

Downscaling of Daily Rainfall Occurrence over Northeast Brazil Using a Hidden Markov Model

ANDREW W. ROBERTSON

International Research Institute for Climate Prediction, The Earth Institute at Columbia University, Palisades, New York

SERGEY KIRSHNER AND PADHRAIC SMYTH

School of Information and Computer Science, University of California, Irvine, Irvine, California

(Manuscript received 4 December 2003, in final form 24 April 2004)

ABSTRACT

A hidden Markov model (HMM) is used to describe daily rainfall occurrence at 10 gauge stations in the state of Ceará in northeast Brazil during the February–April wet season 1975–2002. The model assumes that rainfall occurrence is governed by a few discrete states, with Markovian daily transitions between them. Four “hidden” rainfall states are identified. One pair of the states represents wet-versus-dry conditions at all stations, while a second pair of states represents north–south gradients in rainfall occurrence. The estimated daily state-sequence is characterized by a systematic seasonal evolution, together with considerable variability on intraseasonal, interannual, and longer time scales. The first pair of states are shown to be associated with large-scale displacements of the tropical convergence zones, and with teleconnections typical of the El Niño–Southern Oscillation and the North Atlantic Oscillation.

A nonhomogeneous HMM (NHMM) is then used to downscale daily precipitation occurrence at the 10 stations, using general circulation model (GCM) simulations of seasonal-mean large-scale precipitation, obtained with historical sea surface temperatures prescribed globally. Interannual variability of the GCM’s large-scale precipitation simulation is well correlated with seasonal- and spatial-averaged station rainfall-occurrence data. Simulations from the NHMM are found to be able to reproduce this relationship. The GCM-NHMM simulations are also able to capture quite well interannual changes in daily rainfall occurrence and 10-day dry spell frequencies at some individual stations. It is suggested that the NHMM provides a useful tool (a) to understand the statistics of daily rainfall occurrence at the station level in terms of large-scale atmospheric patterns, and (b) to produce station-scale daily rainfall sequence scenarios for input into crop models, etc.

1. Introduction

One of the major challenges in tailoring seasonal climate forecasts to meet societal needs, is that the potential users of climate information are often concerned with the characteristics of high-frequency weather at a particular location. Unfortunately, the statistics of local weather are generally poorly represented in the coarse-resolution general circulation models (GCMs) that are typically used to make seasonal climate forecasts. Moreover, the seasonal predictability of high-frequency local information is often in serious doubt. As part of the endeavor to produce useful seasonal climate forecasts, an important task is to understand, on a regional basis, just what aspects of the daily “weather-within-climate” can be obtained from GCM simulations made with prescribed historical or predicted sea surface temperature (SST) distributions (Wilks 2002). While the determin-

istic predictability limit of weather is on the order of 10 days, some boundary-forced longer-term predictability of weather *statistics* may often exist—termed predictability of the “second kind” by Lorenz (1963).

Northeast (NE) Brazil is a region with a high potential for seasonal rainfall predictability during the February–April rainy season, due to strong teleconnections with the El Niño–Southern Oscillation (ENSO) and with variability of the tropical Atlantic Ocean (Hastenrath and Heller 1977; Moura and Shukla 1981; Nobre and Shukla 1996). At the seasonal-mean scale, some of this forecast potential has been realized in the seasonal forecasts made at the International Research Institute for Climate Prediction (IRI) during the 1998–2001 period (Goddard et al. 2003). Some of the highest rainfall skills on the globe occurred over NE Brazil during this short time interval since the advent of routine IRI seasonal forecasts. The state of Ceará in NE Brazil is also the target of a water-resource management project¹ whose goal is

Corresponding author address: Andrew W. Robertson, IRI–Monell 230, 61 Route 9W, Palisades, NY 10964.
E-mail: awr@iri.columbia.edu

¹ Information can be found online at <http://iri.columbia.edu/america/>.

to incorporate climate forecast information into sectorial decision making, for which the “downscaling” potential from GCM simulations to local daily rainfall is an important issue.

The first goal of this paper is to examine the probability distribution of local daily rainfall occurrence in the state of Ceará on time scales of intraseasonal to interannual, and to identify relationships with large-scale atmospheric circulation patterns. Our approach is based on fitting a hidden Markov model (HMM) to rainfall records on a network of 10 rain gauge stations for the period 1975–2002. The HMM is a doubly stochastic model in which the daily probability of local rainfall occurrence (typically on a spatial network of stations) is conditioned on a small number of discrete (“hidden”) weather states, with Markovian daily transitions between them. It is a data-driven model, in which the parameters defining each state and the estimated daily sequence of states are derived from a historical record of daily rainfall occurrence. The concept of discrete weather states has a long history in synoptic (Bauer 1951) as well as theoretical midlatitude meteorology (Charney and DeVore 1979). No such literature exists for the Tropics, however. We seek to establish a physical basis for the predictability of local daily rainfall statistics over NE Brazil by constructing relationships between the HMM’s hidden rainfall states and the large-scale modes of atmosphere–ocean variability associated with ENSO and tropical Atlantic variability (TAV).

For the downscaling of rainfall, Hughes and Guttorm (1994a) pioneered the use of the *nonhomogeneous* HMM (NHMM) in which the transition probabilities between states are not held fixed, as they are in the classical “homogeneous” HMM, but are allowed to evolve in time according to a small number of large-scale atmospheric exogenous predictor variables. The NHMM links the local rainfall at a network of stations to large-scale atmospheric variables, using the hidden weather states as intermediaries. In the context of downscaling of seasonal forecasts, it may be most meaningful to allow the daily transition probabilities between states to be a function of the GCM’s seasonal-average simulations. This would be tantamount to assuming that the information content of the GCM is limited to the seasonal average; it provides a benchmark against which more complex hypotheses can be tested involving atmospheric forcing variables derived from the GCM’s daily output. A related downscaling approach has been taken by Wilks (2002), in which the parameters of a weather generator are conditioned on the seasonal forecast.

Northeast Brazil is located near the boundary between two large-scale tropical convergence zones—the Atlantic intertropical convergence zone (ITCZ) to the north, and the South American monsoon system (SAMS) centered to the west—and the influence of the South Atlantic subtropical anticyclone situated to the southeast. The rainfall over NE Brazil is highly seasonal and is

sensitive to anomalies in the extent of any one of these three phenomena (Hastenrath and Heller 1977). It is semiarid for much of the year, falling under the influence of the western fringe of the South Atlantic subtropical anticyclone. The principal rainfall season occurs in February–March–April (FMA), when the Atlantic ITCZ reaches its southernmost position and directly overlies NE Brazil, merging with the SAMS near the mouth of the Amazon, as the SAMS (which peaks in austral summer) retreats northward. The time lag in the seasonal migration of the maritime ITCZ is caused by the large thermal heat capacity of the underlying ocean, together with the impact of the continental convection associated with the SAMS which tends to weaken the ITCZ during FMA.

Interannual variability of rainfall is large, and takes the form of droughts that occur when the usual southward seasonal migration of the ITCZ fails to occur, and large rainfall when the latter is amplified. The interannual behavior of the Atlantic ITCZ during boreal spring is closely tied to two interrelated factors: (a) anomalies in the Walker circulation associated with ENSO events over the tropical Pacific, and (b) changes in the meridional SST gradient in the equatorial Atlantic—the so-called Atlantic *meridional mode*—that may or may not be associated with ENSO (Chiang et al. 2002). The largest ITCZ displacements occur in ENSO years in which preexisting Atlantic SST anomalies are such that they amplify the direct impact of ENSO (Giannini et al. 2004). The North Atlantic Oscillation (NAO)—which peaks in boreal winter—is another factor that influences the tropical Atlantic meridional mode. The NAO is an intrinsic mode of the atmosphere with a time scale of about a week (Feldstein 2000). The influence of North Atlantic SST anomalies on the NAO is generally found to be small (Hurrell et al. 2002), although some seasonal NAO predictability may arise from stratospheric coupling (Thompson et al. 2002) or antecedent autumn snow cover (Gong et al. 2003). The meridional mode is influenced from the south as well, and Rossby wave energy on submonthly time scales emanating from the South Pacific storm track can reach NE Brazil even during the austral summer (Liebmann et al. 1999).

GCM forecasts of the seasonal-average precipitation and near-surface temperature are currently made on a routine basis at IRI and at other centers, issued as 3-month seasonal averages on the grid-scale of the atmospheric GCMs (typically about 300 km). Experimental downscaled seasonal forecasts are now also being made on an ongoing basis for NE Brazil using high-resolution regional dynamical models (60-km grid; Sun et al. 2005). Preliminary results are encouraging. However, these high-resolution dynamical models also suffer from model biases and are computationally expensive, motivating the use of statistical approaches.

Having established that certain HMM rainfall states are strongly tied to the large-scale atmospheric circu-

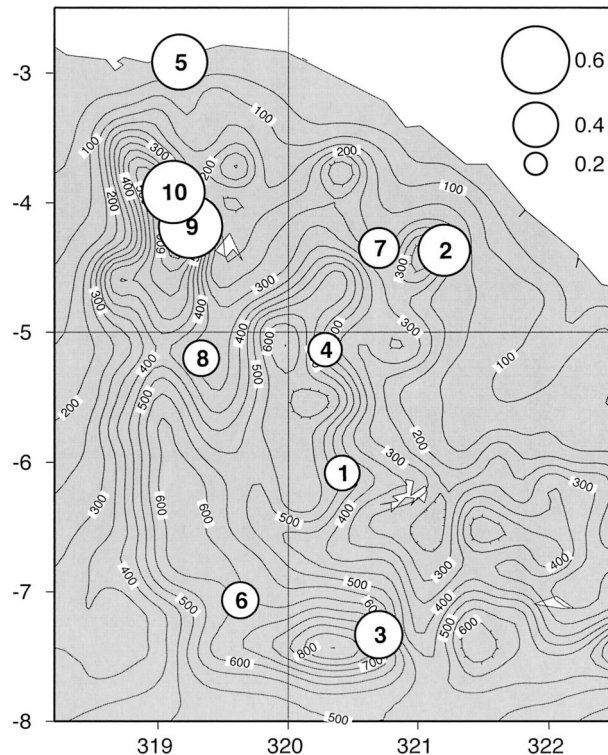


FIG. 1. Rainfall station locations with topographic contours (m). Circle radius denotes the Feb–Apr climatological daily rainfall probability 1975–2002. The stations are (1) Acopiara (317 m), (2) Araçoiaba (107 m), (3) Barbalha (405 m), (4) Boa Viagem (276 m), (5) Camocim (5 m), (6) Campos Sales (551 m), (7) Caninde (15 m), (8) Crateus (275 m), (9) Guaraciaba Do Norte (902 m), and (10) Ibiapina (878 m). One degree of lon/lat corresponds to about 110 km at the equator.

lation, the paper's second goal is to determine to what extent the nonhomogeneous HMM (NHMM) can be used to downscale atmospheric GCM simulations, with prescribed historical SSTs, over Ceará. The remainder of the paper is laid out as follows. Section 2 describes the observed daily rainfall dataset and the GCM used in the downscaling experiment. The homogeneous HMM is described briefly in section 3. The hidden states of rainfall occurrence derived from the HMM are presented in section 4 and discussed in terms of concurrent atmospheric conditions. The results of the GCM downscaling experiment using an NHMM are described in section 5. A summary and conclusions are presented in section 6.

