Simulation of extreme precipitation in the Rhine basin by nearest-neighbour resampling

T. Brandsma, T. A. Buishand

Hydrology and Earth System Sciences, 2(2–3), *195–209* (1998) © EGS

Hydrology & Earth System Sciences

Simulation of extreme precipitation in the Rhine basin by nearest-neighbour resampling

Theo Brandsma and T. Adri Buishand

Royal Netherlands Meteorological Institute (KNMI), P.O. Box 201, 3730 AE De Bitt, The Netherlands

Abstract

The use of the nonparametric nearest-neighbour resampling technique is studied for generating time series of daily rainfall and temperature for seven stations in the German part of the Rhine basin. The emphasis is on the reproduction of extreme N-day precipitation amounts. The daily temperatures are used to determine snow accumulation and melt in winter. Two versions of the resampling method, conditional on the atmospheric circulation and unconditional, show comparable results. For precipitation, the autocorrelation properties are well reproduced, whereas for temperature the autocorrelation coefficients are systematically underpredicted. The distributions of the N-day annual maximum precipitation amounts are adequately preserved. Despite the systematic underprediction of the temperature autocorrelation, the distributions of N-day maximum snowmelt are well reproduced. A 1000-year simulation for the seven stations shows that unprecedented rainfall situations can be generated.

Introduction

There is a growing interest in the simulation of synthetic time series of weather variables like precipitation and temperature. Part of the interest stems from dissatisfaction with current methods of estimating floods, like frequency analysis of observed peak river discharges or running a design storm through a catchment model. Improved flood estimates are expected from the use of synthetic sequences of weather variables, in combination with a physically-based model of the river basin. A stochastic weather generator can also be a powerful tool to translate the large-scale information of General Circulation Models (GCMs) into representative local time series data ('downscaling') in order to assess the hydrologic impacts of possible future changes in climate.

Several types of model have been used to generate synthetic sequences of daily rainfall. The most popular approach is to describe the occurrence of wet and dry days by a two-state Markov chain or an alternating renewal process and to represent the distribution of the precipitation amounts on the wet days by a gamma distribution or an other positively skewed distribution; see Woolhiser (1992) for a review. A recent development is to link model parameters to properties of the atmospheric circulation in order to improve the reproduction of persistence of daily rainfall or to assess the effects of systematic changes in the atmospheric circulation, e.g. resulting from increased greenhouse gas concentrations (Bárdossy and Plate, 1991; Katz and Parlange, 1993).

The majority of the literature is on the generation of single rainfall sequences. However, hydrologic studies generally require more meteorological input, e.g. temperature or rainfall at several sites in a river basin. Methods for generating such multivariate data generally assume an underlying normal distribution, because its multivariate extension is straightforward and well-known. Richardson (1981) used a multivariate normal first-order autoregressive process to generate maximum and minimum temperature and solar radiation conditional on a synthetic sequence of daily rainfall. The same type of model has been explored for a stochastic description of multi-site daily rainfall (Richardson, 1977; Bárdossy and Plate, 1992; Hutchinson, 1995). Unfortunately, the assumption of an underlying multivariate normal distribution is a questionable one. Although transformations can be useful to obtain the marginal distributions that are more or less normal, this does not ensure that the dependence structure corresponds to that of a multivariate normal distribution. The latter requires that the relations between (transformed) variables are linear with constant variance.

Recently, nonparametric resampling procedures have been proposed to avoid the specification of multivariate distributions (Hughes *et al.*, 1993; Lall and Sharma, 1996; Zorita *et al.*, 1995). For the joint simulation of daily

precipitation and other weather variables, resampling from nearest-neighbours or analogues is perhaps the most promising technique. The method does not require the partitioning of continuous variables into discrete states. Furthermore, nearest-neighbour resampling easily allows for linkage with the atmospheric circulation.

Resampling from analogues or nearest-neighbours is rather new in hydrological and meteorological literature. Zorita *et al.* (1995) used the analogue method to generate multi-site daily precipitation in a climate-change study. The generated amounts for a given day were set equal to the observed amounts for the day in the historical record with the most similar sea-level pressure field. In a later paper (Cubasch *et al.*, 1996), the conditioning was done on both sea-level pressure and 700 hPa temperature. Lall and Sharma (1996) discussed a nearest-neighbour bootstrap for generating hydrologic time series. Resampling was done from the successors to the historical k nearest-neighbours of the values generated for the previous day, rather than taking the observed precipitation for the closest neighbour, in terms of the atmospheric circulation, as in Zorita *et al.* (1995) and Cubasch *et al.* (1996). For a high-elevation site in Utah, Rajagopalan and Lall (1995) compared nearest-neighbour resampling with Richardson's (1981) method for generating multivariate weather data (precipitation, maximum and minimum temperature, solar radiation, dew point temperature and wind speed). Various properties of the daily values (quartiles and skewness of their distributions, lag 0 and lag 1 cross-correlation coefficients, and lag 1 autocorrelation coefficients) were best reproduced by nearest-neighbour resampling.

Different variables have been used to find the analogues in the historical data. In contrast to the papers in the meteorological literature, Rajagopalan and Lall (1995) did not consider the atmospheric circulation. It is also not known how far nearest-neighbour resampling can reproduce properties of extreme rainfall, which is important for hydrologic design. In this paper, the nearest-neighbour resampling technique of Rajagopalan and Lall is studied for generating single-site values of precipitation and temperature in the German part of the Rhine basin. Four different sets of variables are compared for the selection of analogues. The emphasis in this comparison is on the reproduction of the autocorrelation structure and the distribution of multi-day maximum precipitation amounts. Snow accumulation and melt are also dealt with.

The paper is organised as follows. Firstly, the necessary background of the nearest-neighbour method is presented. Secondly, the results for seven German stations are discussed. Finally, some overall conclusions are given.

Methodology

NEAREST-NEIGHBOUR RESAMPLING

The principle of nearest-neighbour resampling is presented in Rajagopalan and Lall (1995) and Lall and Sharma (1996). The main points are repeated here. A distinction is made between unconditional simulation of weather variables and conditional simulation of weather variables given the atmospheric circulation.

The generation of precipitation P and temperature T for day t requires a feature vector $\mathbf{D}_t$ to find analogue situations in the historical data. In the method of Rajagopalan and Lall (1995) for generating multivariate daily weather data, $\mathbf{D}_t$ would contain the values of P and T, generated for day $t-1$. The k nearest-neighbours of $\mathbf{D}_t$, in terms of Euclidean distance, are abstracted from the historical record. Let $t(j), j = 1, \ldots, k$, be the times associated with these nearest-neighbours, such that the distance of $\mathbf{D}_{t(j)}$ to $\mathbf{D}_t$ increases with increasing j. The vector of weather variables following $\mathbf{D}_{t(j)}$, the successor to $\mathbf{D}_{t(j)}$, is denoted as $\mathbf{x}_{t(j)}$. One of the successors to the k nearest-neighbours is sampled using a discrete probability distribution or kernel $\{p_j\}$. In Lall and Sharma (1996) the following decreasing kernel was recommended:

$$p_j = \frac{1/j}{\sum_{j=1}^{k} 1/j}, \quad j = 1, \ldots, k \qquad (1)$$

In this paper $k = 20$ is used for resampling from observed records of 30 years. A sensitivity study showed that a reasonable reproduction of extreme-value properties could be obtained with a rather broad range of k values (Brandsma and Buishand, 1997).

Figure 1 presents four methods to find analogues in the historical record. Methods 1 to 3 are examples of conditional simulation of weather variables given the atmospheric circulation. The vector $\mathbf{C}$ in Fig. 1 consists of indices that characterise the atmospheric circulation. In method 1 there are only observed circulation indices for day t in the feature vector $\mathbf{D}_t$. Analogues of the observed circulation for the day of interest form also the starting point in the study of Zorita *et al.* (1995). In order to improve the reproduction of persistence of daily rainfall, method 2 also considers the wet/dry status generated for the previous day. The search for analogues is here restricted to days with the same wet/dry status as day $t-1$. The same restriction is imposed in the resampling procedure of Hughes *et al.* (1993) and Conway *et al.* (1996). In this paper, methods 1 and 2 are extended by including the observed circulation indices on day $t-1$ and the values of P and T that were generated for day $t-1$ (method 3). Resampling occurs from the observed precipitation and temperatures on the days $t(j)$ in the nearest neighbourhood.

Method 4 is an example of unconditional simulation. This method extends the method of Rajagopalan and Lall (1995) by adding atmospheric circulation indices to the weather variables in $\mathbf{D}_t$.

