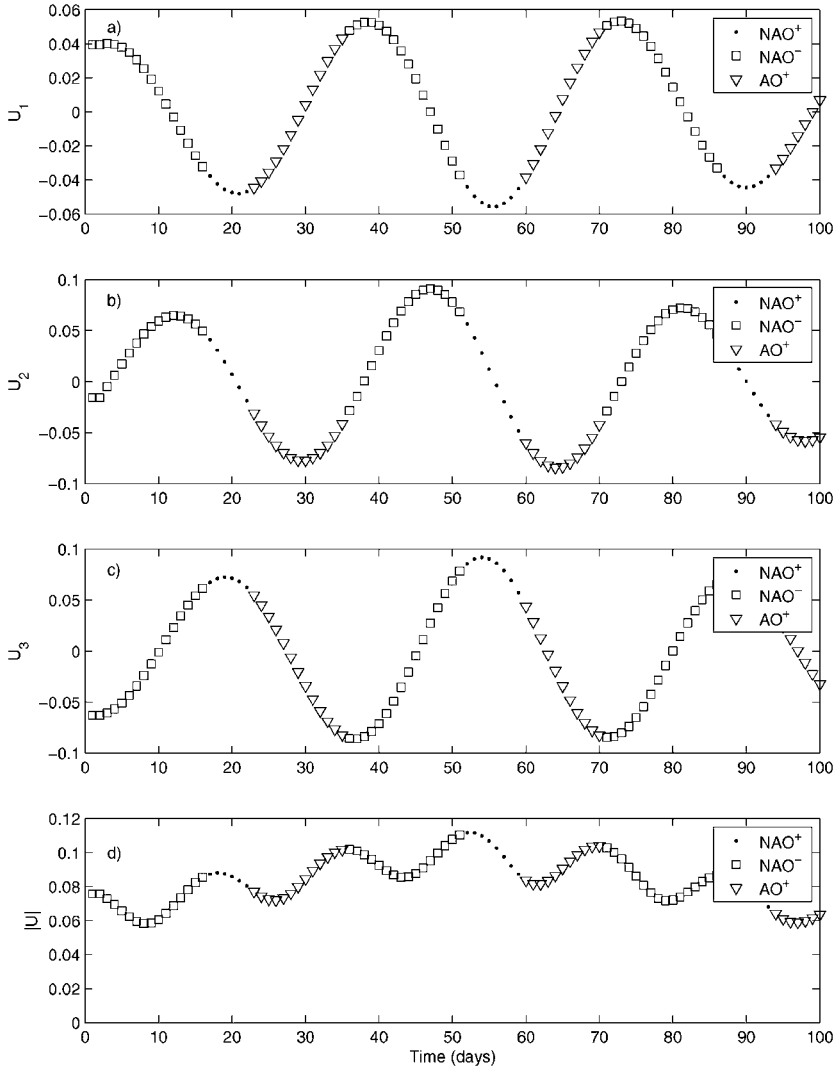

FIG. 8. Phase velocity in a space of three leading EOFs along the trajectory of RCs 8–9 and 15–16 during two high-amplitude oscillatory spells: (a), (b), (c) three velocity components (nondimensional units); (d) absolute value for RCs 8–9; (e)–(h) the corresponding figures for RCs 15–16. Symbols mark the clusters that the trajectory traverses at that moment.

nents, and Fig. 8d presents the speed. Likewise, the three velocity components and the speed for a 150-day-long oscillatory spell with a 19-day period (RCs 15–16) are plotted in Figs. 8e–h.

Each velocity component is highly periodic, with the periods of 37 and 19 days, respectively, and the speed goes through phases of significant acceleration and deceleration. The period of the speed (Figs. 8d,h) seems to equal half the period of the velocity components, in the case of both unstable orbits. Alternate minima of the speed are associated with NAO^+ and AO^+, for the 39-day cycle (Fig. 8d), as well as for the 19-day cycle (Fig. 8h).

6. Concluding remarks

a. Summary

We have examined the low-frequency variability (LFV) of the NH midlatitude atmosphere in the intermediate-complexity QG3 model of Marshall and Molteni (1993). The model is based on the QG equations for the conservation of potential vorticity and it has three levels in the vertical and T21 truncation in the horizontal (see section 2.) Using fairly realistic topography, land–sea contrast and residual forcing (see Corti et al. 1997), it produces respectable climatology and LFV.

