A technique for generating regional climate scenarios using a nearest neighbor bootstrap

David Yates[1,2], Subhrendu Gangopadhyay[2], Balaji Rajagopalan[2],
and Kenneth Strzepek[2,3]

[1]National Center for Atmospheric Research*
Boulder, Colorado

[2]Department of Civil Engineering
The University of Colorado
Boulder, Colorado

[3]International Water Management Institute
Colombo, Sri Lanka

Submitted to
Water Resources Research

February 2002

*The National Center for Atmospheric Research is sponsored by the National Science Foundation

Corresponding author:

David N. Yates
NCAR/RAP
PO Box 3000
Boulder, Colorado 80307-3000
---
p:303/497-8394
f: 303/497-8401
e: yates@ucar.edu

Abstract

A *K*-nearest neighbor (*K-nn*) re-sampling 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, etc. 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 the fact that the variables do not have to be standardized nor is there a requirement to pre-assign weights to variables. An adaptable, moving window is used to identify candidate neighbors. The model is applied to two sets of station data in climatologically diverse areas of the US, including the Rocky Mountains and the North Central US and is shown to reproduce synthetic series that largely preserve important cross and autocorrelations. Likewise, the adapted *K-nn* algorithm is used to generate alternative climate scenarios based upon prescribed conditioning criteria.

1 Introduction

Integrated assessment (IA) studies link biophysical and socio-economic 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 agro-ecosystem 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; and Yates and Strzepek 1996 and 1998a].

Due to their detailed characterization of biophysical processes, many of these IA models require high-resolution climate data (typically, this is daily meteorologic/weather data such as precipitation, temperature, solar radiation etc.) 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 datasets are used as forcing variables [Yates and Strzepek 1998a,b; Yates et al. 2000; and 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,b; 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 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 GCM's 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 CO_2 and aerosol changes.

Stochastic weather generators address this issue with their ability to simulate plausible climate scenarios and have themselves been used as downscaling techniques in global 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 can be used in understanding the sensitivity of IA models and their individual components to climate variability and climate change is critical, an issue that has not received its due attention [Morgan and Dowlatabadi 1996; Kunkel et al. 1998; Mearns et al. 1996 and 1997; Adams 1999; Mingkui and Woodward 1998; Risbey and Stone 1996; Smith and Pitts 1997; Mitchell et al. 1999].

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 this *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

Stochastic weather generators are routinely used in water, agricultural and erosion control management [Skidmore and Tatarko, 1990; Wilks 1997; and 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 non-stationarity (i.e. seasonality). There are several drawbacks to this parametric approach, including – (i) the MAR implicitly assumes a normal distribution of the variables, which is difficult to satisfy, consequently, non-normal features in the data such as bimodality etc. cannot be captured (ii) only linear relationships between the variables can be reproduced, (iii) A large number of parameters have to be fitted, as models are

fit to each season separately. If the simulations are to be conditioned – e.g. conditioned upon the state of ENSO, then separate models have to be fit for each ENSO state which further increases the number of parameters, (iv) 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.

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 three to five days) of the current day's weather, and one of these was then randomly selected and used as the next day's weather. This approach mostly preserves the cross correlation between the variables, but Young also noted biases (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 re-sampling to multivariate data (i.e. daily weather). This approach is very similar to Young's method but with two main differences – (i) the discriminant function is not used, (ii) *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 etc.).

n 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 following section.

3 Methodology

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. An adaptable, moving window that can be used to identify the pool of nearest neighbors around the current day is also implemented, allowing for better short-term dependency (Sharma and Lall 1999). For example, in some areas in the US, springtime represents a period of rapid transition, thus a smaller window would be used; while a larger window would be used during more climatologically stable periods such as late summer. The user can specify the window width depending on the location within the year, while in this study a constant 14-day window was chosen. Finally, a strategy for giving priority to certain years from which to resample is presented, which can be used to derive alternative scenarios with differing statistical attributes.

3.1 The K-nn algorithm

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)), then:

1. The weather on day *t* is randomly chosen for the corresponding day from the set of *N* years, which includes all variables *p*. This is the feature vector $\boldsymbol{D}_t^m$. The objective of the algorithm is to select the next day *t* + 1 starting from day *t*. The simulation continues by:
2. Computing the mean vector of the *p* variables across the stations corresponding to year *t* as,

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

where,

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

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

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

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

3. 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 January 1 (of a randomly selected year from set *N*) then the window-of-days consists of all days between December 25 and January 8, excluding January 1 for all *N* years. January 1 is excluded since it includes the vector, $\bar{x}_t$. Thus, there are (*w* * *N)* days that are potential neighbors to day *t*. As mentioned earlier this temporal window width is adaptable (e.g. the window width can change from day-to-day), although we have chosen a constant, 14-day moving window (*w* = 14) in this study.
4. Mean vectors, $\bar{x}_i$ are computed for each day *i, where i* =1 to (w * *N).* These mean vectors are computed using the equations from step 2.
5. A *pxp* covariance matrix, $\boldsymbol{S_t}$ is computed for day *t* using the data block (w * *N) x p* described in step 3 (Note: this step is done only once, since $\boldsymbol{S_t}$ is computed from the historic data and can be used repeatedly).
6. Mahalanobis distances, d_i^2 are computed between the mean vector of the current day's weather, $\bar{x}_t$ and the mean vectors, $\bar{x}_i$ as

$$d_i^2 = \left(\bar{x}_t - \bar{x}_i\right) S_t^{-1} \left(\bar{x}_t - \bar{x}_i\right) \quad (5)$$

for all *i* = 1 to *(w * N)*.

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, $K = \sqrt{w*N}$ generally performs well as seen by them and also by Rajagopalan and Lall [1999]. Thus, if fifty-eights years of data are used with a 14-day window then $K = \sqrt{14*58} \approx 28$.
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-neighbors. The neighbor with the shortest distance gets a high weight, while the neighbor with the smallest distance (i.e. the K^{th} 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.

8. The $t + 1$ day is selected as one of the K-neighbors using the probability metric, p_j leading to a new feature vector, $\boldsymbol{D}_{t+1}^{m}$. 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 other cases, $t + 1$ corresponding to d_q is selected based on a linear interpolation of u and the corresponding proximity of the p_j's.
9. Steps 1 through 8 are repeated for as many days and years of required simulations. Since each year is generated independently, the number of years to be simulated is not relevant to the length of the historic record.

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, changes in variability, etc.) as will be seen in the next section

3.2 Strategic Resampling

When all (w * *N)* days are given equal probability of being selected using 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 El-Nino, Southern Oscillation Index (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 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.

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* n)* 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* sub-set 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 $\boldsymbol{n_t}$ subsets from all *N* years. These $\boldsymbol{n_t}$ subsets can change both inter-annually and intra-annual according to year-to-year, weekly deviations (or anomalies) observed in the historic record.

