

Stochastic Characterization of Regional Circulation Patterns for Climate Model Diagnosis and Estimation of Local Precipitation

EDUARDO ZORITA

Max-Planck-Institut für Meteorologie, Hamburg, Germany

JAMES P. HUGHES

Department of Statistics, University of Washington, Seattle, Washington

DENNIS P. LETTEMAIER

Department of Civil Engineering, University of Washington, Seattle, Washington

HANS VON STORCH

Max-Planck-Institut für Meteorologie, Hamburg, Germany

(Manuscript received 10 August 1993, in final form 12 September 1994)

ABSTRACT

Two statistical approaches for linking large-scale atmospheric circulation patterns and daily local rainfall are applied to GCM (general circulation model) climate simulations. The ultimate objective is to simulate local precipitation associated with altered climate regimes. Two regions, one in the Pacific–American sector (western region) and one in the American–Mid-Atlantic sector (eastern region), are explored.

The first method is based on Classification and Regression Trees (CART) analysis. The CART method classifies observed daily sea level pressure (SLP) fields into weather types that are most strongly associated with the presence/absence of rainfall at selected index stations. After applying this method to historical SLP observations, precipitation simulations associated with GCM SLP output were validated in terms of probability of occurrence and survival time of the weather states identified by the CART analysis. Daily rainfall time series were then generated from weather classes derived by application of CART to both daily SLP fields derived from historical observation and from GCM simulations. While the mean rainfall and probability distributions were rather well replicated, the precipitation generator based on this version of the CART technique had two important deficiencies: the generated dry periods were too short, on average, and the identification of weather states may be not invariant under coordinate rotations.

The second rainfall generator is based on the analog method and uses information about the evolution of the SLP field from several previous days. It considers a pool of past observations for the circulation patterns closest to the target circulation. It is similar to the CART method and in certain aspects it performs better, although some downward bias in the simulated rainfall persistence was still present. Applying both methods to the output of a $2 \times \text{CO}_2$ GCM simulation produced only small changes in simulated precipitation, which is due to the small sensitivity of this variable to greenhouse forcing. The selection characteristics of the analogs are similar for observations, a control run, and a $2 \times \text{CO}_2$ run, indicating that analogs for possible altered climates can be found in the historical record.

1. Introduction

One of the largest uncertainties in climate simulations produced by the present generation of general circulation models (GCMs) is the hydrological cycle at the land surface (Chahine 1992). The physical processes that contribute to the atmospheric and surface hydrologic cycle, such as cloud formation, precipita-

tion, infiltration, evaporation, and runoff production, evolve over a much smaller scale than the resolution of today's GCMs, which are limited by computational considerations to a typical range of 200 to 1000 km. Therefore, these processes have to be incorporated into the GCMs by means of parameterizations, which may introduce additional errors in the GCM simulations (Thomas and Henderson-Sellers 1991). On the other hand, changes in the hydrological cycle caused by an increase of atmospheric greenhouse gases could have a considerable societal impact (Rind et al. 1992), so that there is a need to assess the potential effects of

Corresponding author address: Dr. Eduardo Zorita, LOYDC, Université Pierre and Marie Curie, 4 Place Jussieu, Tour 14, Paris 75252, France.

climate change at scales that cannot be resolved by current GCMs (Grotch and MacCracken 1991).

Three general strategies have been suggested to overcome this scale mismatch (e.g., Giorgi 1991). The first is to develop finer-resolution regional climate models that are driven by boundary conditions simulated by global GCMs at coarser scales (Giorgi 1990). In theory, these nested models should be able to replicate the physical processes operating at regional scales and can take into account orographic features that are partially or totally absent in a GCM and that may be important for regional climates. However, this approach is computationally costly, and at present the resolution attainable using this approach requires that some processes still must be parameterized. Another problem with this approach is that feedbacks from the regional model into the GCM are not usually incorporated. While alternative numerical schemes such as the adaptive multigrid method could allow such feedbacks to be modeled (see, e.g., Barros and Lettenmaier 1993), these schemes have not yet been applied to GCMs. Another strategy that has been recently developed is the use of time-slice GCM experiments. In such experiments a high-resolution atmospheric GCM is forced by the boundary conditions for the atmosphere generated in a coupled integration of a low-resolution atmosphere-ocean GCM.

Another possibility is to derive statistical models from the observed relationships between the large-scale atmospheric fields, such as sea level pressure (SLP) or geopotential heights, and local variables, such as precipitation or surface temperature. Once the statistical model parameters are estimated from a training set of large-scale and local observations, the models may be used to infer changes in the local variables due to changes in the large-scale fields simulated by GCM sensitivity experiments. For instance, multiple regression equations linking the 700-mb geopotential heights and precipitation (Klein and Bloom 1988) and geopotential heights and fire weather elements (Klein and Whistler 1990) have been used in the United States. Wigley et al. (1990) used, among other variables, large-scale spatial averages of near-surface temperature and correlated them with local temperature time series. With a slightly different strategy, Karl et al. (1990) identified statistical relationships between a set of free atmosphere variables as predictors and near-surface temperature and precipitation as predictands. Von Storch et al. (1993) used canonical correlation analysis to relate local monthly precipitation to the large-scale SLP field. Hewitson (1994) constructed regression equations between the atmospheric circulation and local surface temperature, allowing for nonlinear interactions between different atmospheric regimes. These studies have all concluded that local climate change inferred directly from GCM simulations interpolated to the local scale may differ markedly from local simulations derived from the statistical approach.

All of the statistical techniques noted above essentially make use of correlations between the time series of the large-scale and local variables. However, there are some important variables that are discontinuous in time, like daily precipitation, which are not suitable for statistical techniques such as regression-based methods. On the other hand, there are whole families of sector models (for prediction of agricultural production, hydroelectric power production, surface water supply, terrestrial and aquatic ecosystems, to mention a few), which require as input local precipitation amounts at daily or near-daily timescales. In this case statistical models have to be based upon other techniques, such as the use of weather classification schemes applied to an altered climate, in which the GCM-simulated large-scale fields are classified into weather states (types) and the local observations are sampled from days belonging to a particular weather state. Bardossy and Plate (1992) made use of such a strategy with the classification scheme traditionally used by the German Weather Service. Wilson et al. (1992) defined weather states through a combined Principal Components Analysis of sea level pressure, 850-mb temperature, and 850-mb geopotential height. Hughes et al. (1993) applied Classification and Regression Tree (CART) analysis to identify the weather types that were most related to occurrence or absence of precipitation.

At least three assumptions underlie this type of statistical strategy. First, the GCMs are assumed to simulate realistically the large-scale atmospheric features that give rise to the observed distribution of regional climates, such as the subtropical highs, subpolar lows, and storm tracks. This condition is common to all regionalization techniques, either statistical-empirical or based on high-resolution nested models, and obviously has to be taken as given. If the GCMs themselves fail to reproduce reasonably well the large-scale climate any downscaling approach is doomed to fail. The final estimation of regional climate change strongly depends on the degree of confidence put on the GCM simulations. Second, the relationships between the large-scale and local variables are assumed to hold under the altered climate. This condition is almost impossible to check in practice and this difficulty is in some sense equivalent to the assumption that GCMs will also simulate properly altered climates. It can, however, be argued that within the natural variability of the observed climate there exist all kinds of situations that may eventually prevail in an altered climate. Therefore, if the training period for the statistical method is long enough it will be able to identify the most important factors affecting the regional climate. The uncertainties will, however, always remain, although they can be somewhat limited by imposing a third condition, namely, that the statistical procedure to estimate the local variable is assumed to replicate the historical data, or at least important aspects of their statistical behavior, when it is driven with the observed large-scale circu-

lation. This condition is the counterpart of requiring that the GCMs be also able to simulate past climates. Often this condition cannot be fulfilled because of the lack of observations.

This paper fits within the framework of the coupled empirical/statistical approach. Its aim is to check to what extent some of the above assumptions are fulfilled in practice and to help establish the degree of confidence that can be placed in these procedures. The applications of the method are for daily rainfall at selected stations in two North American regions at midlatitudes: the Columbia River basin, located in the Pacific–North American sector (western region), and the middle-Atlantic region of the eastern United States (eastern region). The choice of the variable that will represent the atmospheric circulation deserves some discussion. It is widely assumed that geopotential height at 500 or 700 mb is strongly related to local rainfall and, therefore, should be a strong candidate. However, a few important considerations have to be taken into account. It is desirable that the observed time series be as long as possible to increase the reliability of the statistical analysis and to allow checking of the statistical relationships on an independent dataset. This dataset should be as separated in time as possible from the training set, so that it could be considered a “different” climate (von Storch et al. 1993). Another question is related to the application of the method to greenhouse-gas experiments. Due to global tropospheric warming, geopotential heights rise globally in a $2 \times \text{CO}_2$ experiment, but this rise is not necessarily bound to changes in the circulation. Therefore, geopotential heights contain information about both temperature and circulation changes and the interaction between the two, the former not being directly related to rainfall. If this effect is not corrected, unrealistically large rainfall changes are to be expected in the statistical downscaling. We decided to use SLP in our study instead of the geopotential height. It offers the advantage that the available time series are globally about 100 yr long and it is marginally affected by temperature effects. Furthermore, SLP has been found to be an acceptable predictor for rainfall (von Storch et al. 1993) and represents well extratropical cyclones (Jones and Simmons 1993) in near-coastal areas, where topography is not so important as in the interior.

With this goal in mind, we first analyze two present-climate simulations of the general circulation models of the Geophysical Fluid Dynamics Laboratory (GFDL), and the Max-Planck-Institute für Meteorology (MPI), respectively. In section 2 the regional performance of the control runs of two GCMs is examined. For this purpose the simulated long-term mean SLP field, its standard deviation, and its coherent patterns of variability (empirical orthogonal functions) are compared to the corresponding patterns derived from observations.

One of the statistical models used to generate daily rainfall at the selected stations is based on a circulation-type classification by CART analysis that was used by Hughes et al. (1993). The CART analysis classifies the daily circulation into weather types objectively, based on the values of several circulation indices. These indices may be the value of SLP at certain grid points or, as in this paper, the amplitude of spatial patterns selected a priori, for instance the leading empirical orthogonal functions. In this paper the CART analysis is applied to historical SLP and rainfall data in the Columbia River basin and middle-Atlantic regions and is used to identify the circulation types that are most strongly related to rainfall at the selected stations. To assess the quality of the downscaling procedure, it is of interest to validate the two GCM control runs in terms of the circulation types by computing quantities such as probability of occurrence and lifetimes of each weather state and observing their change in a $2 \times \text{CO}_2$ experiment performed with the MPI model. This is described in sections 4 and 5 of the paper.