2. Observed rainfall data and GCM

We use daily rainfall data collected at 10 stations from the state of Ceará in NE Brazil over the years 1975–2002, provided by Fundaco Cearense de Meteorologia (FUNCME; Fig. 1). These 10 stations have the longest and most complete reliable records. A wet day is defined as having nonzero rainfall. The mean seasonal cycle of rainfall occurrence is plotted in Fig. 2. We select the

90-day period beginning 1 February (FMA), corresponding to the peak rainy season over NE Brazil, and retain the seasonal cycle of rainfall occurrence within the FMA season. The years 1976, 1978, 1984, and 1986 were omitted because of missing data for certain months at one or more stations, yielding 24 complete 90-day seasons (2160 days).

The climatological values of daily rainfall occurrence probability are plotted at each station on a map of topography in Fig. 1. The largest values occur in the northwest decreasing southward, consistent with the large-scale rainfall pattern associated with the ITCZ and the SAMS. Local topography also appears to play a role in the large probabilities at stations 9 and 10, and perhaps 3 and 6.

The GCM is the ECHAM 4.5 model (Roeckner et al. 1996), for which an ensemble of 24 integrations was available, with historical SSTs prescribed at the lower boundary from the National Centers for Environmental Prediction–National Center for Atmospheric Research (NCEP–NCAR) reanalysis dataset (Kalnay et al. 1996). Each ensemble member uses the identical version of the GCM, and differs only in its initial condition.

3. The homogeneous hidden Markov model (HMM)

Adopting similar notation to that in Hughes et al. (1999), let $\mathbf{R}_t = (R_t^1, \dots, R_t^M)$ be a multivariate random vector of rainfall occurrences for a network of M rain stations. Let the observed value $r_t^i = 1$ if rain is observed on day t at station i and $r_t^i = 0$ if it is dry. Let S_t be the hidden rainfall state for day t , taking on one of the K values from 1 to K . By $\mathbf{R}_{1:T}$ and $S_{1:T}$ we will denote daily sequences of precipitation occurrences and hidden rainfall states. An HMM for rainfall data makes two conditional independence assumptions (e.g., see Hughes and Guttorp 1994a). The first assumption is that the multivariate precipitation observations \mathbf{R}_t at time t are independent of all other variables in the model up to time t , conditional on the weather state S_t at time t , that is,

$$P(\mathbf{R}_t | S_{1:t}, \mathbf{R}_{1:t-1}) = P(\mathbf{R}_t | S_t). \quad (1)$$

The second assumption is that the hidden state process, $S_{1:T}$, is first-order Markov:

$$P(S_t | S_{1:t-1}) = P(S_t | S_{t-1}), \quad (2)$$

and that this first-order Markov process is homogeneous in time, that is, the $K \times K$ transition probability matrix for Eq. (2) does not change with time. Later in the paper we develop a nonhomogeneous extension of this model where the transition probabilities are allowed to vary over time. The conditional independence assumptions are easily visualized as edges in a directed graph of the HMM, as shown in Fig. 3.

For $P(\mathbf{R}_t | S_t)$, we make the simplifying assumption that the rainfall observation at each station at time t is




FIG. 2. The mean seasonal cycle of daily rainfall occurrence frequency at each station over Cear .

independent from observations at other stations at time t , conditional on the hidden state:

$$P(\mathbf{R}_t = \mathbf{r} | S_t = s) = \prod_{m=1}^M P(R_t^m = r | S_t = s) = \prod_{m=1}^M p_{smr},$$

where $r \in \{0, 1\}$, each $p_{smr} \in [0, 1]$, and $p_{sm0} + p_{sm1} = 1$. Note that this *conditional* independence assumption given the states does not imply spatial independence of the rainfall process (which would be unreasonable). Spatial dependence is captured implicitly via the state variable, as we will see later in the paper.

Techniques for parameter fitting and prediction using homogeneous HMMs are well known in the statistical literature and, thus, need not be elaborated on here—the reader is referred to standard references such as Rabiner (1989) and MacDonald and Zucchini (1997) for details.

4. Hidden states of daily rainfall occurrence

a. Model performance

In this section, we use the homogeneous HMM to construct states of daily rainfall occurrence from the 10-




FIG. 3. Graphical model representation of a hidden Markov model.

gauge daily record. As a baseline we also evaluate the performance of a model with no hidden states, where independent Markov chains were fit to the historical data from each station using maximum likelihood. The stateless model is equivalent to the classical single-station “weather generator,” commonly used to model daily rainfall occurrence (e.g., Wilks and Wilby 1999). However, more sophisticated multistation weather generators have also been developed (e.g., Wilks 1998), and the stateless model should be interpreted as a very basic benchmark.

We used cross-validation to evaluate the quality of the fitted HMMs as a function of K the number of states. HMMs for different values of K were fit to the training data leaving out different one-fourths (six consecutive years) of the data at a time and then calculating the log-probability of the observed data for the left-out years. The resulting normalized out-of-sample values of the cross-validated log-likelihood for each model are given in Table 1 for $K = 2, 3, 4, 5, 6$, together with a normalized Bayes information criterion (BIC) for each model (see Appendix).

The *normalized* cross-validated log-likelihood is defined to be the negative of the original total (not averaged over years) cross-validated log-likelihood divided by the total number N of binary events used in evaluating the model. Here $N = 24 \times 90 \times 10$, the number of individual rainfall events across all years, days, and stations. Scaled in this fashion the normalized log-likelihood corresponds to a form of predictive entropy or uncertainty of the model in bits (for base-2 logarithms).

TABLE 1. Performance of HMMs with different numbers of states K compared to a model with independent Markov chains.

Model	Out-of-sample normalized cross-validated log-likelihood	Normalized BIC score	Out-of-sample average absolute difference in spatial correlation (per station pair)	Out-of-sample average absolute difference in rainfall persistence
Markov chains	0.6314	0.6325	0.2453	0.0641
$K = 2$	0.5856	0.5844	0.0627	0.0793
$K = 3$	0.5743	0.5743	0.0511	0.0678
$K = 4$	0.5709	0.5723	0.0465	0.0635
$K = 5$	0.5687	0.5725	0.0460	0.0644
$K = 6$	0.5684	0.5737	0.0450	0.0651

The resulting normalized scores are somewhat more interpretable than unnormalized log-likelihood scores since they lie on a scale between 0 and 1. The BIC scores in Table 1 are scaled in a similar manner (Appendix).

Not surprisingly in Table 1 the HMMs have lower predictive uncertainty and lower normalized BIC scores than the independent chains model (where lower scores correspond to better predictive ability). The cross-validated out-of-sample normalized log-likelihood decreases substantially from $K = 2$ to $K = 4$, and then levels off. The normalized BIC score reaches its minimum at $K = 4$. Given that both scores indicate that $K = 4$ is a reasonable choice for K , this is the value we chose in the following.

Measures of the models' ability to reproduce both (a) observed spatial correlations between stations and (b) rainfall persistence (the probability of rain given rain the previous day for a particular station) are also included in the table. Each number corresponds to the average (over the six cross-validation runs) of the absolute difference between the statistic calculated from (a) 3000 yr of simulated data from a model trained on

18 yr of data, and (b) actual data from the remaining six out-of-sample years. In this manner the numbers in Table 1 indicate the out-of-sample predictive power of the model in terms of both spatial correlation and rainfall persistence. The spatial correlation between a pair of stations is computed as Pearson's correlation coefficient of their respective daily rainfall binary occurrence time series; its average value over the observations is 0.2479. Rainfall persistence is defined by the ratio of number of pairs of consecutive rainy days to the total number of rainy days, with a station-averaged observed value of 0.5556.

As a benchmark, the average absolute error for a model which has no spatial correlation is 0.2453, which is also the error of the independent chains model since this model has (by definition) zero spatial correlation in its simulations (over the long run). All of the HMMs reduce this error by roughly a factor of 4, or equivalently can model roughly 80% of the total daily spatial correlation in the data. The HMM models achieve this via the hidden states which can implicitly capture marginal spatial dependence. In practice, an appropriate operational comparison of spatial dependence would be with a spatially coherent set of independent Markov models, from which sophisticated weather generators have been constructed (Wilks 1998).

There is less difference between the HMMs and the independent chains model in terms of persistence. We might expect the independent Markov chains model to outperform any HMM since wet and dry spells in daily rainfall data are often well-modeled as first-order Markov (Wilks and Wilby 1999). Empirically, in Table 1 the HMMs provide similar persistence numbers to that of the independent chains, out-of-sample. The specific value of K for the HMMs does not seem to have a large impact on the spatial correlation or persistence error numbers, apart from a slight flattening out in error after $K = 4$.

