

A technique for generating regional climate scenarios using a nearest-neighbor algorithm

David Yates¹

National Center for Atmospheric Research, University of Colorado, Boulder, Colorado, USA

Subhrendu Gangopadhyay, Balaji Rajagopalan, and Kenneth Strzepek

Department of Civil, Environmental and Architectural Engineering, University of Colorado, Boulder, Colorado, USA

Received 14 October 2002; revised 17 March 2003; accepted 15 April 2003; published 31 July 2003.

[1] A K -nearest neighbor (K - nn) resampling scheme is presented that simulates daily weather variables, and consequently seasonal climate and spatial and temporal dependencies, at multiple stations in a given region. A strategy is introduced that uses the K - nn algorithm to produce alternative climate data sets conditioned upon hypothetical climate scenarios, e.g., warmer-drier springs, warmer-wetter winters, and so on. This technique allows for the creation of ensembles of climate scenarios that can be used in integrated assessment and water resource management models for addressing the potential impacts of climate change and climate variability. This K - nn algorithm makes use of the Mahalanobis distance as the metric for neighbor selection, as opposed to a Euclidian distance. The advantage of the Mahalanobis distance is that the variables do not have to be standardized nor is there a requirement to preassign weights to variables. The model is applied to two sets of station data in climatologically diverse areas of the United States, including the Rocky Mountains and the north central United States and is shown to reproduce synthetic series that largely preserve important cross correlations and autocorrelations. Likewise, the adapted K - nn algorithm is used to generate alternative climate scenarios based upon prescribed conditioning criteria. *INDEX TERMS:* 1854 Hydrology: Precipitation (3354); 1869 Hydrology: Stochastic processes; 1899 Hydrology: General or miscellaneous; 1630 Global Change: Impact phenomena; *KEYWORDS:* bootstrap, resampling, climate data, climate change

Citation: Yates, D., S. Gangopadhyay, B. Rajagopalan, and K. Strzepek, A technique for generating regional climate scenarios using a nearest-neighbor algorithm, *Water Resour. Res.*, 39(7), 1199, doi:10.1029/2002WR001769, 2003.

1. Introduction

[2] Integrated assessment (IA) studies link biophysical and socioeconomic models for studying the effects of climate change and other anthropogenic stressors on both natural and human systems [Cohen, 1986; Dowlatabadi and Morgan, 1993; Mearns *et al.*, 1996]. They do this by predicting, for example, local patterns of spatial change in agroecosystem boundaries, soil carbon storage, changes in soil moisture, and water resource availability, and then, in conjunction with various development policies, simultaneously address the implications of these local impacts at broader regional and national scales [Rosenzweig and Parry, 1994; Smit *et al.*, 1996; Yates and Strzepek, 1996, 1998a].

[3] Because of their detailed characterization of biophysical processes, many of these IA models require high-resolution climate data (e.g., precipitation, temperature, wind, solar radiation, and so on) at relatively fine spatial and temporal (at least daily) scales to drive these

process models. To calibrate these models and evaluate their performance, observed meteorological data sets are used as forcing variables [Yates and Strzepek, 1998a, 1998b; Yates *et al.*, 2000; Kenny, 2000]. To study the likely effects of climate change for IA analysis, climate scenarios are generated through downscaling techniques that involve developing statistical relationships between historic meteorological observations and outputs from regional and/or global climate models [Wilks, 1992; Robock *et al.*, 1993; Easterling, 1999; Hewitson and Crane, 1996; Semenov, 1997; Wilby *et al.*, 1998; Sailor and Li, 1999a, 1999b; Mearns *et al.*, 1999]. While these approaches for simulating climate scenarios for IA analysis are useful, they do have limitations. For example, a climate change scenario based on output from a general circulation model (GCM) simulation is a single realization of many possible climatic futures, while an ensemble of climate scenarios that could rigorously explore the decision space of IA models would be more useful. Also, in many cases GCMs do not adequately replicate the historic climate of a region, so there is a great deal of uncertainty regarding the regional GCM output under future scenarios of increasing CO₂ and aerosol changes.

[4] Stochastic weather generators address the issue with their ability to simulate plausible climate scenarios and have themselves been used as downscaling techniques in global

¹Also at Department of Civil, Environmental and Architectural Engineering, University of Colorado, Boulder, Colorado, USA.

change studies [Wilks, 1992]. Typically, a stochastic weather generator is developed based on the historically observed data at a location, and it can then be used to simulate climate scenarios consistent with the global change scenarios. Furthermore, this technique can be used to investigate the sensitivity of IA models and their individual components to climate variability and climate change, an issue that has not received its due attention [Morgan and Dowlatabadi, 1996; Kunkel et al., 1998; Mearns et al., 1996, 1997; Adams, 1999; Mingkui and Woodward, 1998; Risbey and Stone, 1996; Smith and Pitts, 1997; Mitchell et al., 1999].

[5] In this paper we develop and demonstrate the utility of a modified version of the K -nearest-neighbor (K -nn) based nonparametric stochastic weather generator [Buishand and Brandsma, 2001; Rajagopalan and Lall, 1999] to simulate climate scenarios on a regional scale. The paper first gives a brief background on stochastic weather generators. Section 3 describes the methodology of the K -nn regional weather generator, including the methodology to adapt it to simulate daily weather sequences conditioned upon alternative climate scenarios. Section 4 then applies the algorithm to data from two regions of the United States, generating synthetic series that preserve historic statistics and series with alternative statistical attributes compared with historic. A summary of the techniques and results concludes the paper.

2. Background

[6] Stochastic weather generators are routinely used in water, agricultural, and erosion control management [Skidmore and Tatarko, 1990; Wilks, 1997; Dubrovsky et al., 2000]. Weather generators based on parametric statistical techniques typically use precipitation as the driving variable in a number of models [see Jones et al., 1972; Nicks and Harp, 1980; Richardson, 1981], where precipitation occurrence and amount are generated independently, and the other variables are then generated based on the stochastically generated precipitation. For example, a precipitation occurrence and amount model (e.g., a two-state Markov model, with exponentially distributed rainfall amount) can be used to generate the sequence of dry and wet days and precipitation amount. The other variables are simulated using a lag 1 multivariate autoregressive model with exogenous precipitation input (MAR-1). Furthermore, a model is fit to each month (or season) separately to capture the nonstationarity (i.e., seasonality). There are several drawbacks to the parametric approach, including the following: (1) The MAR implicitly assumes a normal distribution of the variables, which is difficult to satisfy; consequently, nonnormal features in the data such as bimodality cannot be captured; (2) only linear relationships between the variables can be reproduced; (3) a large number of parameters are separately fit to each season. If the simulations are to be conditioned, e.g., conditioned upon the state of the El Niño-Southern Oscillation Index (ENSO), then separate models have to be fit for each ENSO state, which further increases the number of parameters; (4) by simulating other variables conditioned on precipitation, only that part of the dependency that is related to precipitation is captured. Furthermore, Katz [1996] points out that modifying the parameters of a stochastic model can lead to unanticipated effects. For example, modifying the

probability of daily precipitation occurrence using a stochastic weather generator [Richardson, 1981] changed not only the mean of daily temperature, but its variance and autocorrelation as well.

[7] Nonparametric methods based on simulating from kernel-based multivariate probability density estimators [Rajagopalan et al., 1997] and K -nearest-neighbor (K -nn) bootstrap methods [Young, 1994; Rajagopalan and Lall, 1999; Buishand and Brandsma, 2001] can improve upon the parametric models. The nearest neighbor methods are based on the classic bootstrapping techniques [Efron, 1979; Härdle and Bowman, 1988; Zucchini and Adamson, 1989; Yakowitz, 1993]. Young [1994] employed a multiple discriminant function to identify K -nn (or K unique days, where K was 3–5 days) of the current day's weather, and one of these was then randomly selected and used as the next day's weather. Young's model mostly preserves the cross correlation between the variables, but biases were also noticed (e.g., reduced persistence and underestimation of the fraction of dry months) in the generated series. Rajagopalan and Lall [1999] adapted the K -nn bootstrap method for time series resampling to multivariate data (i.e., daily weather). Rajagopalan and Lall's algorithm is very similar to Young's method but with two main differences: (1) The discriminant function is not used, and (2) K -nearest neighbors are obtained to current day's weather and one of the neighbors is selected based on a probability metric that depends on the closeness of the neighbor. This approach preserved the persistence (i.e., lag-1 correlations) and also seasonal statistics. Buishand and Brandsma [2001] extended this approach to simulate weather scenarios at multiple locations preserving the dependencies across the locations. These nonparametric methods provide a flexible framework, are parsimonious, make no assumptions of the underlying marginal or joint probability distribution of the variables, have the ability to reproduce any arbitrary functional relationship, and can easily be modified to do conditional simulation (e.g., conditioned upon ENSO).

[8] In this paper we use a modification of the K -nn method of Lall and Sharma [1996], Rajagopalan and Lall [1999], and Buishand and Brandsma [2001] to generate synthetic climate data and generate alternative climate scenarios at several sites. The general algorithm used for these purposes is described in the next section.

3. Methodology

[9] The K -nn algorithm employed here uses the Mahalanobis distance metric [Davis, 1986], which has an operational advantage over a Euclidean distance approach by not requiring the explicit weighting of variables and does not require the variables to be standardized. A strategy is introduced that gives priority to certain years from which to resample, which can be used to derive alternative scenarios with differing statistical attributes.

3.1. The K -nn Algorithm

[10] Suppose that historic daily weather data are available at m stations for N years and for p variables (here $p = 3$ including precipitation (PPT), maximum temperature (TMX), and minimum temperature (TMN)), and that generation begins on day t corresponding to 1 January. The algorithm proceeds according to the following steps:

[11] 1. Regional means of the p variables across the stations corresponding to day t are given as

$$\bar{x}_t = \begin{bmatrix} \overline{PPT}_t \\ \overline{TMN}_t \\ \overline{TMX}_t \end{bmatrix} \quad (1)$$

where

$$\overline{PPT}_t = \frac{1}{m} \sum_{j=1}^m PPT_{j,t} \quad (2)$$

$$\overline{TMN}_t = \frac{1}{m} \sum_{j=1}^m TMN_{j,t} \quad (3)$$

$$\overline{TMX}_t = \frac{1}{m} \sum_{j=1}^m TMX_{j,t} \quad (4)$$

where \overline{PPT}_t , \overline{TMN}_t , and \overline{TMX}_t are the regional means of precipitation, minimum temperature, and maximum temperature computed for day t from all m stations, respectively.

[12] 2. All days within the temporal window of width w centered on day t are selected as potential candidates for day $t + 1$. For example, if a 14-day temporal window is chosen, and t is 1 January (of a randomly selected year from set N), then the window of days consists of all days between 25 December and 8 January, including 1 January for all N years but excluding day t . Thus there are $[(w + 1) * N] - 1$ days that are potential neighbors to day t . A fixed length, 14-day moving window ($w = 14$) was used in this study.

[13] 3. Mean vectors \bar{x}_i are computed for each day i , where $i = 1$ to $[(w + 1) * N] - 1$. These mean vectors are computed using the equations from step 1.