To account for the systematic annual cycle in the various weather variables, the search for the k nearest-neighbours of the feature vector is restricted to days in a

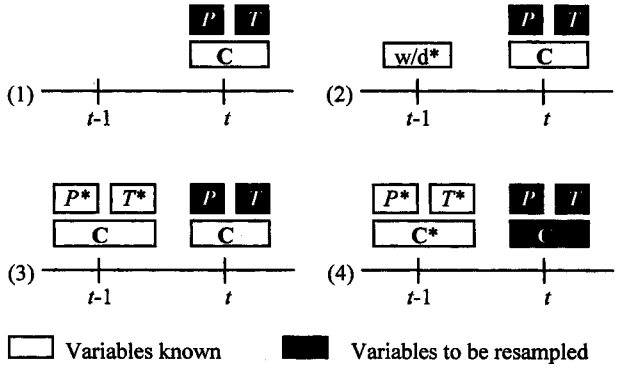

Fig. 1. *Four methods for the generation of new variables (solid squares) using different sets of known variables (open squares). P refers to precipitation, T to temperature, C is a vector of circulation indices and w/d indicates whether the day concerned is wet or dry. The asterisks indicate that the corresponding variables are resampled values of the previous time step.*

specified moving window of width W_{mw} days, centred at the day of interest (Fig. 2). The use of a moving window, instead of fixed seasons, prevents sharp transitions between seasons. Here $W_{mw} = 61$ days is used (Brandsma and Buishand, 1997). For a historical time series of 30 years, the Euclidean distances for a specific day are then calculated for $61 \times 30 = 1830$ days.

A further reduction of seasonal variation can be achieved by working with standardised variables. In Rajagopalan and Lall (1995) standardisation was done by subtracting the calendar day's mean m_d and dividing by the calendar day's sample standard deviation s_d:

$$\tilde{x}_t = (x_t - m_d)/s_d,\ t = 1,\ldots,\ n;\ d = (t-1) \bmod 365 + 1 \quad (2)$$

where x_t and $\tilde{x}_t$ are the original and standardised variable for day t, respectively, and n is the total number of days in the time series. For variables with a normal or almost normal distribution, $\tilde{x}_t$ usually takes values between −3 and +3. However, for daily precipitation the range of $\tilde{x}_t$ is quite different. For a dry day, $\tilde{x}_t = -m_d/s_d \oplus -0.5$ at a lowland station in the Rhine basin, whereas for days with heavy precipitation $\tilde{x}_t \oplus 10$. After resampling a vector of standardised values for day t, the inverse of Eqn. (2) is applied to each element to get the simulated weather variables $x_{t,\mathrm{sim}}$:

$$x_{t,\mathrm{sim}} = m_d + s_d\, \tilde{x}_{t(j)} \quad (3)$$

where $\tilde{x}_{t(j)}$ is the corresponding resampled value.

For daily precipitation, $x_{t,\mathrm{sim}}$ can be negative because the mean and standard deviation for day t generally differ from those for day $t(j)$. Therefore, instead of using Eqn. (2), the observed daily precipitation was in this study standardised by:

$$\tilde{x}_t = x_t / m_{d,\mathrm{wet}} \quad (4)$$

where $m_{d,\mathrm{wet}}$ is the calendar day's mean precipitation for wet days. Such a division by the mean is popular in hydrology to standardise non-negative variables. For dry days $\tilde{x}_t = 0$ and for the most extreme wet days $\tilde{x}_t \oplus 10$.

To reduce the effect of sampling variability, smooth approximations of m_d, $m_{d,\mathrm{wet}}$ and s_d were used instead of the raw values. Smoothing was done with the so-called supersmoother (Härdle, 1990).

Through the standardisation, the elements v_{ti} of the feature vector $\mathbf{D}_t$ are dimensionless quantities. The weighted Euclidean distance between two vectors $\mathbf{D}_t$ and $\mathbf{D}_u$ is given by:

$$\delta_{tu} = \sqrt{\sum_{i=1}^{q} w_i (v_{ti} - v_{ui})^2} \quad (5)$$

where q is the number of variables in $\mathbf{D}_t$ and $\mathbf{D}_u$, and w_i is the weight associated with the ith variable. Table 1 lists the weights used in this paper.

The full resampling procedure comprises the following steps:

1. Determine the composition of the feature vector.
2. Standardise observed variables using Eqns. (2) and (4).
3. Generate data for $t = 1$.
4. Form a feature vector from the most recent generated variables (method 4) and standardised observed circulation indices (method 3). In methods 1 and 2, the feature vector only contains the standardised observed circulation indices for the day of interest.
5. Determine the k nearest-neighbours of the feature vector within a specified window, using the weighted Euclidean distance in Eqn. (5).
6. Sample one of the k nearest-neighbours (conditional methods 1–3) or one of their successors (method 4), and return to step 4 if more simulated values are needed.
7. Retransform the resampled standardised variables to their original scale, using Eqn. (3) or $x_{t,\mathrm{sim}} = m_{d,\mathrm{wet}}\, \tilde{x}_t$ for precipitation.

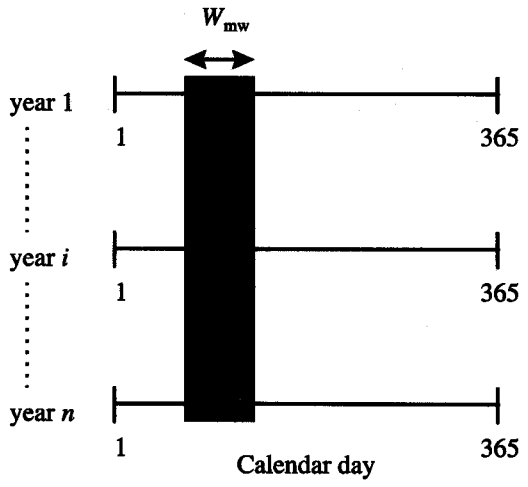

Fig. 2. *Moving window.*

Table 1. Weights applied to the components of the feature vector $\mathbf{D}_t$. The w_i values for the circulation apply to all three components $\tilde{Z}$, $\tilde{W}$ and $\tilde{S}$ of the vector $\tilde{\mathbf{C}}$. The asterisks in $\mathbf{D}_t$, indicate that the corresponding variables are resampled values from the previous time step and the tilde refers to standardised values.

Method	$\mathbf{D}_t$	w_i
1,2	$(\tilde{\mathbf{C}}_t)$	(1)
3	$(\tilde{\mathbf{C}}_t, \tilde{\mathbf{C}}_{t-1}, \tilde{P}^*_{t-1}, \tilde{T}^*_{t-1})$	(1/2,1/2,1,1)
4	$(\tilde{\mathbf{C}}_{t-1}, \tilde{P}^*_{t-1}, \tilde{T}^*_{t-1})$	(1,1,1)

The initialisation in step 3 can be done by random sampling a day within the window for 1 January (method 4) or by resampling a day from the k nearest-neighbours of the observed circulation on the first day (conditional methods). Alternative initialisations are also possible, e.g. starting with the observed weather for $t = 1$.

In the unconditional method 4, the closest neighbour in step 5 (with $\delta_{tu} = 0$) is always the day in the historical record from which the most recent values were resampled. The probability is therefore quite high (⊕0.28 for the kernel used here) that the next day in the historical record is resampled in step 6. Resampling of pairs of successive days in the historical record occurs less frequently in the conditional methods, in particular in method 3.

An important characteristic of a resampling technique is its inability to generate larger 1-day precipitation amounts than the largest historical value (small deviations may occur due to the standardisation of the variables and the subsequent inversion). Multi-day precipitation amounts larger than those observed are a result of other sequences of the historical observations.

From the presentation above, it is clear that there are various options in the nearest-neighbour method. The most important of these is the construction of the feature vector $\mathbf{D}_t$. Once $\mathbf{D}_t$ is constructed, the model can be tuned using the following factors: (1) the number k of nearest-neighbours used for resampling; (2) the width W_{mw} of the running window; (3) the type of kernel used for attaching probabilities p_j to the k nearest-neighbours; (4) the weights w_j in the calculation of distances; and (5) the method used to standardise the variables in the feature vector. Brandsma and Buishand (1997) studied the sensitivity to these factors. This paper focuses on the comparison of feature vectors.

DATA DESCRIPTION

For the research described in this paper, precipitation and temperature data for Essen, Kahler Asten, Trier, Frankfurt, Bamberg, Freudenstadt and Stuttgart were analysed for the period 1961–1990. The stations are situated in the German part of the Rhine basin (Fig. 3). The data were made available by the Deutscher Wetterdienst via the 'International Commission for the Hydrology of the Rhine Basin' (CHR/KHR). Table 2 presents the mean annual temperature and precipitation of these stations, together with the station elevation.

To incorporate atmospheric flow characteristics, daily mean sea level pressure (MSLP) data from the UK Meteorological Office were considered on a 5° latitude by 10° longitude grid. These data extend back to December 1880. For a grid centred at the Rhine basin (see Fig. 4), the following three daily air-flow indices are calculated: (1) total shear vorticity Z; (2) strength of the westerly flow W; and (3) strength of the southerly flow S (see also Jones *et al.*, 1993). These three indices form the elements of the vector $\mathbf{C}$ in Fig. 1.

Before resampling, the data were standardised using the smoothed values of the calendar day's mean, m_d or $m_{d,\mathrm{wet}}$, and standard deviation, s_d. Figure 5 presents m_d and s_d for the relevant variables together with their smooth approximations. Before calculating the smooths, the values for $d = 336, \ldots, 365$ were inserted for $d < 1$ and the values for $d = 1, \ldots, 30$ for $d > 365$ to harmonise the smoothed values at the beginning and end of the year. Smoothing is essential here, because most daily statistics show large sampling variability. Note further that the largest standard deviations of the flow indices (vorticity, strength of the flow) are found in winter. The mean westerly flow is also relatively large in that season. The largest mean wet-day precipitation amounts are found in summer, which is due to the influence of convection (summer showers).