These weekly anomalies are computed from the historic data for each variable (e.g. Ψ^i_{jt} is PPT^i_{jt}, TMN^i_{jt}, or TMX^i_{jt}; where *i* is the year, *j* is the station, and *t* is the day in the week) by computing the regional weekly mean. For simplification, assume the year can be broken into 52, seven-day periods, thus the regional weekly means are given as,

$$\overline{\Psi}_w = \frac{1}{N*m*7} \sum_{i=1}^{N}\sum_{j=1}^{m}\sum_{t=1}^{7}\Psi^i_{j,t} \tag{7}$$

while the regional weekly means for each year are given as,

$$\Psi^i_w = \frac{1}{m*7}\sum_{j=1}^{m}\sum_{t=1}^{7}\Psi^i_{j,t} \tag{8}$$

Regional, weekly anomalies are computed as $\delta^i_w = \Psi^i_w - \overline{\Psi}_w$, which are then ranked by week according to their relative magnitude. For example, regionally minimum temperature anomalies

could be ranked from lowest to highest for all weeks and years such as, $\{\delta_1^i$: (1945 -7.8, 1962 -7.2, .. , 1967 +0.5, 1957 +4.5, 1987 +8.0); δ_2^i : (1946 -6.8, 1945 -5.0,.., 1896 +4.0, 1937 +6.8) ; .. δ_w^i .. (,) **:** δ_{52}^i (1956 -6.8, 1945 -6.1, .., 1978 +6.8)}. Here, the coldest first week of January was 1945 (-7.8ºC below the week 1 mean), while the warmest first week of January was 1987 (+8.0ºC above the week 1 mean). The sorted list is then assigned a rank index I_1^i ,..,I_1^N ;..; I_{52}^i ,..,I_{52}^N according to position and corresponding ranked probabilities are given as $1/I_1^i$,..,$1/I_1^N$;..; $1/I_{52}^i$; $1/I_{52}^N$. To compute the $\boldsymbol{n_t}$ year sub-sets, scaled $u \sim U[0,1]$ random numbers are generated that correspond to the ranked probabilities ($1/I_w^i$). These are mapped to the I_w^i index that corresponds to a year in the ranked lists. With this technique if N, $U[0,1]$ random numbers are selected there is equal probability that any year is part of the $\boldsymbol{n_t}$ sub-set. Likewise, there is the probability that years could be repeatedly selected and appear multiple times in the $\boldsymbol{n_t}$ sub-set, and are thus more likely to contribute to the daily weather simulations at any time t. This is the main difference from when resampling is done where $\boldsymbol{n_t}$ is prescribed to be all N years.

Paired Ranking

The above applies for single variables, while paired ranking is slightly more difficult. 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, etc.) relative to "normal" conditions. This is done by first standardizing the weekly mean of each variable,

$$\mathrm{Z}_w^i = (\Psi_w^i - E[\Psi_w^i]) / \{E[\Psi_w^i - E[\Psi_w^i]\}^{1/2} \tag{9}$$

for all years i =1 to N; then computing the standardized weekly mean $\overline{Z}_w$ according to 7, and then computing the distance as,

$$\delta_w^i = \sqrt{(\mathrm{Z}(1)_w^i - \overline{\mathrm{Z}}(1)_w)^2 + (\mathrm{Z}(2)_w^i - \overline{\mathrm{Z}}(2)_w)^2} \tag{10}$$

where the number within the parenthesis (e.g. Z(1) and Z(2)) indicates the appropriate variable (e.g. PPT, TMN, TMX). Ranking is done similarly to the single variable case described above.

Biasing the ranked list

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 $\boldsymbol{n_t}$ sub-set. However, it would be advantageous to select the index in such a way as to bias certain years over others. The integer function,

$$I_w^i = \text{INT}[N-(U[0,1]^{\lambda_w^i} * (N-1))] \quad (11)$$

is used to select the index, and hence years from which to derive the $\boldsymbol{n_t}$ sub-sets, where λ_w^i is the shape parameter of the function. A value of λ_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 $u \sim U[0,1]$. 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 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 11 is ~33 when repeatedly queried with $U[0,1]$ random numbers, suggesting a bias towards the selection of warm years. Likewise, a value of 0.66 would bias cold years with an expected value of ~17 (Figure 1).

Introducing Inter and Intra Annual Variability

The shape of the curve corresponding to equation 11 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 (inter-annual change). The shape parameter can be altered according to,

$$\lambda_w^i = \lambda_{w-1}^{i-1} (1+\rho_i)(1+\beta_w) \quad (12)$$

where β_w is the intra-annual trend parameter and ρ_i is the inter-annual trend parameter, which is 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 (Equ. 11) is queried, it's 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 9 and 50 through 52 (corresponding to 10 December to 26 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.

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 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 $\beta_{k,}$ for the weeks of December, January and February), would likely yield greater temperature changes for January when compared with December and February.

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

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 US. 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 (Fig. 2). Region 4 had 34 stations, while Region 7 had seven. These two regions were chosen for the diversity of their climate. Daily observational data for a fifty-eight 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.

Thirty series, each comprised of 30 realizations were generated, leading to a total of 900 independent synthetic series. Suites of statistics for each month of the year, such as mean, standard deviation, coefficient of skew, cross-correlations between the variables, cross-correlation between stations were computed from the realizations of each series, and are summarized with boxplots. Although the weather simulations were made on a daily time scale, the statistics from the daily data have been aggregated to the monthly time scale. In all the boxplot figures shown below, the following holds: The box in the boxplots (e.g. Fig. 3) indicates the inter quartile range of the simulations and the whiskers show the 5^{th} and 95^{th} 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 a 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

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 $\lambda_w^i = 1$ for all w and i. Thus, all ($w * N$) days have an equal probability of being 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 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 for each month [Richardson, 1981]. Given this, the performance was impressive.

The upper right graph of Figures 3 and 4 are the monthly standard deviations for these two stations and the simulations adequately reproduced this statistic, which was greatest in the spring (114198) and in the spring and late summer (52281). Skewness, an important attribute of daily precipitation that describes the symmetry of the distribution of rainfall, is plotted in the lower left graph of Figures 3 and 4 for these two stations. Typically rainfall 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: Jan., Apr., May, Jun., Sept., Dec.; Fig 4) generally corresponds to months where the model had a tendency to over select dry days leading to less precipitation (Station 52281: Jan., Apr., Jun., Sept., and Dec.; Fig. 4).

To illustrate a measure of persistence, the average number of days between precipitation events for each month is illustrated in the lower right graph 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 complex seasonal variability comprised of an early spring minima in March and a summer minima in July and August corresponding to the North American monsoon in the southwestern US. Station 114198 in Region 4 exhibits increased occurrence of wet days, given as fewer days between precipitation events in the spring.

Figure 5 shows the boxplots of monthly mean temperatures of the daily data (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 monthly mean temperatures for these two 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.

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. 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.

Correlation among variables at a station

Since the same days weather are 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 anti-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 (Fig. 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 (Fig. 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 auto correlation 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 auto correlations of temperature at a station were well preserved (not shown).

Correlation among stations

To investigate how well the spatial dependencies of precipitation were captured by the model, the auto correlations across stations are summarized with scatter plots in Figure 7 (Wilks 1998). Both Lag-0 (Corr[$\boldsymbol{D}_t^k$, $\boldsymbol{D}_t^l$; $k \neq l$]) and Lag-1 (Corr[$\boldsymbol{D}_t^k$, $\boldsymbol{D}_{t+1}^l$; $k \neq l$]), for the collection of stations (k,l = 1 to m) are presented and are stratified by the winter (Oct., Nov., Dec., Jan., Feb., Mar.) and summer seasons (Apr., May, Jun., Jul., Aug. Sep). There are m = 35 stations of Region 4 and m = 7 stations of Region 7, yielding (m*(m-1)/2)*2*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, Fig. 7).