[14] 4. A pxp covariance matrix S_t is computed for day t using the data block $([(w + 1) * N] - 1) \times p$ described in step 2.

[15] 5. The weather on the first day t (e.g., 1 January) comprising all p variables at m stations is randomly chosen from the set of all days of the historic record of N years (e.g., all 1 January are candidate days with equal probability of selection) and includes all p variables. This is the feature vector F_t^i and represents the stochastically generated weather for day t of year i given for each station. The algorithm continues with the objective of selecting the next day, $t + 1$.

[16] 6. Mahalanobis distances d_i are computed between the mean vector of the current day's weather, \bar{x}_t and the mean vectors \bar{x}_i as

$$d_i = \sqrt{(\bar{x}_t - \bar{x}_i)^T S_t^{-1} (\bar{x}_t - \bar{x}_i)} \quad (5)$$

for all $i = 1$ to $[(w + 1) * N] - 1$, where T is the transpose of the vector.

[17] 7. The distances d_i are sorted and the first K -nearest neighbors are retained. *Lall and Sharma* [1996] provide an

objective method to choosing K , but the heuristic choice, $K = \sqrt{[(w + 1) * N] - 1}$, generally performs well as seen by them and also by *Rajagopalan and Lall* [1999]. Thus if 58 years of data are used with a 14-day window, then $K = \sqrt{[(14 + 1) * 58] - 1} \approx 29$.

[18] 8. A probability metric with a weight function given as

$$p_j = \frac{1/j}{\sum_{i=1}^K 1/i} \quad (6)$$

for all $j = 1$ to K is constructed. The weight function assigns weights to each of the K -nearest neighbors. The neighbor with the shortest distance gets a high weight, while the neighbor with the smallest distance (i.e., the K th nearest neighbor) gets the least weight. *Lall and Sharma* [1996] argue the efficacy and simplicity of the weight function in equation (6). Furthermore, the probability metric does not depend on the d_i distances so it need not be computed at each time step.

[19] 9. The $t + 1$ day is selected as one of the K neighbors using the probability metric p_j , leading to a new feature vector, F_{t+1}^i . A random number, $u \sim U(0,1)$ is first generated, and if $u \geq p_1$, then the $t + 1$ day corresponding to distance d_1 is selected. If $u \leq p_K$, then the $t + 1$ day corresponding to d_K is selected. For $p_1 < u < p_K$, the day $t + 1$ corresponding to d_j is selected for which u is closer to p_j . For example, with 29 nearest neighbors, then from equation (6), $p_1 = 0.2524$; $p_2 = 0.1262$; ...; $p_{29} = 0.0087$. Now, if $u = 0.1250$, which is closest to p_2 , then the day corresponding to d_2 is selected.

[20] 10. Steps 6 through 9 are repeated for as many days of required simulation as are necessary. If multiple years of independent simulation are desired, then the algorithm starts at step 5, randomly selecting the first day from the N years of data.

[21] Since the daily weather at the stations within a region is sampled simultaneously, dependence between the stations and among variables at each station should be largely preserved. Further, this being a bootstrap of the historical data, non-Gaussian features in the probability density functions of the variables are retained [*Lall and Sharma*, 1996; *Rajagopalan and Lall*, 1999]. Finally, the algorithm can be used to perform strategic simulation to derive new daily weather series with altered attributes as compared with historically observed attributes (such as shifts in the mean or changes in variability), as will be seen in the next section.

3.2. Strategic Resampling

[22] When all K days constitute the candidate days for selection with the K -nn algorithm, the goal of the simulation is to produce synthetic series that reproduce the statistics of the historic data. However, many decision models (e.g., watershed models and reservoir operation models) require weather sequences to be conditioned upon some large-scale climate signal such as the ENSO so that the sensitivity of the system to alternative climate scenarios can be tested. Likewise, it would be advantageous to test IA models with climate data that consists of divergent patterns relative to historically observed climate, such as warmer and wetter springs or a gradual warming trend over a certain period of

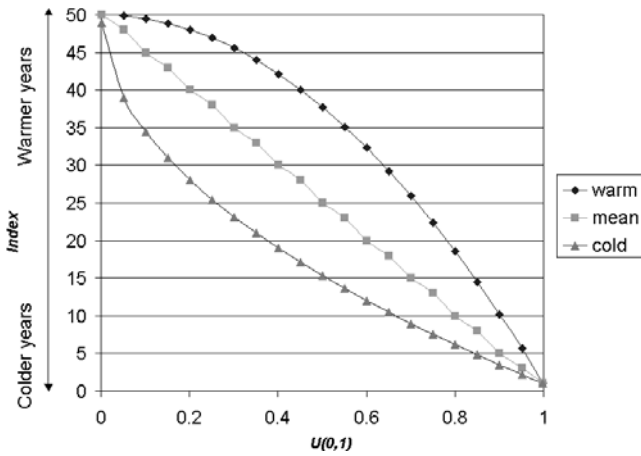


Figure 1. Plot of the index function given by equation (12). A $\lambda_w^i = 1.0$ corresponds to the straight line (mean), while a $\lambda_w^i = 2.0$ introduces curvature into equation (12), biasing warm years and $\lambda_w^i = 0.6$ value will yield a cold bias.

time and region. This would allow for an improved understanding of the complex interactions and system response of these models. Strategies to adapt the K - nn algorithm for developing such alternative climate scenarios are described below.

[23] Strategic resampling simply implies that a subset of years, $n \in N$ will be used in the K - nn algorithm based on some conditioning criteria. The simplest criteria would be some large-scale climate signal such as the annual ENSO index, which influences regional climate. Thus, to generate daily weather with similar ENSO characteristics, only those years with a particular ENSO index would comprise the $n \in N$ subset, and the K - nn algorithm would be constrained to only resample from the $[(w + 1) \times n] - 1$ days of daily data. This approach would work fine for short weather sequences, such as a year; but for longer sequences, it would be desirable to dynamically select the n subset from the larger population of N years so that longer weather sequence would reflect different climate regimes. To do this, a temporal, probabilistic resampling scheme is introduced that generates subsets of years n_w^τ for each week w . The superscript τ represents the fact that the set of years from which to resample is not fixed, but can change both interannually and intra-annually according to year-to-year, weekly deviations (or anomalies) observed in the historic record.

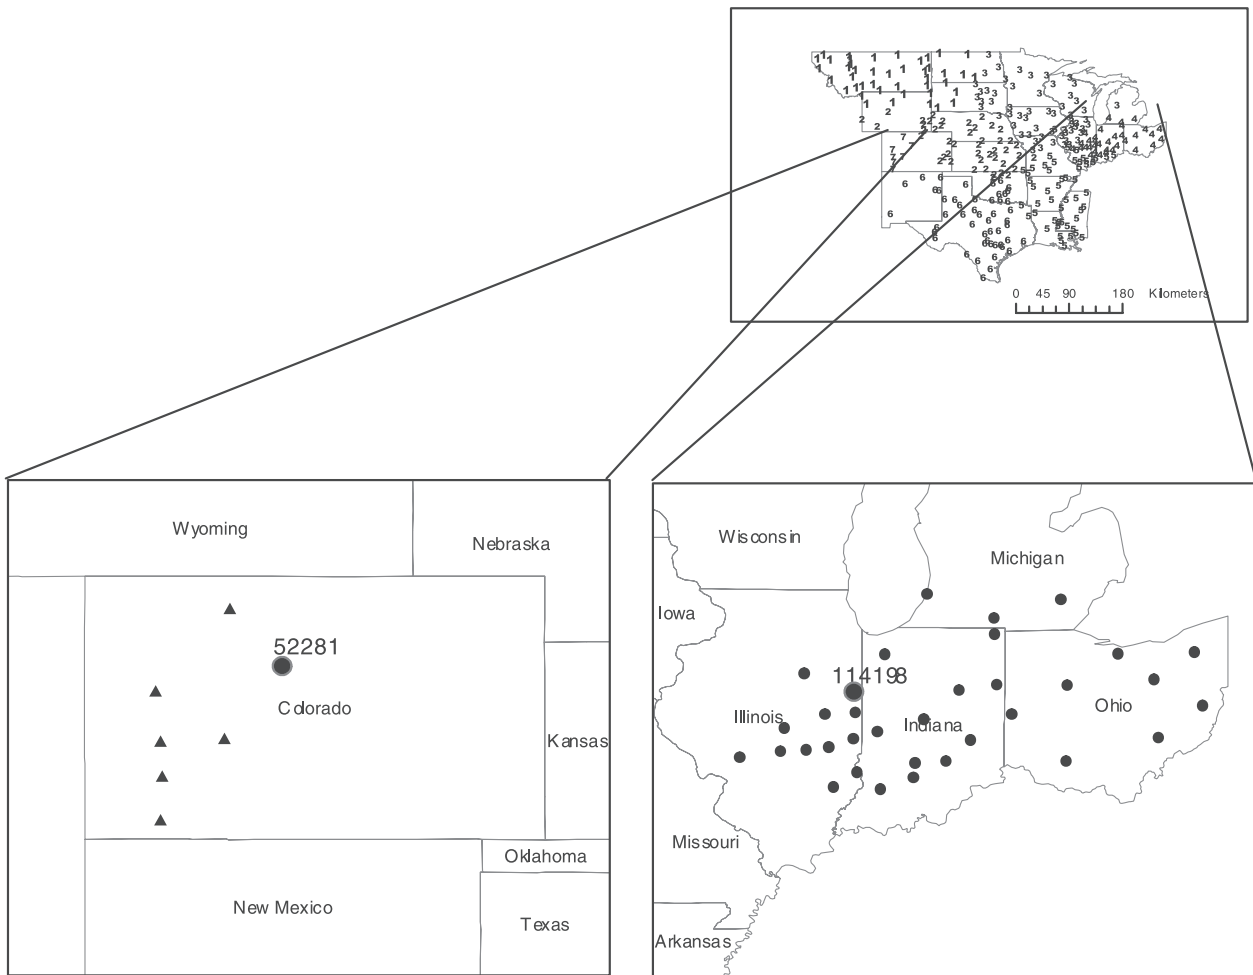


Figure 2. Map depicting the 21-state area of interest in this study (top right graph), and the two focus areas (regions 4 and 7) and their corresponding stations (bottom graphs). The numbers indicate stations grouped by region. The two stations of focus (114198 in region 4 and 52281 in region 7) are highlighted.