Table 2. Characteristics of the stations that have been used in the study (mean annual values for the period 1961–1990).

Station	Altitude (m above m.s.l.)	Mean annual temperature (°C)	Mean annual precipitation (mm)
Essen	152	9.6	931
Kahler Asten	839	4.9	1476
Trier	265	9.2	784
Frankfurt	112	9.7	658
Bamberg	239	8.5	634
Freudenstadt	797	6.7	1681
Stuttgart	373	8.8	719

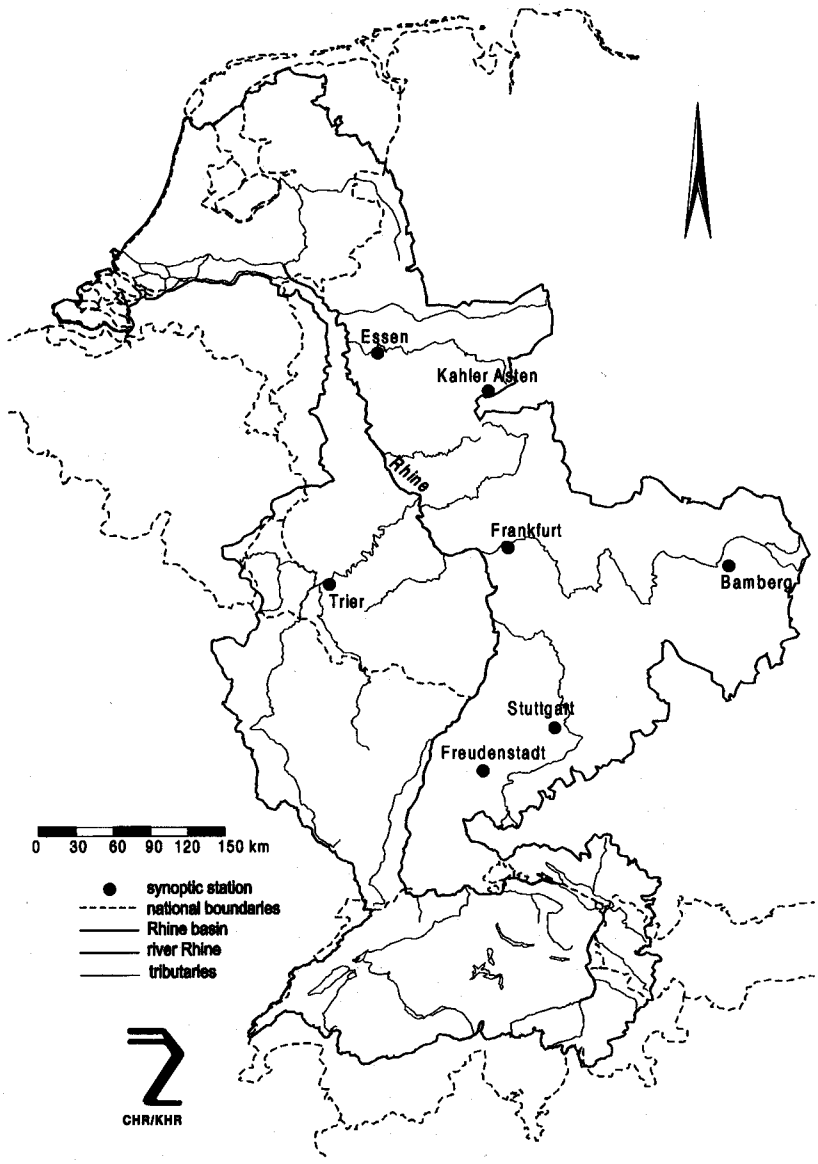

Fig. 3. *Location of the seven German stations used in this study in the drainage basin of the river Rhine.*

Results

For all four methods in Fig. 1, ten 30-year simulations were carried out. For the unconditional simulation, these 30-year runs are independent through the use of different random number seeds. For the conditional methods, there is some correlation because each run is related to the observed circulation. Unless specified otherwise, the results in this section apply to all ten simulations.

The results for Stuttgart are used to show that methods 1 and 2 are inferior to methods 3 and 4 with respect to the reproduction of the autocorrelation coefficients and the extreme-value distributions. The results for the other stations are, therefore, presented only for methods 3 and 4. The reproduction of autocorrelation properties is discussed first. Then, the N-day maximum precipitation and snowmelt amounts are addressed. Finally, simulated long-duration series of 1000 years are considered.

AUTOCORRELATION AND VARIABILITY OF MONTHLY VALUES

The reproduction of the occurrence of extreme multi-day precipitation amounts requires that not only the lag 1 autocorrelation coefficient is preserved, but also the higher order autocorrelation coefficients. Figure 6 compares for Stuttgart the lag 1, 2 and 3 autocorrelation coefficients of daily precipitation and temperature for the historical data, with those for methods 1 to 4 (Fig. 1). For the historical data the standard error se is also presented. The autocor-

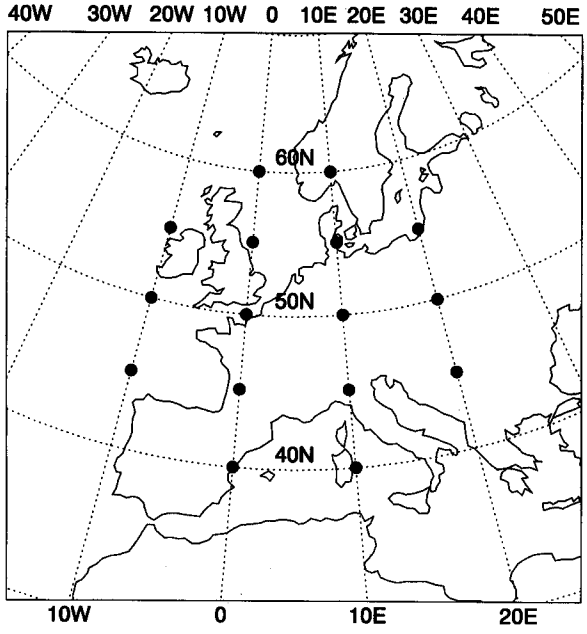

Fig. 4. *Grid points of mean sea-level pressure used for the calculation of the air-flow indices over the Rhine basin.*

relation estimates for the simulated data have a much smaller se because they are based on ten 30-year simulations rather than a single 30-year record. Both the autocorrelation coefficients and their se were estimated with the jackknife method of Buishand and Beersma (1993). For precipitation, the figure shows a clear seasonal cycle of the autocorrelation coefficients. In winter, precipitation is dominated by widespread frontal rainfall resulting in relatively large autocorrelation coefficients. In contrast, the large contribution of convective precipitation (showers) to summer rainfall results in a relatively weak autocorrelation of the daily values in that season. For temperature (all 3 lags) and precipitation (lag 1), it is immediately seen that the autocorrelation is seriously underestimated in methods 1 and 2. Making allowance for the wet/dry status of the previous day in the search for analogues, gives only a small improvement in the lag 1 precipitation autocorrelation. The reproduction of the autocorrelation is much improved by using $\tilde{P}^*_{t-1}$ and $\tilde{T}^*_{t-1}$ in $\mathbf{D}_t$ (methods 3 and 4). Although $\mathbf{D}_t$ does not contain circulation indices of day t in case of unconditional simulation of P and T (method 4), it gives similar results as the most advanced method of conditional simulation (method 3).

For the historical and simulated data of Stuttgart, Table 3 presents the mean lag 1, 2 and 3 autocorrelation coefficients

$$\bar{r}(l) = \sum_{m=1}^{12} r_m(l)/12, \quad l = 1,2,3 \tag{6}$$

with $r_m(l)$ the estimated lag l autocorrelation coefficient for month m as shown in Fig. 6. The table also lists the mean standard deviation of the monthly precipitation totals and monthly mean temperatures

$$\bar{s} = \sum_{m=1}^{12} s_m/12 \tag{7}$$

For the simulated data, the monthly standard deviations tend to be too small if their autocorrelation is not adequately reproduced. In contrast to the comparisons for each individual lag, a test on the monthly standard deviation considers all lags simultaneously. This test is useful to discover a too rapid decay of the autocorrelation coefficients with increasing lag in the simulated data. Such a departure has often been observed with unconditional simulation of daily rainfall using simple parametric models (Buishand, 1978; Katz and Parlange, 1996).

In order to judge the statistical significance of the differences between the historical and simulated values, the standard errors of $\bar{r}(l)$ and $\bar{s}$ were calculated for the 30-year historical record. The sampling variability of these estimates would be the main source of differences between the historical and simulated values, if the resampling method preserves the autocorrelation structure. The se of $\bar{r}(l)$ was obtained by the jackknife method of Buishand and Beersma (1993) and se($\bar{s}$) was obtained by a bootstrap technique (Appendix A). A criterion of 2 × se is used in Table 3 to indicate significant differences between historical and simulated values. Neglecting sampling variability

Table 3. Mean lag 1, 2 and 3 autocorrelation coefficients of daily values and mean standard deviation of monthly values for the historical record of Stuttgart (1961–1990) and simulated data (ten runs of 30 years for each method). Estimates in italics for the simulated data indicate that they differ more than 2 × se from the corresponding estimate for the historical data.