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, Fig. 7). The summer months are characterized by more localized, convective events reflected in lower station cross correlations. Note the model bias, which tended to over-estimate those stations with low cross correlations, while under-estimating 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 auto correlations (bottom right graph, Fig. 7), and the model adequately produced these Lag-1 correlations with the above noted biases.

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 multi-station meteorological data that largely reproduced the spatial and temporal statistics of the historic data, given values of $\lambda_w^i = 1$ for all weeks and years.

4.2 Cases of λ_w^i <> 1: Generating alternative scenarios

To test the ability of the algorithm to generate alternative climate scenarios using different values of λ_w^i, Region 4 was selected (Fig. 1) and a simple increasing average temperature (TAV) scenario was contrived. Average temperature was computed simply as a weighted average of minimum and maximum temperatures. For all stations in Region 4, 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.

Again, 30 synthetic series each comprised of 30 realizations (a total of 900 independent years of synthetic data) were generated but with $\lambda_w^i = 0.6$, 2.0, and 3.0 for all w,i. These λ_w^i values were used in equation 11 to generate indices I_w^i from which the ranked TAV list was sampled. Thus, a scenario corresponding to $\lambda_w^i = 0.6$ should result in cooler regional temperatures, while the $\lambda_w^i = 2.0$ and 3.0 scenarios should produce warmer regional temperatures.

The Table 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 $\lambda_w^i <> 1.0$ did lead to changes in the average regional

temperature. Although the biasing was only based on the regional 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 λ_w^i = 2.0 and 3.0 resulted in regional TMN and TMX increases, while the synthetic series corresponding to λ_w^i = 0.6 resulted in TMN and TMX regional temperature decreases (see Table). Regional average precipitation decreased by 6 percent for the λ_w^i = 0.6, while it increased by 7 percent for the λ_w^i = 2 case and by 12 percent for the λ_w^i = 3.0 case.

Boxplots, similar to those presented for the λ_w^i = 1 case are presented for both the λ_w^i=1 and λ_w^i = 3 scenarios for station 114198, which includes both total monthly precipitation (top row) and average monthly temperature (bottom row) based on all 900 series. These plots clearly show that the strategic resampling can produce alternative scenarios with different statistical attributes. Warming has occurred for all months, while precipitation has increased most notably for the winter months. For these same scenarios, the station cross correlations are summarized with scatter plots (Fig. 9) similar to those presented in Figure 7 but only for Region 4, and includes both the λ_w^i=1 scenario (left column) and the λ_w^i = 3 scenario (right column). There was some difference in the Lag-0 cross correlations between the λ_w^i = 1 and λ_w^i = 3 cases implying some change in the Lag-0 cross correlation of the new scenario. Likewise, there was some shift in the Lag-1 cross correlation for the λ_w^i = 3 scenarios.

4.3 Case of λ_w^i <> 1 : varying β_W and ρ_i by week and year

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 Integrated Assessment 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. This ranked list was then strategically sampled to produce $\boldsymbol{n_t}$ sub-sets from which to apply the *K-nn* algorithm.

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 12. Thus, when equation 11 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 (Fig. 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 inter-annual trend parameter, ρ_i in an attempt to mimic a gradual increase of warmer-drier springs throughout the region.

The 100-year precipitation series are given in Figure 11 (left graph) for select months that include the late winter, spring, and early summer months, 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. Not that although the scenario did not bias the selection of any particular year in January, there was a modest decline in January precipitation on over the 100 year simulation (Fig. 11, left graph).

Figure 11 includes plots of the 100-year time series of monthly maximum temperatures for the 100-year series (right graph). The long-term trend of increasing maximum temperature for April and May is clear from this figure, with a last decade increase of +2.5ºC and +2.3ºC for April and May, over the 100-year time period. There was little to no trend observed in the January and June temperatures.

Summary and Conclusions

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 Euclidan 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 the covariance matrix in the Mahalanobis distance computation (Equation 4) weights the variables as per their covariance (or correlation) – i.e. strongly correlated variables get less weight. This modified *K-nn* also incorporates an adaptable, moving window centered on the day of the year providing a realistic neighborhood – e.g., weather on February 1 is more likely to be closer to weather on days closer to this day of the year.

The *K-nn* algorithm was demonstrated by simulating weather sequences at multiple sites over two regions in the US. The performance of this technique in reproducing the temporal and spatial statistics was good. Capturing the spatial and temporal dependencies is very important in using the 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.

A common shortcoming of this and other bootstrap method 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 auto-regressive 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.

ACKNOWLEDGMENTS

This research was supported through a research grant from the National Institute for Global Environmental Change (NIGEC) of the US Department of Energy; and the United States Environmental Protection Agencies Office of Research and Development, Global Change Research Program (CX 82876601); and the Research Applications Program at the National Center for Atmospheric Research. The rigorous review by two anonymous reviewers greatly enhanced the message and content of the paper, and was greatly appreciated.

References

Adams, R., On the search for the correct economic assessment method, *Climatic Change*, 41(3/4) 363-370, 1999.

Buishand, T.A, 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 CO2-induced climatic change on water resources in the Great Lakes Basin, *Climatic Change*, 8, 135-153, 1986.

Davis, J., *Statistics and Data Analysis in Geology*, John Wiley and Sons, New York, 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, *Climatic Change* 46, 447- 472, 2000.

Easterling, D., Development of regional climate scenarios using a downscaling approach, *Climatic Change*, 41(3), 615-634, 1999.

Efron, B., Bootstrap methods: Another look at the Jackknife, *Ann. Stat*., 7, 1–26, 1979.

Härdle, W., and A. W. Bowman, Bootstrapping in nonparametric regression: Local adaptive smoothing and confidence bands, *JASA J. Am. Stat. Assoc.*, 83, 102–110, 1988.

Hewitson, B.C. and Crane, R.G., Climate downscaling: techniques and application, *Climate Research*, 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, *Climate 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, *Climatic 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 conterminous United States', In: Proceedings of the 10th Conference on Applied Climatology, 20-24 October 1997, Reno, NV. American Meteorological Society, Boston, 1997.

Kunkel, K., S. Changnon, T. Croley II, and F. Quinn, Transposed climates for study of water supply variability on the Laurentian Great Lakes, *Climatic 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, *Climatic 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, *Climatic 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. of Geophysical Research*, 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, W. J. Ingram, Towards the construction of climate change scenarios, *Climatic Change,* 41(3/4), 547-581, 1999.

Morgan, M. G., and H. Dowlatabadi, Learning from integrated assessments of climate change, *Climatic Change,* 34(3-4), 1996.

Nicks, A. D., and J. F. Harp, Stochastic generation of temperature and solar radiation data, *J. Hydrol*., 48, 1–7, 1980.

Rajagopalan, B., U. Lall, D. G. Tarboton, and D. S. Bowles, Multi-variate nonparametric resapling scheme for generation of daily weather variables, *Stochastic Hydrol. Hydraul.*, 11(1), 523–547, 1997.