Finally, in section 6, the circulation types identified by CART analysis are used for the generation of daily precipitation time series at individual stations. For this purpose the historical SLP fields, as well as the ones simulated by the GFDL and MPI GCMs, are used to simulate some important statistical properties of the rainfall time series. One important deficiency of this method, as noted by Hughes et al. (1993), is that precipitation sequences generated from the observed SLP fields were not as persistent in terms of the occurrence or absence of precipitation as sequences in the observations. Hughes et al. (1993) explored a modified model that included dependency on precipitation in the previous day, which improved their results. However, in the context of climate change assessment, this modification is not conceptually very satisfactory because it requires the ad hoc assumption that this dependency will remain unchanged in a new climate. In view of these problems, CART analysis is compared to another, simpler, rainfall generator based on an analog method. In the analog method a pool of historical observations is compared to the target pattern and the closest one is chosen as an analog circulation. The rainfall amounts observed simultaneously with the analog circulation are then ascribed to the target pattern. This basic strategy can be modified by defining the similarity between circulation patterns in different ways. Here, similarity is defined by comparison of patterns not only at the current time step, but also in some number of previous days. Therefore, the analog pattern has a similar circulation evolution over several previous days.

2. Data and statistical techniques

Results from three climate simulations were used in this study. The first was a 10-yr control run of the

GFDL GCM performed with prescribed sea surface temperatures and interactive clouds. The GFDL model uses a spectral formulation with R30 resolution, approximately equivalent to a regular grid of 3.7° long \times 2.2° lat. The second experiment was a control run of the MPI coupled ocean–atmosphere model (Cubasch et al. 1992). The atmospheric portion of this climate model (Roeckner et al. 1989) is also a spectral model with a T21 resolution (about $5.6^\circ \times 5.6^\circ$). The third experiment was a $2 \times \text{CO}_2$ run with the MPI model. In this experiment the atmospheric greenhouse gas concentration was doubled from its value in the control run. Daily means of the SLP pressure field between years 76 and 100 of these two MPI simulations were used.

Twenty years of daily SLP analyses from the U.S. National Meteorological Center (NMC) for the period 1965 to 1984 were used in this study. These data were retrieved from a CD prepared by the Department of Atmospheric Sciences of the University of Washington (Mass et al. 1987) and interpolated to a rectangular latitude–longitude grid from the NMC octagonal grid, using software developed at NCAR. Infrequent missing data were filled in by interpolation between the previous and following days. All data were interpolated to the lowest resolution T21 to avoid possible inconsistencies in the subsequent statistical analysis.

Daily station precipitation data in the period 1965–1984 were retrieved from U.S. National Climatic Data Center records. Gaps in these data were removed by using information from nearby stations via a prorating method (Wallis and Lettenmaier 1991). The station positions are indicated in Fig. 1. Two multivariate statistical techniques were applied to the datasets. Empirical Orthogonal Function analysis (EOF; Preisendorfer 1988) is often used in climatology to reduce the number of degrees of freedom of large-scale anomaly fields by identifying a limited number of variables that can describe most of the variance. Mathematically this is achieved by diagonalizing the cross-covariance matrix calculated between anomalies at grid nodes. The eigenvectors of this matrix are called EOF loadings and the variance explained by the EOF is given by its associated eigenvalue. Each EOF has an associated time series (also known as scores) that describes the time evolution of the EOF and that can be calculated at each time step by projecting the EOF onto the field. Since the normalizing constant for each EOF is not defined (as for all eigenvectors of any matrix), we choose this constant in such a way that the associated time series has standard deviation unity (with no physical units), so that the physical units are carried by the EOF. One important property of the scores is that the correlation between any pair of them is zero.

The second statistical technique used was Classification and Regression Tree (CART) analysis (Breiman et al. 1984). This technique provides an objective way to define a scheme to classify the daily circulation into




FIG. 1. Positions of the stations used in this paper.

a small number of weather classes, relevant for the precipitation occurrence or absence at a certain set of stations (precipitation is considered here as a two-valued discrete variable, wet or dry). The daily circulation states are classified by means of a binary decision tree, the nodes of which are split depending on the values of one of the input variables, such as, for instance, the value of the SLP field at a certain grid node or some other large-scale circulation index. The algorithm begins by examining all possible splits of the data based on the input variables and choosing the one that gives the maximal separation of the precipitation occurrence distribution. That is, if there are 16 possible rainfall patterns (corresponding to rain/no rain at four stations, say), then the split that minimizes

$$\sum_{i \neq j} p(i)p(j) \quad (1)$$

[where $p(i)$ and $p(j)$ are the probabilities of the i th and j th rainfall patterns in the two daughter nodes] is chosen. Since there are a finite number of observations there are a finite number of such splits to examine. Tree construction continues until a balance is reached between the "cost" (the sum of the cost function over all terminal nodes) and the complexity (the number of terminal nodes) is reached. Additional details can be found in Breiman et al. (1984). Each terminal node of the decision tree corresponds to a weather state and each day can be classified into one of the weather states by following the binary decision tree. Typically, the splits are made on individual input variables. Thus, the boundaries of the weather classes are defined by CART as inequalities on the input variables (e.g., $a(1) < z_1$ and $a(3) > z_2$, where $a(1)$ and $a(3)$ are input variables and z_1 and z_2 are constants). Graphically, these inequalities can be represented as hyperplanes parallel to the axes in the input variables space (e.g., the EOF space). It is conceptually possible, but much more computationally demanding, to allow the splits to be based on linear combinations of the input variables (once again, for a finite number of observations there are a finite number of such linear combinations that uniquely partition the data). Therefore, the input variables for a CART analysis should be chosen with care, which may not always be an easy task. In our analysis, we compared (for western region winter data) the weather states that resulted from CART runs that alternately did or did not allow for rotations of the input variables. We found no qualitative difference; therefore, we report the unrotated results only.

3. Regional model validation

In this section some basic statistical parameters derived from the control runs of the two GCMs are compared with the same parameters derived from the NMC data. This comparison is restricted to sea level pressure in the two regions of interest, the west region (30° – 70° N, 115° W– 180°) and east region (30° – 70° N, 45° – 100° W). Both regions were investigated for winter and summer. All of the analyses are based on daily data, in the case of the NMC observations, taken at 0000 GMT.

Figure 2 shows the long-term mean of the SLP field in January and July for the observations and the two model runs for the western region. In January, both models produce a stronger Aleutian Low than is observed in the historical data. In the GFDL run the low is centered northeast of its correct position. In July the MPI control run correctly reproduces the summer high pressure cell over the Pacific Ocean, whereas in the GFDL run the high pressure is much stronger than in the observations. The long-term standard deviations of the SLP fields calculated from the daily means for

January and July are shown in Fig. 3. The historical data for January show a broad variability maximum centered over the Aleutian islands. In the MPI run this maximum is fairly well simulated. The GFDL model tends to be much more variable than the observations, with its maximum extending eastward into the continent. In July the variability of the MPI model is smaller than in the observations, and it also misses the variability maximum over the Pacific Ocean, whereas the GFDL model replicates quite well the variability distribution.

The coherent patterns of the SLP variability have been identified using EOF analysis, separately for the winter (DJF) and summer (JJA) months. Figure 4 depicts the first four (five for the observations) EOFs of the SLP anomaly field, along with their relative explained variances. In winter both models are able to reproduce the pattern of the first EOF, although the explained variance is underestimated by the MPI and overestimated by the GFDL model. Lower-order EOFs are, in general, well reproduced by the MPI run, whereas in the GFDL model the second and third EOFs are interchanged. In summer, the variance accounted for by the first EOF is again overestimated by the GFDL model and underestimated by the MPI model, which also produces a pattern that deviates from the observations. Other higher-order EOFs are satisfactorily reproduced by both models.

A similar comparison between observations and the model simulations was carried out for the eastern American coast. For the observed mean SLP field in January (Fig. 5) the edge of the Icelandic low can be seen, which is also present in the GFDL run with the right location and strength. In contrast, in the MPI simulation, this low is shifted too far southwestward. In July both models produce correctly the position of the quasi-permanent anticyclone over the North Atlantic, but in the GFDL run, its strength is overestimated, as was also the case for the North Pacific, while the MPI model underestimates SLP over North America. With respect to variability (Fig. 6), the NMC analysis for January shows a broad maximum of the SLP standard deviation over Greenland. In the MPI run this maximum is displaced somewhat to the southwest, whereas in the GFDL run this maximum is very distorted and elongated into the continent. In July, the observed SLP variability presents a marked zonal symmetry; this feature is also found in the MPI and GFDL simulations, although the models tend to be more variable. With respect to the EOFs in the winter months (Fig. 7), the patterns simulated by the GFDL model, except for EOF 2, do not resemble, even qualitatively, the ones derived from the historical data. The performance of the MPI model is much better in this respect. In contrast, in the summer, the GFDL model achieves much better results, most notably in replicating the first EOF, whereas the MPI results are not as good as in winter, especially for the lower-order EOFs.




FIG. 2. Mean SLP fields (mb) in the Pacific-North American sector as derived from the NMC analysis, the GFDL model, and the MPI model: (a) January; (b) July.

4. Weather state classification

In this section, we describe an application of CART analysis to identify the weather states that are most closely related to the occurrence or absence of precipitation in four selected stations in the Columbia River basin of the northwestern United States. Based on the CART analysis, a stochastic model linking the weather states with the presence/absence of precipitation, and precipitation amounts, at selected stations was developed. In implementing this approach for the conditional simulation of station precipitation associated with GCM scenarios, two problems must be addressed. First, the long-term mean climate simulated by a GCM

is usually not the same as in the observations. This problem can be avoided in sensitivity experiments with GCMs by considering only the differences between a control run and an anomaly run. This approach is justified only when these differences are small, so that the simulated climate change may be similar to the real climate change. If these differences are large, the use of the GCM is meaningless and no information about future climates can be inferred from it, either at large or at small scales. This approach is followed here and in the CART classification scheme by performing the analysis using only anomalies of the SLP field. The CART procedure also requires for computational reasons that the number of input variables be limited, but




FIG. 3. SLP standard deviation (mb) in the Pacific-North American sector as derived from the NMC analysis, the GFDL model, and the MPI model: (a) January; (b) July.

at the same time, we are interested in capturing the essentials of the large-scale atmospheric circulation. For these reasons the input variables used in the CART procedure are the time series associated with the most important EOFs of the SLP anomaly field. To account for lags between the large-area pressure fields and local precipitation, we used the EOF scores for the present and the previous day. In the Columbia River basin the stations selected for the CART analysis were Arrowrock Dam, Ellensburg, Stampede Pass, and Wickiup. These stations provide a reasonable coverage of the basin and have a minimum of missing data during the period under study.