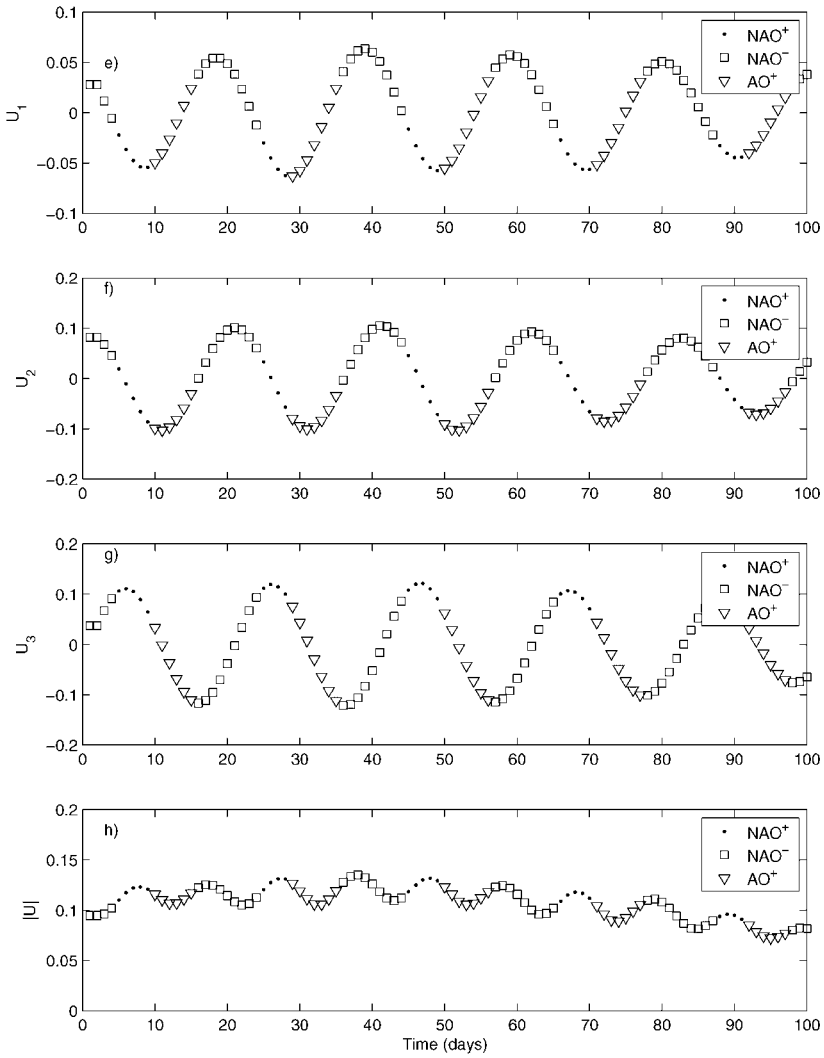

FIG. 8. (*Continued*)

1) MARKOV CHAIN

We analyzed a 54 000-day-long perpetual-winter integration of this model from the point of view of coarse-graining its phase space (Ghil 1987; Haines 1994; Ghil and Robertson 2002). The first step in this approach is to verify the presence of clusters in the dataset that corresponds to multiple weather regimes. These represent, in principle, regions of higher probability density function (PDF) on the model's attractor.

Using both the *k*-means method of Michelangeli et al. (1995) and the Gaussian mixture modeling of Smyth et al. (1999), and optimizing the result with respect to both methods (see Tables 1–4 here and Robertson and Ghil 1999), we obtained four NH clusters (Fig. 1) in the subspace of the four leading EOFs of the simulation's 500-hPa streamfunction field. The projections of the cluster centroids of the fourth and subsequent EOFs are very small and so the remainder of the analysis was carried out in a 3D subspace (see Fig. 1).

The four clusters (see Fig. 2) are NAO^+ and NAO^-, that is, the zonal and blocked phases of the North Atlantic Oscillation (NAO), and AO^+ and AO^-, that is, the high- and low-index phases of the Arctic Oscillation (AO). These four regimes are not identical to but in fairly good conceptual agreement with the observational results of Cheng and Wallace (1993), using hierarchical clustering, and those of Smyth et al. (1999), using a Gaussian mixture model for a slightly different NH dataset.

As discussed by Smyth et al. (1999) and Ghil and Robertson (2002), the four regimes that are best supported by synoptic experience, as well as statistical analysis of the upper-air data for the past half-century, are the zonal and blocked phases of westerly flow in the

Atlantic–Eurasian and Pacific–North American (PNA) sector, respectively. Seeing the PNA-sector patterns replaced in our classification by the hemispheric AO^+ and AO^- ones is thus not too surprising: the latter two overlap largely in the Atlantic–Eurasian sector with NAO^+ and NAO^-, respectively, from which they differ mostly by very strong centers of action in the PNA sector (see again Fig. 2).

The second step in the multiple-regime approach to LFV is to determine the Markov chain of transitions between regimes (Ghil 1987; Mo and Ghil 1988; Molteni et al. 1990; Kimoto and Ghil 1993b). In the language of dynamical systems theory, this corresponds to a particular representation of the ''skeleton of the attractor,'' which we take here to mean its most interesting and robust part.

To get the most reliable results for the transition probabilities from this long run, we used ellipsoidal clusters based on the Gaussian mixture model that had semiaxes ranging from 0.75σ to 1.25σ, with σ the variance in the corresponding principal direction (see Fig. 1 and Tables 5 and 6). Transitions were evaluated for significance with respect to statistical testing for fixed-size ellipsoids at the 95% level, as well as by comparing results obtained with different-size clusters. Excepting self-transitions, we found five transitions to be highly significant (see Table 6): $NAO^+ \rightarrow AO^+$, $NAO^- \rightarrow NAO^+$, $AO^+ \rightarrow NAO^-$, $AO^+ \rightarrow NAO^+$, and $NAO^- \rightarrow AO^+$; the first three were significant independently of cluster size, the last two only for larger clusters (1.25σ).

These five significant transitions can be organized into two cycles. The first one, $NAO^+ \rightarrow AO^+ \rightarrow NAO^- \rightarrow NAO^+$, is significant at the 95% level, independently of cluster size. The second one, $AO^+ \leftrightarrow NAO^+$ is only so for the larger-size clusters. It is remarkable that direct transitions between the opposite phases of the AO are significantly unlikely, at the 95% level, independently of cluster size. Likewise a direct transition from the zonal NAO^+ to the blocked NAO^- is quite unlikely (see again Table 6). As previously argued by Mo and Ghil (1988) and Kimoto and Ghil (1993b), based on observational data, the present model results strongly support the nonlinear origin of these well-established NH regimes.