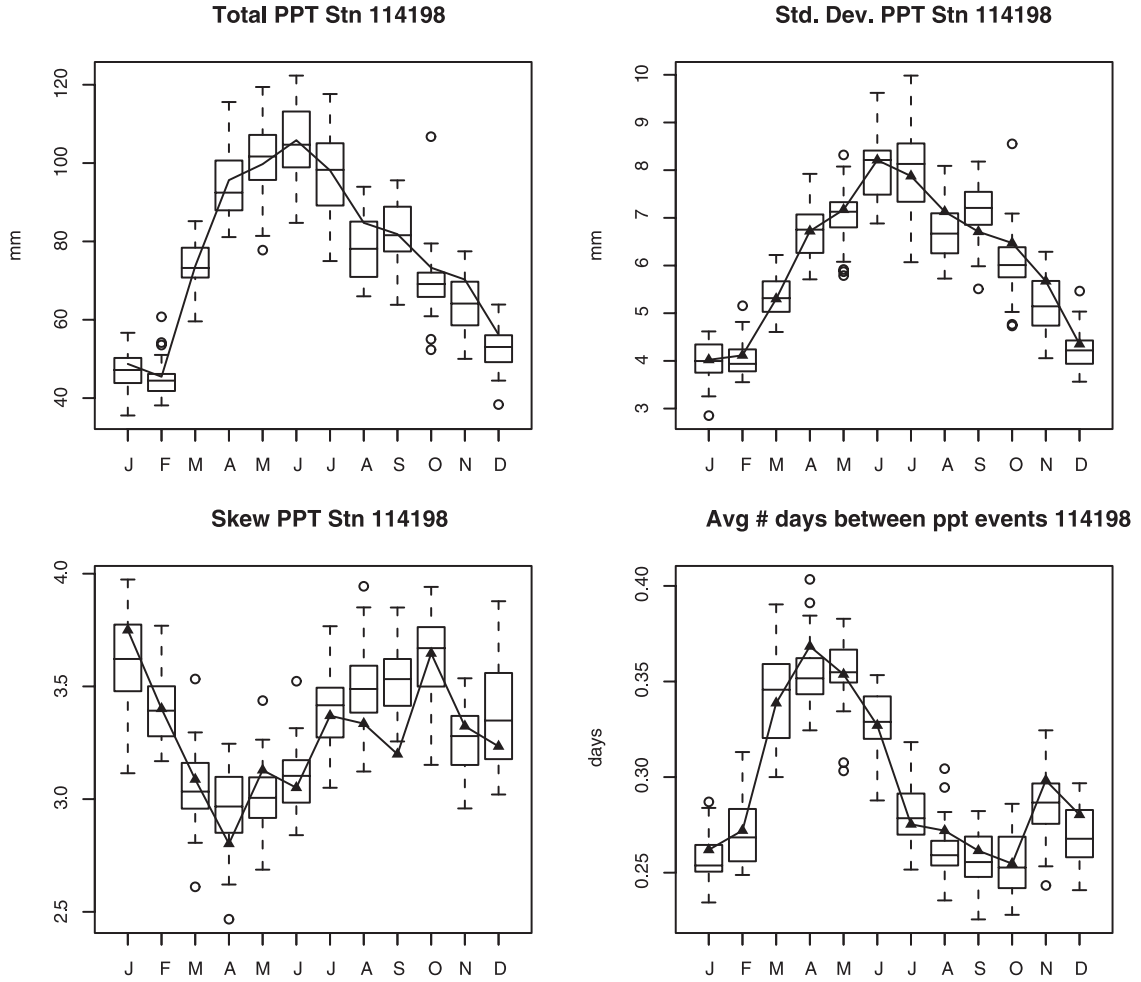


Figure 3. Box plots of precipitation statistics for station 114198 in region 7, with constant $\lambda_w^i = 1.0$. Top left graph is total monthly precipitation, top right graph is the standard deviation of daily precipitation, bottom left graph is the skewness of daily precipitation, and the bottom right graph is the average number of days between precipitation events for the month. The solid line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

[24] The weekly anomalies are computed from the historic daily data for each variable (e.g., Ψ_{jt}^i is PPT_{jt}^i , TMN_{jt}^i , or TMX_{jt}^i , where i is the year, j is the station, and t is the day in the week) first by computing the regional weekly means. For simplification, assume the year can be broken into w , $d(w)$ day periods (e.g., weeks); thus the regional weekly means of the entire record are given as

$$\bar{\Psi}_w = \frac{1}{Nmd} \sum_{i=1}^N \sum_{j=1}^m \sum_{t=1}^d \Psi_{j,t}^i \quad (7)$$

while the regional weekly means for each year are given as

$$\Psi_w^i = \frac{1}{md} \sum_{j=1}^m \sum_{t=1}^d \Psi_{j,t}^i \quad (8)$$

Regional, weekly anomalies are computed as $\delta_w^i = \Psi_w^i - \bar{\Psi}_w$, which are then ranked by week according to their relative magnitude. For example, regional minimum temperature anomalies could be ranked from lowest to highest for all weeks w and N years such as $\mathbf{R} = \{\delta_1^{1:N}; (1945$

$-7.8, 1962 -7.2, \dots, 1967 +0.5, 1957 +4.5, 1987 +8.0)$; $\delta_2^{1:N}; (1946 -6.8, 1945 -5.0, \dots, 1896 +4.0, 1937 +6.8)$;... $\delta_w^{1:N}; (\dots); \delta_{52}^{1:N}; (1956 -6.8, 1945 -6.1, \dots, 1978 +6.8)$. Here the coldest first week of January was 1945 (-7.8°C below the week 1 mean), while the warmest first week of January was 1987 ($+8.0^\circ\text{C}$ above the week 1 mean). The sorted list for each week is then assigned a rank index where the first index I_1^1 , corresponding to δ_1^1 , is given a value of 1; $I_1^2 = 2$; $I_1^3 = 3$; and the last index $I_1^N = N$, which corresponds to δ_1^N in the ranked list. To generate the subset of years, \mathbf{n}_w^T for each of the 52 weeks, $u \sim U(0,1)$ random numbers are used to compute index values between 1 and N according to

$$I_w^{1:N} = INT[N(1-u)] + 1, \quad (9)$$

which directly correspond to certain years in the ranked list, \mathbf{R} . With this technique, if N , $U(0,1)$ random numbers are selected, there is equal probability that any year is part of the \mathbf{n}_w^T subset. Likewise, there is the probability that years could be repeatedly selected and appear multiple times in the \mathbf{n}_w^T subset and are thus more likely to contribute to the daily

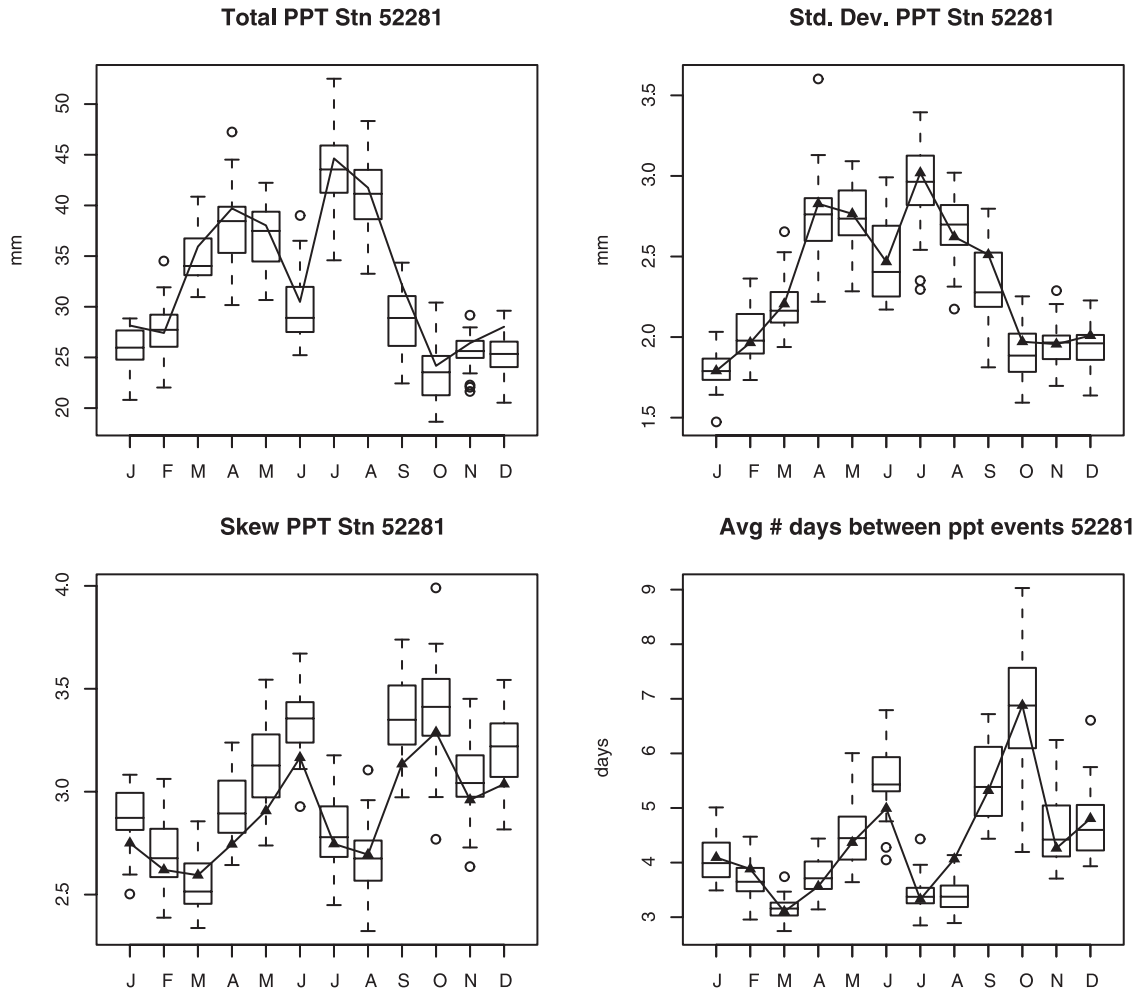


Figure 4. Same as Figure 3, but for station 52281 in region 4.

weather simulations. This is the main difference from when resampling is done where n_w^τ is prescribed to be all N years.

3.3. Paired Ranking

[25] The above applies for single variables, while paired ranking is slightly more involved. Departures from the weekly mean for a pair of climate variables might include minimum temperature and precipitation, maximum temperature and precipitation, or minimum and maximum temperature. A relative distance to the mean for each paired combination is computed which is then used to determine the year of the most extreme pair (e.g., the warmest and wettest, the warmest and driest, the coldest and wettest) relative to “normal” conditions. This is done by first generating standardized weekly anomalies of each variable,

$$Z_w^i = (\Psi_w^i - \bar{\Psi}_w) / \left\{ E \left[(\Psi_w^i - \bar{\Psi}_w)^2 \right] \right\}^{1/2} \quad (10)$$

for all years $i = 1$ to N and weeks $w = 1, 52$. The anomalies for each climate variable are simply summed to give

$$\delta_w^i = Z(1)_w^i + Z(2)_w^i \quad (11)$$

where $Z(1)$ and $Z(2)$ indicate the appropriate variable (e.g., PPT, TMN, TMX). By means of example, if the normalized precipitation anomaly is negative (e.g., a dry year) and

the temperature anomaly is negative (a warm year), then the $\delta_w^i = Z(PPT)_w^i + Z(TMN)_w^i$ would be a large negative number, and would be ranked toward the bottom of the list. Once the rank list is generated from the paired variables, the subset of years n_w^τ can be generated in a manner similar to the single variable case.