The results of the CART analysis for the winter months (DJF) are schematically shown in Fig. 8 and in Table 1. The CART procedure identified three

weather states based on the binary decision scheme illustrated in Fig. 8. The variables that are identified to define the weather states are the one-day-lagged scores of EOF 3 to discriminate between state 1 and the other two, and EOF 5 to discriminate state 2 and 3. For instance, 11 days in which the score of PC 3 is less than -0.085 are classified as belonging to state 1. If not, they may belong to state 2 or 3 according to the value of PC5.

To illustrate the situations that give rise to each of the weather states, SLP composite plots based upon the days classified as belonging to each of the weather states are displayed in Fig. 8. The most probable rainfall occurrence patterns at the four stations for each weather state can be found in Table 1. It can be seen that the first weather state is usually associated with the occur-

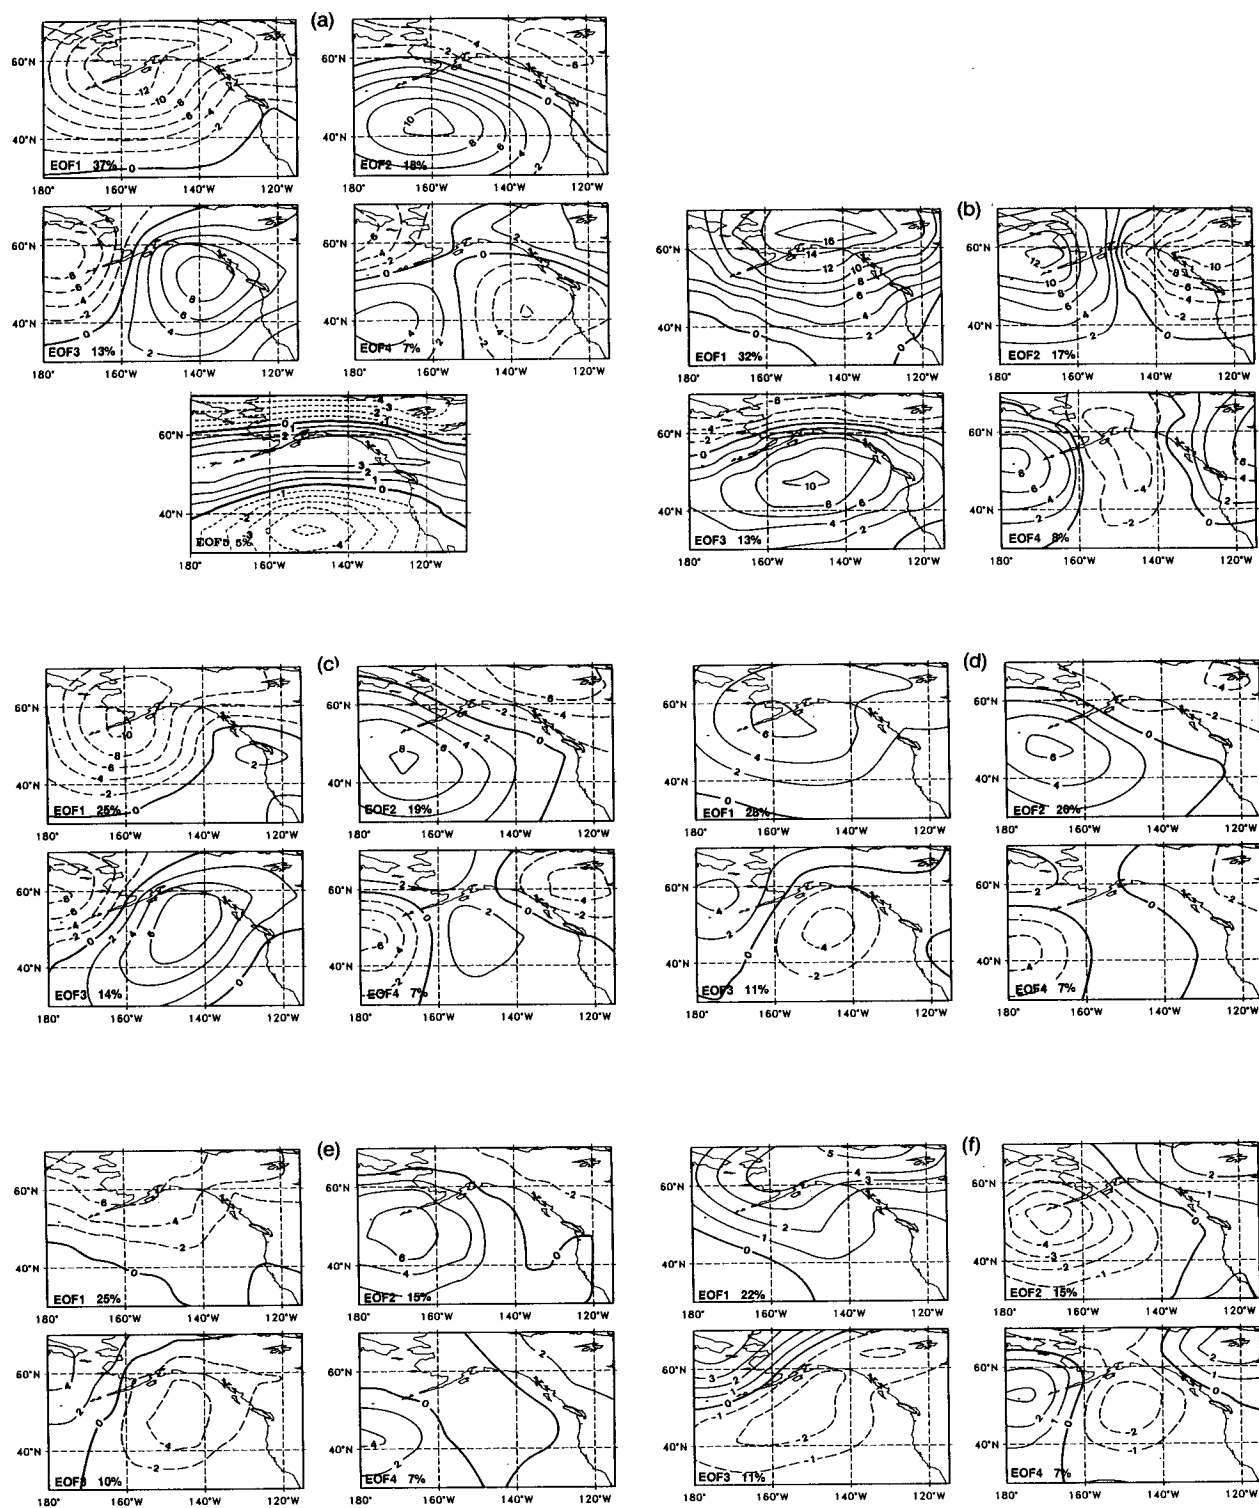


FIG. 4. Leading EOFs of the SLP anomaly field (mb) in the Pacific-North American sector as derived from (a) NMC (DJF); (b) GFDL (DJF); (c) MPI (DJF); (d) NMC (JJA); (e) GFDL (JJA); and (f) MPI (JJA).