The third step in our coarse-graining approach had only been suggested in previous work (Ghil 1987; Ghil and Childress 1987, section 6.4). This step consists of examining whether there are preferential directions in which the system's trajectory in phase space leaves the regimes and it has been carried out here for the first time, to the best of our knowledge.

We performed this step in two different ways. First, we computed, in the 3D subspace spanned by the model's three leading EOFs, the exit PDF on the unit sphere around each regime centroid (shown in Fig. 4 for the larger-size clusters). The solid-angle PDFs for all but two significant transitions exhibit one ($NAO^- \rightarrow AO^+$) or two ($NAO^+ \rightarrow AO^+$ and $NAO^- \rightarrow NAO^+$) sharp peaks; the single peak for $AO^+ \rightarrow NAO^-$ and the double one for $AO^+ \rightarrow NAO^+$ are somewhat broader and thus the PDF maxima, while still present, are less pronounced.

Second, we computed full-map composites of all the regime breaks associated with the five significant transitions. To be included in such a composite, the anomaly map has to be the midpoint between a data point on the full phase-space trajectory that has just left a cluster and the preceding one, just inside the cluster. The difference between this composite exit anomaly and that of the cluster centroid from which exit occurred is shown, for each significant transition, in Fig. 5. It is also projected on the three leading EOFs and then on the unit sphere around each cluster centroid. The results are shown as filled circles in Figs. 4a–e.

For each panel in Fig. 4 we have thus to compare three points: the asterisk that shows the direction of the straight line passing through the two regime centroids involved in the transition, the filled triangle that indicates the global maximum of the exit-angle PDF, and the filled circle, as just explained. The difference in locations between the filled circle and the filled triangle corresponds to the fact that the former is computed using only the values of PCs 1–3 of each map, while the latter is based on the actual map (including, as it were, all PCs). This difference is only sizable for Fig. 4a.

Clearly, in the case of the exit-angle PDFs with a single sharp peak (Fig. 4c) the filled triangle marks this peak. For the double-peaked PDFs the preferred mean exit angle (filled circle) lies roughly between the two peaks (Figs. 4b,e). Only for the transition $AO^+ \rightarrow NAO^+$ (Fig. 4e) does the straight line that passes through the two centroids (asterisk) lie close to the preferred mean exit angle (filled triangle or filled circle); see also Figs. 1a,b. It follows that the actual preferred paths for transitions between pairs of clusters lean overwhelmingly away from the line segment that unites the corresponding pair of centroids.

2) INTRASEASONAL OSCILLATIONS

A complementary way of describing the skeleton of our model's attractor is by using the least-unstable periodic orbits (Ott et al. 1994, and references therein). We carried out multichannel SSA (M-SSA) of the 54 000-day-long trajectory (Plaut and Vautard 1994; Ghil et al. 2002) in the same subspace spanned by the three leading EOFs as for the Markov chain. Two pairs of oscillatory modes are statistically significant at the 95% level (Fig. 6); they have periods of 37 days (modes 8–9) and 19 days (modes 15–16). The spatial patterns associated with a half-cycle of the 37-day oscillation are shown in Fig. 7; they are $NAO^+ \rightarrow AO^+ \rightarrow PNA \rightarrow NAO^-$. The Markov-chain cycle $NAO^+ \rightarrow AO^+ \rightarrow NAO^- \rightarrow NAO^+$ thus appears embedded in the 37-day oscillation.

In fact, the average time taken by all the trajectory

segments that complete the above cycle of three residence times in NAO^+, AO^+, and NAO^-, as well as the three transitions indicated, equals about 19 days. The composite patterns for the 19-day oscillation resemble very well those of the Fig. 7, but the two oscillations do not seem to be phase locked.

The spatiotemporal patterns of the oscillations present certain similarities with the Branstator (1987)–Kushnir (1987) wave, on the one hand, and the oscillatory topographic instability of Ghil and associates (Ghil 1987; Ghil and Robertson 2002, and further references therein), on the other. A closer evaluation of these similarities, as well as of the differences, goes beyond the purpose of the present paper.

b. Discussion

The preferred transition paths between the two regimes found in a stochastically forced Lorenz (1963a) model (see Figs. 9 and 10) also point away from the line segment uniting the two centroids. In this simple case, the dynamic cause of this statistical result is well known: it is associated with the geometry of the unstable and stable manifolds of the original, deterministic model version's unstable fixed points (Guckenheimer and Holmes 1983; Ghil and Childress 1987, section 5.4) and with the fact that the deformation of these manifolds, as model parameters change, leads to the existence of hetero- and homoclinic ''explosions'' (Sparrow 1982) that generate its deterministic chaos. The stochastic forcing (Sutera 1980) merely plays the role of the small-scale, high-frequency components of the QG3 model and smears out the PDF associated with the deterministic Lorenz model's chaotic behavior.

Hetero- and homoclinic orbits are well understood to play a major role in the nonlinear behavior of dynamical systems in general, in spite of their elusiveness (Wiggins 1988). A homoclinic orbit arises when a fixed point's unstable manifold returns as its stable manifold. A heteroclinic connection involves the unstable manifold of a fixed point becoming the stable manifold of another fixed point. The importance of such orbits in the irregular behavior of atmospheric and oceanic models, and maybe of the actual atmosphere and oceans, has recently been discussed more widely. We sketch, therefore, a brief history of these applications and how it might relate to the results of this paper.