The correlation and persistence numbers in Table 1 summarize the difference in expected values for the statistics over multiple years. We find that the HMM trained on the data can also recover most of the year-to-year variability of the statistics. Figures 4 and 5 show for a selected pair of stations (for correlation) and for one station (for persistence) the observed 24 yearly values of the statistic. Also shown for comparison are 10




FIG. 4. Observed correlations between stations 1 and 8 for all 24 yr of data compared with correlations between the same stations from the data simulated 10 times (plotted stacked above each other) from the four-state HMM trained on all 24 yr. The histogram of the simulated distribution was computed from 12 000 90-day seasons.




FIG. 5. Observed persistence for station 5 for all 24 yr of data compared with persistence for the same station from the data simulated from the four-state HMM trained on all 24 yr. Details of simulations as in Fig. 4.

HMM simulations of 24 seasonal values of the same statistic produced from an HMM trained on the observed data. Statistics from the simulated data appear to be similar to the ones calculated from the observed data. Moreover, the variability in statistics for the observed years generally agrees with the empirical probability distribution of simulated statistics (using 12 000 sequences or years). Other choices of station yield similar results.

In general, the distributions of wet- and dry-spell lengths are captured quite well by the HMM, as illustrated in Fig. 6. The spell-length distributions (both simulated and observed) follow approximately geometric distributions (near-straight lines on the semilog plots in Fig. 6), while the HMM tends to underestimate the spell durations at some stations. The geometric distribution is a characteristic of the Markov model. However, the HMM treats the state sequence as a Markov chain, rather than the rainfall. Thus, while the states in data simulated from an HMM will have run lengths whose distribution is geometric, the observed precipitation may be more “bursty,” leading (for example) to possible underestimation of rainfall persistence. The within-state probability of precipitation is constant.

The state transition-probability matrix is given in Table 2. The self-transitions are relatively large indicating that the states are persistent, with states 1 and 2 being the most so, and state 4 being the least persistent. Direct transitions between states 1 and 2 are rare, with states 3 and (especially) 4 playing the role of intermediaries. There are no very clear transition directions, though state 1 tends to follow state 4 ($p = 0.22$) rather than precede it ($p = 0.11$).

The rainfall probabilities for each state are plotted in Fig. 7, along with the number of days assigned to each

state. The four states fall roughly into two pairs with states 1–2 characterized by wet or dry conditions at all stations and states 3–4 describing anomalous north–south gradients (see also Fig. 10 which shows the probabilities as anomalies from climatology). State 1 (state 2) is characterized by increased (decreased) probability of rain at all stations, compared to the climatological probabilities in Fig. 1; state 2 has near-zero probability of rain everywhere except on the coast. State 3 (state 4) has anomalously small (large) probabilities in the south and slightly increased (reduced) probabilities in the north, with near-climatological probabilities in the center of the domain. It is notable that the rainfall probabilities for states 1 and 2 are more spatially uniform than in the climatology, while the opposite is true of states 3 and 4 which are characterized by larger meridional gradients.

The most likely state sequences, calculated using the Viterbi algorithm (e.g., Rabiner 1989), are shown in Fig. 8. The sequence exhibits considerable variability on in-traseasonal, as well as interannual time scales. We now examine systematic seasonal and interannual variability.

b. Physical interpretation

The mean seasonality of state occurrence is shown in Fig. 9. The frequency of state 2 decreases from early February to mid-March, while the prevalence of state 1 maximizes in March, indicating the peak of the wet season at all stations. State 4 decreases in prevalence toward the end of the rainy season, while state 3 tends to become more frequent, indicating a contraction of the wet season northward. The ratios of minimum/maximum frequencies (from 24-yr averages of 10-day running means) plotted in Fig. 9 are 0.32, 0.25, 0.36, and 0.28 for states 1–4, respectively; that is, the state frequency varies by a factor of 3–4 within the season. This substantial seasonality represents an implicit nonhomogeneity in the HMM.

We now examine the meteorological characteristics associated with each rainfall state, by compositing anomalies of National Oceanic and Atmospheric Administration (NOAA)-interpolated outgoing longwave radiation (OLR) and NCEP–NCAR reanalysis winds over the days assigned to each state. Figure 10 shows anomaly composites of OLR and 850-hPa winds for the tropical South American–Atlantic sector, along with the state rainfall probabilities displayed as anomalies from the FMA-mean climatology. Note that the anomalies in winds and OLR are computed here as deviations from the mean seasonal cycle. States 1 and 2 are clearly associated with strongly contrasting large-scale anomalies in OLR and the cross-equatorial flow. Together these represent north–south displacements in the Atlantic ITCZ, as well as zonal contractions or expansions of the SAMS core region over Amazonia.

State 1 represents an anomalously southward-displaced ITCZ, together with an eastward expansion of

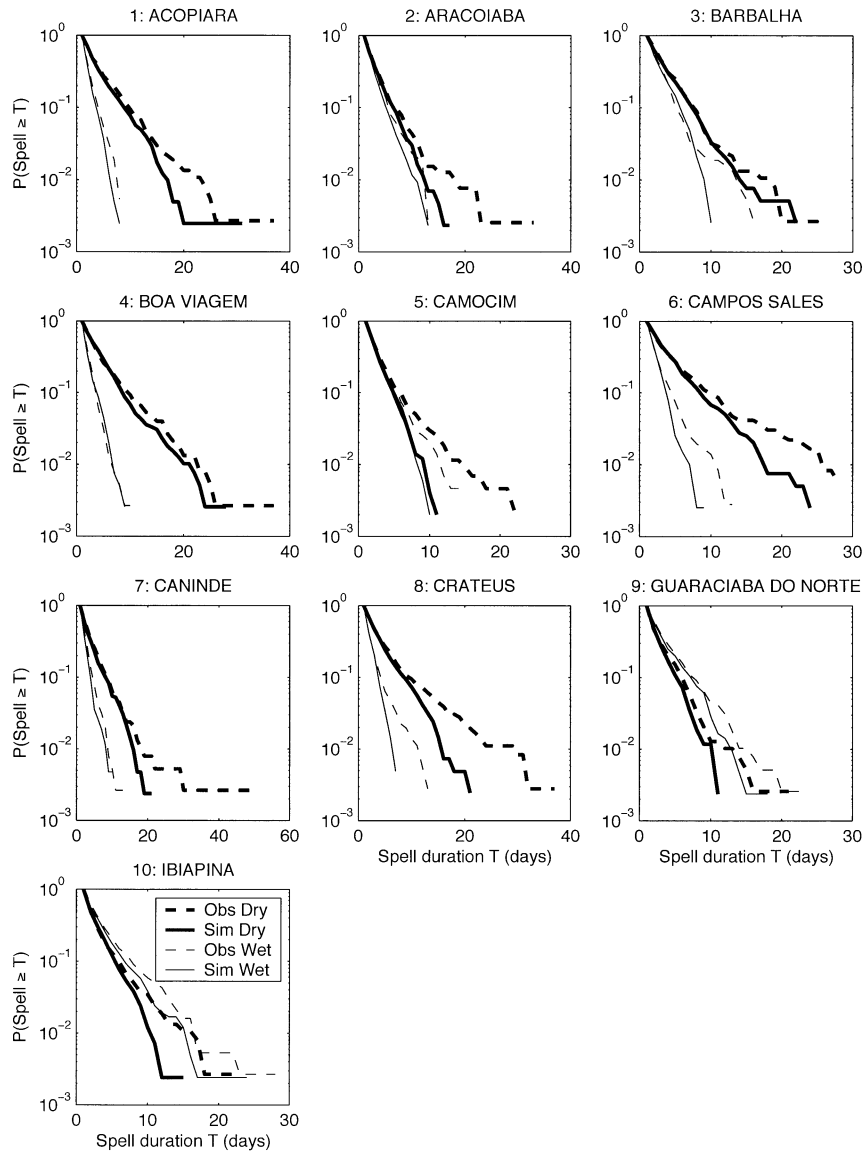


FIG. 6. Distribution of spell lengths for wet (thin) and dry (thick) spells at each station. Solid lines are computed from the four-state HMM simulations, with the observed data shown as dashed lines. Plotted is the probability of a spell exceeding a particular duration. A geometric distribution would plot as a straight line on these semilog plots.

TABLE 2. Transition probabilities for four-state HMM.

	To state			
	1	2	3	4
From state				
1	0.70	0.01	0.18	0.11
2	0.02	0.68	0.16	0.13
3	0.18	0.14	0.61	0.08
4	0.20	0.20	0.12	0.48

SAMS. Thus, the ITCZ and SAMS merge to a greater extent than in the FMA climatology, consistent with state 1 tending to correspond to the peak of the seasonal evolution (Fig. 9). In this configuration, Ceará comes entirely under the influence of the large-scale convection zones, and rainfall probabilities are high at all 10 stations. The contrasting situation characterizes state 2, in which SAMS and the ITCZ become more separated, reminiscent of the DJF climatology, with Ceará largely dry. It is worth emphasizing that states 1 and 2 are associated with sizeable atmospheric anomalies with respect to the mean seasonal cycle (i.e., those plotted in Fig. 10), so that these states are not merely aspects of