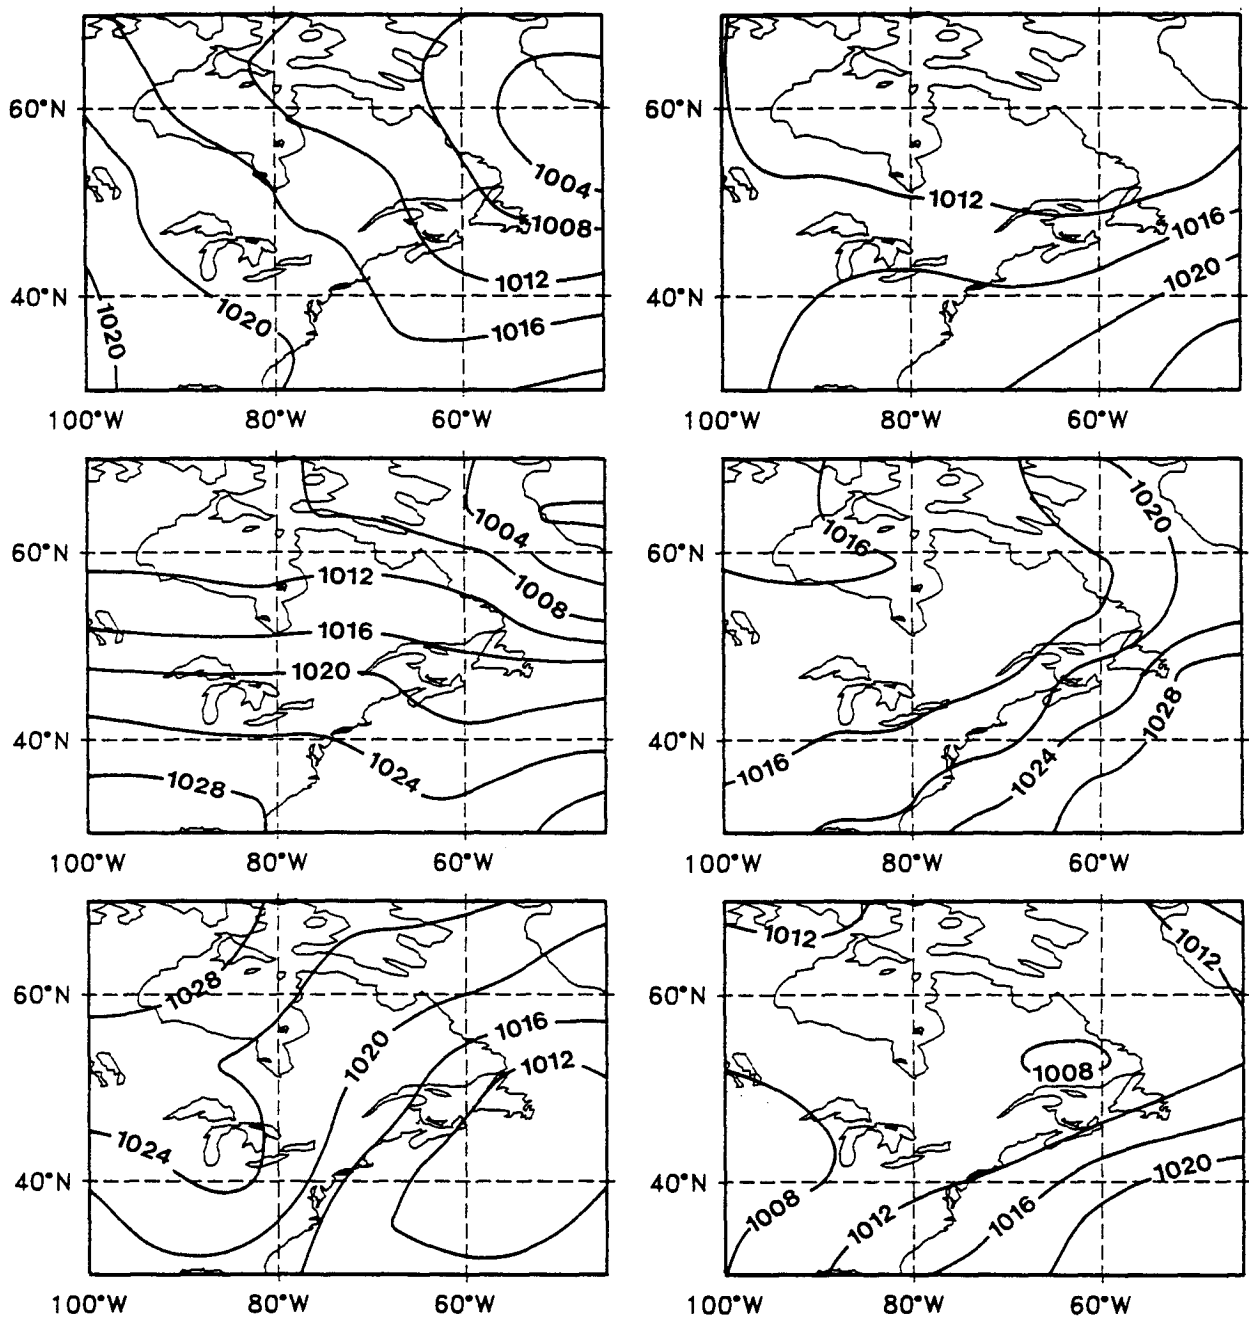


FIG. 5. Mean SLP fields (mb) in the Atlantic–North American sector as derived from the NMC analysis, the GFDL model, and the MPI model: (a) January; (b) July.

rence of precipitation at all stations. The third state is mainly related to no precipitation at any of the stations, whereas the second weather state may be accompanied by no rain, rain at the most western stations, or rain at all of them. The interpretation of these results may be as follows: when the score of EOF 3 is negative (see Fig. 4) a lower-than-normal pressure cell sitting off the northwest coast of the United States advects humid air

from the southwest into the continent. When the EOF 3 score is positive this mechanism does not operate but the CART procedure identifies another way by which rain may reach some or all of the stations. EOF 5 (Fig. 4) is associated with anomalous geostrophic zonal wind at approximately the latitude of the index stations. When the score of EOF 5 is positive, the zonal (westward) circulation is enhanced and may bring

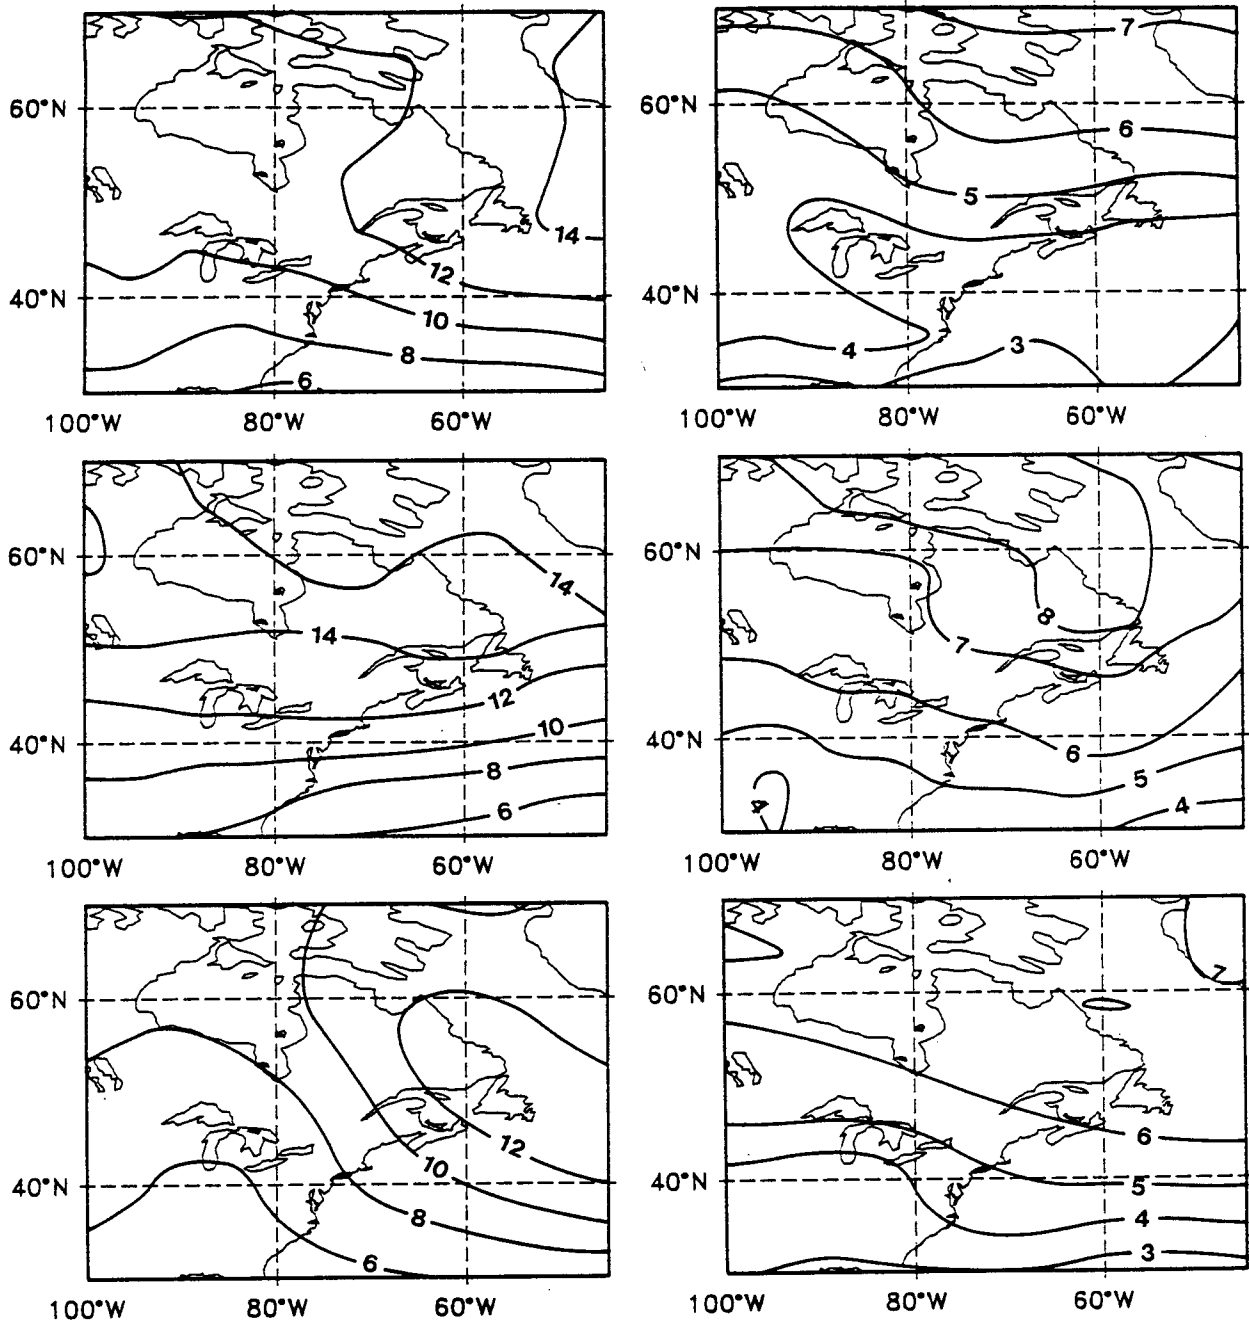


FIG. 6. SLP standard deviation (mb) in the Atlantic-North American sector as derived from the NMC analysis, the GFDL model, and the MPI model: (a) January; (b) July.

some rain with it. Depending on the relative strength of EOF 5 and EOF 3, the rain may reach all stations, only the most western stations, or none of them (state 2). If the EOF 5 score is negative, both mechanisms work jointly to produce dry days at all the stations (state 3). That the splitting variables are the one-day lagged scores should be roughly due to the time needed for the large-scale weather patterns over the Pacific to reach the stations in the continent.

A relevant question in this context is whether the weather states identified by using data from these four stations are stable with respect to changes in the number or location of the stations. A similar CART analysis was also performed retaining just two stations in the set: one located near the coast (either Ellensburg or Stampede Pass) and the other located in the interior (either Arrowrock Dam or Wickiup). It turned out that the weather states remained essentially the same.