3.4. Biasing the Ranked List

[26] If the ranked lists are sampled with a uniform random number to generate the I_w^i index, then each year has an equal probability of making up the n_w^τ subset. However, it would be advantageous to select the index in such a way as to bias certain years over others. The integer function of equation (9) becomes

$$I_w^{1,N} = INT \left[N \left(1 - u^{\lambda_w^i} \right) \right] + 1 \quad (12)$$

where λ_w^i is a shape parameter of the function. A value of $\lambda_w^i = 1.0$ gives equal probability that any year could be selected for the given week w . Different values of the shape parameter bias certain index values above or below the mean when queried with a $u \sim U(0,1)$ random number. If years are ranked from coldest (with an index value of 1) to warmest (with an index value of N), then shape parameter values greater than 1.0 tend to bias the selection of warm years, while shape parameter values less than 1.0 tend to bias the

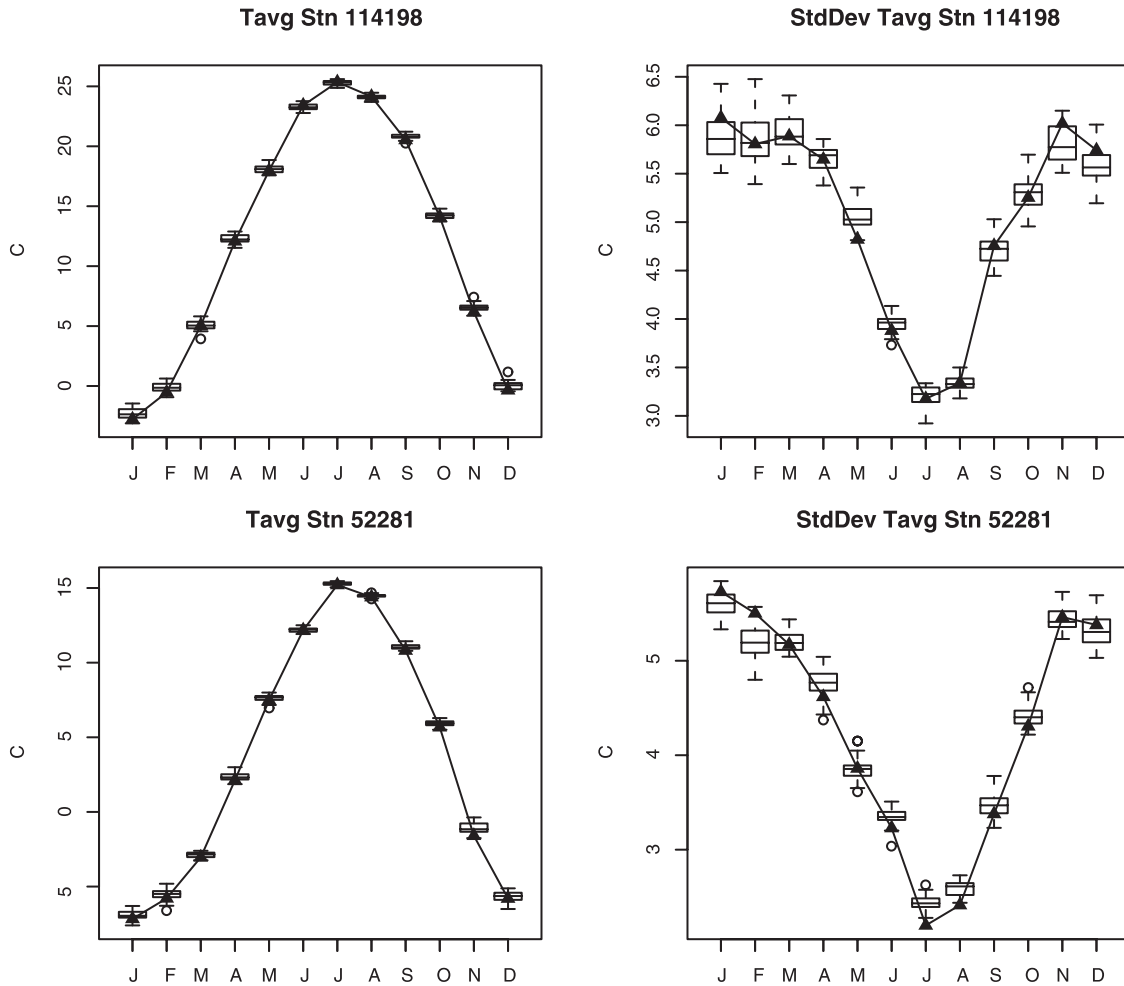


Figure 5. Box plots of daily average temperature for station 114198 (top left graph) and station 52281 (bottom left graph), and the standard deviation of the daily average temperature for the same two stations (right graphs), given constant $\lambda_w^i = 1.0$. The solid line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

selection of cold years. Figure 1 describes an example using a 50-year record of average temperatures ranked from coldest (with an index value of 1) to warmest (with an index value of 50). If a λ_w^i value of 2.0 is used, then the expected value of the function described in equation (12) is approximately 33 when repeatedly queried with $U(0,1)$ random numbers, suggesting a bias toward the selection of warm years. Likewise, a value of 0.66 would bias cold years with an expected value of approximately 17 (Figure 1).

3.5. Introducing Interannual and Intra-annual Variability

[27] The shape of the curve corresponding to equation (12) can be modified through a manipulation of the λ_w^i parameter according to the week in the year (intra-annual change) and along years (interannual change). The shape parameter can be altered according to

$$\lambda_w^i = \lambda_{w-1}^{i-1} (1 + \rho_i) (1 + \beta_w) \tag{13}$$

where β_w is the intra-annual trend parameter and ρ_i is the interannual trend parameter, which are adjusted for each

week and for each year, respectively. Choosing different values of β_w and ρ_i for certain weeks and years produces values of λ_w^i yielding changes in the shape of the response curve, and when the integer function (equation (12)) is queried, its index values are used to select particular years from the ranked list. In the case of the three variables used here, the various ranked lists (e.g., TMN, TMX, PPT, PPT-TMN, PPT-TMX, and TMN-TMX) can be assigned to specific weeks in the year. If a scenario of warmer-wetter winters for weeks 1 through 8 and 50 through 52 (corresponding to 10 December to 25 February) were to be generated, then these weeks would be resampled from the warmer-wetter ranked list. Because values of λ_w^i can be uniquely defined for each week, then the middle weeks (weeks 3 and 4, or mid-January) could have higher values of λ_w^i relative to the other winter months, which would likely produce warmer and wetter days in mid-January relative to the other winter weeks.

[28] Of course, this depends on the attributes of the underlying data. It is likely that a region’s climate is such that there is greater variability during certain times of the year, which is reflected in the historic climate data. If the

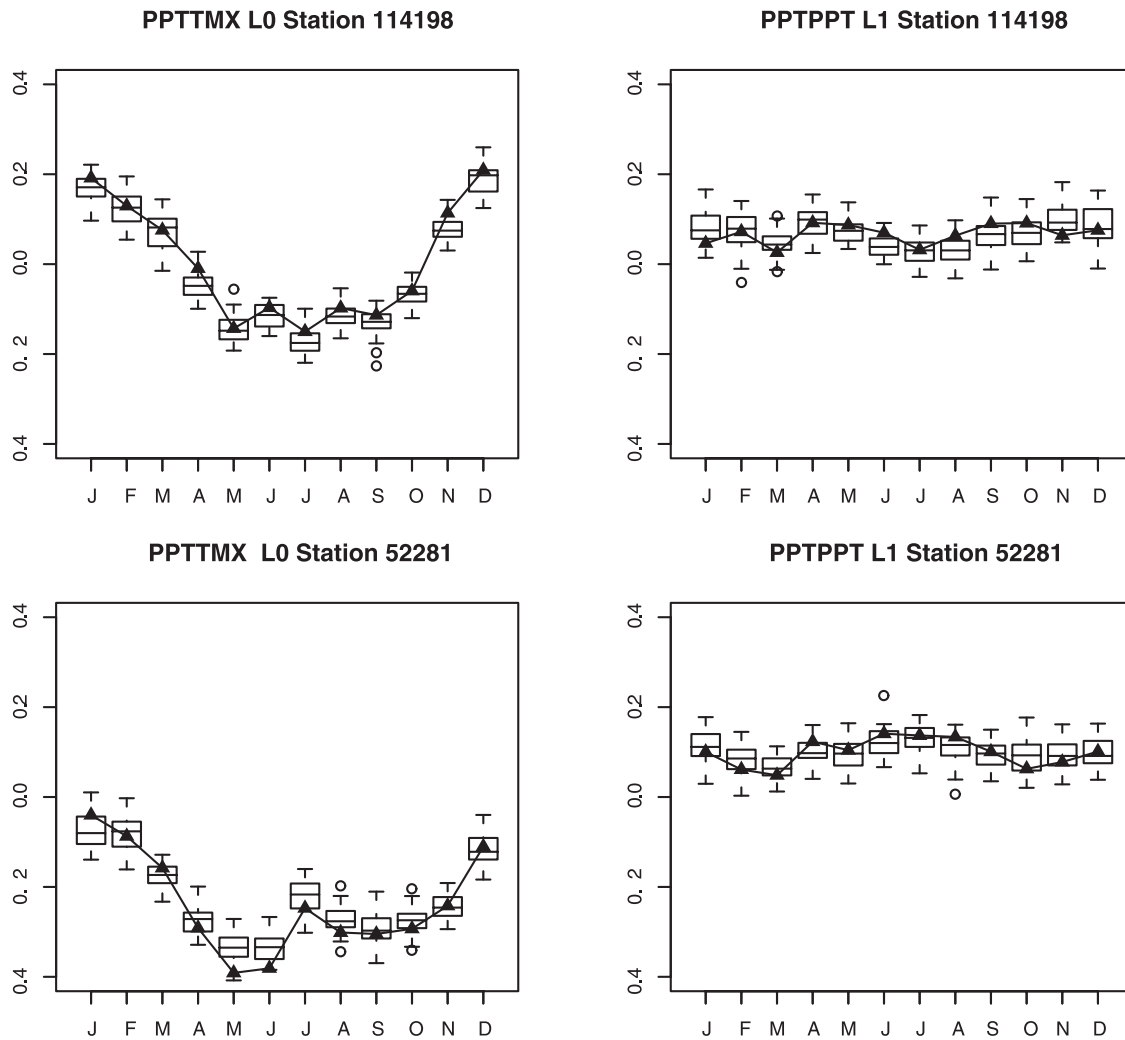


Figure 6. Box plots of lag-0 correlation between precipitation and the daily average maximum temperature at station 114198 and 52281 (left graphs); and the lag-1 auto correlation of precipitation for the same two stations (right graphs), given constant $\lambda_w^i = 1.0$. The solid line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

standard deviation of temperature is greater in January than it is in December and February, then a conditioned set generated with the same scenario function (or the same values of ρ_j and β_k , for the weeks of December, January, and February) would likely yield greater temperature changes for January when compared with December and February.