Lorenz (1963b) was the first, to the best of our knowledge, to mention the potential role of a heteroclinic orbit in transition to chaos for an atmospheric model. Ghil and Childress (1987, sections 5.3 and 5.5) discussed at length this role in connection with the Shilnikov (1965) ''route to chaos.'' These authors also discussed, at some length, the role of unstable equilibria and a homoclinic orbit (see Fig. 6.12 and section 6.4 there) in the genesis of regimes, as well as in their persistence, slow onset, and rapid break. Both hetero- and homoclinic orbits are hard to compute exactly, because of their existence at isolated parameter values only. Still, they can strongly influence the behavior of the system at neighboring parameter values. A simple illustration is given in Ghil and Childress (1987, Fig. 12.10 and section 12.2) for an oscillatory paleoclimate model.

Kimoto and Ghil (1993b) and Weeks et al. (1997) mentioned the possible role of a heteroclinic orbit in giving rise to regime transitions. Hetero- and homoclinic orbits were shown to play an important role in the interannual and interdecadal LFV of the midlatitude ocean's wind-driven circulation (Meacham 2000; Chang et al. 2001; Nadiga and Luce 2001; Simmonet et al. 2003a,b). Crommelin (2002, 2003) explored in detail the role of such orbits in a low-order and an intermediate-order model of the extratropical atmosphere, respectively.

It is plausible that the statistically significant transitions between our model's four regimes are related to the presence, at nearby parameter values, of heteroclinic connections between unstable fixed points that lie near the cluster centroids. This conjecture is supported by the results of D'Andrea and Vautard (2001) and D'Andrea (2002), who found a good correspondence between their low-order model's quasi-stationary states and the QG3 model's regimes. The presence of heteroclinic connections is especially likely in the case of the transitions that lean quite noticeably away from the straight-line segments between pairs of centroids. Work in progress (D. Kondrashov et al. 2003, unpublished manuscript) using a lower-order, stochastically forced model is pursuing further the possible connections between preferred transitions and nonlinear dynamics in the QG3 model.

The model's closed Markov chain cycle and its 37-day oscillatory mode both exhibit composite 500-hPa maps that resemble the NAO^+, AO^+, and NAO^- spatial patterns in this order. This strong similarity is another illustration of the interesting, but elusive, relationship between the episodic, multiple-regime description of atmospheric LFV and the oscillatory one, based on intraseasonal oscillations (Kimoto and Ghil 1993b; Plaut and Vautard 1994; Ghil and Robertson 2002; Koo and Ghil 2002). It is quite possible, although still to be confirmed (see the previous paragraph), that these three regimes, all of which lie in the PC-1 ≤ 0 half of the model's phase space, are connected by heteroclinic orbits. Not all three of these need coexist at exactly the same set of parameter values; it suffices that the appropriate parameter values lie close to each other and to those that are most realistic for matching the model behavior to NH atmospheric LFV.

Last, but not least, the original motivation for undertaking this study was our interest in exploiting satellite data, including satellite images, for the early identification of regime breaks. The preferred transition patterns shown in Figs. 5a–e have synoptic features that are easily identifiable and can foretell an abrupt break in a given regime, while the day's map lies still within

a fairly small distance of that regime's centroid. This identification might be done numerically, within the evolution of an operational forecast-analysis cycle that uses an advanced numerical weather prediction model.

It might also be possible to identify these transition patterns by using directly advanced image processing methods on cloud imagery or other fields produced by satellite-based instruments, with or without the use of model assimilation of these fields. It is true that spatio-temporal LFV patterns captured in observational datasets only reflect a few tens of percent of the total variance of the data (Kimoto and Ghil 1993a,b; Brunet 1994). Still, advanced data mining and knowledge discovery tools from computer science and nonparametric statistics (Hand et al. 2001) can help extract this information from the high level of weather noise, due to its regularities in space or time.

Acknowledgments. We are grateful to Fabio D'Andrea for providing a code of the QG3 model and information on its performance. Andrew W. Robertson and Padhraic J. Smyth provided the codes for the *k*-means method and the Gaussian mixture model, respectively; it is a pleasure to thank them for the codes and their advice on using these. We also wish to thank Sergey Kravtsov and Ricardo Todling for helpful discussions, and Daan Crommelin for sharing preprints of his recent work. Gilbert Brunet and an anonymous reviewer provided constructive comments on a previous version of the paper. Our study was supported by NASA Grant NAG5-9294 and NSF Grant ATM00-8213. This is publication 5825 of UCLA's Institute of Geophysics and Planetary Physics.