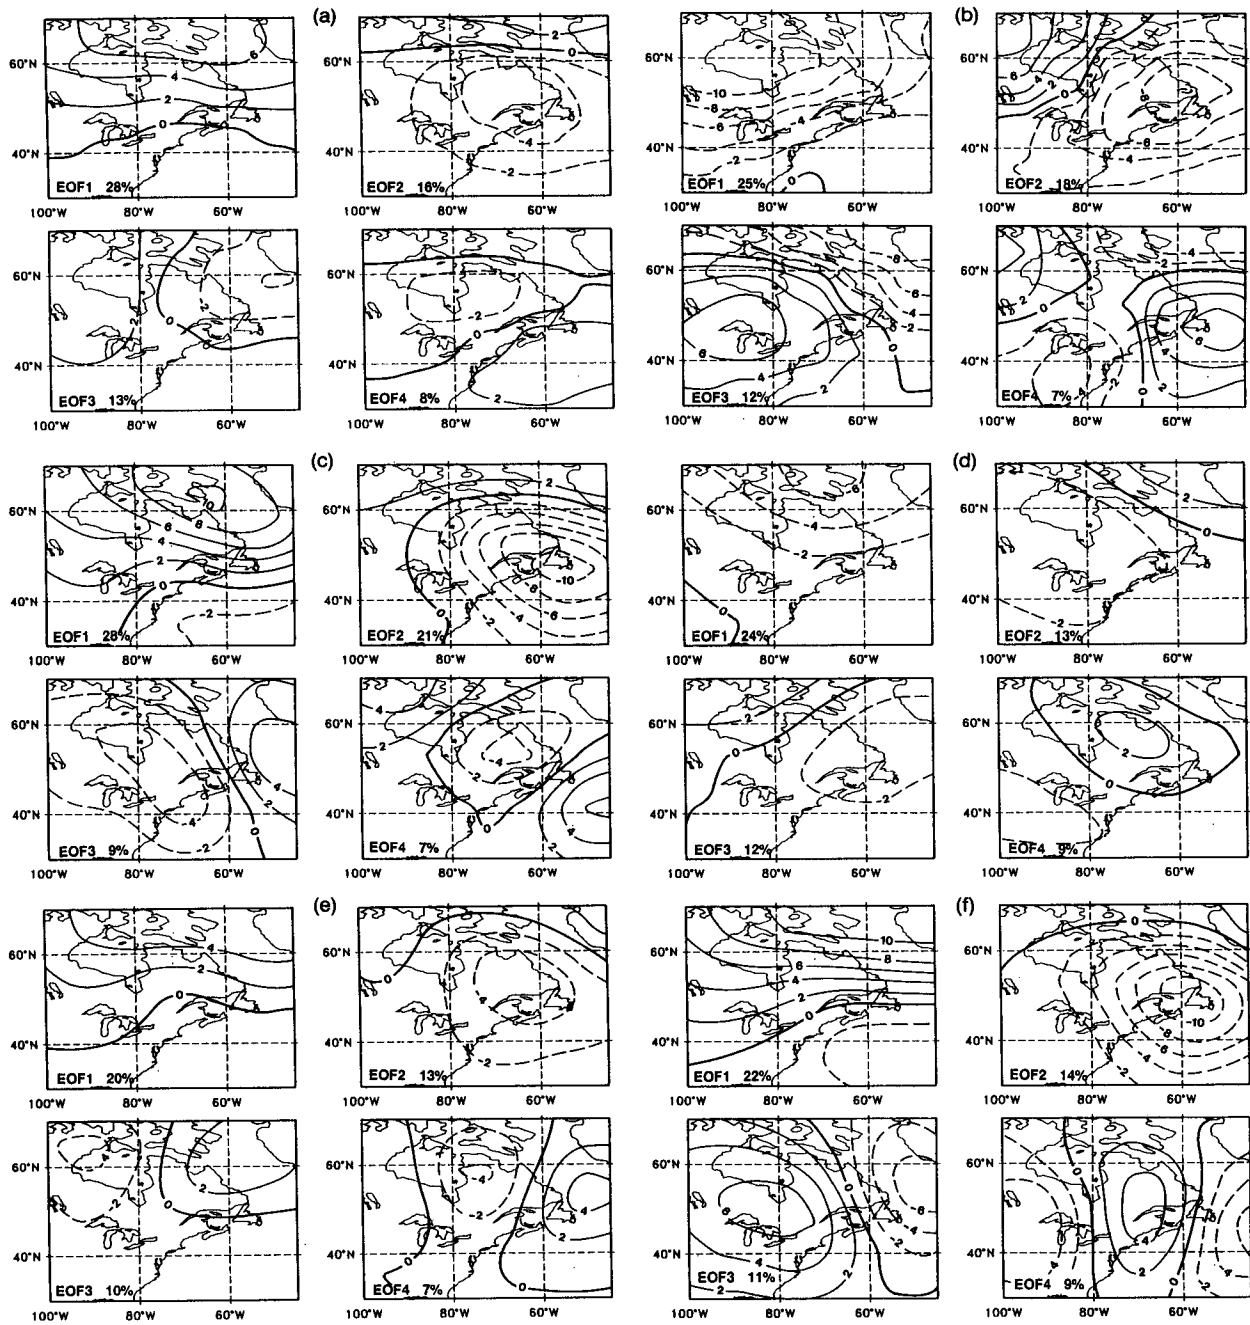


FIG. 7. First four EOFs of the SLP anomaly field (mb) in the Atlantic-North American sector as derived from (a) NMC (DJF); (b) GFDL (DJF); (c) MPI (DJF); (d) NMC (JJA); (e) GFDL (JJA); and (f) MPI (JJA).

Furthermore the same weather states were identified when the data were split into two half-periods of ten years each. We also applied the CART analysis to individual stations in the Columbia River basin. The discriminant variable between dry and wet states was always the score of EOF3, indicating that this statistical association is quite stable. We turn to this question again when presenting the results for the eastern region.

In the summer months (JJA), the CART procedure was not able to produce any decision tree when all four stations were used, perhaps indicating that precipitation at each of the stations is not so closely related to one another in this season as it is in winter. This hypothesis was checked by removing the two interior stations (Arrowrock Dam and Wickiup) from the set and performing the CART analysis with the remaining two

near-coast stations. In this case two weather states were found, defined by the classification scheme of Fig. 9 and Table 2. The precipitation occurrence distribution (Table 2) indicates that these two states are mainly associated with dry or wet days in both stations simultaneously, respectively. The only node in the tree is the split based on a relatively large (negative) value of PC5.

Composites calculated from the days belonging to each state are also shown in Fig. 9, which can be interpreted using the same ideas as in winter: for days in which PC5 is negative a high-pressure zone is located northwest of the two stations, preventing moist air from reaching the coastal areas and giving rise to the dry state. The reverse reasoning is valid for the wet state. If either of the interior stations is included in the analysis, the CART procedure is not able to find any classification of weather states. We believe that the reason is that summer precipitation is produced by processes of smaller scale than in winter.

The classification of weather states for the eastern region turned out to be more difficult. Initially a set of four stations was selected (Delaware, Ohio; Hatteras, North Carolina; Hightstown, New Jersey; and Newport, Tennessee) with a geographical separation

TABLE 1. Most frequent rainfall patterns for each of the CART weather states in the Pacific-North American sector in DJF and their relative occurrence probability (conditional on the weather state). Here A = Arrowrock Dam, E = Ellensburg, S = Stampede Pass, Wi = Wickiup; W = Wet, and D = Dry.

Rainfall pattern	State 1	State 2	State 3
A E S Wi			
WWWW	0.33	0.15	0.06
WWWD	0.23	0.22	0.11
WWD D	0.16	0.15	0.09
D D D D	0.03	0.14	0.39

roughly of the same order as those in the Columbia River basin. However, no simple classification tree could be created for these stations by the CART procedure, either in winter or in summer. We applied the CART analysis to three stations located closer to one another (Baltimore, Maryland; Hightstown, New Jersey; and Lexington, Virginia), again with negative results for both seasons. An attempt was also made to identify weather states by using a smaller region in the EOF calculation, with the hope of identifying a smaller-scale process that could be relevant for all the stations simultaneously but this also proved to be unsuccessful.

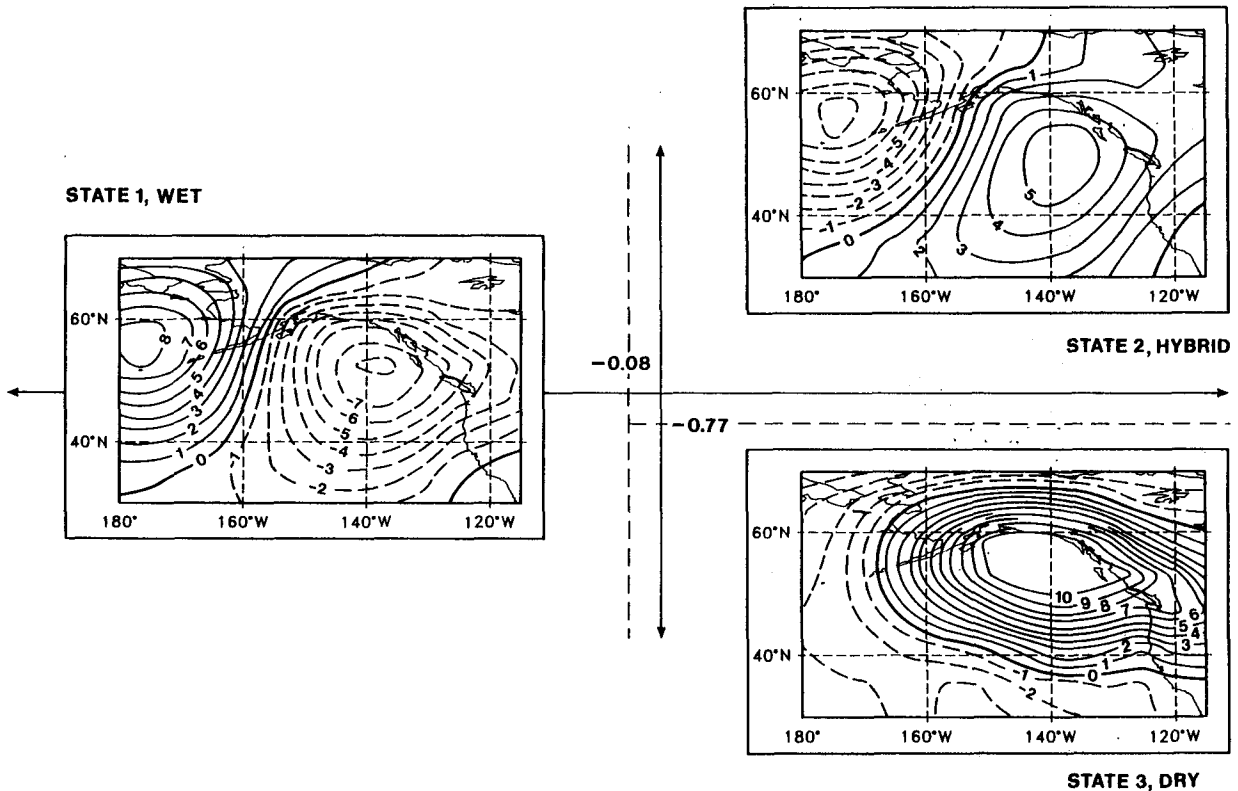


FIG. 8. Results of the CART analysis for the SLP field in the Pacific-American sector in stations in the Columbia River basin. The discriminant variables are PC3 and PC5 (principal components are normalized). For instance all days in which PC3 is less than -0.085 belong to state 1. Plots correspond to the composite SLP fields calculated from the days belonging to each weather state (see also Table 1).




FIG. 9. Results of the CART analysis for the SLP field in the Pacific-American sector in JJA, derived from the NMC analysis and four stations in the Columbia River basin. The discriminant variable is PC5 (principal components are normalized). Plots correspond to the composite SLP fields calculated from the days belonging to each weather state (see also Table 2).