	Precipitation				Temperature			
	$\bar{r}(1)$	$\bar{r}(2)$	$\bar{r}(3)$	$\bar{s}$ (mm)	$\bar{r}(1)$	$\bar{r}(2)$	$\bar{r}(3)$	$\bar{s}$ (°C)
Stuttgart	0.205	0.066	0.032	30.4	0.789	0.571	0.434	1.74
method 1	*0.043*	*0.025*	0.018	*25.3*	*0.288*	*0.190*	*0.146*	*1.02*
method 2	*0.102*	*0.045*	0.022	*26.7*	*0.307*	*0.183*	*0.136*	*1.01*
method 3	0.181	0.062	0.024	28.9	*0.721*	*0.506*	*0.368*	*1.45*
method 4	0.187	0.061	0.018	27.8	*0.750*	*0.531*	*0.373*	*1.51*

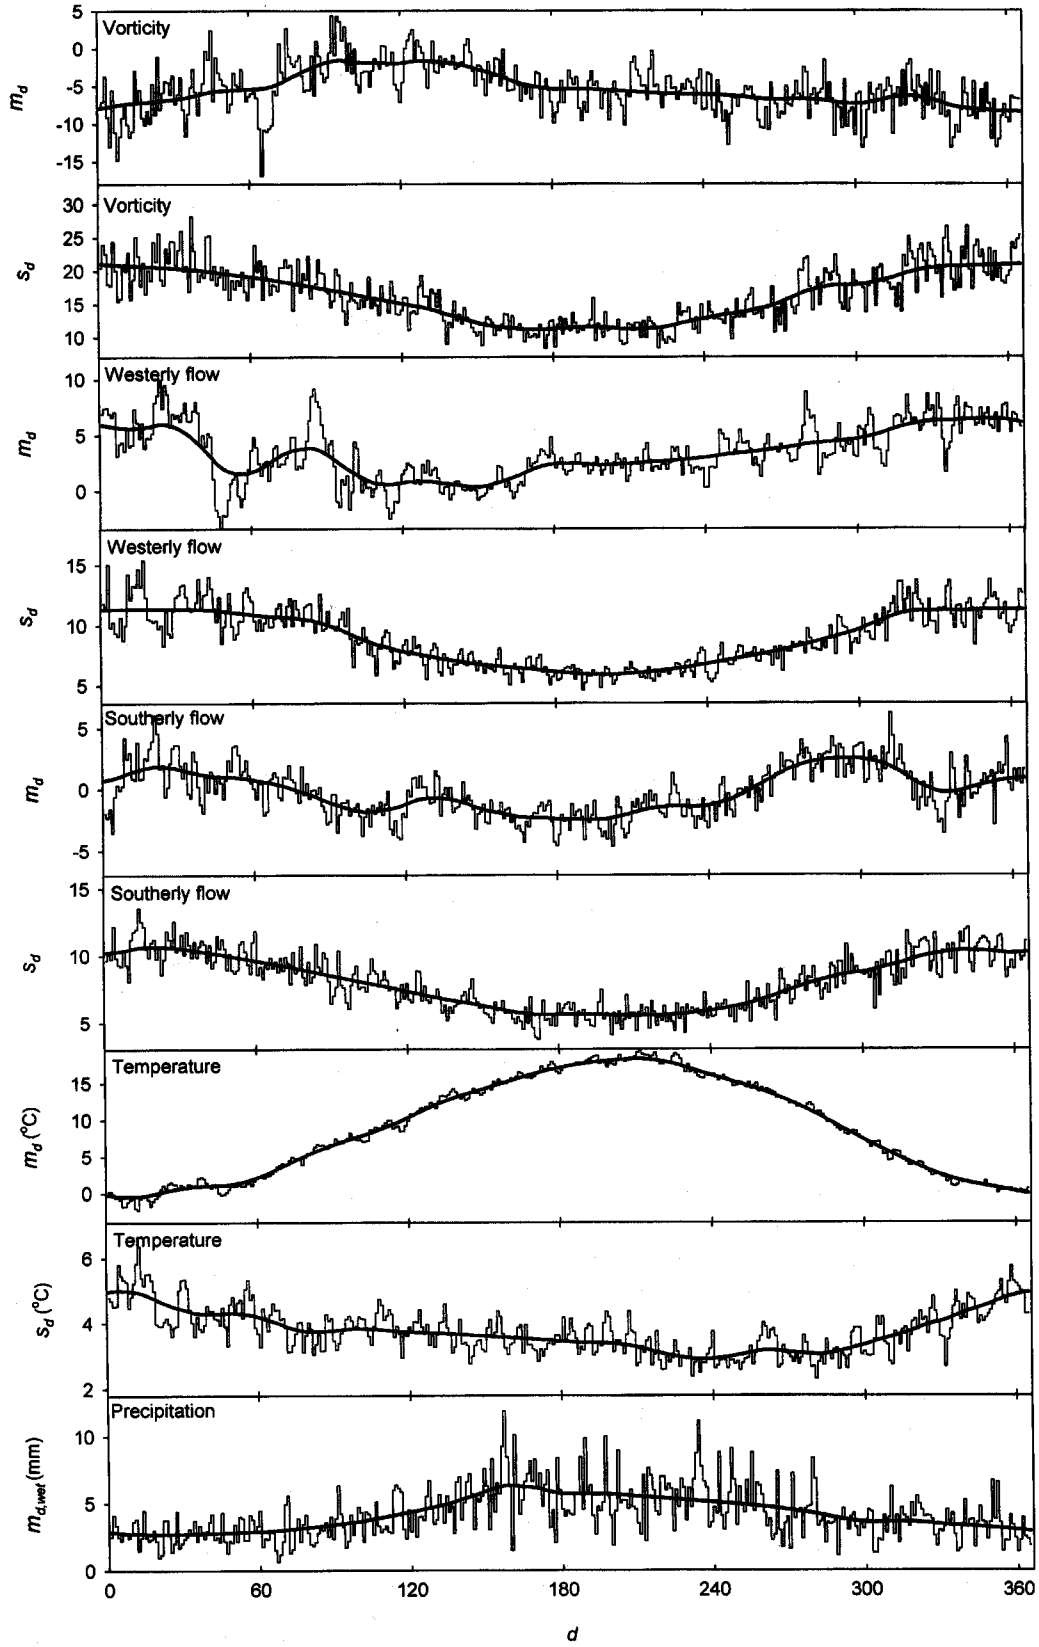

Fig. 5. *Values of m_d and s_d for vorticity (Z), westerly flow (W), southerly flow (S), temperature (T) and $m_{d,wet}$ for precipitation (P) together with their smooth approximations (P and T values apply to Stuttgart) as a function of calendar day d for the period 1961–1990. Vorticity and flow units are geostrophic, expressed as hPa per 10° latitude at 50°N (1 unit is equivalent to 0.73 × $10^{-6}s^{-1}$ and 0.65 ms^{-1} for vorticity and flow, respecively). The smooth curves are based on the supersmoother (Härdle, 1990).*

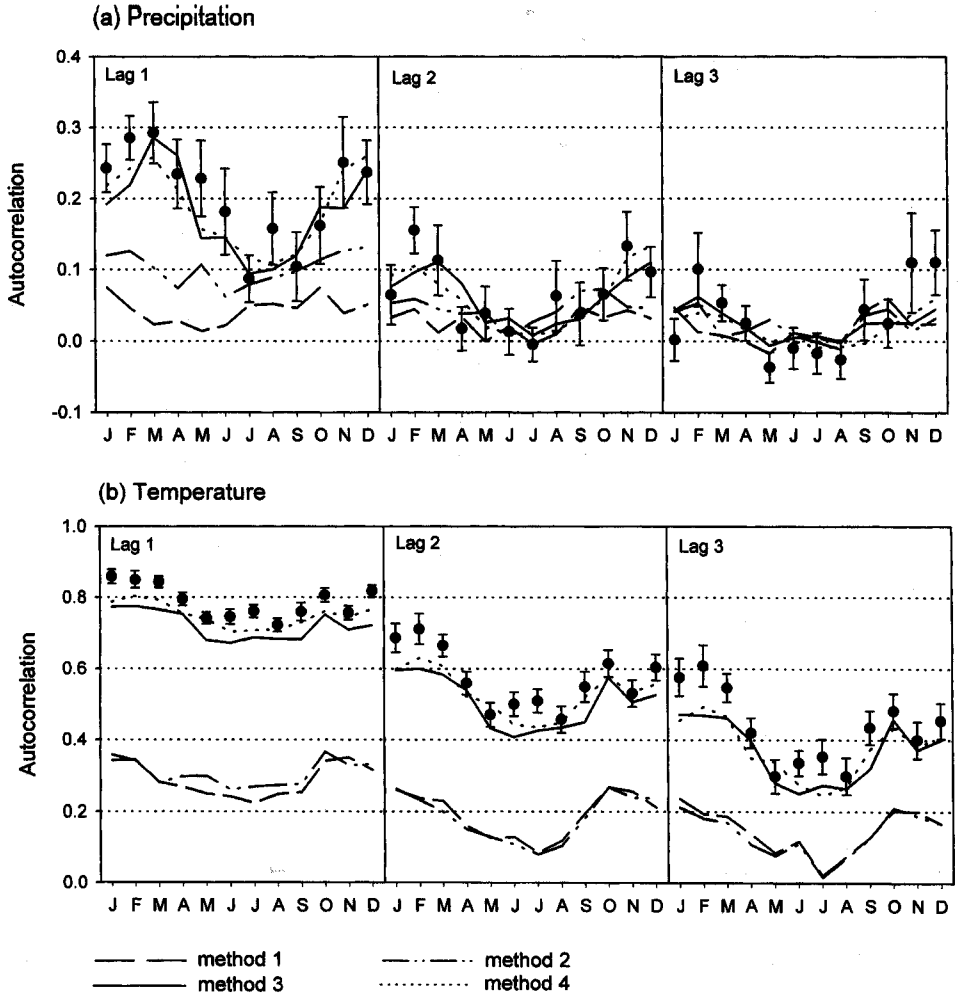

Fig. 6. *Lag 1, 2 and 3 autocorrelation coefficients of observed daily precipitation and temperature (dots plus 1 × se bars) at Stuttgart and simulated precipitation and temperature (average estimates for ten runs for each method) for each month of the year.*