[29] With this technique, different weeks in the year can be assigned to the different ranked list types, leading to any of number of climate change scenarios with particular attributes. One arguable drawback of this climate scenario generating technique is that because all variables are grouped, the choice to select, for example, higher minimum temperatures will likely produce trends in the other variables. For when a particular site experiences above average minimum temperatures, this likely corresponds to increases or decreases in precipitation, which will then be a characteristic of a “higher minimum temperature” scenario. Thus true partial derivative experiments (e.g., generating climate scenarios with higher minimum temperatures but that hold the other climate variables constant) are not possible;

however, the derived scenarios will be realistically consistent with the observed climate relationships. Along the same lines, if there is a positive correlation between maximum temperature and precipitation in the observed data during winter, it may be difficult to produce scenarios with an increase in maximum temperature and a decrease in precipitation for that season.

4. Generating Precipitation and Temperature Fields

[30] The K - nn algorithm described above was used to simulate daily minimum and maximum temperature and precipitation at several stations in the central Midwest and Rocky Mountain region of the United States. These two regions are shown as regions 4 and 7, respectively, based on the data requirements of another research project focused on the Central Great Plains region of the United States (Figure 2). Region 4 had 34 stations, while region 7 had seven. These two regions were chosen for the diversity of

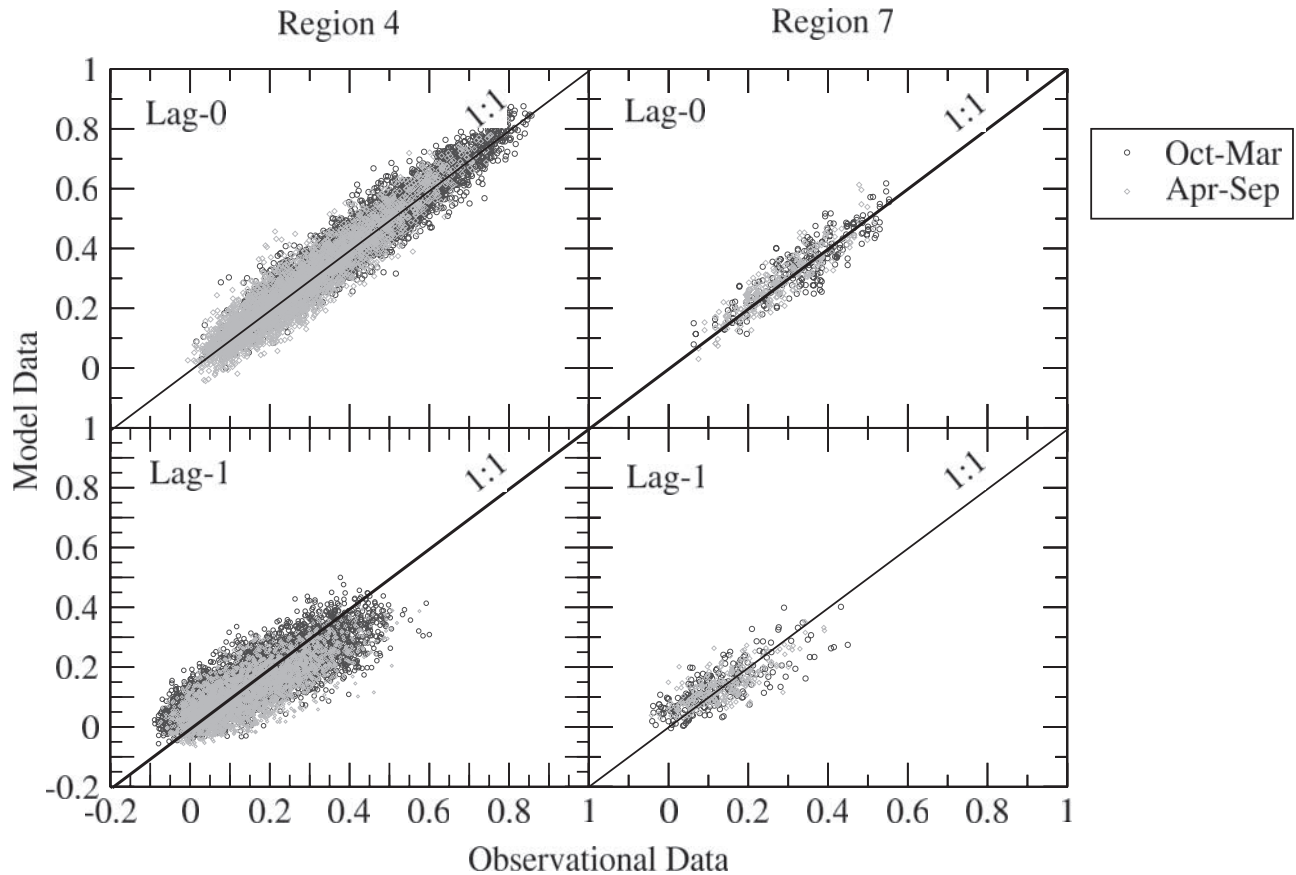


Figure 7. Scatterplots of lag-0 cross correlations of precipitation for region 4 (upper left graph) and region 7 (upper right graph) stratified by season; and the lag-1 cross correlations for region 4 (lower left graph) and region 7 (lower right graph) also stratified by season. Simulations were for $\lambda_w^i = 1.0$ constant values.

their climate. Daily observational data for a 58-year period spanning 1930 to 1987 were taken from the VEMAP archive [Kittel *et al.*, 1997], consisting of quality-controlled meteorological data corresponding to the stations indicated in Figure 2.

[31] A total of 900 independent years of weather data (PPT, TMN, TMX) were generated and a set of statistics computed for each month of the year based on this synthetic daily data. The statistics include means, standard deviations, skewness coefficients, cross correlations between the variables, and cross correlations between stations, and are summarized with box plots. Although the weather simulations were made on a daily timescale, the statistics from the daily data have been aggregated to the monthly timescale. In all the box plot figures shown below, the following holds: The box in the box plots (e.g., Figure 3) indicates the interquartile range of the simulations, and the whiskers show the 5th and 95th percentile of the simulations, while the dots indicate values outside this range. The horizontal lines within the box indicate the median value, and the solid lines join the values of the statistic from the observed data. Typically, if the statistics of the observed data fall within the box, it indicates that the simulations adequately reproduced the statistics of the historical data. For illustrative purpose, the results at station 114198 in

region 4 and station 52281 in region 7 are highlighted (Figure 2).

4.1. The Special Case of $\lambda_w^i = 1$: Series Yielding Statistics of the Historic Data

[32] The first experiment was to reproduce the statistics of the historic data, and to do this, fixed values of β_w and ρ_i for all weeks and years were prescribed, resulting in $\lambda_w^i = 1.0$ for all w and i . Thus all $[(w + 1) \times N] - 1$ days are likely to be part of the pool of neighbor candidates. The upper left graph of Figure 3 is the box plot of total monthly precipitation for station 114198 in region 4, while the upper left graph of Figure 4 is the same for station 52281 in region 7. It can be seen that the simulations adequately reproduced the historical monthly precipitation, also suggesting that the annual cycle in the precipitation was captured. There was a slight underestimation of precipitation at station 114198, particularly for April and September, but overall the performance was good. Note the two peaks of precipitation in the spring and late summer at station 52281, with a spring snow-regime maxima and a later summer rainfall maxima associated with the North American monsoon. It should be noted that the algorithm was not designed to explicitly model monthly statistics, unlike the parametric models that are fitted separately

Table 1. The Region 4 Annual Mean of TMN, TMX, and PPT Computed From the 34 Stations From the 900 Independent Series, and for Different Values of λ_w^i , Using an Average Temperature Ranked List Developed From Region 4 Historic Data^a

$\lambda_{k,j}$	TMN, °C	TMX, °C	PPT, mm
Historic	5.2	16.6	941
0.6	4.2; $-\Delta 1.0$	15.6; $-\Delta 1.0$	882; -6%
1.0	5.3; $+\Delta 0.1$	16.7; $+\Delta 0.1$	921; -2%
2.0	6.9; $+\Delta 1.7$	18.4; $+\Delta 1.8$	1007; $+7\%$
3.0	7.8; $+\Delta 2.6$	19.3; $+\Delta 2.7$	1058; 12%

^aAlso included are the region 4 historic averages for these same variables. For the scenario rows, the first numbers are the absolute values, and the second numbers are the deviations from the historic (precipitation is the percent difference).

for each month [Richardson, 1981]. Given this, the performance was impressive.

[33] The upper right graphs of Figures 3 and 4 are the standard deviations of daily precipitation for each month, and the simulations adequately reproduced this statistic, which was greatest in the spring at station 114198 and in the spring and late summer at station 52281. Skewness, an

important attribute of daily precipitation that describes the symmetry of the distribution of rainfall, is plotted in the lower left graphs of Figures 3 and 4 for these two stations. Typically precipitation data are positively skewed. The simulations generally capture the historical skew, although there is a slight overestimation for a couple of months at both stations. Note those months where the model over-predicted the magnitude of the skew (particularly for station 52281: January, April, May, June, September, December; Figure 4) generally corresponded to months where the model had a tendency to over-select dry days leading to less precipitation (station 52281: January, April, June, September, and December; Figure 4).

[34] To illustrate a measure of persistence, the average number of days between precipitation events (e.g., days in the historic record where precipitation is greater than zero) for each month is illustrated in the lower right graphs of Figures 3 and 4. This statistic gives insight into the ability of the $K-nn$ technique to reproduce dry and wet spell lengths. This statistic is significant in several decision models, especially in agricultural yield models and also in watershed and reservoir operation models. There is significant seasonal variability in the average number of days between rain events at both stations. Station 52281 in region 7 shows