Only when two stations (Baltimore and Hightstown) situated fairly close to each other were used could a relatively simple tree be created for the winter months; however, even in this case no tree was found for the summer. CART analysis was also applied to individual stations. For some of them no simple classification scheme could be found, suggesting that the large-scale EOFs cannot discriminate between dry or wet days at these stations. For others the discriminant variables were different for each station (except for Baltimore and Hightstown in winter). The reason for the failure of CART in the eastern region is not completely clear. It may happen that other variables play a more important role than in the western region or that the SLP EOFs in this region do not define well the synoptic regimes that are associated with rainfall (see section 2). In any case, it seems clear that the problem is much more complicated in the eastern than in the western region. It is intended to investigate this point further with more sophisticated versions of CART. For this reason no results of the CART analysis for the eastern region will be presented.

5. Weather states in the GCM simulations

Once a relevant set of weather states has been identified it is interesting to investigate whether these states can be realistically simulated in the GCM runs. One difficulty arises from the fact that the weather states are defined in terms of the EOFs derived from the observations, which may deviate from those derived from the GCM simulations. For instance, the second and third EOF in the GFDL run for the western region appear interchanged with respect to the observations. Therefore, the scores cannot be directly used to classify the daily circulations into weather states. One consistent way to overcome this problem is to project the simulated SLP anomalies onto the same EOF patterns obtained from the historical data and use the resulting time series $y_j(t)$ in the classification scheme:

$$y_j(t) = \sum_{i=1}^N f^i(t)g_j^i/\sigma_j^2, \quad (2)$$

where N is the number of grid points, $f^i(t)$ is the simulated SLP anomaly at time t and grid point i , g_j is the j th EOF derived from the SLP NMC analysis, and σ_j^2 is the variance explained by EOF j . Since the area covered in this study is relatively small, there is no need to account for the shrinking gridpoint separation with latitude. For the GFDL and MPI control runs the anomalies were calculated by subtracting the respective long-term mean from the simulated data; therefore, the possible bias of the models is removed. For the $2 \times \text{CO}_2$ MPI experiment, however, the long-term mean of the control run was used to retain the possible climate signal due to CO_2 doubling. The following weather state analysis can be used to validate the GCM's ability to simulate the behavior of observed regional synoptic situations. Our attention will be focused on two aspects: the probability of occurrence of each of the weather states and the log-survivor function of each state. This function is defined as

$$L(t) = \ln(1 - F(t)), \quad (3)$$

where $F(t)$ is the cumulative distribution function of the state lifetime. Thus, the log-survivor function gives the log probability that a certain state will last for more than t consecutive days. Tables 3 and 4 give the prob-

TABLE 2. Rainfall patterns for each of the CART weather states in the Pacific-North American sector in JJA and their relative occurrence probability (conditional on the weather state). Here E = Ellensburg, S = Stampede Pass; W = Wet, and D = Dry.

Rainfall pattern	State 1	State 2
E S		
W W	0.21	0.43
W D	0.11	0.11
D W	0.10	0.14
D D	0.58	0.32

TABLE 3. Probabilities of occurrence of the three weather states identified by CART analysis in the Western American Sector in DJF as observed and in the different model runs.

	State 1	State 2	State 3
Observations	0.46	0.42	0.12
MPI	0.45	0.48	0.07
MPI (2 × CO ₂)	0.44	0.45	0.11
GFDL	0.52	0.35	0.12

ability of occurrence of each state derived from the observations and from the three model runs in winter and summer, respectively. The differences between these parameters are in general small. This result is to some extent due to fact that the scores have by construction zero mean (by taking anomalies the model bias has been removed), so that when a node in the decision tree is split upon a very small (absolute) value of one variable it is likely that half of the days will be classified to follow one branch and the other half to follow the other one, regardless of the data origin. However, this will not be, in general, the case if the node is split upon a nonzero value of a variable, since the standard deviations of the scores are in general different in the models than in the observations.

Figure 10 depicts the log-survivor functions for the weather states derived from the NMC analyses and both GCMs. These functions are nearly straight lines, which should be the case if the transitions between states follow a purely Markov process. In that case the absolute value of the derivative would be the inverse of the exponential parameter t . For the western region in winter, the GFDL model tends to produce longer series of state 1 (wet state) and state 2 (hybrid) and shorter series of state 3 (dry state) than were observed. The MPI model simulates correctly the lifetime of state 1 but underestimates that of state 2 and overestimates the life time of state 3. The doubling of atmospheric CO₂ causes a slight increase of the lifetimes of all the weather states in the MPI model. A plausible explanation might be increased atmospheric stability due to a reduced meridional temperature gradient in the 2 × CO₂ climate. In summer, both models replicate properly the behavior of the dry weather state, whereas the GFDL generates longer, and the MPI model shorter, series of the wet state.

6. Precipitation generation for GCM experiments

The weather state classification scheme can be used for the stochastic generation of precipitation time series based on the daily circulations simulated by the GCMs. Since the present generation of GCMs lacks the necessary resolution to simulate precipitation reliably at specific stations, the precipitation generator is an alternative method to transfer the results of the GCM experiments to the local scale.

In this section two different stochastic precipitation generators are described. The first is based on the CART classification scheme. Although in general this method behaves satisfactorily, the precipitation generator based on CART proved to be unsuccessful in replicating some aspects of the precipitation processes, such as the storm interarrival times. This quantity is fairly important when precipitation time series are used by hydrologists to drive catchment (precipitation runoff) models. Furthermore, CART proved to be not always successful in identifying the relevant weather states. Therefore, a second type of precipitation generator, based on a modified analog method, was also explored in an attempt to resolve some of these deficiencies.

a. Precipitation generator based on CART

The first method for generating precipitation sequences is based on the CART classification scheme. This method consists of applying the CART tree to each day of the historical record, which is assigned the value of the resulting weather class. The precipitation values for each day of the historical record are then assigned to "bins" associated with each of the corresponding weather states. Given a sequence of weather states (e.g., resulting from the application of CART to a GCM simulation run), the simulated sequence of precipitation states is determined by sampling, with replacement, from the appropriate bin.

This procedure was applied to the three GCM experiments (GFDL and MPI base climate, and MPI 2 × CO₂). Initially, though, the performance of this method was checked by analyzing the results produced when the input weather classes were taken from the observation series (NMC data); the resulting conditional simulation of precipitation was then compared with the precipitation that actually occurred (in the sense of precipitation presence/absence). Because the procedure is not restricted to generating precipitation for the stations used to "train" the CART algorithm (Arrowrock Dam, Ellensburg, Stampede Pass, and Wickiup), we illustrate the simulated results for an independent station, Forks (located along the Pacific Coast). We will focus on three properties of the simulated precipitation that are of concern to hydrologic modelers: mean precipitation, storm interarrival times,

TABLE 4. Probabilities of occurrence of the two weather states identified by CART analysis in the western American sector in JJA.

	State 1	State 2
Observations	0.24	0.76
MPI	0.19	0.81
MPI (2 × CO ₂)	0.22	0.78
GFDL	0.24	0.76

and the daily precipitation probability distribution. It should be noted that, because the precipitation is re-sampled from the historical record, there is no point in comparing simulated and observed quantities such as the absolute maximum of daily precipitation or spatial precipitation correlation: these quantities are essentially the same in the simulated and observed records.

The precipitation time series generated from the GCM simulations cannot, of course, be compared directly with the observations; the comparison must, of necessity, be with the long-term mean precipitation (Table 5). It can be seen that the differences between the real and simulated mean precipitation are well within the observed interannual variability for both models. However, it should be noted that a completely random resampling from the observations would have yielded the right mean precipitation; therefore, these numbers alone do not demonstrate the validity of the precipitation generator. With respect to the $2 \times \text{CO}_2$ MPI run, no significant deviations were found from the $1 \times \text{CO}_2$ experiment in terms of mean precipitation. This can be traced to the small changes in the mean SLP field between the $2 \times \text{CO}_2$ and the control run in the MPI model, as in other regions in the same experiment (von Storch et al. 1993).

The probability density functions of daily precipitation for Forks (western region) as derived from the observations and generated from the observed and simulated SLP fields are shown in Fig. 11. All proba-

TABLE 5. Mean rainfall (mm day^{-1}) in Forks (Columbia River basin) as observed and simulated by the CART and analog methods from the observed SLP field and from the different model runs in winter and summer.

Observations	Winter 15.1		Summer 2.3	
	CART	Analog	CART	Analog
NMC	14.9	14.4	2.6	2.1
MPI	14.6	15.6	2.4	1.5
MPI ($2 \times \text{CO}_2$)	14.2	15.9	2.4	1.5
GFDL	16.1	14.9	2.1	1.7

bility distributions are in good agreement but again this is not necessarily due to the skill of the method.

Figure 12 shows the log-survivor functions of the storm interarrival times for Forks derived from the historical precipitation data, by applying the precipitation generator to the daily SLP fields observed (NMC) and simulated by the three GCM runs. The log-survivor functions show that dry periods tend to last longer in the observations than in all model runs, in winter as well as in summer. Since there is no systematic under or overestimation of the weather state lifetimes in the model simulations (see Fig. 10), this result indicates that there is some persistence in the precipitation process that is not properly captured by this precipitation generator. This reasoning is supported by the fact that the storm interarrival times generated with the historical SLP data as input variables are also shorter than




FIG. 10. Log-survivor functions of the CART weather states in the Pacific-North American sector in DJF and JJA for the NMC analysis and in the different model runs. The log-survivor function is the log probability that a certain state lasts for more than t consecutive days: — NMC; --- GFDL; ... MPI; — MPI ($2 \times \text{CO}_2$).




FIG. 11. Daily rainfall (mm day^{-1}) probability distribution (percent) as observed and generated by the CART and ANALOG methods from the NMC analysis and from different model runs in winter (DJF) and summer (JJA): (a) Forks (Pacific sector); (b) Hightstown (Atlantic sector).

the actual storm interarrival times, consistent with the findings of Hughes et al. (1993). In winter, for all datasets the storm interarrival times are generally shorter at Forks than at West Glacier (not shown), a station located in the interior. Since weather state 1 (wet state) and weather state 2 (dry state) are mainly associated with wet or dry days, respectively, at all stations simultaneously, this difference in the behavior of the log-survivor functions can be traced to weather state 3 (hybrid state). This also explains why the GFDL model produces shorter storm interarrival times than the MPI model: the GFDL model simulates shorter series of weather state 3 (hybrid state), and weather state 2 (dry state) has a fairly low probability in both models.