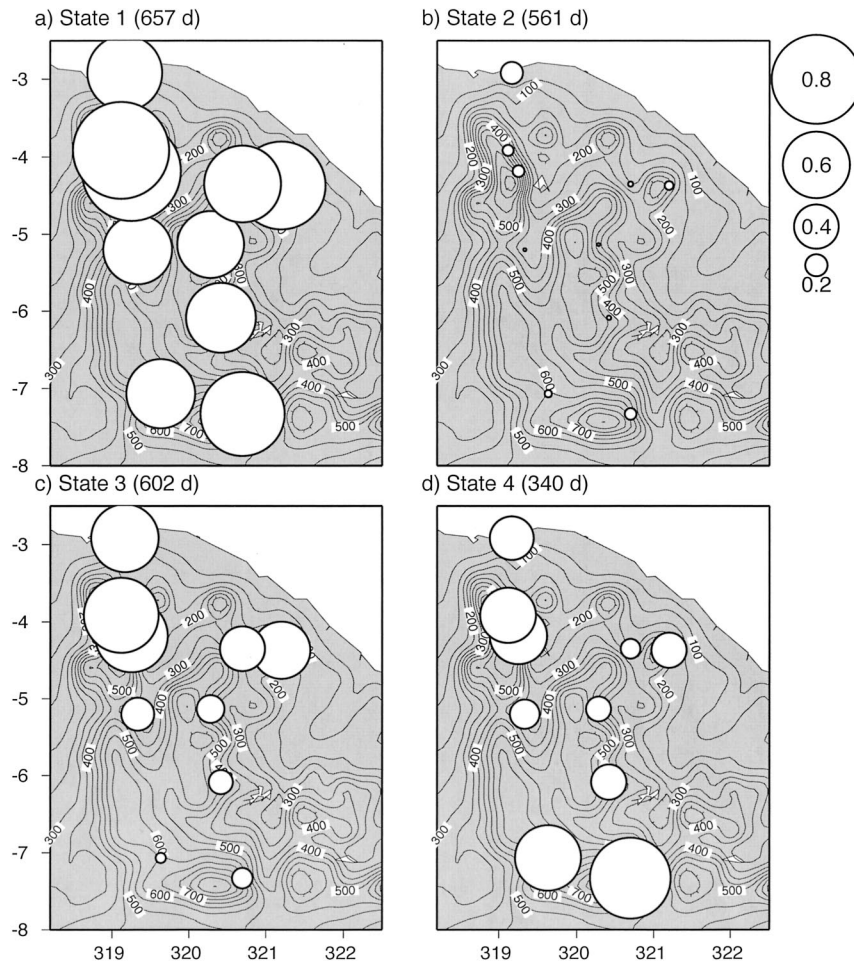


FIG. 7. Four-state HMM rainfall probabilities (circle radius) together with topographic contours. The number of days in each state is given in brackets.

the latter. Indeed the anomalous winds associated with states 1 and 2, plotted in Fig. 10, change little if anomalies are computed with respect to the FMA long-term mean.

States 3 and 4 are associated with smaller anomalies in rainfall probability, and this is reflected in much smaller anomalies in OLR and wind. The north–south gradients in the OLR anomalies are consistent with those in the gauge–rainfall HMM probabilities. However, the wind anomalies are not very coherent.

In order to identify any large-scale atmospheric teleconnections, Fig. 11 shows composites of 850- and 200-hPa wind anomalies of states 1 and 2 over a larger domain. Over NE Brazil, the direction of the wind anomalies reverses with height, typical of the first baroclinic mode of tropical atmospheric dynamics. Over the tropical Pacific, states 1 and 2 are characterized by anomalous Walker circulations, and these are consistent with OLR anomalies over the equatorial Pacific (not shown). Over the Atlantic, they suggest opposite polarities of the Atlantic meridional mode, with NAO-like wind anomalies and anomalies in the NE winds. These

are recognized to be the two mechanisms that influence NE Brazil rainfall on interannual time scales, and both leave their imprint on the daily rainfall states.

Similar pictures for states 3 and 4 (not shown) do not highlight any coherent circulation patterns that might be remotely forcing these rainfall states. However, there is a hint, especially for state 3, that Rossby wave activity propagating from the South Pacific may be a contributing factor (cf. Liebmann et al. 1999).

Interannual variability of state occurrence is plotted in Fig. 12 in terms of the number of days assigned into each state. Large interannual variations do occur, especially in states 1 and 2, which vary in opposition to each other consistent with their rainfall and meteorological characteristics seen in Figs. 10 and 11. The occurrence frequency of state 3 appears to show an upward trend from the 1980s onward, while state 4 shows little interannual variation. La Niña (El Niño) years tend to be associated with more (less) of state 1 compared to state 2.

To probe the interannual variability further, composites of seasonal-mean SST anomalies are shown in Fig.




FIG. 8. The estimated state sequence.

13 for FMA seasons in which the interannual state-frequency (Fig. 12) is in the top 15% of years (i.e., 4 or 5 yr). The shading depicts local significance at the 90% level. A clear ENSO SST anomaly signature is




FIG. 9. Mean seasonal variation of four HMM daily state occurrence, smoothed with a 10-day running mean prior to averaging.

present for years in which states 1 or 2 are highly prevalent, but statistical significance is only high for state 1 (La Niña). States 1 and 2 are also associated with Atlantic SST anomalies characteristic of the Atlantic meridional mode, but these are weak. States 3 and 4 are not associated with appreciable SST anomalies.

Finally, to check whether these four specific states might be sensitive to our choice of model (i.e., the HMM), we also ran the well-known *K*-means clustering algorithm (e.g., Jain and Dubes 1988) on the data with *K* = 4 clusters. Here the input was a set of 10-dimensional vectors with binary components corresponding to the daily rainfall occurrence measurements. The four *K*-means clusters were found to match closely those derived from the HMM. Specifically, the 10-dimensional means for each cluster from *K* means (real-valued numbers between 0 and 1) were compared to the conditional probability vectors for each state from the HMM, and the composite wind and OLR maps resulting from the most likely assignments of each day to each cluster (for both *K* means and HMMs) were also visually compared. The means and maps obtained from the two different methods were found to be quite similar (not shown). The fact that an alternative clustering meth-

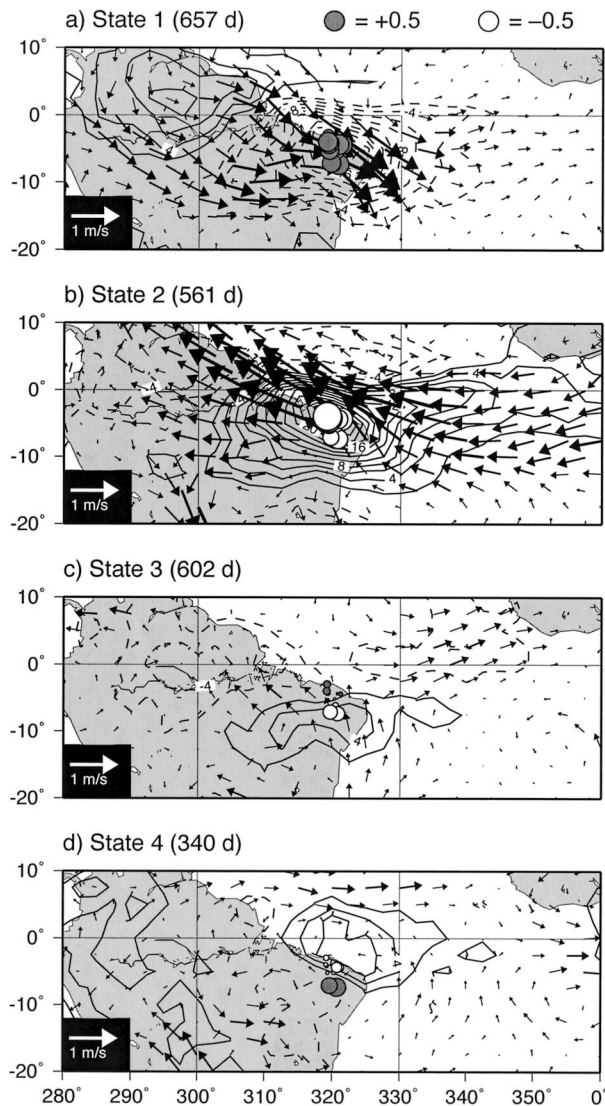


FIG. 10. HMM state anomaly composites of 850-hPa wind (vectors) and OLR (contours), together with rainfall anomaly-probabilities at each station (circles). The wind and OLR anomalies are deviations from the mean seasonal cycle; the latter was computed by 10-day low-pass filtering and then averaging all years together. The composites are defined on the days assigned to each state, with the number of days given in the caption. Contour interval: 2 W m^{-2} , with negative contours dashed.

odology such as K means, that uses no information about temporal ordering of the rainfall measurements, produces state descriptions that are qualitatively similar to those produced by the HMM, suggests that these states are an inherent property of the data and insensitive to the particular modeling methodology being used.

In summary, daily rainfall states 1 and 2 identified by the HMM are associated with well-known patterns of interannual variability in winds, OLR, and SST. These associations provide a basis for the downscaling of seasonal GCM simulations, and this is pursued in the following section.

5. A nonhomogeneous HMM downscaling prototype

The NHMM generalizes the homogeneous HMM in that the transition probabilities in Eq. (2) are allowed to vary with time. In particular, for downscaling applications the transition probabilities between states are allowed to vary as a function of external inputs. Hughes and Guttrop (1994a) introduced this model in the context of modeling rainfall occurrence. The NHMM used in this paper is based on this original work of Hughes and Guttrop, with some minor modifications.

In this section we illustrate the ability of an NHMM to downscale atmospheric GCM simulations over NE Brazil. It is found that introducing atmospheric input variables does not visibly change the appearance of the state composites, nor appreciably change the rainfall probabilities. Thus, a four-state model is chosen for consistency with the HMM in the previous section.

For demonstration purposes and for consistency with IRI's current seasonal-forecast scheme, we define the inputs to the NHMM from the GCM's simulated seasonal-mean rainfall anomaly. The daily values needed as inputs to the NHMM are derived by simply repeating the seasonal-mean input value for each day within the FMA season.

a. The nonhomogeneous hidden Markov model

Let \mathbf{X}_t be a D -dimensional column vector of predictors for day t , derived, for example, from a GCM. By $\mathbf{X}_{1:T}$ we will denote the sequence $\mathbf{X}_1, \dots, \mathbf{X}_T$. We now replace Eq. (2) in the homogeneous HMM with:

$$P(S_t | S_{1:t-1}, \mathbf{X}_{1:T}) = P(S_t | S_{t-1}, \mathbf{X}_t), \quad (3)$$

so that the hidden state on day t depends both on the predictor vector \mathbf{X}_t for day t and the value of the hidden rainfall state S_{t-1} on day $t-1$. Because \mathbf{X}_t can vary in time, this results in transition probabilities between states that can vary in time in response to changes in \mathbf{X}_t , that is, an inhomogeneous model. The role of \mathbf{X}_t is clearly visualized in the corresponding graphical model, shown in Fig. 14.