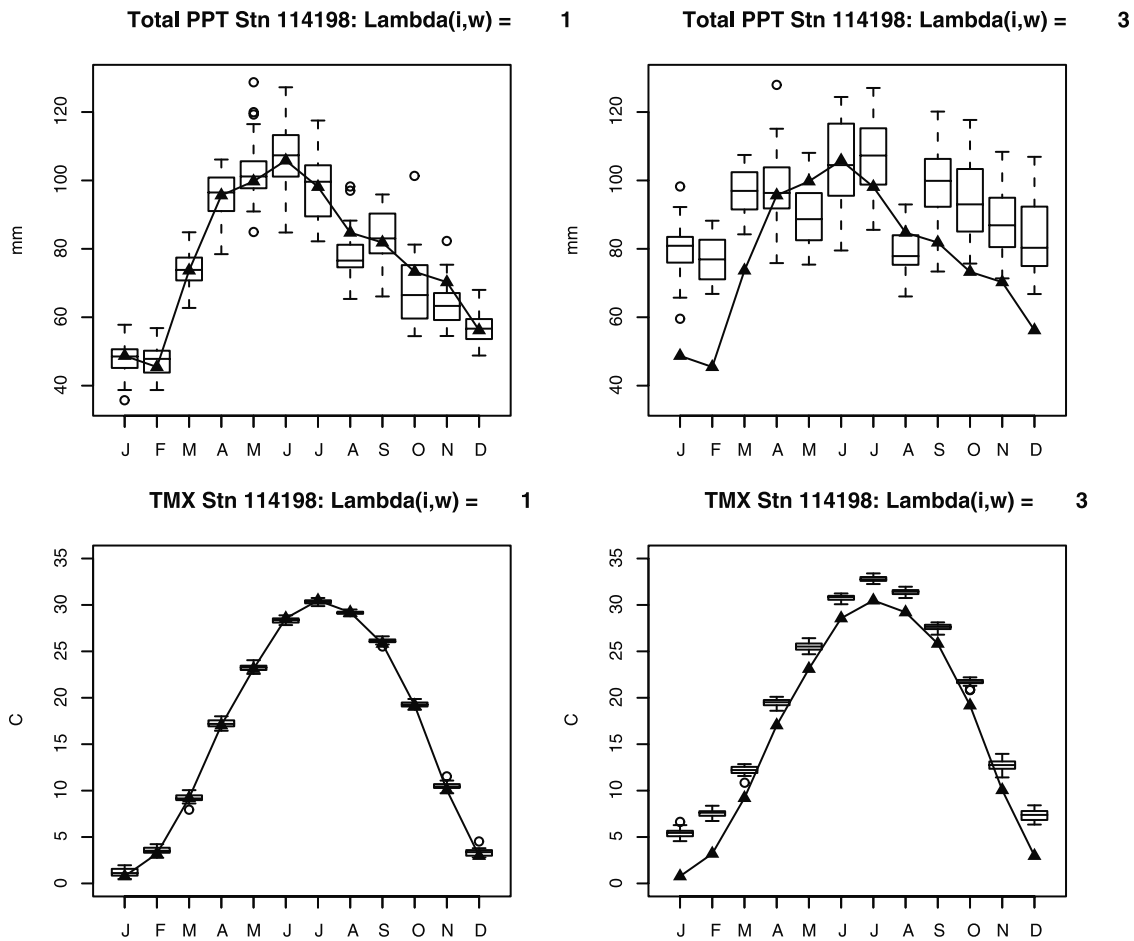


Figure 8. Box plots of total monthly precipitation (top left graph) and daily average temperature (bottom left graph) for station 114198, given $\lambda_w^i = 1.0$ constant values. Right graphs are the same statistics, but for constant values of $\lambda_w^i = 3.0$. The solid line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

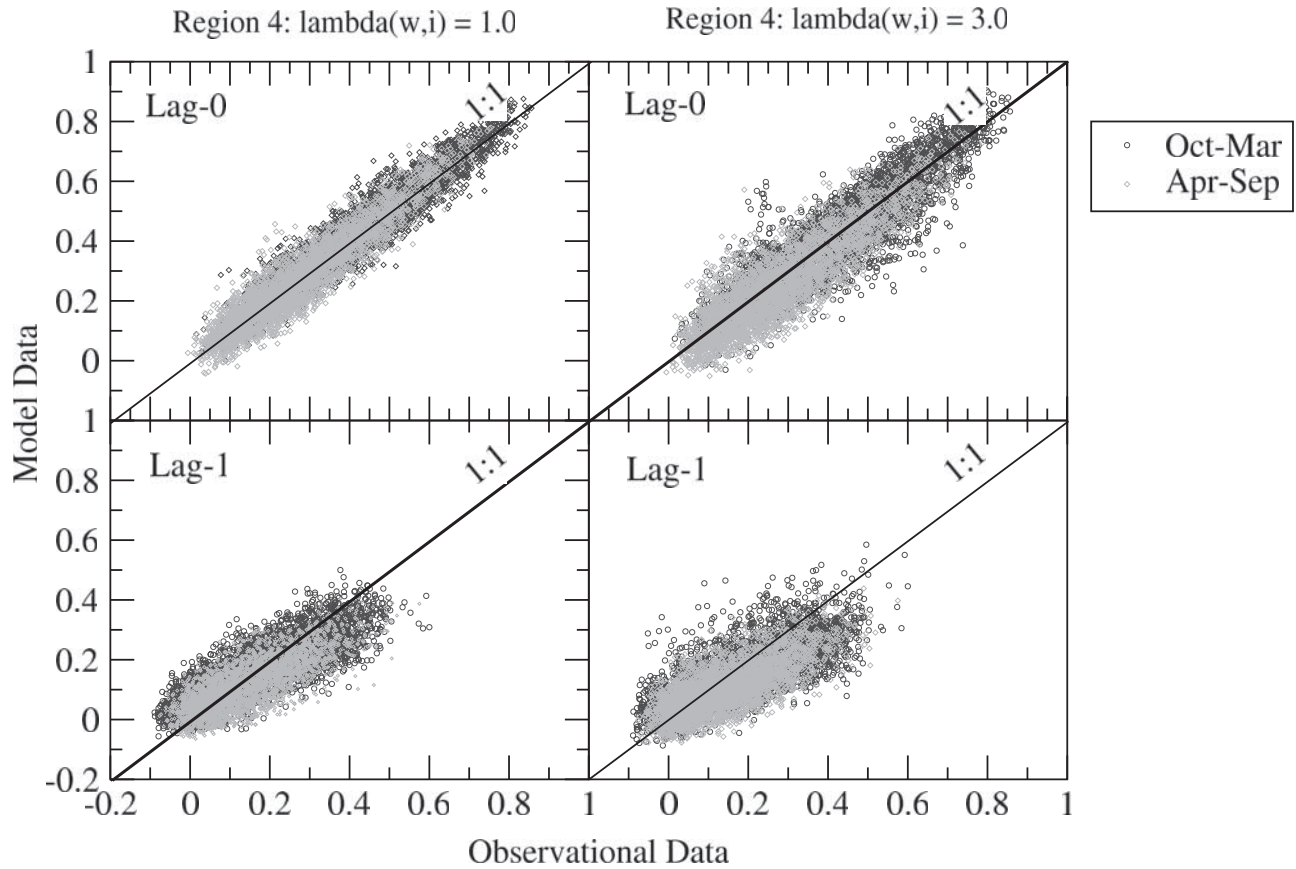


Figure 9. Scatterplots lag-0 (upper left graph) and lag-1 (lower left graph) cross correlations of precipitation for region 4 with constant $\lambda_w^i = 1.0$ values stratified by season; and similar scatterplots of lag-0 (upper right graph) and lag-1 (lower right graph) cross correlations of precipitation for region 4 with constant $\lambda_w^i = 3.0$.

complex seasonal variability composed of an early spring minima in March and a summer minima in July and August corresponding to the North American monsoon in the southwestern United States. Station 114198 in region 4 exhibits increased occurrence of wet days, given as fewer days between precipitation events in the spring.

[35] Figure 5 shows the box plots of daily average temperature ($TAV = 0.6TMAX + 0.4TMN$) for each month (left column) for stations 114198 and 52281, derived as a weighted average of the minimum and maximum temperatures. Figure 5 also includes the standard deviation of the daily temperature for each month of both stations (right column). The simulations reproduce both of these statistics very well. Both stations exhibit higher temperature variability in the winter months, and the derived series were able to capture these historic observations.

[36] To investigate how well the spatial and temporal dependence is captured in our model, we present lag-0 and lag-1 correlations between the variables at a station and spatial correlations between the variables at different stations (cross correlations). For the correlation box plots, the y-axis has been kept constant (-0.4 to 0.4) to better compare the relative strength of the correlation statistics to one another.

4.1.1. Correlation Among Variables at a Station

[37] Since the same day's weather is selected for all stations, we would expect strong lag-0 correlations among

variables at a station. The lag-0 correlations (e.g., PPT-TMN, PPT-TMX, TMN-TMX, and TAV-PPT) were well preserved by the model, with only the correlations between precipitation and maximum temperature presented. One might expect some degree of negative correlation between precipitation and temperature, as precipitation likely occurs with lower temperatures. This seems to be true for the PPT and TMX correlations at station 52281, particularly for the summer months (Figure 6, bottom left graph). Interestingly, station 114198 shows a positive correlation between maximum temperature and precipitation in the winter, which transitions and becomes negative in the summer (Figure 6, upper left graph). In the Central Great Plains cold season, maximum temperatures are often observed to be higher on rainy days, while during the warm season just the opposite is true. Figure 6 also includes lag-1 autocorrelation for precipitation (right column). Somewhat surprisingly, there is little lag-1 correlation in precipitation at both stations; the simulations adequately capture this trend. Lag-1 autocorrelations of temperature at a station were well preserved (not shown).

4.1.2. Station Cross Correlations

[38] To investigate how well the spatial dependencies of precipitation were captured by the model, the station cross correlations are summarized with scatterplots in Figure 7 [Wilks, 1998]. Both lag-0 ($\text{corr}[D_t^k, D_t^l; k \neq l]$) and lag-1 ($\text{corr}[D_t^k, D_{t+1}^l; k \neq l]$), for the collection of

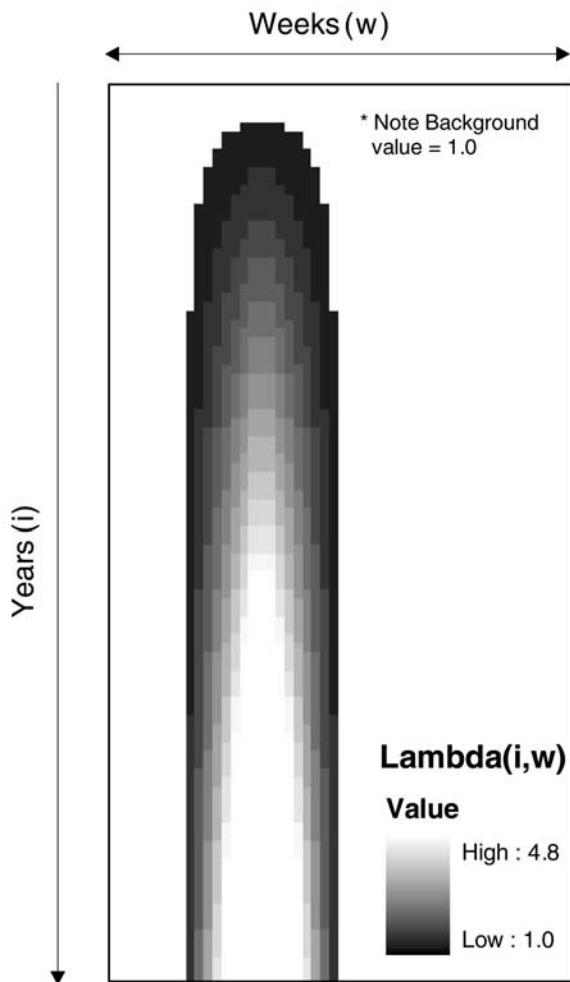


Figure 10. Plot of the temporally varying (by year i and month w) λ_w^i surface for the “warmer-drier” spring scenario.

stations ($k, l = 1$ to m) are presented and are stratified by the winter (October, November, December, January, February, March) and summer seasons (April, May, June, July, August, September), where D_t^m is the day’s weather at each station. There are $m = 35$ stations in region 4 and $m = 7$ stations in region 7, yielding $(m(m - 1)/2) \times 2 \times 12$ combinations of station pairs, each both leading and lagging for all months (e.g., region 4: 13464; region 7: 504). Not surprising, since the K - nn algorithm chooses the same $t + 1$ day for all stations, the lag-0 cross correlations were strongly preserved for both regions (top two graphs, Figure 7).