In the summer months the storm interarrival times are similar in all stations (only results for Forks are shown in Fig. 12). This should be due to the fact that in summer only two weather states have been identified, an "all-dry" state and an "all-wet" state. Furthermore, both models behave nearly equally with respect to the storm interarrival times. This is easily explained since the two models simulate very similar lifetimes of the dry weather state (Fig. 10).

b. Precipitation generator based on analog method

CART analysis is a sound classification scheme that also provides a physical interpretation of the resulting weather types. However, the failure of the rainfall generator based on CART to replicate the observed storm interarrival times could limit the application of this method. Possible reasons for this shortcoming may lie in the dependence of rainfall occurrence on the precipitation amounts in previous days (Hughes et al. 1993; Bardossy and Plate 1992). Another possibility is that precipitation events depend more strongly on the evolution of the atmospheric circulation in previous days than assumed in the CART analysis. In this section an alternative precipitation generator will be presented to investigate this second possibility.

Given a SLP field simulated by a GCM experiment, this method finds in the daily observations the closest possible SLP field and then takes for the simulated precipitation the observed precipitation on the day so identified. This is the well-known analog method described in the literature (Lorenz 1969), but we use it here with a slight modification proposed by Barnett and Preisendorfer (1978). Instead of comparing directly the simulated and the possible analog SLP fields, both are projected onto the five leading EOFs derived from the observations. Thus, each circulation pattern corresponds to a point in a five-dimensional space, the coordinates of which are its projections onto the first five EOFs. The "distance" between them is then calculated based on their coordinates in this new basis by the usual definition of Euclidean distance, where N is the number of EOFs, x and y are the EOF coordinates of the two daily circulations, and a is the metric used.

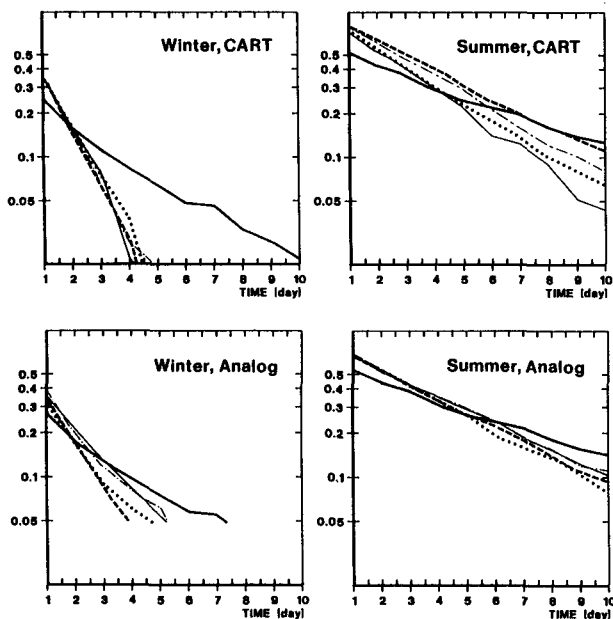


FIG. 12. Log-survivor functions of the storm interarrival times in Forks (Pacific sector) in winter (DJF) and (JJA) as observed and generated with the CART and ANALOG methods from the NMC analysis and the different model runs. The log-survivor function is the log probability that a dry period lasts for more than t consecutive days: — Observations; ··· NMC; — GFDL; - - - MPI; - · - · MPI ($2 \times \text{CO}_2$).

In this way it is expected that a part of the noise present in the daily SLP fields will be filtered out. Another advantage is that anomaly fields may be used, thus removing the possible bias in the mean circulation simulated by the GCMs.

As with the CART-based precipitation generator, this method was also checked using the observed SLP fields as input. (In this case, to avoid any artificial skill, the “analog” circulation was selected from years other than the one being simulated.) The generated storm interarrival times (not shown) tended still to be too short when compared with the historical precipitation data. As a further refinement to this method, the analog circulation was sought not just by comparing single SLP fields but comparing the evolution of the circulation in the previous four days; that is, the circulation in the present day and in the previous four days was compared to all the historical five-day segments and the most similar of them (also in terms of the EOFs coordinates) was taken as the analog. Then, the observed precipitation for the fifth day was ascribed to the circulation pattern in question. For brevity, only the results of this analog-based method for Forks (western region) and Hightstown (eastern region) are presented. While the generated mean precipitation (Tables 5 and 6) and probability density functions (Fig. 11) do not deviate significantly from the results of the “one-segment” analog method, the log-survivor func-

tions of the storm interarrival times (Figs. 12 and 13) are more similar to the ones derived from the observed precipitation data.

Part of this apparent improvement might be spurious. Assume for simplicity that there is no time autocorrelation in the atmospheric circulation and that the analog for day d was found to be day k using both methods. Then, the probability that the analog for day $d + 1$ lies in the surroundings of day k is greater in the “five-day-segment” method (since only 20 percent of the information is new), than in the “one-day-segment” method. Therefore, the simulated storm interarrival times would artificially resemble those observed without necessarily implying that the method is representing the rainfall process better. To check how strong this problem can be in our case, consider Fig. 14. It shows the probability that the analogs of day d and day $d + n$ are separated exactly by n days, for the “one-day segment”, “five-day segment,” and “ten-day segment” methods in DJF in the North Pacific–American sector. It can be seen that these probabilities grow with increasing segment length, as expected, but they are only moderately high for a separation of one day.

We believe, based on this analysis, that the persistence of the precipitation process may be better captured by weather generators that take into account the evolution of the daily SLP fields and that CART analysis using information from the SLP in the roughly five previous days would improve the results. However, this would require considerably longer computer times, since the computational requirement grows exponentially with the number of input variables.

In many applications of the analog method the question arises of how close each situation is to its analog. In our case a measure of this similarity can be given by the patterns of local time correlation between each SLP field and its analog: Fig. 15 shows the result of this calculation for observations, the MPI control run, and the MPI greenhouse experiment for wintertime in the western region. Similar results are obtained for summer and for the eastern region (not shown). It can be seen that the correlations are high where the observed variability is high. This is not surprising since the search for analogs was performed in the space

TABLE 6. Mean rainfall (mm day^{-1}) in Hightstown (Atlantic sector) as observed and simulated by the analog method from the observed SLP field and from the different model runs in winter (DJF) and summer (JJA).

	Winter 2.9	Summer 3.4
Observations	Analog	Analog
NMC	2.8	3.7
MPI	3.0	3.7
MPI ($2 \times \text{CO}_2$)	2.2	3.4
GFDL	2.8	3.9

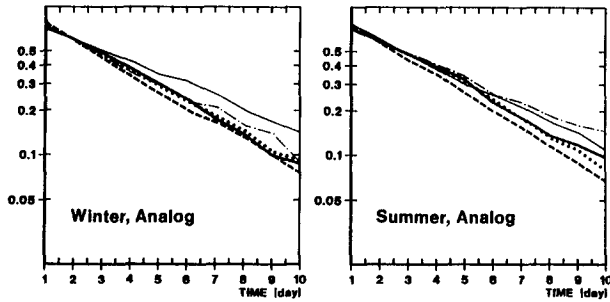


FIG. 13. Log-survivor functions of the storm interarrival times in Hightstown (Atlantic sector) DJF and JJA as observed and generated with the ANALOG method from the NMC analysis and from the different model runs. The log-survivor function is the log probability that a dry period lasts for more than t consecutive days: — Observations; ··· NMC; — GFDL; --- MPI; - - - MPI ($2 \times \text{CO}_2$).

spanned by the leading EOFs. Interestingly, the correlations are high and of the same order for the three datasets, which means that on average it is possible to find analogs of similar quality in the pool of observations, for the observations, for the control run, and for the $2 \times \text{CO}_2$ run.

7. Conclusions

The ability of two GCMs to replicate selected regional features of the lower atmosphere circulation has been investigated and compared with historical data. In general, both of the GCMs are reasonably successful in simulating the major features of the observed mean circulation. For instance, both models reproduce qual-

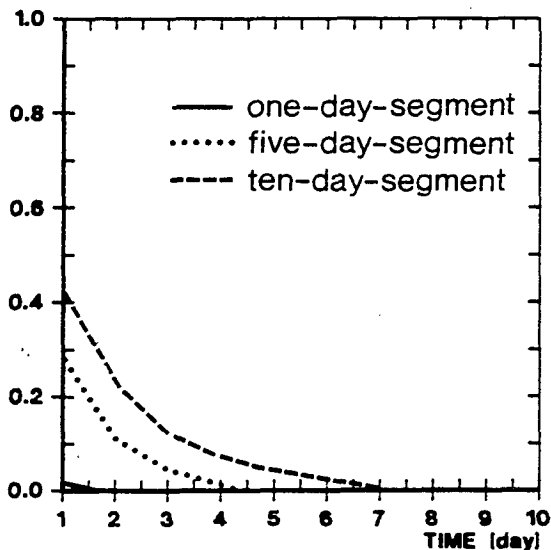


FIG. 14. Probability that the analogs chosen for day d and day plus n (n from 1 to 10 days) are separated by exactly n days in DJF in the Pacific-North American sector for the “one-day segment,” “five-day segment,” and “ten-day segment” methods.

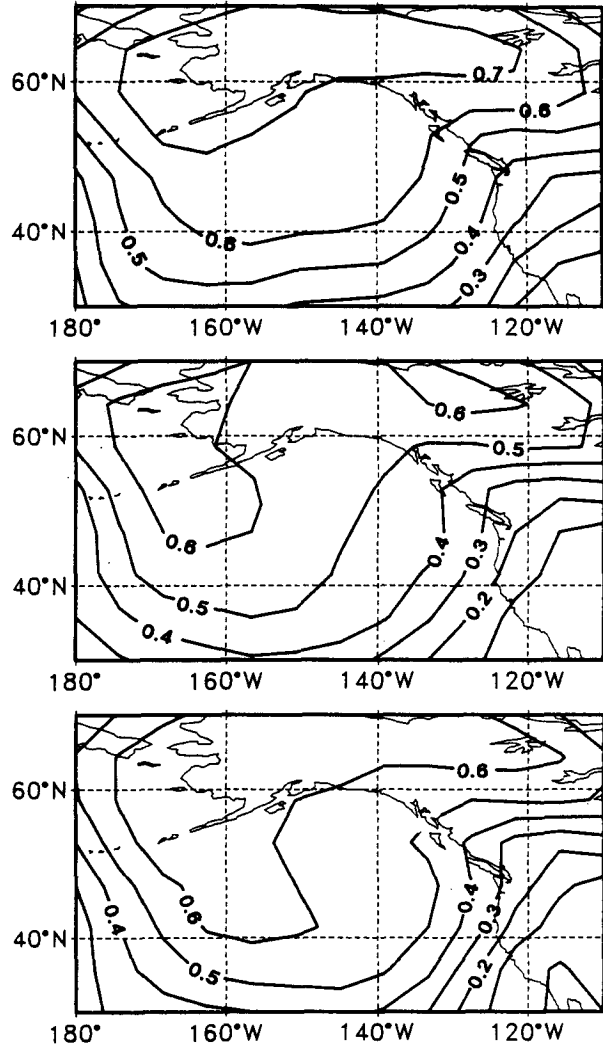


FIG. 15. Time-correlation patterns between daily circulation and their analogs from the pool of historical observations: (a) for observations (analog is chosen in years different than the circulation analyzed); (b) MPI control run; and (c) MPI $2 \times \text{CO}_2$ run.

itatively the quasi-permanent anticyclones over the Pacific and Atlantic Oceans in summer and the Aleutian and Icelandic Lows in winter, although in some cases the exact position and intensity of these features are misrepresented. The models also reproduce quite well the leading patterns of daily SLP variability on these regional scales, although some improvement is still necessary. In general the GFDL model shows a greater variability than the observations, whereas the MPI model is usually less variable. It is not obvious if this fact is related to the different ocean representation in these models, since other authors have found no significant differences in a GCM extratropical variability in experiments performed with fixed or varying SSTs (Chervin 1986).