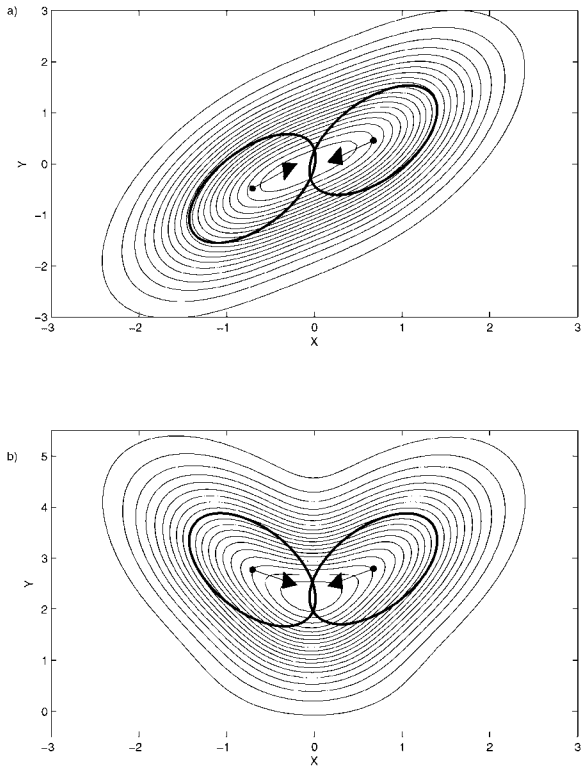

FIG. A1. PDF distribution of the stochastically perturbed Lorenz model, as estimated by a mixture model with $k = 2$, projected onto (a) the (X, Y) plane and (b) the (X, Z) plane. All three time series have been normalized by the standard deviation of $X(t)$ and filled circles correspond to the two unstable steady states W_1 and W_2 (see text for details). Contours as in Fig. 1 and the heavy solid ellipses correspond to 1.0σ. Arrows indicate the preferred transition paths obtained by the method described in section 4b.

APPENDIX

Regime Transition Angles in the Lorenz Model

The Lorenz (1963a) system of three coupled, nonlinear, ordinary differential equations is obtained by spectral truncation of the Boussinesq equations for 2D Rayleigh–Bénard convection. We use this system subject to stochastic forcing, following Sutera (1980):

$$\dot{X} = -sX + sY + \epsilon, \tag{A1a}$$

$$\dot{Y} = -XZ + rX - Y + \epsilon, \tag{A1b}$$

$$\dot{Z} = XY - bZ + \epsilon, \tag{A1c}$$

where ϵ is a white-noise process (see below). In the original Lorenz system (for $\epsilon = 0$) the parameters $s = 10$, $r = 28$, and $b = 8/3$ lead to chaotic internal variability caused by the system's instabilities and nonlinearity (see also Guckenheimer and Holmes 1983; Ghil and Childress 1987).

There are two unstable steady states at $W_1 = (6\sqrt{2}, 6\sqrt{2}, 27)$ and $W_2 = (-6\sqrt{2}, -6\sqrt{2}, 27)$. Each of these states is stable in one direction, and unstable in the other two. The third steady state $W_0 = (0, 0, 0)$ is stable along the Z axis and in one additional direction, and unstable in the third. Due to the oscillatory instability of W_1 and W_2, the trajectories spiral out around either one of these two points with increasing amplitude. As each such spiral reaches a threshold value imposed by the system's quadratic nonlinearity, it leaves the neighborhood of the unstable point it started from, crosses near the axis of mirror symmetry given by $X = Y = 0$, and lands near the opposite unstable point. This behavior then repeats itself in the opposite direction so that the trajectory switches irregularly from the neighborhood of W_1 to that of W_2 and back. These repetitions occur irregularly in time and give rise to two regimes in the Lorenz system, each of which consists of the spiral segments of orbit around W_1 and W_2, respectively.

To make the analogy with the QG3 model more realistic, we stochastically force the Lorenz system with a white-noise term $\epsilon = \epsilon(t)$, having variance $\sigma_\epsilon^2 = 2$. This forcing provides a crude approximation of the effect of the smaller scales and higher frequencies on the LFV of the QG3 model and the atmosphere. More pre-