[39] The region 4 domain spans a larger east-west transect, and there is a tendency for synoptic-scale precipitation systems to move from west to east across the region, which is reflected in the slightly higher lagged cross correlations, particularly for the winter months (bottom left graph, Figure 7). The summer months are characterized by more localized, convective events reflected in lower station cross correlations. Note the model bias, which tended to overestimate those stations with low cross correlations, while underestimating those

stations with higher lag-1 cross correlations in both regions. The region 7 domain spans a smaller geographic region, and there is little difference in the seasonal lag-1 cross correlations (bottom right graph, Figure 7), and the model adequately produced these lag-1 cross correlations with the above noted biases.

[40] In summary, the use of Mahalanobis distance proved to be an adequate methodological approach within this K - nn resampling framework, as the algorithm led to new realizations of multistation meteorological data that largely reproduced the spatial and temporal statistics of the historic data, given values of $\lambda_w^i = 1.0$ for all weeks and years.

4.2. Cases of $\lambda_w^i \neq 1$: Generating Alternative Scenarios

[41] To test the ability of the algorithm to generate alternative climate scenarios using different values of λ_w^i , region 4 (Figure 2) was selected and a simple increasing average temperature (TAV) scenario was contrived. Weekly average temperatures were computed as a weighted average of daily minimum and maximum temperatures to generate the ranked list. For all stations in region 4, the average weekly temperature means and anomalies were computed for each year, which were then ranked from coldest (index of 1) to warmest (index of 58) according to the method of section 3.2.

[42] Again, 900 independent years of synthetic data were generated for each $\lambda_w^i = 0.6, 2.0,$ and 3.0 for all w, i scenarios. Recall that the $p = 3$ variables used in the K - nn algorithm are PPT, TMN, and TMX, while TAV was only used to rank the list of years from which to resample. Synthetic series corresponding to $\lambda_w^i = 0.6$ should result in cooler regional temperatures, while synthetic series generated using values of $\lambda_w^i = 2.0$ and 3.0 should yield warmer regional temperatures.

[43] Table 1 includes the annual historic average temperature and total precipitation, and the annual average regional temperature and precipitation from all scenarios ($\lambda_w^i = 0.6, 1.0, 2.0,$ and 3.0) and shows that the sets with $\lambda_w^i \neq 1.0$ did lead to changes in the average regional temperature. Although the biasing was only based on the regional average temperature, this also led to changes in regional precipitation. Recall from Figure 6 that maximum temperature and precipitation are slightly negatively correlated for the summer months and positively correlated for the winter months; thus this result was not surprising. Resampling from the TAV ranked list with $\lambda_w^i = 2.0$ and 3.0 resulted in regional TMN and TMX increases, while the synthetic series corresponding to $\lambda_w^i = 0.6$ resulted in TMN and TMX regional temperature decreases (see Table 1). Regional average precipitation decreased by 6 percent for the $\lambda_w^i = 0.6$ case, while it increased by 7 percent for the $\lambda_w^i = 2.0$ case and by 12 percent for the $\lambda_w^i = 3.0$ case.

[44] Box plots similar to those presented for the $\lambda_w^i = 1.0$ case are presented for both the $\lambda_w^i = 1.0$ and $\lambda_w^i = 3.0$ scenarios for station 114198, which includes both total monthly precipitation (top row) and average maximum daily temperature (bottom row) based on all 900 series (Figure 8). These plots clearly show that the strategic resampling can produce alternative scenarios with different statistical attributes. Maximum temperature has increased for all months (Figure 8, bottom graphs), while precipita-

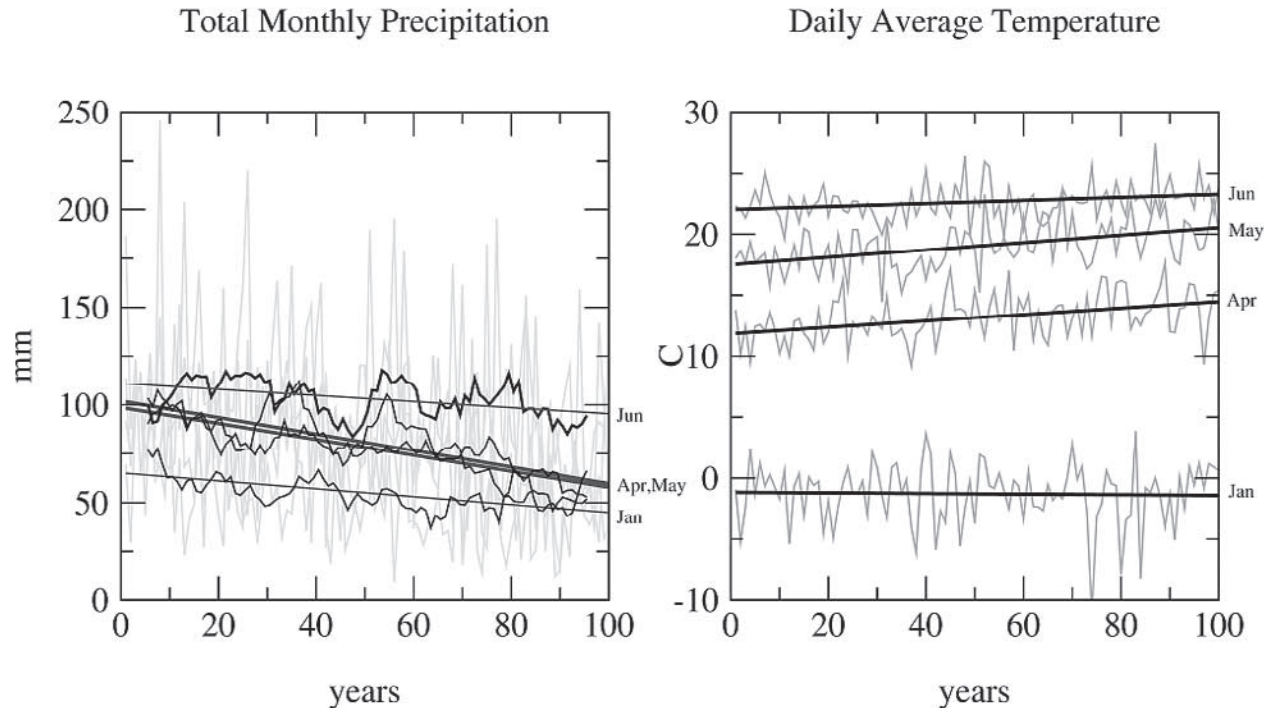


Figure 11. The 100-year times series of total monthly precipitation (left graph) for the “warmer-drier” spring scenario for select months. The shaded lines are the regionally averaged time series, while the solid lines are the 10-year moving averages, and the straight lines are the linear trends for January, April, May, and June. The right graph shows the daily average temperature per month and includes the regionally averaged time series (shaded lines) and the linear trend of this series for the 100 years.

tion has increased most notably for the winter months (Figure 8, top graphs). In the summer months, where precipitation has a slight negative correlation with maximum temperature (Figure 6, top left graph), precipitation decreases for May and August for the $\lambda_w^i = 3.0$ case relative to the $\lambda_w^i = 1.0$ case. Surprisingly, precipitation in the months of June, July, and September did not decline, despite the historic observation of negative correlation with temperature. Recall that the ranked list and thus this biasing is only based on average temperature. The negative correlation between precipitation and maximum temperature is not strong, so there are likely years in which this negative correlation did not occur, from which the neighbors were selected.

[45] For these same scenarios, the station cross correlations are summarized with scatterplots (Figure 9) similar to those presented in Figure 7 but only for region 4, and includes both the $\lambda_w^i = 1.0$ scenario (left column) and the $\lambda_w^i = 3.0$ scenario (right column). There was some difference in the lag-0 cross correlations between the $\lambda_w^i = 1.0$ and $\lambda_w^i = 3.0$ cases, implying some change in the lag-0 cross correlation of the new scenario, although slight. Likewise, there was some shift in the lag-1 cross correlation for the $\lambda_w^i = 3.0$ scenarios.

4.3. Case of $\lambda_w^i \neq 1$: Varying β_w and ρ_i by Week and Year

[46] The second test of the scenario-generating algorithm was to produce an increasing trend of warmer-drier springs throughout region 4 for a single realization of a 100-year

time series. This is the kind of scenario whose data could be directly used in an IA model. To generate such a scenario, a colder-wetter/warmer-drier ranked list was generated according to the method of section 3.2, where higher index values corresponded to warmer-drier years. Average temperature (TAV) was used as the temperature variable, and thus the paired variables TAV- PPT were used to generate the ranked list. This ranked list was then strategically sampled to produce n_w^i subsets from which to apply the K - nn algorithm.

[47] In this case, the β_w and ρ_i were varied by week and by year, resulting in different values for λ_w^i according to equation (13). Thus when equation (12) was randomly sampled with these different λ_w^i values, different indices I_w^i were generated that biased the selection of springtime, warm-dry years. Figure 10 is a plot of the resulting λ_w^i surface for this scenario and shows the increasing values of λ_w^i for the spring months, with late April and early May corresponding to the peak of the λ_w^i values in an attempt to produce a scenario with the maximum increases in temperature and decreases in precipitation to occur during this period. Recall that precipitation shows a spring transition from being positively correlated with temperature in the winter to being negatively correlated with temperature in the summer (Figure 6, top left graph). Thus it should be fairly straightforward to generate a “warmer-drier” spring scenario. Also, a long-term trend of increasing values of λ_w^i was imposed by adjusting the interannual trend parameter ρ_i in an attempt to mimic a gradual increase of warmer-drier springs throughout the region.

[48] The 100-year precipitation series are given in Figure 11 (left graph) for select months that include the late winter, spring, and early summer periods, computed as the average of the 34 stations in region 4 using the above scenario. The changes in the April and May precipitation total over the 100-year series was the greatest, with a long-term decline in average precipitation (first decade average monthly total, 90 mm and 103 mm; last decade average monthly total, 53 mm and 66 mm, respectively). The June decline was more modest, with a first decade average total of 100 mm and a last decade average total of 95 mm. Note that although the scenario did not bias the selection of any particular year in January, there was a modest decline in January precipitation over the 100 year simulation (Figure 11, left graph).

[49] Figure 11 includes plots of the 100-year time series of daily average temperature as a weighted average of minimum and maximum temperature (right graph). The long-term trend of increasing temperatures for April and May is highlighted in this figure, with a last decade increase of approximately $+2.7^{\circ}\text{C}$ and $+3.0^{\circ}\text{C}$ for April and May over the 100-year time period. There was little to no trend observed in the January and June temperatures.