The hidden state transitions in Eq. (3) are modeled by a polytomous (or multinomial) logistic regression:

$$P(S_t = i | S_{t-1} = j, \mathbf{X}_t = \mathbf{x}) = \frac{\exp(\sigma_{ji} + \boldsymbol{\rho}'_i \mathbf{x})}{\sum_{k=1}^K \exp(\sigma_{jk} + \boldsymbol{\rho}'_k \mathbf{x})}. \quad (4)$$

For the specific case of S_1 , the first hidden state in the sequence,

$$P(S_1 = i | \mathbf{X}_1 = \mathbf{x}) = \frac{\exp(\lambda_i + \boldsymbol{\rho}'_i \mathbf{x})}{\sum_{k=1}^K \exp(\lambda_k + \boldsymbol{\rho}'_k \mathbf{x})}. \quad (5)$$

All σ s and λ s are real-valued parameters while the $\boldsymbol{\rho}$ s are D -dimensional real-valued parameter vectors, where

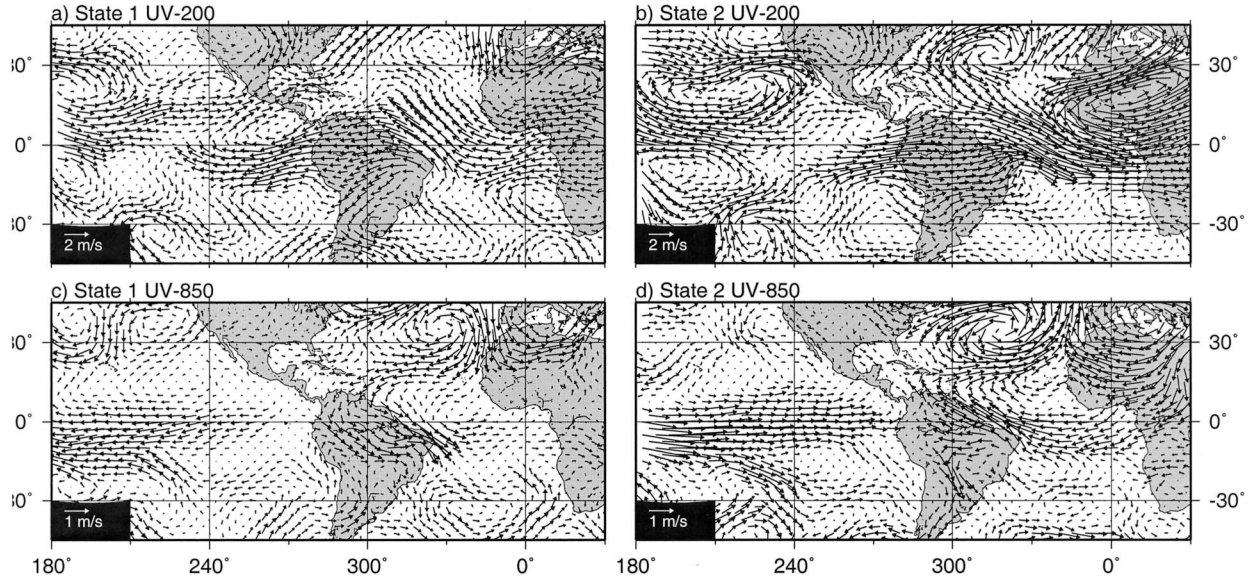


FIG. 11. HMM state anomaly composites of wind over a larger domain for states 1 and 2: (a)–(b) 200-hPa winds, (c)–(d) 850-hPa winds. Details as in Fig. 10.

the prime denotes the vector transpose. This parameterization can be shown to be equivalent to the one defined in Hughes and Guttrop (1994a) and in Hughes et al. (1999) in which the baseline transition matrix is multiplied by a function of the atmospheric predictors:

$$\begin{aligned}
 P(S_t = i | S_{t-1} = j, \mathbf{X}_t) \\
 &\propto P(S_t = i | S_{t-1} = j) P(\mathbf{X}_t | S_{t-1} = j, S_t = i) \\
 &= \gamma_{ji} \exp \left[-\frac{1}{2} (\mathbf{X}_t - \boldsymbol{\mu}_{ji})' \mathbf{V}^{-1} (\mathbf{X}_t - \boldsymbol{\mu}_{ji}) \right]
 \end{aligned}$$




FIG. 12. Interannual variability of HMM state-occurrence frequency. The letters E and L denote El Niño and La Niña years, respectively, defined as an anomaly of SST $\geq 0.6^\circ\text{C}$ over the Niño-3.4 region during FMA in the NOAA Extended Reconstructed SST dataset (Smith et al. 1996).

$$\propto \exp \left[\left(\ln \gamma_{ji} - \frac{1}{2} \boldsymbol{\mu}'_{ji} \mathbf{V}^{-1} \boldsymbol{\mu}_{ji} \right) + \boldsymbol{\mu}'_{ji} \mathbf{V}^{-1} \mathbf{X}_t \right]. \quad (6)$$

Here, $\boldsymbol{\mu}_{ji}$ is the mean of the atmospheric predictor-vector associated with transitions from state j at day $t - 1$ to state i at day t . If we make the simplification that $P(\mathbf{X}_t | S_t, S_{t-1}) = P(\mathbf{X}_t | S_t)$, then $\boldsymbol{\mu}_{ji} = \boldsymbol{\mu}_i$, and the parameterization in Eq. (6) becomes equivalent to the one in Eq. (4). This can be shown by setting $\lambda_i = \ln \gamma_{ji} - (1/2) \boldsymbol{\mu}'_i \mathbf{V}^{-1} \boldsymbol{\mu}_i$ and $\boldsymbol{\rho}_i = \mathbf{V}^{-1} \boldsymbol{\mu}_i$. The parameters λ_i , σ_{ji} , and $\boldsymbol{\rho}_i$ are set to zero to guarantee the identifiability of the transition parameters. Note that the homogeneous transition matrix [(Eq. (2))] can be viewed as a special case where $\boldsymbol{\rho}_i = \mathbf{0}$ for all $i = 1, \dots, K$.

An example of the transition probabilities obtained with this parameterization for a four-state model with a univariate normalized input is shown in Fig. 15. This case corresponds approximately to the downscaling example to be presented later, and demonstrates how the transition probabilities from state 2 to states 1–4 are modulated by the value of a univariate X_t . Each curve can be interpreted as part of a logistic “S-shaped” curve, with the value of $X = 0$ corresponding to the homogeneous HMM, and the central portion of the plot being the most relevant.

Given a fixed number of hidden states K , we learn the parameters Θ of the NHMM by searching for parameters that best fit the observed data. To do this, we employ the commonly used maximum likelihood principle. Specifically, we search for Θ that maximizes the conditional probability of the observed data as a function of Θ —this conditional probability function is referred to as the likelihood:

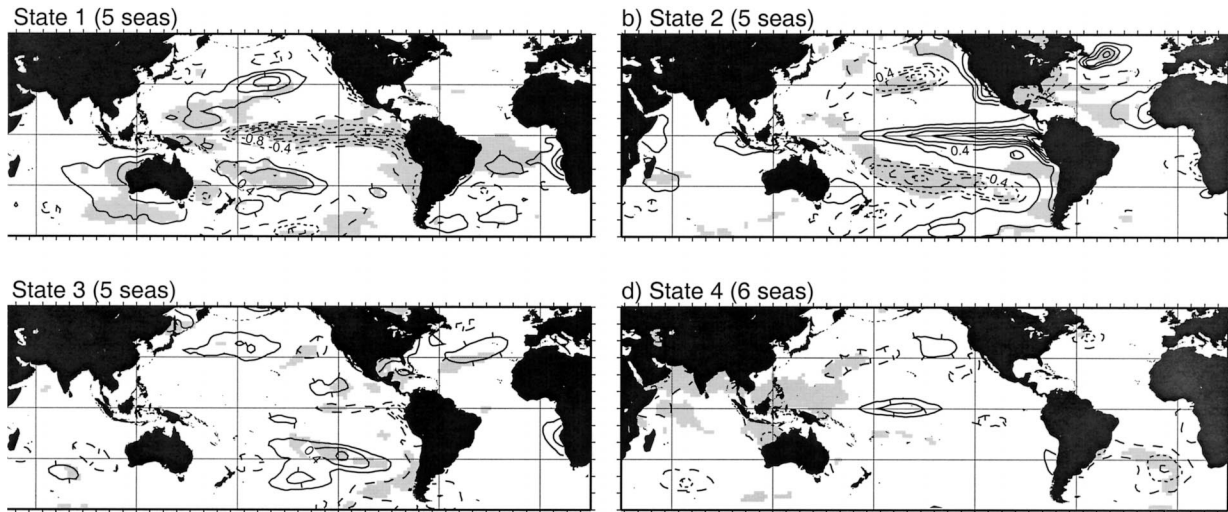


FIG. 13. Composites of anomalous SST for years (Feb–Apr) when HMM states are most prevalent, defined as years in the upper 15% of the interannual distribution of state frequency. The number of years in each composite is given in brackets. Shading represents 90% statistical significance according to a two-sided Student *t* test. Negative contours are dashed, and the zero contour is omitted. Contour interval: 0.2°C.

$$\begin{aligned}
 L(\Theta) &= P(\mathbf{R}_{1:T} | \mathbf{X}_{1:T}, \Theta) \\
 &= \sum_{S_{1:T}} P(S_1 | \mathbf{X}_1) \prod_{t=2}^T P(S_t | S_{t-1}, \mathbf{X}_t) \prod_{t=1}^T P(\mathbf{R}_t | S_t).
 \end{aligned}
 \tag{7}$$

The set of parameters Θ that maximize $L(\Theta)$ can be obtained using the well-known Baum–Welch algorithm (e.g., Rabiner 1989), a variation of an iterative expectation–maximization (EM) algorithm (Dempster et al. 1977) for obtaining maximum likelihood parameter estimates for models with hidden variables and/or missing data. In all the HMM and NHMM results in this paper we start the EM algorithm 10 times from different random starting positions in parameter space, run EM to convergence for each starting point, and choose as our solution the parameter vector Θ that has the maximum likelihood over all 10 runs (this helps avoid poor local maxima that EM can sometimes converge to). Full details on the specific EM procedure used in this paper for NHMM parameter estimation can be found in Robertson et al. (2003).