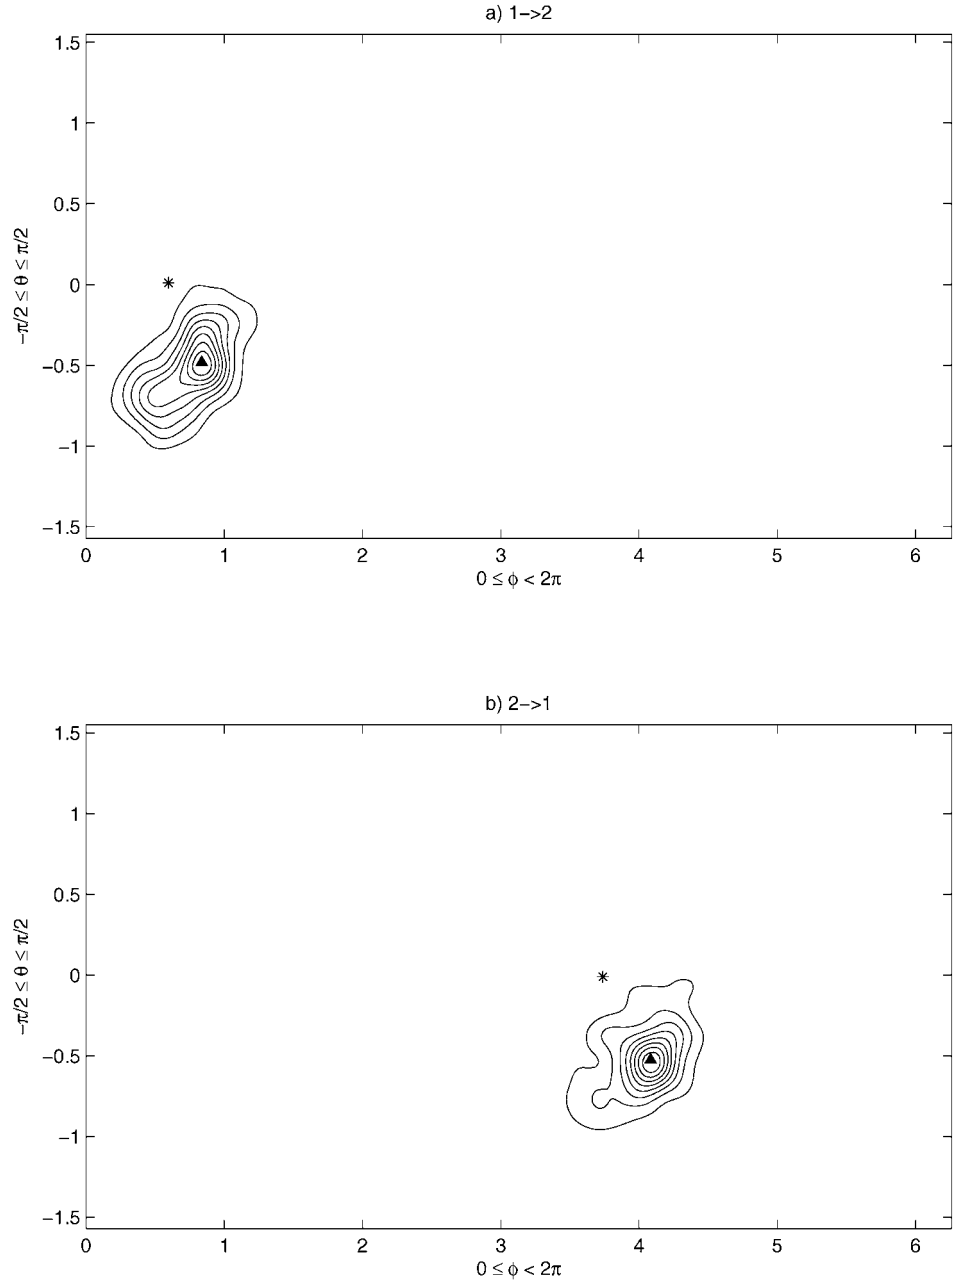

FIG. A2. PDF distribution of regime exit angles for the Lorenz model's two clusters. Asterisks correspond to the line connecting the cluster centroids, and filled triangles correspond to PDF maxima.

cisely, we transform the system (A1) into its proper stochastic differential form:

$$d\mathbf{X} = \mathbf{f}(\mathbf{X})dt + \sigma_\epsilon d\beta, \tag{A2}$$

where $\mathbf{X} = (X, Y, Z)^T$, and $\mathbf{f}$ is the nonlinear, deterministic right-hand side of (A1), while $\beta(t)$ is the appropriate Wiener process whose ''derivative'' is the white noise $\epsilon(t)$. We use a Euler scheme with a time step of 0.01 (Kloeden and Platen 1992) to obtain a long realization of the stochastic process governed by (A2) for 50 000 time steps.

The strong non-Gaussian nature of the Lorenz attractor leads to a systematic bias toward a higher number of clusters when a cross-validated log-likelihood procedure is used (Hannachi and O'Neill 2001). Since we know already the true equations generating the data and the regimes in our system, we apply the mixture model with two clusters to the time series of (X, Y, Z). In the

original Lorenz system ($\epsilon = 0$) the trajectories are bound away from the points W_1 and W_2 at this value of s, r, and b. Stochastic forcing fills in these gaps, as seen in Figs. A1a,b of the mixture model's PDF, projected on the X–Y and X–Z planes, respectively. The arrows in Fig. A1 are projections of vectors that lie along preferred transition paths between regimes that are obtained by the same method as described in section 4.

The regime exit angles ϕ and θ and their PDF distributions for the Lorenz system's two clusters are computed with the same procedure as for the QG3 model, and the results are shown in Fig. A2. The locations of the two sharp maxima are explained by the orbit's crossing from one regime into the other in the vicinity of the Z axis. When the orbit's projection on the X–Y plane spirals out within one of the clusters, it approaches the origin along the stable direction of the W_0 state; this direction does not coincide with the straight line that passes through W_1 and W_2. The orbit then leaves the neighborhood of the origin along the unstable direction of the W_0 state, thus entering the other cluster. The values of ϕ at the PDF maxima are therefore higher than the value that corresponds to the $W_1 \rightarrow W_2$ line segment. The negative but small value of θ at the PDF maximum is caused by the orbit proceeding in the negative direction along the Z axis while exiting either cluster.

Ideally, the symmetry of the dynamics about the Z axis should lead to the symmetry of the exit-angle PDF between the two regimes; this symmetry would imply that Figs. A2a and A2b should be identical, modulo a shift of π in the ϕ direction. The length of the dataset here, however, has been chosen to match approximately the length of the QG3 model integration, in appropriate units. Therefore the mismatch between Figs. A2a and A2b gives an idea of the sampling-error effect in Figs. 4a–e.

The connection between the detailed geometry of regime transitions in the Lorenz model, as captured in Figs. A1 and A2, and the structure of the stable and unstable manifolds of the unstable fixed points W_0, W_1, and W_2 should be apparent by inspecting Figs. 5.10 and 5.11 of Ghil and Childress (1987). Sparrow (1982) has investigated thoroughly the role of homoclinic and heteroclinic orbits that arise from these manifolds, given changes in the Lorenz (1963a) model's parameters s, r, and b, in generating the model's chaotic behavior.

REFERENCES

Allen, M. R., and A. W. Robertson, 1996: Distinguishing modulated oscillations from coloured noise in multivariate datasets. *Climate Dyn.,* **12,** 775–784.