Rajagopalan, B., U. Lall, D., A *K*-nearest neighbor simulator for daily precipitation and other variables, *Wat Res. Res.,* 35(10), 3089-3101, 1999.

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. Climate*, 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, *Climatic 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 Resource Review*, 11, 64-75, 1999a.

Sailor, D.J. and X. Li, A semi-empirical downscaling approach for predicting regional temperature impacts associated with climatic change, *Journal of Climate*, 12, 103-114, 1999b.

Semenov, M.A., Use of a stochastic weather generator in the development of climate change scenarios, *Climatic Change*, 35(4): 397-414, 1997.

Sharma, A. and U. Lall, A nonparametric approach for daily rainfall simulation, *Mathetmatics and Computers in Simulation*, 48, 361-371, 1999.

Skidmore, E.L., and J. Tatarko, Stochastic wind simulation for erosion modeling. *Trans. ASAE, 33,6, 1893-1899,* 1990.

Smit, B., D. McNabb, J. Smithers, Agricultural adaptation to climatic variation, *Climatic Change,* 33(1), 7-29, 1996.

Smith, J. and G. Pitts, Climate change scenarios for vulnerability and adaptation assessments. *Climatic Change,* **36(1/2),** 3-21, 1997.

Wilby, R.L., Wigley, T.M.L., Conway, D., Jones, P.D., Hewitson, B.C., Main, J. and Wilks, D.S., Statistical downscaling of general circulation model output: A comparison of methods, *Water Resour. Res.*, (42) 2995-3008, 1998.

Wilks, D.S, Adapting stochastic weather generation algorithms for climate change studies, *Climatic Change*, 22, 67-84, 1992.

Wilks, D.S., Forecast Value in Prescriptive Decision Studies. In: R.W. Katz and A.H. Murphy, eds., Economic Value of Weather and Climate Forecasts. Cambridge University Press, 109-145, 1997.

Wilks, D.S., Multisite generation of a daily stochastic precipitation generation model, *J. of Hydrology*, 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 Modeling and Assessment,* 1, 119-135, 1996.

Yates, D. and K. Strzepek, An assessment of integrated climate change impacts on the agricultural economy of Egypt, *Climatic Change,* 38, 261-287, 1998a.

Yates, D. and K. Strzepek, Modeling the Nile Basin Under Climatic Change, *J. of Hydrologic Engineering*, 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, *Climatic 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.

Figure 1. Plot of the index function given in Equ. 7. A $\lambda_{k,j} = 1.0$ corresponds to the straight line (mean), while a $\lambda_{k,j} = 2$ introduces curvature into Equ. 7, biasing warm years, while a $\lambda_{k,j} = 0.6$ value will yield a cold bias.

Figure 2. Map depicting the 21-state area of interest in this study (top right), 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.

Fig 3. Box plots of precipitation statistics for Station 114198 in Region 7, with constant $\lambda_{k,j} = 1.0$. Top left graph is total monthly precipitation, top right graph is the standard deviation of precipitation, bottom left graph is the skewness of daily precipitation, and the bottom right graph is the number of days between precipitation events for the month. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

Fig 4. Same as Figure 3, but for Station 52281 in Region 4.

Fig 5. Box plots of average temperature statistics for Station 114198 (top left graph) and Station 52281 (bottom left graph), and the standard deviation of temperature for the same two stations (right graphs), given constant $\lambda_{k,j} = 1.0$. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

Fig 6. Box plots of Lag-0 correlation between precipitation and 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 $\lambda_{k,j} = 1.0$. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

Fig 7. Scatter plots of station-to-station Lag-0 auto correlation of precipitation for Region 4 (upper left graph) and Region 7 (upper right graph) stratified by season; and the station-to-station Lag-1 auto correlations for Region 4 (lower left graph) and Region 7 (lower right graph) also stratified by season. Simulations were for $\lambda_{k,j} = 1.0$ constant values.

Fig 8. Box plots of total monthly precipitation (top left graph) and average monthly temperature (bottom left graph) for Station 114198, given $\lambda_{k,j} = 1.0$ constant values. Right graphs are the same statistics, but for constant values of $\lambda_{k,j} = 3.0$. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

Fig 9. Scatter plots of station-to-station Lag-0 auto correlation of precipitation for Region 4 (upper left graph) and Region 7 (upper right graph) stratified by season; and the station-to-station Lag-1 auto correlations for Region 4 (lower left graph) and Region 7 (lower right graph) also stratified by season. Simulations were for $\lambda_{k,j} = 1.0$ constant values (Same as Fig. 7) and for $\lambda_{k,j} = 3.0$ (right graphs).

Figure 10. Plot of the temporally varying $\lambda_{k,j}$ surface for the "warmer-drier" spring scenario.

Figure 11. The 100-year times series of total monthly precipitation (left graph) for the 'warmer-drier' spring scenario for select months. The light lines are the actual data, while the dark lines with variability 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 actual data from the scenario (light lines) and the linear trend line for the 100-year times series of average monthly temperature.

Table. The Region 4, annual mean of TMN, TMX, and PPT computed from the 34 stations for 30 annual series, each comprised of 30 realizations for a total of 900 independent series, and for different values of $\lambda_{k,j}$, using an average temperature ranked list developed from Region 4 historic data. Also included are the Region 4 historic averages for these same variables. For the scenario rows, the left numbers are the absolute values, while the right numbers are the deviations from the historic (precipitation is the percent difference).

$\lambda_{k,j}$	**TMN (C)**	**TMX (C)**	**PPT (mm)**
Historic	5.2	16.6	941
0.6	4.2; -Δ1.0	15.6; -Δ1.0	882; -6%
1.0	5.3; +Δ0.1	16.7; +Δ0.1	921; -2%
2.0	6.9; +Δ1.7	18.4; +Δ1.8	1007; +7%
3.0	7.8; +Δ2.6	19.3; +Δ2.7	1058; 12%