of the average estimates for the simulated data, this criterion corresponds to a two-sided test at the 5% level because $\bar{r}(l)$ and $\bar{s}$ are approximately normally distributed.

Table 3 confirms what was observed in Fig. 6. For precipitation, methods 1 and 2 are unable to reproduce $\bar{r}(l)$ (lag 1 and 2) and $\bar{s}$, whereas for temperature all methods fail. Though the absolute differences between the historical and simulated temperature autocorrelation coefficients for methods 3 and 4 are small, the differences are statistically significant because of the small standard errors. The differences can be reduced by leaving out the circulation indices in $\mathbf{D}_t$, that is by unconditional joint simulation of P and T. However, in this case, the reproduction of the autocorrelation properties of daily rainfall deteriorates (Brandsma and Buishand, 1997), in agreement with experiences with simple parametric stochastic rainfall models. It should further be noted, that the autocorrelation coefficients of the vorticity indices appear to be too low in the case of unconditional simulation (method 4). The systematic departure from the observed lag 1 autocorrelation coefficient is about 0.05, which is comparable to that for the simulated daily temperatures. The autocorrelation coefficients of the simulated vorticity indices also show a too rapid decay with increasing lag.

Table 4 presents values of $\bar{r}(l)$ and $\bar{s}$ for all seven stations for methods 3 and 4. For temperature there is a small, but significant, underprediction of $\bar{r}(l)$ and $\bar{s}$ for all stations. The results are better for precipitation, where only a weak tendency for underprediction of $\bar{r}(l)$, in particular for method 3, and $\bar{s}$ is found. The better reproduction of the lower-order autocorrelation coefficients of daily precipitation in method 4, can partly be ascribed to the fact that this method often selects successive days in the historical record. If the day from which the most recent values were sampled is excluded in the search for nearest-neighbours, method 4 still performs slightly better than method 3 with respect to these autocorrelation coefficients.

Table 4. Mean lag 1, 2 and 3 autocorrelation coefficients of daily values and mean standard deviation of monthly values for the historical records of seven stations (1961–1990) and simulated data (ten runs of 30 years for each method). Estimates in italics for the simulated data indicate that they differ more than 2 × se from the corresponding estimate for the historical data.

Station	Precipitation				Temperature			
	$\bar{r}(1)$	$\bar{r}(2)$	$\bar{r}(3)$	$\bar{s}$ (mm)	$\bar{r}(1)$	$\bar{r}(2)$	$\bar{r}(3)$	$\bar{s}$ (°C)
Essen	0.222	0.113	0.074	36.9	0.796	0.580	0.443	1.76
method 3	*0.193*	*0.090*	0.059	35.3	*0.718*	*0.501*	*0.360*	*1.44*
method 4	0.216	0.122	0.069	35.0	*0.753*	*0.533*	*0.376*	*1.47*
Kahler Asten	0.359	0.192	0.104	61.8	0.790	0.563	0.420	1.81
method 3	*0.309*	*0.145*	0.089	*55.4*	*0.723*	*0.502*	*0.356*	*1.50*
method 4	0.338	0.182	0.102	57.4	*0.747*	*0.519*	*0.359*	*1.54*
Trier	0.235	0.101	0.062	33.2	0.800	0.585	0.447	1.74
method 3	0.216	0.088	0.057	32.0	*0.726*	*0.511*	*0.370*	*1.39*
method 4	0.229	0.109	0.054	31.4	*0.761*	*0.545*	*0.389*	*1.54*
Frankfurt	0.221	0.087	0.051	31.2	0.792	0.579	0.446	1.73
method 3	*0.187*	0.075	0.041	*27.9*	*0.725*	*0.514*	*0.373*	*1.39*
method 4	0.202	0.085	0.042	*27.9*	*0.749*	*0.533*	*0.380*	*1.43*
Bamberg	0.198	0.080	0.065	28.3	0.784	0.576	0.445	1.80
method 3	0.175	0.068	*0.036*	26.0	*0.722*	*0.510*	*0.368*	*1.42*
method 4	0.196	0.076	0.044	*25.8*	*0.740*	*0.524*	*0.372*	*1.48*
Freudenstadt	0.339	0.168	0.079	80.3	0.781	0.551	0.411	1.84
method 3	*0.296*	*0.129*	0.070	*70.7*	*0.717*	*0.495*	*0.352*	*1.58*
method 4	0.322	0.164	0.090	76.1	*0.743*	*0.507*	*0.342*	*1.48*
Stuttgart	0.205	0.066	0.032	30.4	0.789	0.571	0.434	1.74
method 3	0.181	0.062	0.024	28.9	*0.721*	*0.506*	*0.368*	*1.45*
method 4	0.187	0.061	0.018	27.8	*0.750*	*0.531*	*0.373*	*1.51*

N-DAY ANNUAL MAXIMUM PRECIPITATION AMOUNTS

For the seven stations, the *N*-day (N = 1,4,10,20) annual maximum precipitation amounts were abstracted from the historical record and all simulated cases. As an example, Fig. 7 compares for Stuttgart the Gumbel plots of the observed 10-day annual maxima with those of a 30-year simulation for each of the methods 1 to 4. Figure 7 gives, of course, only a rough indication of the differences between the four methods, because it applies to only one of the ten 30-year runs. For methods 3 and 4 the figure shows no systematic differences between the observed and simulated values. The other nine 30-year simulations exhibit similar behaviour. In contrast, for methods 1 and 2 the 10-day annual maxima distribution is only reasonably reproduced in about half of the ten runs, the other runs display systematic underprediction similar to that shown in Fig. 7.

For an objective verification of the reproduction of the *N*-day annual maxima distributions for N = 1,4,10 and 20, the following three quantities are considered:

1. The maximum of the *N*-day annual maxima (highest *N*-day amount in the record).
2. The upper quintile mean QM5 of the *N*-day annual maxima. For 30 years of data this is the mean of the 6 largest annual maxima. This mean value has an average return period of 12.5 years (Appendix B).
3. The median M of the *N*-day annual maxima.

Expressions for the standard errors of these quantities are given in Appendix B.

For Stuttgart, a comparison between observed and simulated quantities (methods 1 to 4) is given in Table 5. Although the observed quantities differ not more than twice their se from the average values in the simulated data, methods 1 and 2 exhibit a marked tendency of underestimation for $N > 1$. This is related to the underestimation of the autocorrelation coefficients in these methods. It is obvious from Table 5 that methods 3 and 4 perform much better. This becomes more apparent if the same quantities are compared for the maxima in the winter half-year (October–March), because the most serious

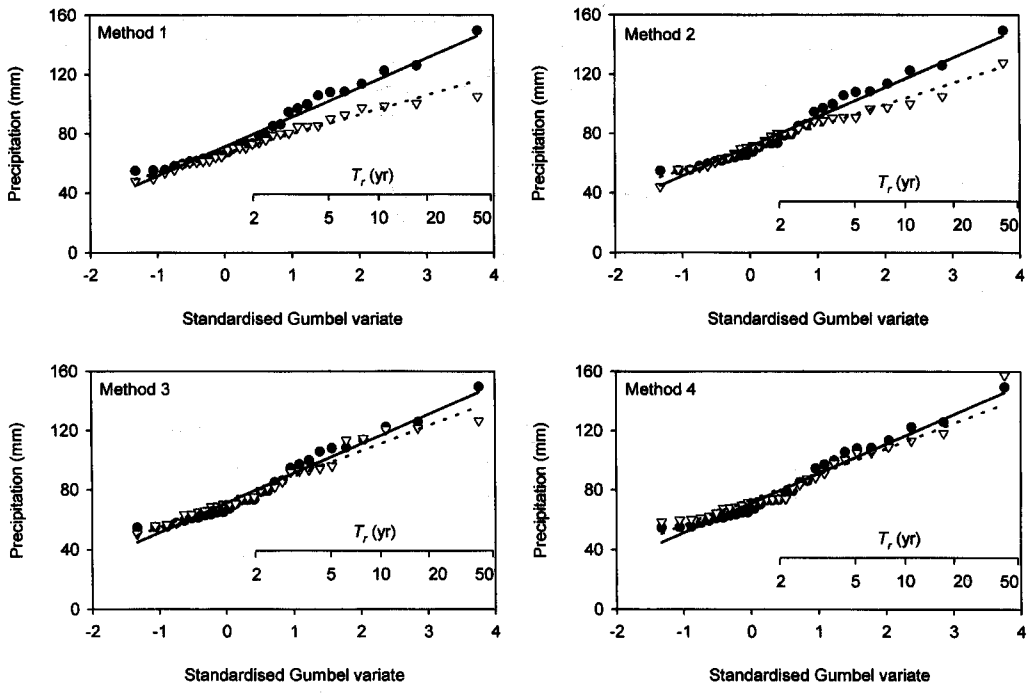

Fig. 7. *Gumbel plots of 10-day annual maxima for observed precipitation at Stuttgart (solid dots, solid lines) and simulated precipitation for each of the methods 1 to 4 (open triangles, dashed lines). The solid and dashed lines are probability-weighted moment fits (Landwehr et al., 1979) to the historical and simulated data, respectively. The annual maxima are plotted using the median plotting position.*

underestimation of the autocorrelation occurs in that part of the year (Fig. 6).