FIG. 14. Graphical model representation of a nonhomogeneous hidden Markov model.

b. Predictor selection: Canonical correlation analysis

The choice of input (“predictor”) variables is non-trivial. Here, canonical correlation analysis (CCA) between historical seasonally averaged rainfall occurrence probabilities at each of the 10 stations, and the GCM’s simulated ensemble-mean seasonal-mean rainfall amount at grid points within the region (21°S–7°N, 80°W–0°W) was used to define the input variables. This domain was chosen so as to include the large-scale South American monsoon over the continent, and the




FIG. 15. Transition probabilities into each state from state 2 (“dry”), as a function of a univariate input variable for a four-state NHMM. The input variable is defined by the ensemble-average GCM simulation of seasonal-average precipitation, averaged spatially over the region of NE Brazil (centered and normalized by its standard deviation). The values of X_t used to train the model are plotted along the abscissa.




FIG. 16. Canonical correlation analysis of GCM gridded seasonal-mean precipitation amount (contours) and Ceará station daily rainfall probability (circles) for concurrent seasons. The leading CCA mode is shown, computed from each half of the dataset in turn: (a) 1975–90, (b) 1991–2002. The amplitudes correspond to a one standard deviation anomaly, with contours every 0.2 mm day⁻¹. All station-rainfall anomalies are positive. Details of the CCA: the optimal PCA truncations are (a) 2 GCM PCs, 3 station PCs, and (b) 5 GCM PCs and 1 station PC, for the two halves of the dataset, respectively. The associated time series correlations are (a) 0.95 and (b) 0.79.

intertropical convergence zone over the Atlantic, whose mean position is close to the equator in February–April. The CCA yields a calibration of the GCM’s simulation of rainfall amount with respect to observed rainfall occurrence probabilities. It is used as a preprocessing step, independent of the NHMM, to define the most appropriate GCM predictors. The 24-member ensemble mean of the GCM’s normalized precipitation field was used, driven by historical estimates of SSTs. The analysis proceeds by first expanding each field into principal components (PCs), and then performing the CCA in the reduced subspace of the two sets of PCs. The GCM precipitation was normalized by its local interannual standard deviation prior to the CCA. The optimal truncations for each PC subspace, as well as the optimal number of CCA modes were determined by 1) computing the CCA for a given choice of truncation and 2) summing the out-of-sample correlations exceeding 0.3 between (a) the GCM predictor(s) so identified and (b) the station rainfall (Tippett et al. 2003).

Since CCA is susceptible to overfitting, a simple but severe cross-validation procedure was employed, consisting of dividing the dataset into two contiguous 12-yr training and validation parts. The CCA modes and NHMM model are derived from the training part of the dataset, using the ensemble-mean of the 24 GCM simulations together with the observed data. The resulting NHMM is then used to simulate rainfall occurrences for the validation part of the dataset, using the input time




FIG. 17. Interannual variability of candidate GCM predictor variables, together with the observed daily rainfall occurrence, averaged over the 10 stations (circles), with the error bars representing their standard deviation. All curves have been normalized by their standard deviation. Key: Leading cross-validated CCA mode (thin solid, $\rho = 0.73$); leading CCA mode without cross-validation (dashed, $\rho = 0.86$); GCM NE Brazil area-average precipitation (dotted, $\rho = 0.66$); NCEP–NCAR reanalysis area-average precipitation (dash-dotted, $\rho = 0.62$). The reported correlation values are with the observed.

series derived by projecting the GCM ensemble-mean simulation from the validation part onto the CCA modes derived from the training part. The procedure is then repeated switching the training and validation parts of the dataset. Thus, two models are fitted and used to simulate the complementary part of the series.

In both cases the CCA yields one statistically significant mode, characterized by correlations between the canonical variate time series of 0.95 and 0.79, for the first and second half of the dataset respectively. Figure 16 shows the structure of the CCA mode for each part of the dataset, in terms of homogeneous covariance maps formed by regressing each field with its respective (normalized) canonical variate time series. Both halves of the dataset are characterized by a correspondence between the GCM’s simulated broad-scale precipitation on the southern flank of the Atlantic ITCZ and the observed station rainfall occurrence frequency over Ceará, with the relationship being stronger from 1975–90 than for 1991–2002.

The resulting (cross-validated) GCM predictor time series is plotted in Fig. 17, together with the station-averaged observed rainfall probability and its standard error; their linear correlation is 0.73. An alternative to using the CCA calibration, is to use a spatial average of the GCM’s seasonal-mean rainfall amount simulations over the region of NE Brazil, obviating the need for cross-validation. Figure 17 also shows these (normalized) spatial averages of NE Brazil rainfall (approximately 8°S–0°N, 44°–35°W) from both the GCM’s ensemble mean ($\rho = 0.66$) as well as the NCEP–NCAR reanalysis ($\rho = 0.62$). It is encouraging that the GCM’s




FIG. 18. Interannual variability of NHMM-simulated rainfall occurrence versus the observed (dashed) averaged over all 10 stations. Plotted is the median of 24 NHMM simulations (solid). The number of rain days per season were summed across the the 10 stations, and then divided by 10. The error bars indicate the entire range of the 24 simulations, with the interquartile range given by the inner ticks.

24-member ensemble mean performs about as well as the reanalysis. The gain achieved by calibrating the GCM with CCA is relatively small in this case. A predictor time series derived using CCA but without cross-validation is also plotted to illustrate the serious problem of overfitting that results ($\rho = 0.86$).

c. NHMM simulations

The generalized EM algorithm was then used to learn the parameters of the four-state NHMM with 10 binary (rainfall) outputs and 1 real-valued (GCM) input, for each half of the dataset separately.

Twenty-four simulations of daily rainfall occurrence were then made in each case, using the GCM ensemble mean input repeated 24 times; use of the individual GCM ensemble members was found to degrade the simulations. Figure 18 shows the median number of rain days per season resulting from the 24 simulations, using the number of rain days averaged over all 10 stations. The linear correlation with the observed value is $\rho = 0.67$, which is similar to the performance of the seasonal-mean input variable in Fig. 17. Thus, the NHMM simulations recover the predictive value of the input variable in this seasonal and station average quantity. Also plotted are the quartiles and extremes of the simulated distribution. The observed curve is inside the simulated interquartile range about 50% of years, indicating that the simulated distribution has a consistent variance. The 24-member simulated distribution also brackets the observed one during all years (but one), and is thus consistent with the NHMM. In other words, the NHMM is capable of generating the station-averaged observed rainfall occurrence time series under

strict cross-validation, except during the 1985 La Niña event.

The interannual performance of the median simulated daily rainfall sequence is plotted for each station individually in Fig. 19. Stations 8 and 10 have correlations with the observed that equal or exceed 0.7, while stations 5 and 6 have near-zero correlation. Stations 7–10 in the north-central part of the domain exhibit the more successful interannual simulations ($\rho \geq 0.50$).

The distribution of dry spells is a parameter that is of particular importance to agriculture. Figure 20 shows the interannual variability of NHMM-simulated dry-spell frequency at each station, in terms of the median of the simulated distribution for each year. Here we define dry spells to be runs of dry days of at least 10 days, with no more than one intervening wet day, and such that the identified dry spells do not overlap. Again, rainfall amount is not considered. The GCM-NHMM is able to simulate interannual variability in dry-spell frequency reasonably well at three stations, with anomaly correlations (between observed and median simulated) equal to or exceeding 0.5. Other studies suggest that rainfall absence is more predictable than occurrence, for example, Wilby (2001).

6. Summary and conclusions

We have used a hidden Markov model (HMM) to analyze daily rainfall occurrence at 10 gauge stations over the state of Ceará in NE Brazil during the rainy season (FMA) 1975–2002. A four-state model is chosen from inspection of the cross-validated log-likelihood of the rainfall data given the model and the Bayes Information Criterion (Table 1), as well as from subjective considerations. Unlike the BIC, the log-likelihood does not reach a peak at $K = 4$ using the leave-six-year-out cross-validation. It may be preferable to omit more years, but the dataset of only 24 complete FMA seasons is a limiting factor.

The HMM is used to estimate the hidden state sequence underlying the observed data, from which seasonal and interannual variability is analyzed. Accompanying meteorological conditions are examined through composites of NCEP–NCAR reanalysis data and NOAA interpolated OLR.

Two of the states are found to correspond to wet or dry conditions at all stations respectively, with similar relative frequencies (Fig. 7). However, the wet state (state 1) tends to be more prevalent during March, with the dry state (state 2) being more prevalent at the beginning of the FMA season (Figs. 8 and 9). Thus, on average, states 1 and 2 describe the seasonal cycle of the “monsoon” over NE Brazil, and this is brought out in the composites of anomalous OLR and low-level winds (Fig. 10). State 1 represents an anomalously southward-displaced ITCZ and an eastward expansion of the SAMS; both convergence zones merge to a greater extent than in the FMA climatology, and Ceará comes