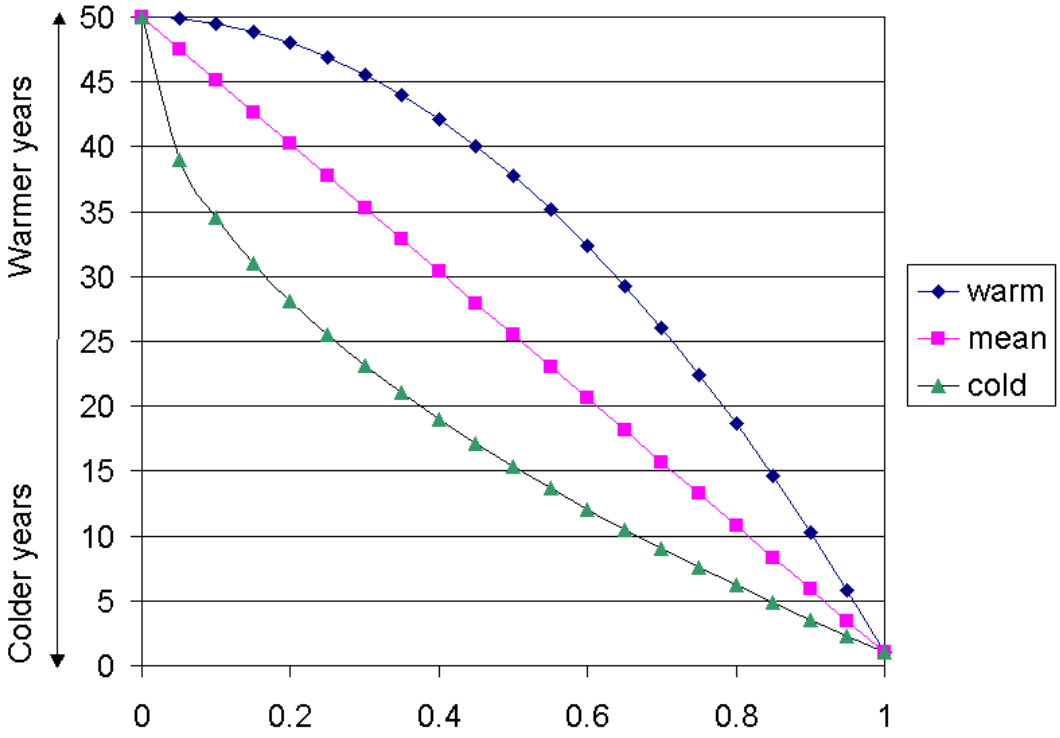

Figure 1. Plot of the index function given in Equ. 7. A $\lambda_{k,j} = 1.0$ corresponds to the straight line (mean), while a $\lambda_{k,j} = 2$ introduces curvature into Equ. 7, biasing warm years, while a $\lambda_{k,j} = 0.6$ value will yield a cold bias.

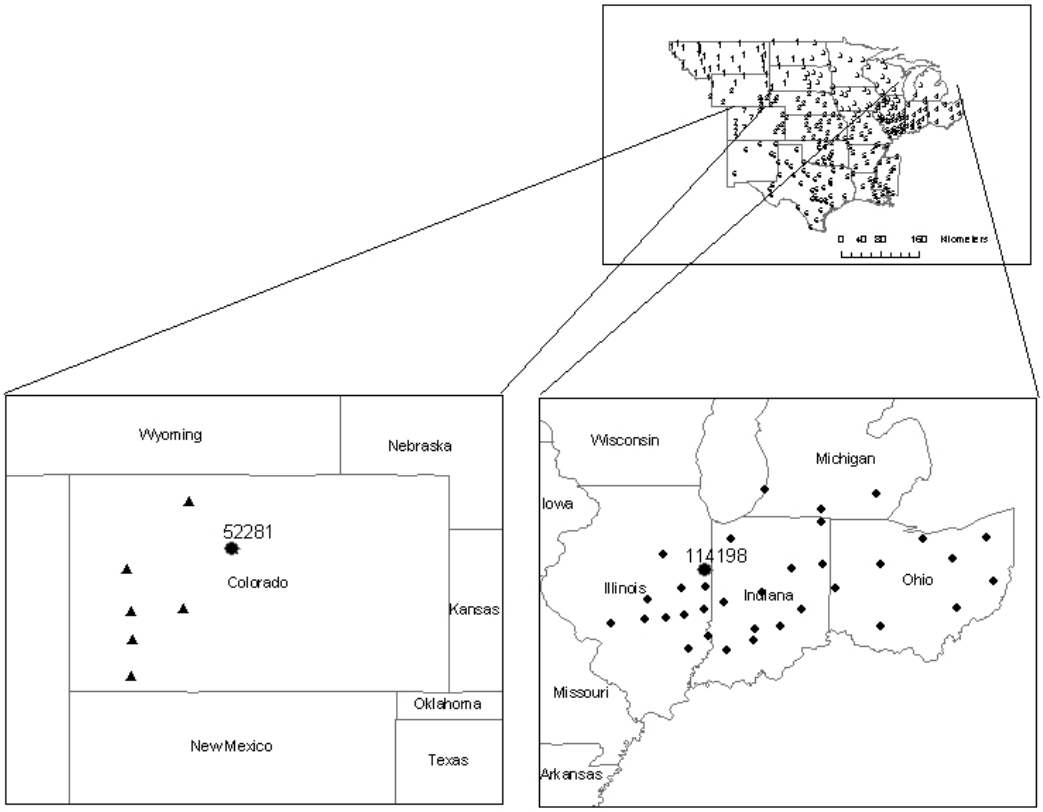

Figure 2. Map depicting the 21-state area of interest in this study (top right), and the two focus areas (Regions 4 and 7) and their corresponding stations (bottom panels). The numbers indicate stations grouped by region. The two stations of focus (114198 in Region 4 and 52281 in Region 7) are highlighted.

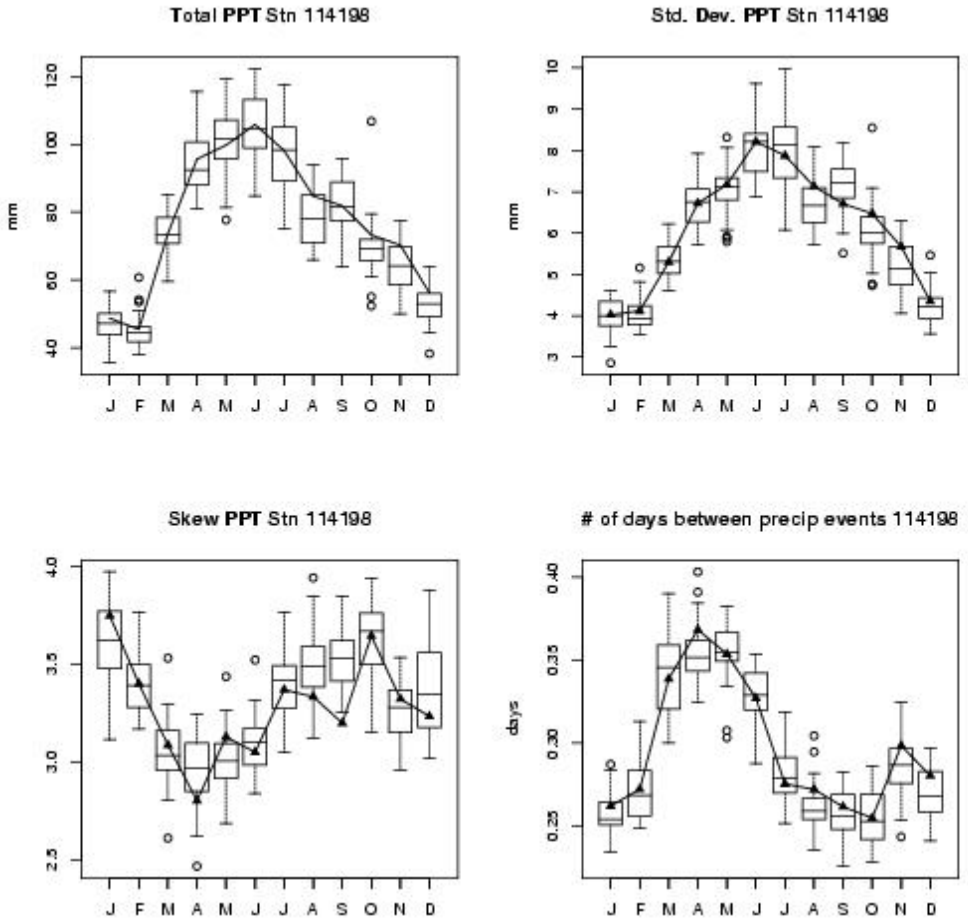

Fig 3. Box plots of precipitation statistics for Station 114198 in Region 7, with constant $\lambda_{k,j}$ = 1.0. Top left panel is total monthly precipitation, top right panel is the standard deviation of precipitation, bottom left panel is the skewness of daily precipitation, and the bottom right panel is the number of days between precipitation events for the month. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