Branstator, G. W., 1987: A striking example of the atmosphere's leading traveling pattern. *J. Atmos. Sci.,* **44,** 2310–2323.

Broomhead, D. S., and G. P. King, 1986a: Extracting qualitative dynamics from experimental data. *Physica D,* **20,** 217–236.

——, and ——, 1986b: On the qualitative analysis of experimental dynamical systems. *Nonlinear Phenomena and Chaos,* S. Sarkar, Ed., Adam Hilger, 113–144.

Brunet, G., 1994: Empirical normal-mode analysis of atmospheric data. *J. Atmos. Sci.,* **51,** 932–952.

Chang, K.-I., M. Ghil, K. Ide, and C.-C. A. Lai, 2001: Transition to aperiodic variability in a wind-driven double-gyre circulation model. *J. Phys. Oceanogr.,* **31,** 1260–1286.

Cheng, X. H., and J. M. Wallace, 1993: Cluster-analysis of the Northern Hemisphere wintertime 500-hPa height field: Spatial patterns. *J. Atmos. Sci.,* **50,** 2674–2696.

Corti, S., A. Giannini, S. Tibaldi, and F. Molteni, 1997: Patterns of low-frequency variability in a three-level quasi-geostrophic model. *Climate Dyn.,* **13,** 883–904.

Crommelin, D. T., 2002: Homoclinic dynamics: A scenario for atmospheric ultralow-frequency variability. *J. Atmos. Sci.,* **59,** 1533–1549.

——, 2003: Regime transitions and heteroclinic connections in a barotropic atmosphere. *J. Atmos. Sci.,* **60,** 229–246.

D'Andrea, F., 2002: Extratropical low-frequency variability as a low-dimensional problem. Part II: Stationarity and stability of large-scale equilibria. *Quart. J. Roy. Meteor. Soc.,* **128,** 1059–1073.

——, and R. Vautard, 2001: Extratropical low-frequency variability as a low-dimensional problem. Part I: A simplified model. *Quart. J. Roy. Meteor. Soc.,* **127,** 1357–1374.

Dickey, J. O., M. Ghil, and S. L. Marcus, 1991: Extratropical aspects of the 40–50 day oscillation in length-of-day and atmospheric angular momentum. *J. Geophys. Res.,* **96,** 22 643–22 658.

Dole, R. M., and N. D. Gordon, 1983: Persistent anomalies of the extratropical Northern Hemisphere wintertime circulation—Geographical distribution and regional persistence characteristics. *Mon. Wea. Rev.,* **111,** 1567–1586.

Ghil, M., 1987: Dynamics, statistics and predictability of planetary flow regimes. *Irreversible Phenomena and Dynamical Systems Analysis in the Geosciences,* C. Nicolis and G. Nicolis, Eds., D. Reidel, 241–283.

——, and S. Childress, 1987: *Topics in Geophysical Fluid Dynamics: Atmospheric Dynamics, Dynamo Theory and Climate Dynamics.* Springer-Verlag, 485 pp.

——, and K.-C. Mo, 1991: Intraseasonal oscillations in the global atmosphere. Part I: Northern Hemisphere and Tropics. *J. Atmos. Sci.,* **48,** 752–779.

——, and A. W. Robertson, 2002: ''Waves'' vs. ''particles'' in the atmosphere's phase space: A pathway to long-range forecasting? *Proc. Natl. Acad. Sci.,* **99** (Suppl. 1), 2493–2500.

——, S. L. Marcus, J. O. Dickey, and C. L. Keppenne, 1991: *AAM the Movie.* Pasadena, CA: NTSC videocassette AVC-91-063. [Available from Caltech/NASA Jet Propulsion Laboratory, Pasadena, CA 91109, also available from author upon request.]

——, and Coauthors, 2002: Advanced spectral methods for climatic time series. *Rev. Geophys.,* **40,** 1003, doi: 10.1029/2000GR000092.

Guckenheimer, J., and P. Holmes, 1983: *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* Springer-Verlag, 453 pp.

Haines, K., 1994: Low-frequency variability in atmospheric middle latitudes. *Surv. Geophys.,* **15,** 1–61.

Hand, D., H. Mannila, and P. Smyth, 2001: *Principles of Data Mining.* MIT Press, 546 pp.

Hannachi, A., and A. O'Neill, 2001: Atmospheric multiple equilibria and non-Gaussian behaviour in model simulations. *Quart. J. Roy. Meteor. Soc.,* **127,** 939–958.

Kalnay-Rivas, E., and L.-O. Merkine, 1981: A simple mechanism for blocking. *J. Atmos. Sci.,* **38,** 2077–2091.

Keppenne, C. L., and M. Ghil, 1993: Adaptive filtering and prediction of noisy multivariate signals: An application to subannual variability in atmospheric angular momentum. *Int. J. Bifurcation Chaos,* **3,** 625–634.