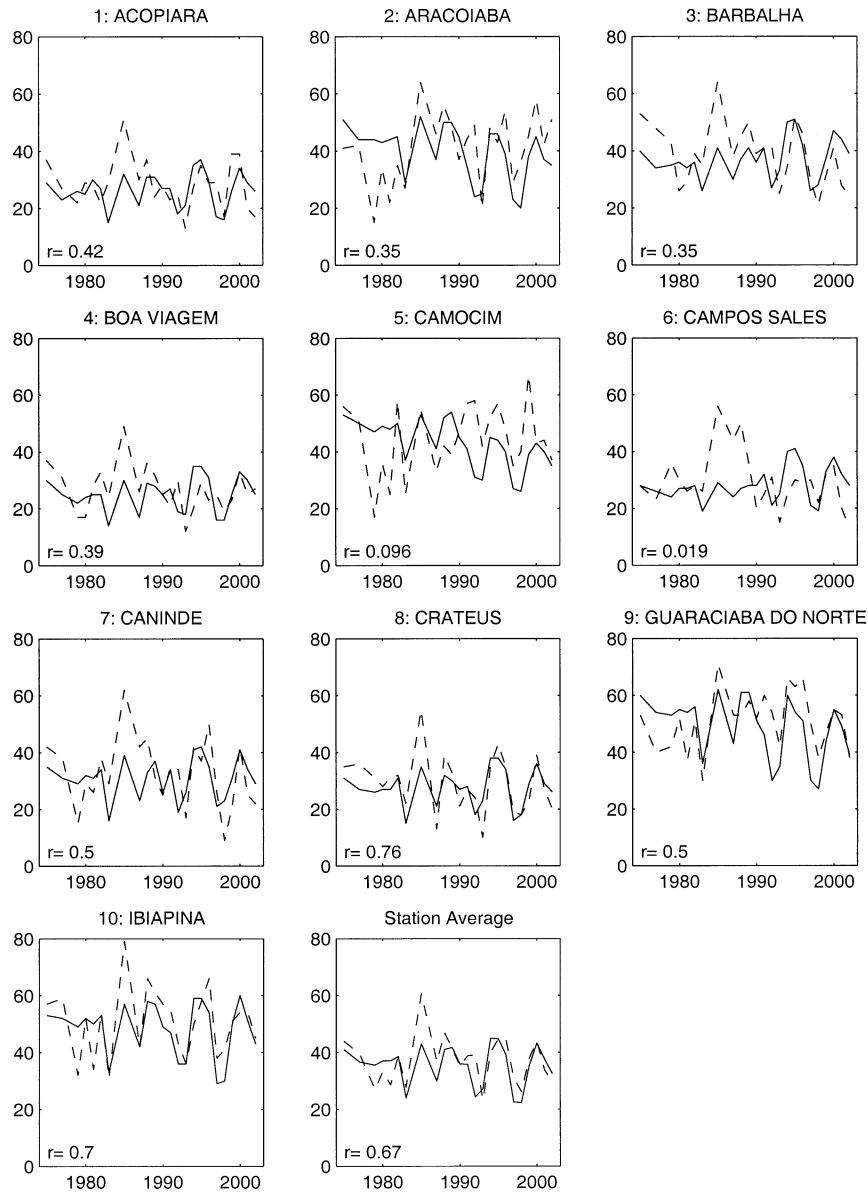


FIG. 19. Interannual variability of NHMM-simulated rainfall occurrence frequency at each station. The median of the 24 simulations is plotted for each year (solid) together with the observed (dashed). The number of rain days per season is plotted on the ordinate.

entirely under their influence. States 3 and 4 are characterized by meridional gradients of rainfall probability that have increased prevalence late in the season as the ITCZ retreats northward. However, the meteorological associations are relatively weak (cf. Fig. 10).

The state-based description of the monsoon over NE Brazil deserves some comment. Daily rainfall occurrence at a single station is binary, and thus it is natural to characterize it by a discrete Markov process. This is the basis of “weather generator” models of rainfall occurrence. The statistics in Table 1 demonstrate that the HMM outperforms a stateless model consisting of independent Markov chains fit to each station. The tem-

poral evolution of the monsoon may actually be better described in terms of a discontinuous weather-state process, than by a continuous one. The state-based model provides a probabilistic description of the onset and end of the rainy season (Fig. 8), together with an average seasonal evolution (Fig. 9). Further work is needed to address this issue, using rain gauge networks covering larger geographical areas that enable the spatial structures of the weather states to be better defined.

Despite the lack of any built-in rainfall persistence in the HMM, the distribution of wet- and dry-spell lengths is generally reproduced surprisingly accurately. Nonetheless, a tendency to underestimate the spell

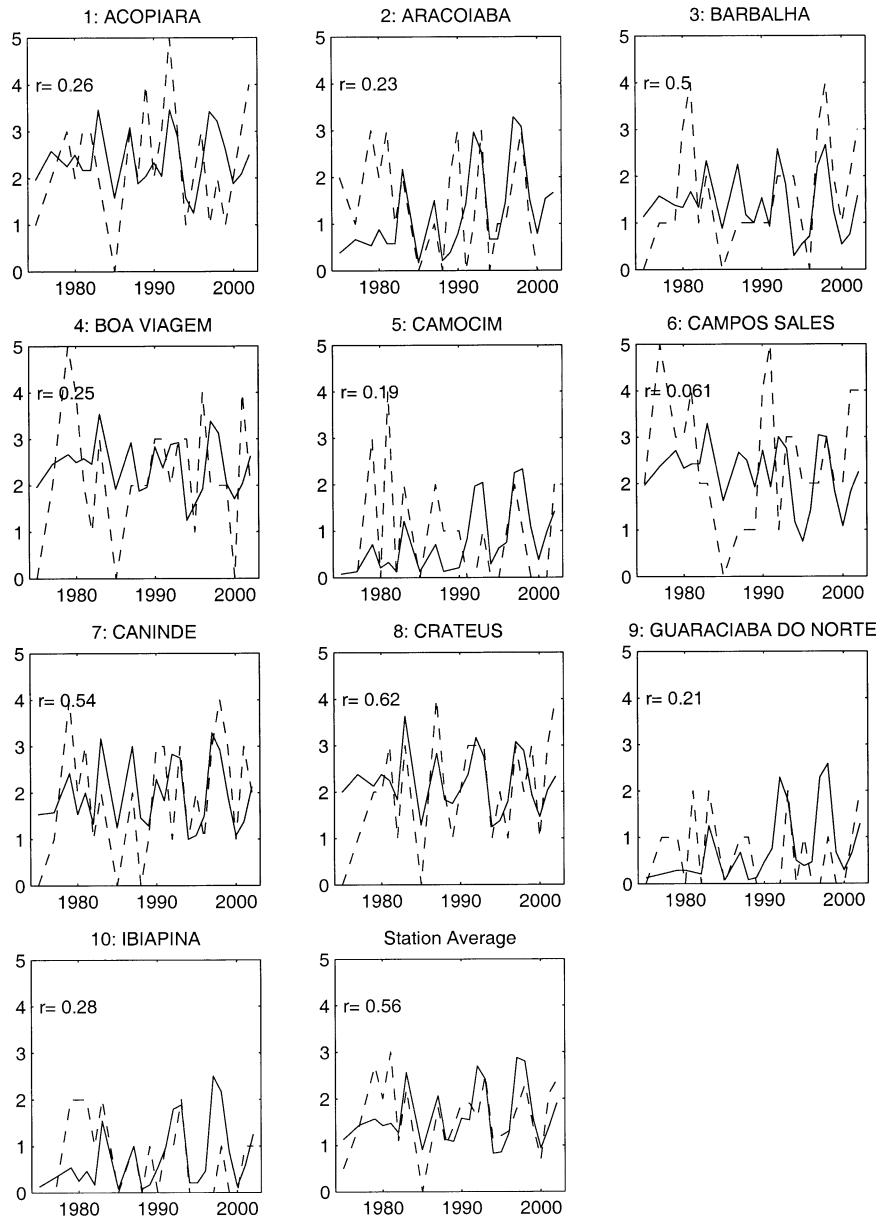


FIG. 20. Interannual variability of NHMM-simulated 10-day dry-spell frequency at each station. The median of the 24 simulations is plotted for each year (solid) together with the observed (dashed). The number of dry spells per season is plotted on the ordinate.

lengths at certain stations is clear. An autoregressive HMM (e.g., Juang and Rabiner 1985) may provide a solution to this defect, by explicitly including arrows between the outputs in Figs. 3 and 14.

The atmospheric composites of states 1 and 2 (Fig. 11) exhibit some well-known characteristics of intra-seasonal (NAO) and interannual (ENSO) teleconnections. These teleconnection patterns influence the NE trade winds in the tropical Atlantic and the position of the ITCZ, with similar characteristics to the mean seasonal evolution. It is not surprising that we find indications of variability on different time scales in the rain-

fall states, because the atmospheric spatial structures are similar. From a regional perspective, several mechanisms can influence the probability of occurrence of the states, and thus the rainfall occurrence. Thus, the hidden states of observed rainfall occurrence appear to correspond to intrinsic weather states, which allow a natural description of rainfall variability across many different time scales.

There are large interannual variations in state frequency, particularly of states 1 and 2, together with some decadal-scale changes (also in state 3) (Fig. 12). The SST anomalies during the years of large anomalies

in state frequency (Fig. 13) bear the hallmarks of ENSO. The La Niña relationship is more statistically significant than for El Niño. This is consistent with the findings of Giannini et al. (2004) who show evidence that the relationship between NE Brazil rainfall and La Niña has been stronger than for El Niño during recent decades, due to preconditioning by tropical Atlantic SST anomalies. From our short 24-yr dataset, we find the relationship between the latter and state frequency to be relatively weak compared to that identified in other studies (e.g., Moura and Shukla 1981).

The second goal of the paper was to utilize GCM simulations of the large-scale circulation to make downscaled simulations of station-scale rainfall. Based on historical SST forcing, the 24-member ECHAM 4.5 GCM ensemble-mean precipitation is found to have considerable cross-validated skill at reproducing interannual variations in 10-station average rainfall occurrence. We find a cross-validated linear correlation of 0.66 with the GCM's precipitation averaged over the NE Brazil region (Fig. 17). This value rises to 0.73 if a canonical correlation analysis is used to find the GCM's pattern of precipitation that is best correlated (under cross validation) with Ceará seasonally averaged rainfall occurrence frequency (Fig. 16).

The GCM's ensemble-mean simulation of seasonal-mean precipitation is used as a univariate input into the NHMM, from which multiple daily sequences of rainfall are then generated at the 10 stations. Validating seasonal-mean rainfall frequency simulated by the NHMM integrated across the 10 stations yields an anomaly correlation skill of 0.67. Thus, the NHMM conveys most of the GCM's large-scale interannual simulation skill to the daily rainfall sequences. At the individual station level, the maximum interannual correlation found between simulation and observed is 0.76, using the median of a 24-member NHMM ensemble of simulations (Fig. 19). Attempting to use the individual GCM ensemble members to make the NHMM ensemble produced an inferior result to that obtained by repeating the GCM's ensemble mean 24 times. Ten-day dry-spell frequency is also hindcast fairly well at certain stations by the GCM-NHMM, with a maximum anomaly correlation skill of 0.62 at station 8 (Fig. 20, Crateus).