For methods 3 and 4, Table 6 presents the results for all seven stations. There is no obvious preference for one of these two methods. The maximum, the upper quintile mean and the median are all well reproduced. The average percentage difference is small for these quantities and the magnitude of the differences between the observed and simulated values for the individual stations is generally not more than that expected from the standard errors.

N-DAY MAXIMUM SNOWMELT

Large river discharges may be partly caused by snowmelt. Therefore, the reproduction of snowmelt is also considered.

Historical estimates and simulated values of snowmelt were derived from daily precipitation and temperature. It was assumed that for $T < 0$ precipitation accumulates on the surface as snow. To calculate the N-day maximum snowmelt, snow was transformed into snowmelt using the degree days method. In that method, the amount of snowmelt on a certain day is proportional to the temperature excess (number of degrees Celsius above freezing point on that day), as long as there is solid precipitation stored on the surface. The constant of proportionality is known as the degree days factor (mm/°C). This factor was set equal to 4, which is an average of the values found in the literature (Linsley *et al.*, 1988; Gray and Prowse, 1993).

Analogous to Table 6, Table 7 shows the results for the

Table 5. Maximum, upper quintile mean and median of the N-day annual precipitation maxima for the historical record of Stuttgart (1961–1990) and simulated data (averages for ten runs of 30 years each).

Station	Maximum (mm)				Upper quintile mean (mm)				Median (mm)			
	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$
Stuttgart	68.3	119.2	149.7	209.1	56.3	94.2	121.5	161.2	36.2	59.6	73.4	111.5
method 1	63.3	92.5	133.2	183.3	53.3	78.3	109.9	154.0	35.1	53.8	76.6	109.7
method 2	65.4	93.9	130.2	175.1	56.0	77.3	107.9	146.8	36.1	51.3	74.3	107.3
method 3	62.5	110.8	158.4	216.8	55.7	87.9	123.3	163.9	37.1	57.1	78.9	112.2
method 4	62.8	108.9	151.0	194.5	53.8	86.0	116.1	158.8	37.3	56.2	78.9	109.6

Table 6. Maximum, upper quintile mean and median of the N-day annual precipitation maxima for the historical records of seven stations (1961–1990) and simulated data (averages for ten runs of 30 years each). The two bottom lines give the percentage difference between the values for the historical and simulated data, averaged over all stations. Estimates in italics for the simulated data indicate that they differ more than 2 × se from the corresponding estimate for the historical data.

Station	Maximum (mm)				Upper quintile mean (mm)				Median (mm)			
	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$
Essen	81.7	133.5	159.8	192.7	51.5	88.3	125.6	178.2	30.4	56.9	84.8	128.4
method 3	65.0	113.2	151.5	202.3	48.9	84.8	120.4	166.6	33.2	56.1	86.8	125.9
method 4	*51.9*	*95.7*	137.0	194.1	44.1	80.1	117.8	167.5	33.2	56.6	85.2	125.7
Kahler Asten	68.8	147.0	232.9	352.3	63.1	129.4	207.9	306.8	46.2	96.2	142.7	214.8
method 3	68.3	144.5	232.8	345.1	60.5	125.5	199.3	295.5	44.2	92.7	145.0	213.0
method 4	68.6	158.2	247.6	348.3	60.3	129.6	204.1	290.5	44.8	93.7	149.8	215.6
Trier	57.8	95.9	113.5	144.7	50.9	80.3	109.3	137.9	33.4	61.4	88.3	123.4
method 3	61.9	98.9	138.6	*183.2*	51.1	81.6	114.8	*154.7*	32.5	*55.4*	*80.4*	*110.9*
method 4	57.9	92.7	126.6	176.0	50.8	75.9	107.7	147.6	33.7	*55.2*	*80.0*	*112.4*
Frankfurt	109.7	159.8	164.6	191.5	66.4	109.6	128.5	158.6	30.8	50.0	78.8	117.2
method 3	85.7	128.1	152.1	186.6	59.0	92.7	115.8	148.8	30.0	51.0	71.9	*99.1*
method 4	81.3	132.7	154.4	182.2	57.0	89.8	112.8	146.7	32.9	51.4	72.2	*100.7*
Bamberg	75.3	93.4	126.2	182.9	57.8	77.0	102.3	145.7	31.5	53.0	69.4	98.7
method 3	69.0	100.6	130.1	163.7	55.2	77.9	104.6	138.6	32.2	50.0	69.9	97.2
method 4	67.6	89.5	118.4	159.9	54.6	76.3	100.9	134.0	31.8	48.3	67.7	95.9
Freudenstadt	112.6	246.4	295.7	443.9	104.0	213.3	276.5	403.2	72.0	143.9	213.3	303.4
method 3	115.0	246.6	335.3	453.8	96.9	196.5	279.0	376.9	69.0	128.1	*188.2*	*256.6*
method 4	114.9	258.2	370.6	530.4	101.0	206.4	309.1	425.4	67.0	135.2	205.5	280.8
Stuttgart	68.3	119.2	149.7	209.1	56.3	94.2	121.5	161.2	36.2	59.6	73.4	111.5
method 3	62.5	110.8	158.4	216.8	55.7	87.9	123.3	163.9	37.1	57.1	78.9	112.2
method 4	62.8	108.9	151.0	194.5	53.8	86.0	116.1	158.8	37.3	56.2	78.9	109.6
% diff (3)	–7.3	–4.7	4.5	3.2	–4.6	–4.8	–1.2	–2.0	0.0	–4.8	–2.5	–6.4
% diff (4)	–11.3	–7.0	2.5	2.3	–6.6	–6.5	–2.3	–2.3	1.5	–4.4	–1.5	–5.3

N-day maximum snowmelt in the winter half-year (October–March). The contrast between the two highest stations, Kahler Asten and Freudenstadt, and the other stations is striking. For Kahler Asten and Freudenstadt, the maximum of the 10-day snowmelt is of the same order of magnitude as the corresponding maximum precipitation amount in Table 6, whereas for the other stations it is less than a third of that value. Despite the systematic underprediction of the autocorrelation of daily temperatures, the correspondence between the historical and simulated values is quite good in Table 7.

LONG-DURATION SIMULATION

Simulation of synthetic rainfall sequences can be useful to study the effect of unprecedented extreme rainfall on river discharges. An interesting question is, therefore, how far nearest-neighbour resampling can provide higher N-day amounts than those observed. The average maxima of the ten 30-year runs in Table 6 are already sometimes larger than those observed in the historical record. To investigate this further, a 1000-year simulation for all seven stations was carried out. Simulation was restricted to method 4 only, because the conditional methods would need an additional model for generating circulation indices for that purpose.

Figure 8 compares for Kahler Asten and Stuttgart the Gumbel plots of the 1-day and 10-day annual maximum precipitation amounts in the historical record with those in the 1000-year simulation. Because the resampled 1-day values are bounded by the observed daily amounts, the plots of the simulated 1-day annual maxima suddenly flattens near the position of the largest historical values. For the peak discharges of large river basins, like the Rhine basin, the 1-day precipitation maxima are of less importance than e.g. the 10-day maxima. The Gumbel plot of the 10-day annual maxima in the 1000-year simulation for Kahler Asten shows values up to 74 mm larger (⊕ 32%) than the largest historical value. A similar exceedance of the largest 10-day values (58 mm, 39%) is found for

Table 7. Maximum, upper quintile mean and median of the N-day maximum snowmelt for the historical records of seven stations (1961–1990) and simulated data (averages for ten runs of 30 years each). The two bottom lines give the percentage difference between the values for the historical and simulated data, averaged over all stations. Estimates in italics for the simulated data indicate that they differ more than 2 × se from the corresponding estimate for the historical data.