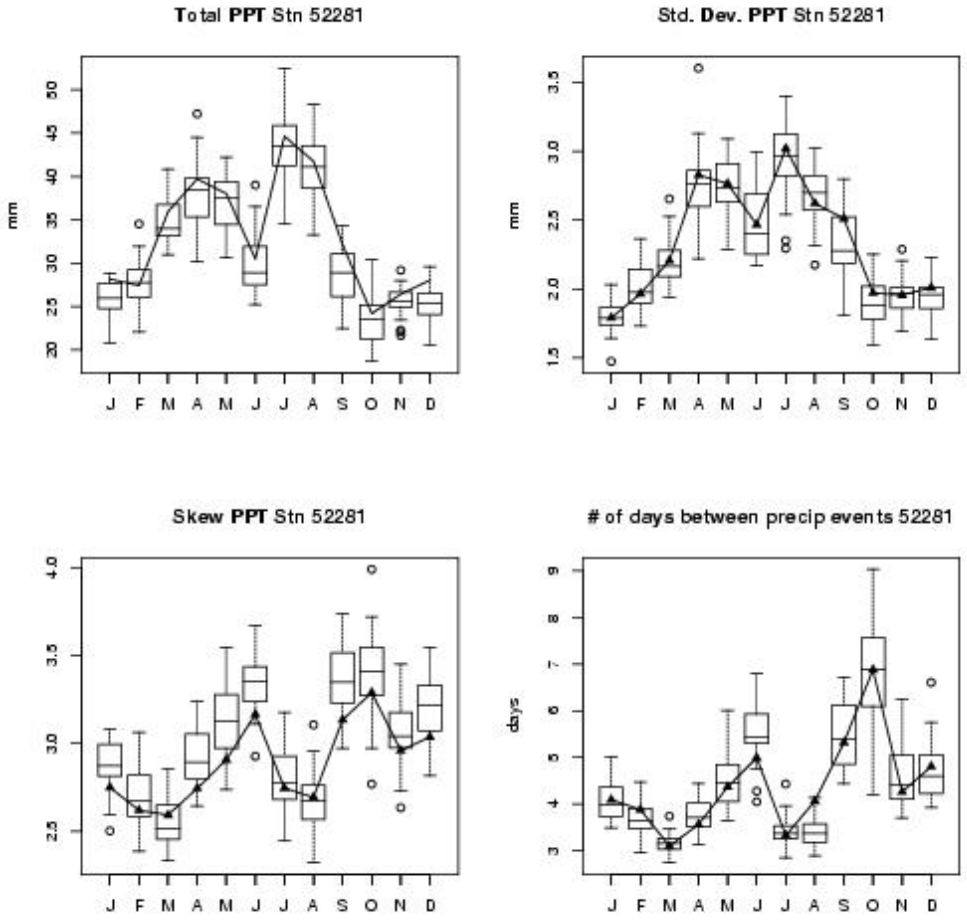

Fig 4. Same as Figure 3, but for Station 52281 in Region 4.

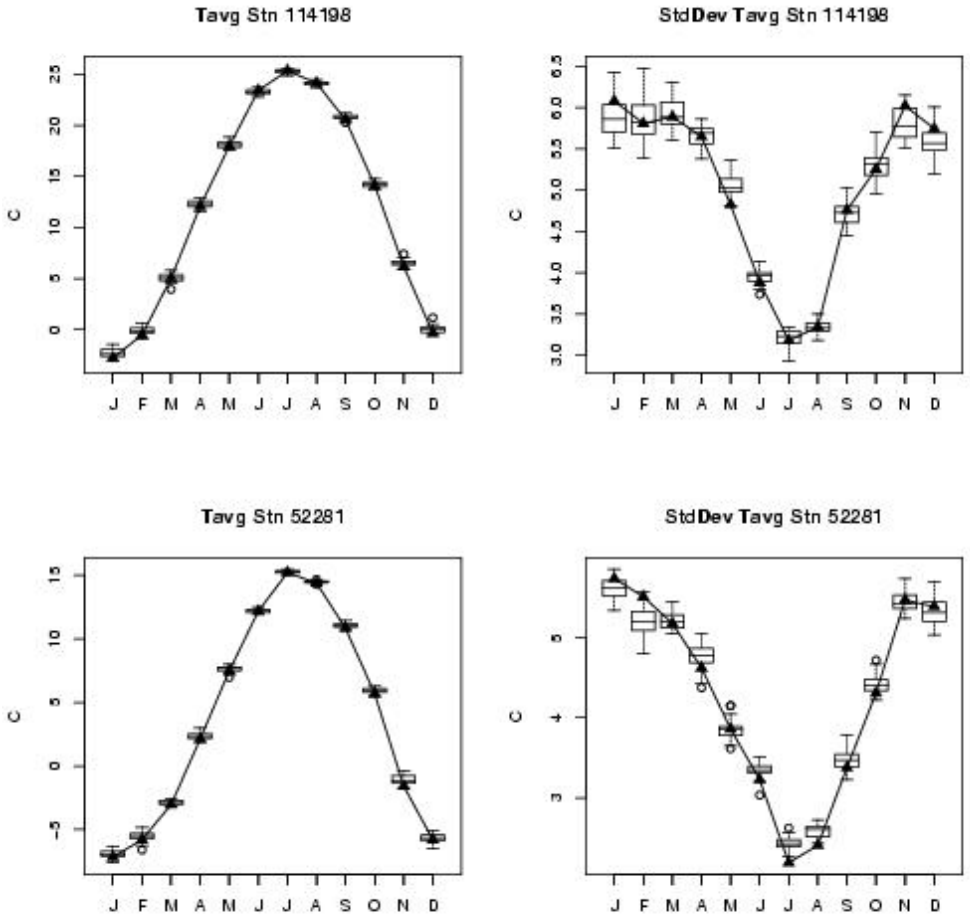

Fig 5. Box plots of average temperature statistics for Station 114198 (top left panel) and Station 52281 (bottom left panel), and the standard deviation of temperature for the same two stations (right panels), given constant $\lambda_{k,j} = 1.0$. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