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

——, and ——, 1993b: Multiple flow regimes in the Northern Hemi-

sphere winter. Part II: Sectorial regimes and preferred transitions. *J. Atmos. Sci.,* **50,** 2645–2673.
Kloeden, P. E., and E. Platen, 1992: *Numerical Solution of Stochastic Differential Equations.* Springer-Verlag, 632 pp.
Koo, S., and M. Ghil, 2002: Successive bifurcations in a simple model of atmospheric zonal-flow vacillation. *Chaos,* **12,** 300–309.
——, A. W. Robertson, and M. Ghil, 2002: Multiple regimes and low-frequency oscillations in the Southern Hemisphere's zonal-mean flow. *J. Geophys. Res.,* **107,** 4596, doi: 10.1029/2001JD001353.
Kushnir, Y., 1987: Retrograding wintertime low-frequency disturbances over the North Pacific Ocean. *J. Atmos. Sci.,* **44,** 2727–2742.
Legras, B., and M. Ghil, 1985: Persistent anomalies, blocking and variations in atmospheric predictability. *J. Atmos. Sci.,* **42,** 433–471.
Lorenz, E. N., 1963a: Deterministic nonperiodic flow. *J. Atmos. Sci.,* **20,** 130–141.
——, 1963b: The mechanics of vacillation. *J. Atmos. Sci.,* **20,** 448–464.
Marshall, J., and F. Molteni, 1993: Towards a dynamical understanding of planetary-scale flow regimes. *J. Atmos. Sci.,* **50,** 1792–1818.
Meacham, S. P., 2000: Low-frequency variability in the wind-driven circulation. *J. Phys. Oceanogr.,* **30,** 269–293.
Michelangeli, P. A., R. Vautard, and B. Legras, 1995: Weather regimes: Recurrence and quasi stationarity. *J. Atmos. Sci.,* **52,** 1237–1256.
Mo, K., and M. Ghil, 1987: Statistics and dynamics of persistent anomalies. *J. Atmos. Sci.,* **44,** 877–901.
——, and M. Ghil, 1988: Cluster analysis of multiple planetary flow regimes. *J. Geophys. Res.,* **93D,** 10 927–10 952.
Molteni, F., and S. Corti, 1998: Long-term fluctuations in the statistical properties of low-frequency variability: Dynamical origin and predictability. *Quart. J. Roy. Meteor. Soc.,* **124,** 495–526.
——, S. Tibaldi, and T. N. Palmer, 1990: Regimes in the wintertime circulation over northern extratropics. 1. Observational evidence. *Quart. J. Roy. Meteor. Soc.,* **116,** 31–67.
Mukougawa, H., 1988: A dynamical model of ''quasi-stationary'' states in large-scale atmospheric motions. *J. Atmos. Sci.,* **45,** 2868–2888.
Nadiga, B. T., and B. P. Luce, 2001: Global bifurcation of Shilnikov type in a double-gyre ocean model. *J. Phys. Oceanogr.,* **31,** 2669–2690.
Ott, E., T. Sauer, and J. A. Yorke, Eds., 1994: *Coping with Chaos: Analysis of Chaotic Data and the Exploitation of Chaotic Systems.* J. Wiley & Sons, 418 pp.
Plaut, G., and R. Vautard, 1994: Spells of low-frequency oscillations and weather regimes in the Northern Hemisphere. *J. Atmos. Sci.,* **51,** 210–236.
Robertson, A. W., and M. Ghil, 1999: Large-scale weather regimes and local climate over the western United States. *J. Climate,* **12,** 1796–1813.
Shilnikov, L. P., 1965: A case of the existence of a denumerable set of periodic motions. *Sov. Math. Dokl.,* **6,** 163–166.
Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis.* Chapman and Hall, 175 pp.
Simonnet, E., M. Ghil, K. Ide, R. Temam, and S. Wang, 2003a: Low-frequency variability in shallow-water models of the wind-driven ocean circulation. Part I: Steady-state solutions. *J. Phys. Oceanogr.,* **33,** 712–728.
——, ——, ——, ——, and ——, 2003b: Low-frequency variability in shallow-water models of the wind-driven ocean circulation. Part II: Time-dependent solutions. *J. Phys. Oceanogr.,* **33,** 729–752.
Smyth, P., K. Ide, and M. Ghil, 1999: Multiple regimes in Northern Hemisphere height fields via mixture model clustering. *J. Atmos. Sci.,* **56,** 3704–3723.
Sparrow, C., 1982: *The Lorenz Equations: Bifurcations, Chaos and Strange Attractors.* Springer-Verlag, 269 pp.
Sutera, A., 1980: Stochastic pertubation of a pure convective motion. *J. Atmos. Sci.,* **37,** 246–249.
Thompson, D. W., and J. M. Wallace, 1998: The Arctic Oscillation signature in the wintertime geopotential height and temperature fields. *Geophys. Res. Lett.,* **25,** 1297–1300.
Vautard, R., 1990: Multiple weather regimes over the North Atlantic: Analysis of precursors and successors. *Mon. Wea. Rev.,* **118,** 2056–2081.
——, and B. Legras, 1988: On the source of low frequency variability. Part II: Nonlinear equilibration of weather regimes. *J. Atmos. Sci.,* **45,** 2845–2867.
——, K. C. Mo, and M. Ghil, 1990: Statistical significance test for transition matrices of atmospheric Markov chains. *J. Atmos. Sci.,* **47,** 1926–1931.
Wallace, J. M., 2000: North Atlantic Oscillation/annular mode: Two paradigms—one phenomenon. *Quart. J. Roy. Meteor. Soc.,* **126,** 791–805.
Weeks, E. R., Y. Tian, J. S. Urbach, K. Ide, H. L. Swinney, and M. Ghil, 1997: Transitions between blocked and zonal flows in a rotating annulus with topography. *Science,* **278,** 1598–1601.
Wiggins, S., 1988: *Global Bifurcations and Chaos: Analytical Methods.* Springer-Verlag, 494 pp.