Station	Maximum (mm)				Upper quintile mean (mm)				Median (mm)			
	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$	$N = 1$	$N = 4$	$N = 10$	$N = 20$
Essen	20.4	30.6	32.2	38.7	17.0	22.1	25.3	28.7	6.9	8.7	10.1	12.3
method 3	21.2	25.1	31.0	37.9	15.9	20.1	24.1	29.1	8.3	10.5	12.4	14.0
method 4	20.9	32.7	37.8	43.3	15.4	22.3	25.8	30.7	8.4	11.1	12.5	13.9
Kahler Asten	51.2	164.0	287.2	314.0	35.5	107.7	191.6	241.5	22.3	63.0	90.2	135.8
method 3	45.8	139.6	243.2	320.4	37.8	108.6	179.7	245.8	25.3	65.1	99.3	140.3
method 4	46.2	142.8	235.6	297.6	38.0	112.2	180.3	246.2	25.3	67.2	98.8	136.6
Trier	21.6	30.3	34.8	37.2	15.8	25.8	31.7	32.8	7.7	10.6	12.5	14.8
method 3	20.5	30.8	36.0	42.8	15.0	21.7	25.0	29.0	8.0	9.9	11.6	13.5
method 4	24.1	41.0	44.8	48.3	17.3	26.2	29.5	33.3	8.1	10.5	11.8	14.0
Frankfurt	18.7	24.6	24.6	24.6	11.5	17.6	18.8	19.9	5.2	6.2	7.1	7.2
method 3	17.9	24.3	26.5	29.8	12.3	15.7	17.9	20.7	5.7	6.7	7.5	8.5
method 4	16.5	20.8	22.8	28.2	12.2	16.3	17.8	20.8	6.0	7.0	8.2	9.2
Bamberg	21.4	30.9	30.9	33.9	15.4	23.7	25.5	29.2	7.8	9.7	12.2	14.1
method 3	18.7	28.0	32.3	40.1	14.4	20.9	24.3	28.9	7.3	10.4	11.7	13.9
method 4	21.3	31.7	34.6	39.3	16.1	23.8	26.5	30.5	8.5	11.5	13.4	15.6
Freudenstadt	42.4	126.8	234.7	243.1	36.8	111.1	184.6	216.6	26.2	68.0	93.7	119.4
method 3	43.5	113.3	165.6	214.7	36.0	92.0	*132.0*	167.6	25.2	58.6	76.7	94.8
method 4	47.0	127.9	191.0	238.3	38.1	103.6	146.1	183.6	25.2	58.4	79.9	104.7
Stuttgart	26.4	42.7	42.7	51.8	20.6	32.6	35.4	38.9	10.8	16.6	19.8	21.8
method 3	23.5	38.4	46.7	54.9	19.0	28.9	33.3	39.3	10.2	14.2	16.3	18.9
method 4	28.1	47.8	51.3	56.8	20.7	32.8	36.8	42.6	10.7	14.8	16.6	20.0
% diff (3)	–5.3	–8.9	–3.3	7.0	–2.3	–10.7	–10.9	–4.0	4.6	0.7	–1.1	–1.1
% diff (4)	1.4	4.2	4.9	10.7	3.0	–1.0	–4.2	2.0	8.8	5.7	3.1	3.6

Stuttgart. The simulated 10-day maxima nicely follow the Gumbel distribution. For the other stations similar results were obtained.

Discussion and conclusions

In this paper the joint simulation of daily precipitation and temperature by nearest-neighbour resampling was explored. Seven stations in the German part of the Rhine basin were considered. Both the precipitation and temperature autocorrelation coefficients were seriously underestimated, if the search for nearest-neighbours was only based on circulation indices for the day of interest or circulation indices extended with the wet/dry status of the previous day. The poor reproduction of the daily precipitation autocorrelation resulted in an underestimation of the quantiles of the multi-day annual maxima distributions. A major improvement was achieved by incorporating the simulated precipitation and temperature for the previous day in the feature vector. The distributions of the multi-day annual precipitation maxima could then be satisfactorily reproduced. For the simulated daily temperatures a small, but significant, underprediction of the autocorrelation coefficients remained. Nevertheless, the distributions of N-day maximum snowmelt derived from these temperatures and the simulated precipitation data were very close to those derived from observed data.

With respect to the reproduction of autocorrelation coefficients and extreme-value distributions, unconditional simulation of precipitation, temperature and circulation indices (method 4) turned out to be at least as good as conditional simulation of precipitation and temperature on circulation indices (method 3). In the case of conditional simulation, the length of a simulation run cannot exceed that of the MSLP-data set (about 120 years). Longer simulation runs need a separate stochastic model for generating circulation indices. It is likely that such a model reproduces better the autocorrelation properties of these indices than the unconditional method presented here. In particular, a significant improvement is expected for vorticity. Furthermore, conditional simulation has potential applications for downscaling GCM results.

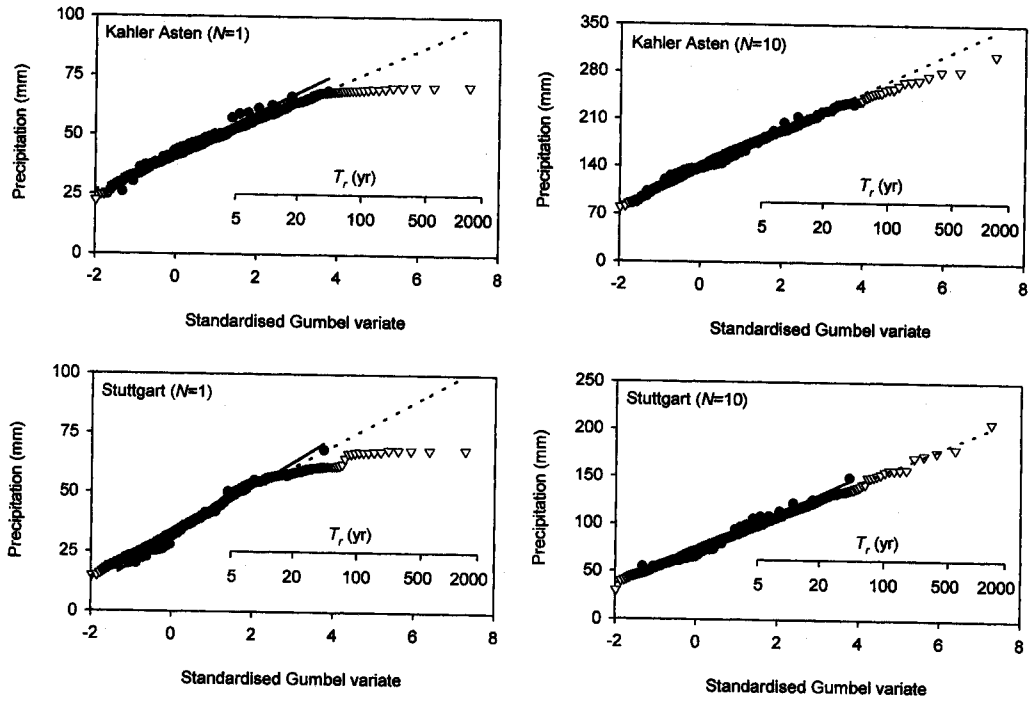

Fig. 8. *Gumbel plots of 1-day and 10-day annual maxima for observed precipitation at Kahler Asten and Stuttgart (solid dots, solid lines) and precipitation in a 1000-year run from resampling method 4 (open triangles, dashed lines). The solid and dashed lines are probability-weighted moment fits (Landwehr et al., 1979) to the historical and simulated data, respectively. The annual maxima are plotted using the median plotting position.*

A difficult question is, how far the length of the observed record puts limitations on the length of a simulation run. For N = 4, 10 and 20 a repetition of large N-day precipitation amounts was not found in the 1000-year simulations in this paper. The distributions of their annual maxima were close to the Gumbel distribution. The most extreme simulated values were generally found to be well above the observed annual maxima.

The application of nearest-neighbour resampling to large river basins requires a multi-site extension. The composition of the feature vector needs further study for such an extension. Although a resampling technique preserves the spatial dependence of the 1-day amounts, this is not necessarily true for the multi-day amounts. The reproduction of the spatial association of large multi-day amounts should therefore be tested.

Large river discharges of the river Rhine in the Netherlands generally occur only during the winter season. The reproduction of extremes during that season is then of particular interest. Methods 3 and 4 are not entirely successful concerning this point (Brandsma and Buishand, 1997). This is related to the weak tendency to underestimate the autocorrelation coefficients and the standard deviation of the monthly totals.

Appendix A: Bootstrap standard errors

The bootstrap estimate of the standard error of a statistic $\hat{\theta}$ is obtained by recomputing $\hat{\theta}$ for a large number of independent bootstrap samples. Each bootstrap sample $X_1^*, X_2^*, \ldots, X_n^*$ is a random sample of size n, drawn with replacement from the original sample $X_1, X_2, \ldots, X_n$ (Efron and Tibshirani, 1993). Let $\hat{\theta}^*(b)$ denote the value of the statistic for the bth bootstrap sample, $b = 1 \ldots, B$. Then the bootstrap estimate of the standard error is the sample standard deviation of the $\hat{\theta}^*(b)$ values:

$$\hat{se}^* = \left\{ \sum_{b=1}^{B} \left[\hat{\theta}^*(b) - \hat{\theta}^*(\circ) \right]^2 / (B-1) \right\}^{\frac{1}{2}} \qquad \text{(A1)}$$

where $\hat{\theta}^*(\circ)$ is the mean of the $\hat{\theta}^*(b)$ values.