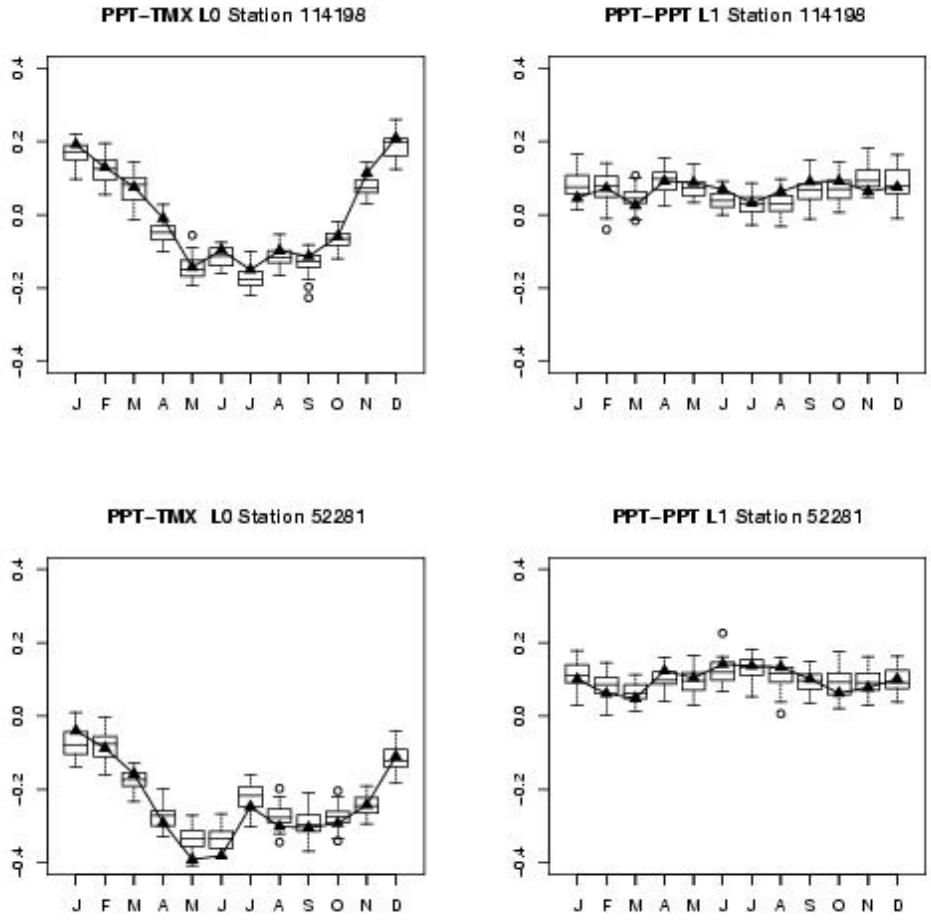

Fig 6. Box plots of Lag-0 correlation between precipitation and maximum temperature at Station 114198 and 52281 (left panels); and the Lag-1 auto correlation of precipitation for the same two stations (right panels), given constant $\lambda_{k,j} = 1.0$. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

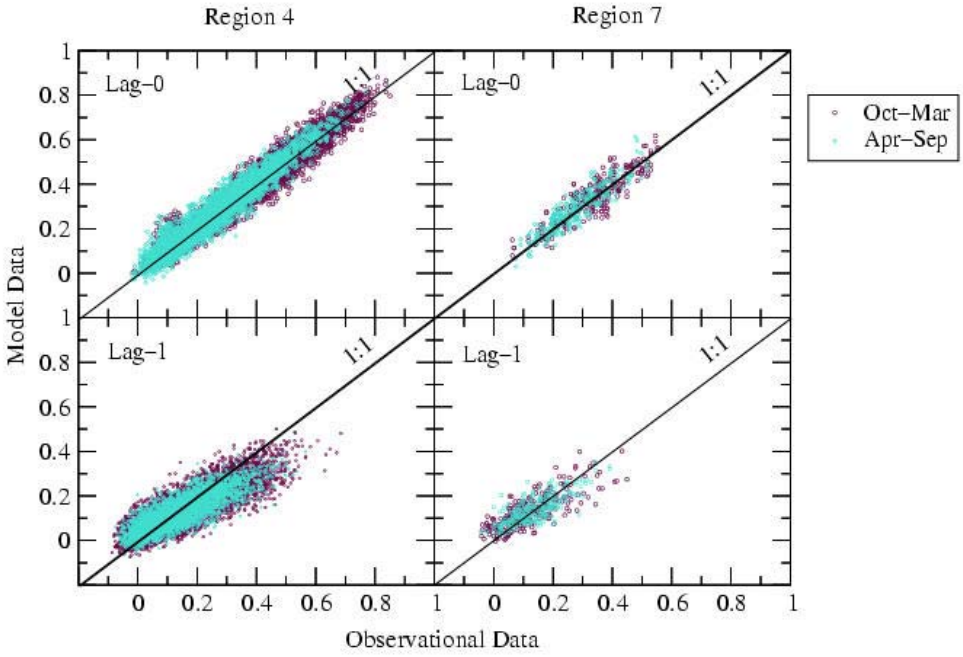

Fig 7. Scatter plots of station-to-station Lag-0 auto correlation of precipitation for Region 4 (upper left panel) and Region 7 (upper right panel) stratified by season; and the station-to-station Lag-1 auto correlations for Region 4 (lower left panel) and Region 7 (lower right panel) also stratified by season. Simulations were for $\lambda_{k,j} =$ 1.0 constant values.