In terms of downscaling, we have shown that the HMM is able to quite accurately capture the characteristics of daily rainfall occurrence in terms of spell lengths (Fig. 6) and (to some extent) spatial interstation correlations (Table 1). The spatial model used in this study assumes that the rainfall stations are conditionally independent of each other, given the rainfall state. More sophisticated spatial models can be used with HMMs, such as the autologistic model (Hughes and Guttorp 1994b), and the Chow-Liu tree model that is the subject of current work (Kirshner et al. 2004).

It is also able to convey an interannual GCM simulations to the local scale. The method thus shows promise as a technique for generating downscaled daily rain-

fall-sequence scenarios for input into crop models that require such inputs, and may be competitive with sophisticated weather-generator models designed for this purpose (Wilks 2002). Seasonally varying predictors could be used to simulate the average seasonality of rainfall, and to investigate the downscalability of the onset of the rainy season. Daily rainfall amounts can be incorporated into the NHMM in a consistent manner (Charles et al. 1999; Bellone et al. 2000). A topic of future research concerns the seasonal predictability of daily rainfall intensity versus occurrence, which will involve extension of the present work to use predictions of SST or antecedent-lagged predictors.

Acknowledgments. We are grateful to two anonymous reviewers for their constructive comments. This work was supported by Department of Energy Grant DE-FG02-02ER63413, by the National Science Foundation under Grant SCI-0225642 for the OptIPuter project, and by NOAA through a block grant to the International Research Institute for Climate Prediction (IRI). NCEP-NCAR Reanalysis data were provided by the NOAA-CIRES Climate Diagnostics Center, Boulder, Colorado, from their Web site available online at <http://www.cdc.noaa.gov>. The daily rainfall data were provided by the Brazilian State of Ceará's Foundation for Meteorology and Water Resources (FUNCEME; <http://www.funceme.br>).

APPENDIX

BIC Scores

The BIC score for an HMM or NHMM model with K states is defined as

$$\text{BIC}_K = 2L(\Theta_K^*) - p \log T,$$

where Θ_K^* is the estimated maximum likelihood parameter vector as found by EM on the training data for a model with K states, $L(\Theta_K^*)$ is the likelihood of the model evaluated at Θ_K^* as in Eq. (7), p is the number of parameters in the K -state model (linear in K), and T is the total number of days of observed data used to train the model. The second term in the BIC expression $-p \log T$ "penalizes" more complex models. BIC can be viewed as a practical approximation to the more ideal (but intractable to compute) Bayes factor for model selection (e.g., see Kass and Raftery 1995). Although not fully justified theoretically for model selection in the context of HMMs and NHMMs [e.g., see Titterton (1990) and Hughes et al. (1999) for further comments], the BIC score can nonetheless provide a useful indication of which models are supported by the data (e.g., Hughes and Guttorp 1994a).

To obtain normalized BIC scores (as in Table 1) that are on roughly the same scale as the normalized log-likelihoods, we replace BIC_K above with $-\text{BIC}_K/2N$, where N is the total number of binary predictions made (here $N = 24 \times 90 \times 10$).

REFERENCES

- Bauer, F., 1951: Extended range weather forecasting. *Compendium of Meteorology*, T. F. Malone, Ed., American Meteorological Society, 814–833.
- Bellone, E., J. P. Hughes, and P. Guttorp, 2000: A hidden Markov model for downscaling synoptic atmospheric patterns to precipitation amounts. *Climate Res.*, **15**, 1–12.
- Charles, S. P., B. C. Bates, and J. P. Hughes, 1999: A spatiotemporal model for downscaling precipitation occurrence and amounts. *J. Geophys. Res.*, **104D**, 31 657–31 669.
- Charney, J. G., and J. G. DeVore, 1979: Multiple flow equilibria in the atmosphere and blocking. *J. Atmos. Sci.*, **36**, 1205–1216.
- Chiang, J. C. H., Y. Kushnir, and A. Giannini, 2002: Deconstructing Atlantic Intertropical Convergence Zone variability: Influence of the local cross-equatorial sea surface temperature gradient and remote forcing from the eastern equatorial Pacific. *J. Geophys. Res.*, **107D**, 4004, doi:10.1029/2000JD000307.
- Dempster, A. P., N. M. Laird, and D. B. Rubin, 1977: Maximum likelihood from incomplete data via EM algorithm. *J. Roy. Stat. Soc.*, **39B**, 1–38.
- Feldstein, S. B., 2000: The time scale, power spectra, and climate noise properties of teleconnection patterns. *J. Climate*, **13**, 4430–4440.
- Giannini, A., R. Saravanan, and P. Chang, 2004: The preconditioning role of Tropical Atlantic variability in the prediction of Nordeste rainfall during ENSO events. *Climate Dyn.*, **22**, 839–855.
- Goddard, L., A. G. Barnston, and S. J. Mason, 2003: Evaluation of the IRI's "net assessment" seasonal climate forecasts: 1997–2001. *Bull. Amer. Meteor. Soc.*, **84**, 1761–1781.
- Gong, G., D. Entekhabi, and J. Cohen, 2003: Modeled Northern Hemisphere winter climate response to realistic Siberian snow anomalies. *J. Climate*, **16**, 3917–3931.
- Hastenrath, S., and L. Heller, 1977: Dynamics of climatic hazards in northeast Brazil. *Quart. J. Roy. Meteor. Soc.*, **103**, 77–92.
- Hughes, J. P., and P. Guttorp, 1994a: A class of stochastic models for relating synoptic atmospheric patterns to regional hydrologic phenomena. *Water Resour. Res.*, **30**, 1535–1546.
- , and —, 1994b: Incorporating spatial dependence and atmospheric data in a model of precipitation. *J. Appl. Meteor.*, **33**, 1503–1515.
- , —, and S. P. Charles, 1999: A non-homogeneous hidden Markov model for precipitation occurrence. *J. Roy. Stat. Soc.*, **48C**, 15–30.
- Hurrell, J. W., Y. Kushnir, G. Ottersen, and M. Visbeck, Eds., 2002: The North Atlantic Oscillation: Climatic Significance and Environmental Impact. *Geophys. Monogr.*, No. 134, Amer. Geophys. Union, 279 pp.
- Jain, A. K., and R. C. Dubes, 1988: *Algorithms for Clustering Data*. Prentice Hall, 320 pp.
- Juang, B. H., and L. R. Rabiner, 1985: Mixture autoregressive hidden Markov models for speech signals. *IEEE Trans. Acoust. Speech Signal Process.*, **33**, 1404–1413.
- Kalnay, E., and Coauthors, 1996: The NCEP/NCAR 40-Year Reanalysis Project. *Bull. Amer. Meteor. Soc.*, **77**, 437–471.
- Kass, R. E., and A. E. Raftery, 1995: Bayes factors. *J. Amer. Stat. Assoc.*, **90**, 773–795.
- Kirshner, S., A. W. Robertson, and P. J. Smyth, 2004: Conditional Chow-Liu tree structures for modeling discrete-valued vector time series. Tech. Rep. ICS-TR 04-04, Information and Computer Science, University of California, Irvine, CA, 26 pp.
- Liebmann, B., G. N. Kiladis, J. A. Marengo, T. Ambrizzi, and J. D. Glick, 1999: Submonthly convective variability over South America and the South Atlantic convergence zone. *J. Climate*, **12**, 1877–1891.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.
- MacDonald, I. L., and W. Zucchini, 1997: *Hidden Markov and Other Models for Discrete-valued Time Series*. Meteor. Monogr., Statistics and Applied Probability, No. 70, Chapman & Hall, 236 pp.
- Moura, A. D., and J. Shukla, 1981: On the dynamics of droughts in northeast Brazil: Observations, theory and numerical experiments with a general circulation model. *J. Atmos. Sci.*, **38**, 2653–2657.
- Nobre, P., and J. Shukla, 1996: Variations of sea surface temperature, wind stress, and rainfall over the tropical Atlantic and South America. *J. Climate*, **9**, 2464–2479.
- Rabiner, L. R., 1989: A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, **77**, 257–286.
- Robertson, A. W., S. Kirshner, and P. J. Smyth, 2003: Hidden Markov models for modeling daily rainfall occurrence over Brazil. Tech. Rep. ICS-TR 03-27, Information and Computer Science, University of California, Irvine, CA, 38 pp.
- Roeckner, E., and Coauthors, 1996: The atmospheric general circulation model ECHAM4: Model description and simulation of present-day climate. Max-Planck-Institut für Meteorologie Rep. 23, 90 pp.
- Smith, T. M., R. W. Reynolds, R. E. Livezey, and D. C. Stokes, 1996: Reconstruction of historical sea surface temperatures using empirical orthogonal functions. *J. Climate*, **9**, 1403–1420.
- Sun, L., D. F. Moncunill, H. Li, A. D. Moura, and F. A. D. S. Filho, 2005: Climate downscaling over Nordeste Brazil using NCEP RSM97. *J. Climate*, in press.
- Thompson, D. W. J., M. P. Baldwin, and J. M. Wallace, 2002: Stratospheric connection to Northern Hemisphere wintertime weather: Implications for prediction. *J. Climate*, **15**, 1421–1428.
- Tippett, M. K., M. Barlow, and B. Lyon, 2003: Statistical correction of central southwest Asia winter precipitation simulations. *Int. J. Climatol.*, **23**, 1421–1433.
- Titterton, D. M., 1990: Some recent research in the analysis of mixture distributions. *Statistics*, **21**, 619–641.
- Wilby, R. L., 2001: Downscaling summer rainfall in the UK from North Atlantic Ocean temperatures. *Hydrol. Earth Syst. Sci.*, **5**, 245–257.
- Wilks, D. S., 1998: Multisite generalization of a daily stochastic precipitation generation model. *J. Hydrol.*, **210**, 178–191.
- , 2002: Realizations of daily weather in forecast seasonal climate. *J. Hydrometeorol.*, **3**, 195–207.
- , and R. L. Wilby, 1999: The weather generation game: A review of stochastic weather models. *Prog. Phys. Geogr.*, **23**, 329–357.