A CART classification scheme based on the daily SLP anomaly field and station precipitation data was applied to identify the weather states that are most strongly associated with the presence/absence of precipitation at those stations. For the winter season, a reasonable set of weather states was identified for the western region. In the summer season only the stations that are located near the coast could be used to classify the daily large-scale SLP fields, indicating that precipitation events in the interior of the Columbia Basin in summer are not as dependent on the large-scale atmospheric circulation as in winter.

For the eastern region of the United States the situation is more complicated. CART analysis, at the level of sophistication used in this paper, is not able to find the circulation regimes associated with rainfall in the index stations. The reason for this failure may lie in the role played by other variables, such as humidity advection or tropospheric temperature, or the input variables, which in this paper are the leading principal components, do not define well the synoptic situations associated with rainfall. If this second possibility is true, CART would try to build a quite complicated classification scheme, resulting in numerous weather states with similar characteristic rainfall patterns. However, the tradeoff between model fit and model complexity upon which CART is based mitigates against such a model, so the algorithm may fail to detect this type of situation. A more sophisticated use of CART analysis should allow linear combinations of discriminant variables that yield the best classification scheme. This possibility, which requires considerably more computer time, will be investigated in a future work. In the cases when a set of weather states could be found, the GCM simulations were quite successful in reproducing the observed weather state occurrence probabilities. These weather states are defined in terms of anomalous circulations deviating from the long-term mean of each dataset, so that the possible bias of the GCMs in simulating the mean state is not reflected in the replication of the weather states. However, the variability of each GCM-simulated dataset is taken fully into account, so that the GCMs do show some skill in simulating the weather types. Some discrepancies were found in the log-survivor functions of the weather states, but there was no systematic bias, indicating that this behavior depends on the particular dynamics of each model.

A stochastic precipitation generator developed by Hughes et al. (1993) and based on CART analysis was applied to the historical and simulated SLP fields. The results so obtained were compared to a simpler stochastic generator based on an analog method. Both methods show strengths and weaknesses. The daily rainfall probability distribution was generally well replicated by both methods when they were driven by the observed daily SLP fields, but the storm interarrival times were in both cases shorter than in the observations. The analog method performed better in terms

of the storm interarrival times. This advantage may stem from the fact that the analog method exploits information from five consecutive days, whereas CART is limited to information of one day because of the computational load. Also, the analog method uses the information available more efficiently, in the sense that it uses the optimal analog from the whole pool of observations, while CART samples at random from the observations that belong to a certain weather state. The price paid by the analog method is that it does not give any physical insight whatsoever about the connection between the large-scale circulation and local rainfall. CART, on the other hand, produces a nice classification of the daily circulations and allows for a physical interpretation in terms of the effect of these circulations regimes on the local rain. The fact that the results of the analog method improve when also the evolution of the SLP field is taken into account may suggest future research to clarify this important point.

The reliability of the application of the rainfall generators to greenhouse gas experiments is not at all clear cut. On one side, the downscaling approach assumes good large-scale simulations by the GCM. This may be approximately true for the present climate, but to what extent it is also true for an altered climate is an open question, since not all GCMs produce the same results even at large scales. A direct way to test the performance of the GCM would be by paleoclimate simulations with all the known difficulties due to the lack of data. A similar test specific for these statistical downscaling methods would be the reconstruction of historical rainfall from the past SLP observations. This has been successfully done for monthly rainfall in previous examples (von Storch et al. 1993). The problem of working with daily data is that there are few time series extending long enough into the past.

With respect to the changes in precipitation caused by a CO₂ doubling, small differences were found between the MPI control run and the 2 × CO₂ experiment. The reasons for this result are probably the low climate sensitivity of the MPI model in comparison with other models (Houghton et al. 1990) and the fact that the signal-to-noise ratio in the SLP is usually lower in altered 2 × CO₂ experiments (Barnett 1991).

REFERENCES

- Bardossy, A., and E. J. Plate, 1992: Space-time model for daily rainfall using atmospheric circulation patterns. *Water Resour. Res.*, **28**, 1247–1259.
- Barnett, T. P., 1991: An attempt to detect the greenhouse signal in transient GCM simulation. *Greenhouse-Gas Induced Climate Change: A Critical Appraisal of Simulations and Observations*. M. E. Schlesinger, Ed., Elsevier, 559–568.
- , and R. Preisendorfer, 1978: Multifield analog prediction of short-term climate fluctuations using a climate state vector. *J. Atmos. Sci.*, **35**, 1771–1787.
- Barros, A. P., and D. P. Lettenmaier, 1993: Multiscale aggregation and disaggregation of precipitation for regional hydroclimatological studies. *Macroscale Modeling of the Hydrosphere*, E. B.

- Wilkonson, Ed., International Association of Hydrological Sciences, Publication 214, 184–194.
- Breiman, L., J. H. Friedman, R. A. Olsen, and J. C. Stone, 1984: *Classification and Regression Trees*. Wadsworth, 391 pp.
- Chachine, M. T., 1992: The hydrological cycle and its influence on climate. *Nature*, **359**, 373–380.
- Chervin, R. M., 1986: Interannual variability and seasonal climate predictability. *J. Atmos. Sci.*, **43**, 233–251.
- Cubasch, U., K. Hasselmann, H. Hoock, E. Maier-Reimer, U. Mikolajewicz, B. D. Santer, and R. Sausen, 1992: Time-dependent greenhouse warming computations with a coupled ocean-atmosphere model. *Climate Dyn.*, **8**, 55–59.
- Giorgi, F., 1990: Simulation of regional climate using a limited area model nested in a general circulation model. *J. Climate*, **3**, 941–963.
- , and L. O. Mearns, 1991: Approaches to the simulations of regional climate change: A review. *Rev. Geophys.*, **29**, 191–216.
- Grotch, S., and M. C. MacCracken, 1991: The use of general circulation models to predict regional climate change. *J. Climate*, **4**, 286–303.
- Hewitson, B., 1994: Regional climates in the GISS General Circulation Model: Air surface temperature. *J. Climate*, **7**, 283–303.
- Houghton, J. T., G. J. Jenkins, and J. J. Ephraums, Eds., 1990: *Climate Change. The IPCC Scientific Assessment*. Cambridge University Press, 364 pp.
- Hughes, J. P., D. P. Lettenmaier, and P. Guttorp, 1993: A stochastic approach for assessing the effect of changes in regional circulation patterns on local precipitation. *Water Resour. Res.*, in press.
- Jones, D., and I. Simmons, 1993: A climatology of Southern Hemisphere extratropical cyclons. *Climate Dyn.*, **4**, 131–145.
- Karl, T. R., W.-C. Wang, M. E. Schlesinger, R. W. Knight, and D. Portman, 1990: A method of relating general circulation model simulated climate to the observed local climate. Part I: Seasonal statistics. *J. Climate*, **3**, 1053–1079.
- Klein, W. H., and H. Bloom, 1988: An operational system for specifying monthly precipitation amounts over the United States from the field of concurrent mean 700-mb heights. *Wea. Forecasting*, **4**, 51–60.
- , and B. Whistler, 1990: Specification of monthly mean anomalies of fire weather elements in the United States. *Agric. For. Meteorol.*, **56**, 145–172.
- Lorenz, E. N., 1969: Atmospheric predictability as revealed by naturally occurring analogs. *J. Atmos. Sci.*, **26**, 636–646.
- Mass, C. F., H. J. Edmon, H. J. Friedman, N. R. Cheney, and E. E. Recker, 1987: The use of compact discs for the storage of large meteorological and oceanographic datasets. *Bull. Amer. Meteor. Soc.*, **68**, 1556–1558.
- Preisendorfer, R., 1988: *Principal Components Analysis in Meteorology and Oceanography*. Elsevier, 425 pp.
- Rind, D., C. Rosenzweig, and R. Goldberg, 1992: Modeling the hydrological cycle in assessment of climate change. *Nature*, **358**, 119–122.
- Roeckner, E., L. Dumenil, E. Kirk, F. Lunkeit, M. Ponater, B. Rockel, R. Sausen, and U. Schlese, 1989: The Hamburg version of the ECMWF model (ECHAM). GARP Report No. 13, WMO, Geneva, WMO/TP No. 332, 7.1–7.4.
- Thomas, G., and A. Henderson-Sellers, 1991: An evaluation of proposed representations of sub-grid hydrologic process in climate models. *J. Climate*, **4**, 898–910.
- von Storch, H., E. Zorita, and U. Cubasch, 1993: Downscaling of climate change estimates to regional scales: Application to winter rainfall in the Iberian Peninsula. *J. Climate*, **6**, 1161.
- Wallis, J. R., and D. P. Lettenmaier, 1991: A daily hydroclimatological data set for the continental United States. *Water Resour. Res.*, **27**, 1657–1663.
- Wigley, T. M. L., P. D. Jones, K. R. Briffa, and G. Smith, 1990: Obtaining sub-grid scale information from coarse-resolution general circulation model output. *J. Geophys. Res.*, **95**, 1943–1953.
- Wilson, L. L., D. P. Lettenmaier, and E. Skillingstad, 1992: A multiple stochastic daily precipitation model conditional on large-scale atmospheric circulation patterns. *J. Geophys. Res.*, **97**, 2791–2809.