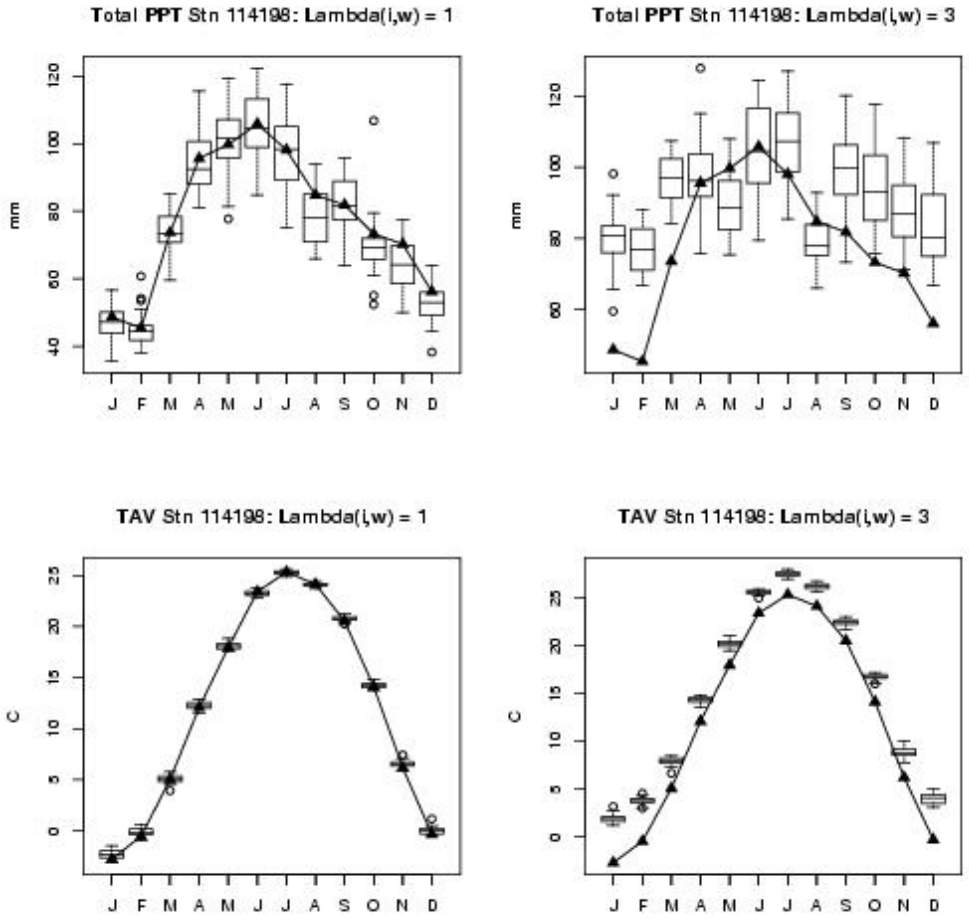

Fig 8. Box plots of total monthly precipitation (top left panel) and average monthly temperature (bottom left panel) for Station 114198, given $\lambda_{k,j} =$ 1.0 constant values. Right panels are the same statistics, but for constant values of $\lambda_{k,j} =$ 3.0. The dark line and marks are the same statistics derived from the historical data for the period 1938 to 1997.

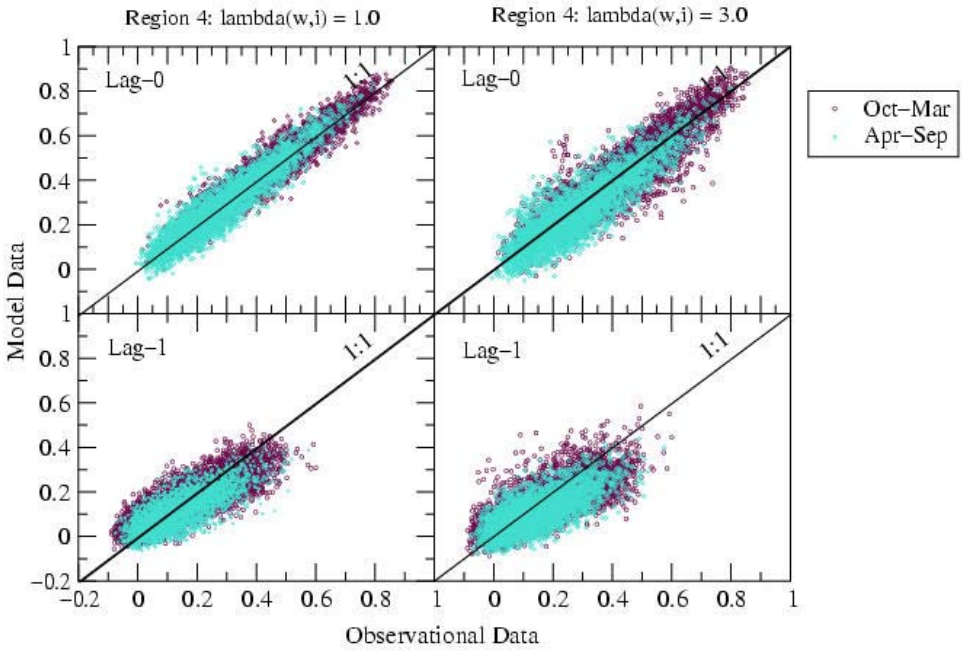

Fig 9. Scatter plots of station-to-station Lag-0 auto correlation of precipitation for Region 4 (upper left panel) and Region 7 (upper right panel) stratified by season; and the station-to-station Lag-1 auto correlations for Region 4 (lower left panel) and Region 7 (lower right panel) also stratified by season. Simulations were for $\lambda_{k,j} =$ 1.0 constant values (Same as Fig. 7) and for $\lambda_{k,j} =$ 3.0 (right panels).

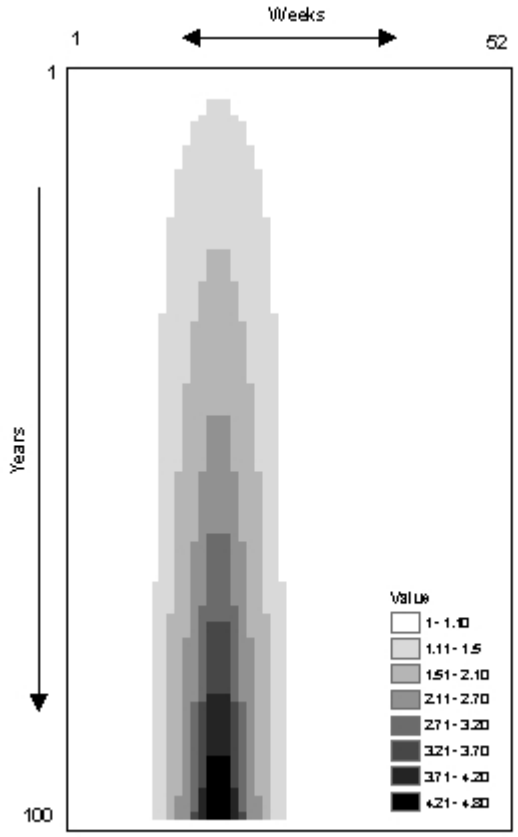

Figure 10. Plot of the temporally varying $\lambda_{k,j}$ surface for the "warmer-drier" spring scenario.

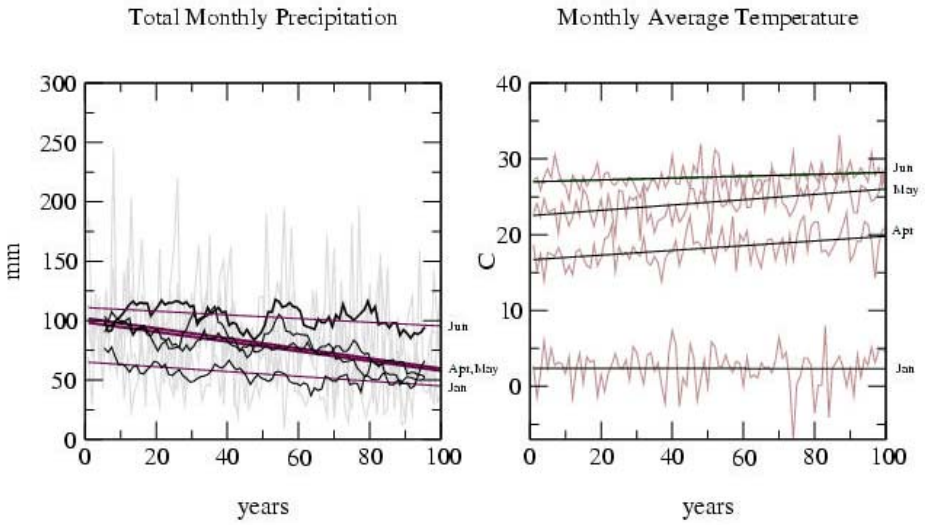

Figure 11. The 100-year times series of total monthly precipitation (left graph) for the 'warmer-drier' spring scenario for select months. The light lines are the actual data, while the dark lines with variability are the 10-year moving averages, and the straight lines are the linear trends for January, April, May and June. The right graph shows the actual data from the scenario (light lines) and the linear trend line for the 100-year times series of average monthly temperature.