5. Summary and Conclusions

[50] A modified K - nn resampling technique was described that simulates regional daily weather, thus largely preserving important spatial and temporal dependencies of the observed climate. This method has an operational advantage over traditional resampling models that make use of the Euclidian distance. In the Euclidian distance approach the variables have to be standardized (i.e., mean removed and divided by the standard deviation) before the distances are computed. This is because variables with higher magnitudes disproportionately influence the neighbor selection. The use of Mahalanobis distance obviates the standardizing of the variables. Likewise, the use of the covariance matrix in the Mahalanobis distance computation (equation (5)) weights the variables as per their covariance (or correlation), i.e., strongly correlated variables get less weight.

[51] The K - nn algorithm was demonstrated by simulating weather sequences at multiple sites over two regions in the United States. The performance of this technique in reproducing the temporal and spatial statistics was good. Capturing the spatial and temporal dependencies is very important when using simulated weather sequences in IA models. Strategies to adapt this K - nn technique for strategic resampling that could be conditioned upon a large-scale climate signal (e.g., ENSO) or some contrived climate change scenarios that could be used in IA models were also developed. These strategies were used to successfully generate several climate change scenarios, e.g., regional temperature increases and increasingly warmer-drier springs for a large geographic region. The nearest-neighbor methodology developed here for simulating regional climate provides a flexible framework with no assumptions of the underlying functional form of the dependencies or the probability density functions of the variables. It is very easy to understand and implement.

[52] A common shortcoming of this and other bootstrap methods is that they do not produce values not seen in

the historical record. However, they produce a rich variety of sequences of daily weather. The variety is important for a majority of decision models and IA models. *Rajagopalan and Lall* [1999] describe a strategy that allows nearest-neighbor resampling with perturbations of the historical data in the spirit of the traditional autoregressive models, i.e., conditional expectation with an added random innovation. Their suggested scheme perturbs the historic data points in the series with innovations that are representative of the neighborhood, and will “fill in” between historic values as well as extrapolate beyond the sample. We plan to implement this scheme in a later version of this model.

[53] **Acknowledgments.** This work was supported through a research grant from the United States Environmental Protection Agency Office of Research and Development, Global Change Research Program (CX 82876601); and the Research Applications Program at the National Center for Atmospheric Research. This research was also supported by the Office of Science (BER), U.S. Department of Energy, through the Great Plains Regional Center of the National Institute for Global Environmental Change under Cooperative Agreement DE-FC03-90ER61010. The rigorous review by two anonymous reviewers greatly enhanced the message and content of the paper, and was greatly appreciated. The National Center for Atmospheric Research is sponsored by the National Science Foundation.

References

- Adams, R., On the search for the correct economic assessment method, *Clim. Change*, 41(3/4), 363–370, 1999.
- Buishand, T. A., and T. Brandsma, Multisite simulation of daily precipitation and temperature in the Rhine basin by nearest-neighbor resampling, *Water Resour. Res.*, 37(11), 2761–2776, 2001.
- Cohen, S., Impacts of CO_2 -induced climatic change on water resources in the Great Lakes Basin, *Clim. Change*, 8, 135–153, 1986.
- Davis, J., *Statistics and Data Analysis in Geology*, John Wiley, Hoboken, N. J., 1986.
- Dowlatabadi, H., and M. G. Morgan, A model framework for integrated studies of the climate problem, *Energy Policy*, 21(3), 209–221, 1993.
- Dubrovsky, M., Z. Zalud, and M. Stastna, Sensitivity of CERES-Maize yields to statistical structure of daily weather series, *Clim. Change*, 46, 447–472, 2000.
- Easterling, D., Development of regional climate scenarios using a downscaling approach, *Clim. Change*, 41(3), 615–634, 1999.
- Efron, B., Bootstrap methods: Another look at the jackknife, *Ann. Stat.*, 7, 1–26, 1979.
- Hårdle, W., and A. W. Bowman, Bootstrapping in nonparametric regression: Local adaptive smoothing and confidence bands, *J. Am. Stat. Assoc.*, 83, 102–110, 1988.
- Hewitson, B. C., and R. G. Crane, Climate downscaling: Techniques and application, *Clim. Res.*, 7, 85–95, 1996.
- Jones, W., R. C. Rex, and D. E. Threadgill, A simulated environmental model of temperature, evaporation, rainfall, and soil moisture, *Trans. ASAE*, 15, 366–372, 1972.
- Katz, R., Use of conditional stochastic models to generate climate change scenarios, *Clim. Change*, 32, 237–255, 1996.
- Kenny, G. J., Investigating climate change impacts and thresholds: An application of the CLIMPACTS integrated assessment model for New Zealand agriculture, *Clim. Change*, 46(1), 91–113, 2000.
- Kittel, T., J. Royle, C. Daly, N. Rosenbloom, W. Gibson, H. Fisher, D. Schimel, L. Berliner, and VEMAP2 participants, A gridded historical (1895–1993) bioclimate dataset for the coterminous United States, in *Proceedings of the 10th Conference on Applied Climatology*, 20–24 October 1997, Reno, NV, pp. 219–222, Am. Meteorol. Soc., Boston, Mass., 1997.
- Kunkel, K., S. Changnon, T. Croley II, and F. Quinn, Transposed climates for study of water supply variability on the Laurentian Great Lakes, *Clim. Change*, 38(4), 387–404, 1998.
- Lall, U., and A. Sharma, A nearest neighbor bootstrap for time series resampling, *Water Resour. Res.*, 32(3), 679–693, 1996.
- Mearns, L., C. Rosenzweig, and R. Goldberg, The effect of changes in daily and interannual climatic variability on CERES-Wheat, *Clim. Change*, 32, 257–292, 1996.

- Mearns, L., C. Rosenzweig, and R. Goldberg, Mean and variances change in climate scenarios: Methods, agricultural applications, and measures of uncertainty, *Clim. Change*, 35, 367–396, 1997.
- Mearns, L. O., I. Bogardi, F. Giorgi, I. Matyasovszky, and M. Palecki, Comparison of climate change scenarios generated from regional climate model experiments and statistical downscaling, *J. Geophys. Res.*, 104, 6603–6621, 1999.
- Mingkui, C., and F. I. Woodward, Dynamic responses of terrestrial ecosystem carbon cycling to global climate change, *Nature*, 393, 249–252, 1998.
- Mitchell, J. B., T. C. Johns, M. Eagles, and W. J. Ingram, Towards the construction of climate change scenarios, *Clim. Change*, 41(3/4), 547–581, 1999.
- Morgan, M. G., and H. Dowlatabadi, Learning from integrated assessments of climate change, *Clim. Change*, 34(3–4), 337–368, 1996.
- Nicks, A. D., and J. F. Harp, Stochastic generation of temperature and solar radiation data, *J. Hydrol.*, 48, 1–7, 1980.
- Rajagopalan, B., and U. Lall, A k -nearest-neighbor simulator for daily precipitation and other variables, *Water Resour. Res.*, 35(10), 3089–3101, 1999.
- Rajagopalan, B., U. Lall, D. G. Tarboton, and D. S. Bowles, Multi-variate nonparametric resampling scheme for generation of daily weather variables, *Stochastic Hydrol. Hydraul.*, 11(1), 523–547, 1997.
- Richardson, C. W., Stochastic simulation of daily precipitation, temperature, and solar radiation, *Water Resour. Res.*, 17(1), 182–190, 1981.
- Risbey, J., and P. Stone, A case study of the adequacy of GCM simulations for input to regional climate change assessments, *J. Clim.*, 9(7), 1441–1467, 1996.
- Robock, A., R. P. Turco, M. A. Harwell, T. P. Ackerman, R. Anderssen, H. S. Chang, and M. Sivakumar, Use of general circulation model output in the creation of climate change scenarios for impact analysis, *Clim. Change*, 23, 293–335, 1993.
- Rosenzweig, C., and M. L. Parry, Potential impacts of climate change on world food supply, *Nature*, 367, 133–138, 1994.
- Sailor, D. J., and X. Li, Tree-structured regression downscaling for regional climate change predictions of precipitation, *World Resour. Rev.*, 11, 64–75, 1999a.
- Sailor, D. J., and X. Li, A semi-empirical downscaling approach for predicting regional temperature impacts associated with climatic change, *J. Clim.*, 12, 103–114, 1999b.
- Semenov, M. A., Use of a stochastic weather generator in the development of climate change scenarios, *Clim. Change*, 35(4), 397–414, 1997.
- Skidmore, E. L., and J. Tatarko, Stochastic wind simulation for erosion modeling, *Trans. ASAE*, 33(6), 1893–1899, 1990.
- Smit, B., D. McNabb, and J. Smithers, Agricultural adaptation to climatic variation, *Clim. Change*, 33(1), 7–29, 1996.
- Smith, J., and G. Pitts, Climate change scenarios for vulnerability and adaptation assessments, *Clim. Change*, 36(1/2), 3–21, 1997.
- Wilby, R. L., T. M. L. Wigley, D. Conway, P. D. Jones, B. C. Hewitson, J. Main, and D. S. Wilks, Statistical downscaling of general circulation model output: A comparison of methods, *Water Resour. Res.*, 34(11), 2995–3008, 1998.
- Wilks, D., Adapting stochastic weather generation algorithms for climate change studies, *Clim. Change*, 22, 67–84, 1992.
- Wilks, D., Forecast value in prescriptive decision studies, in *Economic Value of Weather and Climate Forecasts*, edited by R. W. Katz and A. H. Murphy, pp. 109–145, Cambridge Univ. Press, New York, 1997.
- Wilks, D. S., Multisite generation of a daily stochastic precipitation generation model, *J. Hydrol.*, 210, 178–191, 1998.
- Yakowitz, S., Nearest neighbor regression estimation for null-recurrent Markov time series, *Stochastic Processes Their Appl.*, 48, 311–318, 1993.
- Yates, D., and K. Strzepek, Modeling economy-wide climate change impacts on Egypt: A case for an integrated approach, *Environmental Modell. Assess.*, 1, 119–135, 1996.
- Yates, D., and K. Strzepek, An assessment of integrated climate change impacts on the agricultural economy of Egypt, *Clim. Change*, 38, 261–287, 1998a.
- Yates, D., and K. Strzepek, Modeling the Nile Basin under climatic change, *J. Hydrol. Eng.*, 3(2), 98–108, 1998b.
- Yates, D., T. Kittel, and R. Cannon, Comparing the correlative Holdridge model to mechanistic biogeographical models for assessing vegetation distribution response to climatic change, *Clim. Change*, 44, 59–87, 2000.
- Young, K. C., A multivariate chain model for simulating climatic parameters from daily data, *J. Appl. Meteorol.*, 33, 661–671, 1994.
- Zucchini, W., and P. T. Adamson, Bootstrap confidence intervals for design storms from exceedance series, *Hydrol. Sci. J.*, 34(1), 41–48, 1989.

S. Gangopadhyay, B. Rajagopalan, and K. Strzepek, Department of Civil, Environmental and Architectural Engineering, University of Colorado, Boulder, CO 80309, USA.

D. Yates, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307-3000, USA. (yates@ucar.edu)