Bootstrap estimation of the standard error of the annual mean $\hat{\theta} = \bar{s}$ of monthly standard deviations requires a slight extension of this concept, because each element of the sample is a vector of the twelve monthly precipitation totals of a given year. A bootstrap sample of size n is then obtained by choosing years randomly with replacement, cf. Zwiers and Ross (1991). Every time a particular year is selected the twelve monthly precipitation totals are added to the bootstrap sample.

For each bootstrap sample, the standard deviations $s^*_m(b)$, $m = 1, \ldots, 12$ of the monthly precipitation totals and their annual average $\hat{\theta}^*(b)$ are calculated. The latter is used in Eqn. (A1) to obtain the bootstrap standard error of $\bar{s}$. In this study $B = 500$ bootstrap samples were generated.

Appendix B: Properties of statistics used to compare extreme-value distributions

In this appendix it is assumed that the annual maxima X_i are independent random variables with a Gumbel distribution:

$$\Pr(X_i \leq x) = \exp\{-e^{-(x-\xi)/\delta}\}, \quad -\infty < x < \infty \qquad \text{(B1)}$$

The mean of X_i depends both on ξ and δ:

$$\mathrm{E}(X_i) = \xi + \delta\gamma \qquad \text{(B2)}$$

where $\gamma = 0.5772...$ is Euler's constant. The variance of X_i is given by:

$$\mathrm{var}\, X_i = \pi^2\delta^2/6 \qquad \text{(B3)}$$

The special case $\xi = 0$ and $\delta = 1$ is known as the standard or reduced variable Y_i.

The upper quintile mean QM5 and the median M are simple linear functions of the order statistics $X_{(1)} \leq X_{(2)} \leq \ldots \leq X(n)$:

$$L = \sum_{i=1}^{n} a_i X_{(i)} \qquad \text{(B4)}$$

For QM5 the coefficients are $a_1 = \cdots = a_{24} = 0$ and $a_{25} = \cdots = a_{30} = \frac{1}{6}$, and for the sample median $a_{15} = a_{16} = \frac{1}{2}$, whereas $a_i = 0$ for the other values of i. Writing $X_{(i)} = \xi + \delta Y_{(i)}$, where $Y_{(i)}$ is the ith order statistic in a random sample of size n from the reduced variable, and using the fact that for the two statistics considered here, the sum of the coefficients a_i equals 1, the mean of L becomes:

$$\mathrm{E}(L) = \xi + \delta \sum_{i=1}^{n} a_i \mathrm{E}(Y_{(i)}) \qquad \text{(B5)}$$

The variance of L can be written as:

$$\mathrm{var}\, L = \sum_{i=1}^{n}\sum_{j=1}^{n} a_i a_j \,\mathrm{cov}(X_{(i)}, X_{(j)})$$

$$= \delta^2 \sum_{i=1}^{n}\sum_{j=1}^{n} a_i a_j \,\mathrm{cov}(Y_{(i)}, Y_{(j)}) \qquad \text{(B6)}$$

The means and covariances of the order statistics of the reduced variable can be obtained from tables in Balakrishnan and Chan (1992). For the estimation of var(QM5) and var(M), the scale parameter δ in Eqn. (B6) was replaced by its probability-weighted moment estimate.

For the upper quintile mean QM5, Eqns. (B5) and (B6) reduce to:

$$\mathrm{E(QM5)} = \xi + 2.4840\delta \qquad \text{(B7)}$$

$$\mathrm{var(QM5)} = 0.3130\delta^2 \qquad \text{(B8)}$$

The return period T_r associated with E(QM5) can be obtained as:

$$T_r = \frac{1}{\Pr\{X_i > \mathrm{E(QM5)}\}} = \frac{1}{1 - \exp(-e^{-2.4840})} \approx 12.5 \qquad \text{(B9)}$$

For the sample median M, Eqns. (B5) and (B6) reduce to:

$$\mathrm{E}(M) = \xi + 0.3774\delta \qquad \text{(B10)}$$

$$\mathrm{var}(M) = 0.0671\delta^2 \qquad \text{(B11)}$$

For the return period associated with E(M) it follows $T_r = 2.02$ years. Due to the positive bias of the sample median, this return period slightly differs from the value $T_r = 2$ years for the true median.

A well-known result for the Gumbel variable is that the maximum $X_{\max} = X_{(n)}$ has also a Gumbel distribution with the same scale parameter δ but with location parameter $\xi + \delta \ln n$. Eqn. (B3) therefore also applies to the variance of $X_{\max}$. This result, however, strongly relies on the validity of the Gumbel distribution. Regional analyses of long-duration series of N-day annual maximum precipitation amounts in the United Kingdom and the Low Countries (Dupriez and Demarée, 1988; Dales and Reed, 1989; Buishand, 1991) show that for $N = 1$ the upper tail of the distribution tends to be longer than that of the Gumbel distribution, whereas for large N ($N = 10$ or $N = 20$) the upper tail tends to be shorter. The former implies that the variance of $X_{\max}$ increases with increasing n and will be underestimated if a Gumbel distribution is assumed. The opposite holds if the distribution has a shorter upper tail than the Gumbel distribution. For the German stations used in this study there are, however, no indications of systematic departures from the Gumbel distribution in the upper tail. It should further be noted that the $2 \times$ se criterion should be used with care for the Gumbel distribution. For the Gumbel variable X, it follows from Eqns. (B1), (B2) and (B3) that $\Pr\{X > \mathrm{E}(X) + 2\sigma(X)\} \oplus 0.042$ (compared with 0.023 for the normal distribution) and $\Pr\{X < \mathrm{E}(X) - 2\sigma(X)\} \oplus 0.001$.

Acknowledgements

The authors are grateful to J.J. Beersma for comments on an earlier version of the paper and to B.W.A.H. Parmet and H.C. van Twuiver for discussions during the project. They further thank U. Lall for sending his working paper (co-authored with B. Rajagopalan) on the nearest-neighbour bootstrap for resampling daily precipitation and other weather variables. The UK Meteorological Office gridded MSLP data were kindly provided by P.D. Jones (Climatic Research Unit, University of East Anglia, Norwich). The daily precipitation and temperature data for the German stations were made available by the Deutscher Wetterdienst via the 'International Commission for the

Hydrology of the Rhine Basin' (CHR/KHR). The work was performed in co-operation with the Institute for Inland Water Management and Waste Water Treatment (RIZA), Lelystad.

References

Balakrishnan, N. and Chan, P.S., 1992. Order statistics from extreme value distribution, I: tables of means, variances and covariances, *Commun. Statist.–Simula.*, **21**, 1199–1217.

Bárdossy, A. and Plate, E.J., 1991. Modelling daily rainfall using a semi-Markov representation of circulation pattern occurrence, *J. Hydrol.*, **122**, 33–47.

Bárdossy, A. and Plate, E.J., 1992. Space-time model for daily rainfall using atmospheric circulation patterns, *Wat. Resour. Res.*, **28**, 1247–1259.

Brandsma, T. and Buishand, T.A., 1997. *Rainfall generator for the Rhine basin: single-site generation of weather variables by nearest-neighbour resampling.* KNMI-publication 186–1, KNMI, De Bilt, 47 pp.

Buishand, T.A., 1978. Some remarks on the use of daily rainfall models, *J. Hydrol.*, **36**, 295–308.

Buishand, T.A., 1991. Extreme rainfall estimation by combining data from several sites, *Hydrol. Sci. J.*, **36**, 345–365.

Buishand, T.A. and Beersma, J.J., 1993. Jackknife tests for differences in autocorrelation between climate time series, *J. Climate*, **6**, 2490–2495.

Conway, D., Wilby, R.L. and Jones, P.D., 1996. Precipitation and air flow indices over the British Isles, *Clim. Res.*, **7**, 169–183.

Cubasch, U., von Storch, H., Waszkewitz, J. and Zorita, E., 1996. Estimates of climate change in Southern Europe derived from dynamical climate model output, *Clim. Res.*, **7**, 129–149.

Dales, M.Y. and Reed, D.W., 1989. Regional flood and storm hazard assessment. Report No. 102, Institute of Hydrology, Wallingford, 159 pp.

Dupriez, G.L. and Demarée, G., 1988. *Totaux Pluviométriques sur des Périodes Continues de 1 à 30 Jours: I. Analyse de 11 Séries Pluviométriques de plus de 80 Ans.* Miscellanea Série A, No. 8, Institut Royal Météorologique de Belgique, Bruxelles, 154 pp.

Efron, B. and Tibshirani, R.J., 1993. *An Introduction to the Bootstrap.* Chapman & Hall, New York, 436 pp.

Gray, D.M. and Prowse, T.D., 1993. Snow and Floating Ice. In: *Handbook of Hydrology* (D.R. Maidment, ed.), 7.1–7.58, McGraw-Hill, New York.

Härdle, W., 1990. *Applied Nonparametric Regression.* Cambridge University Press, Cambridge, 333 pp.

Hughes, J.P., Lettenmaier, D.P. and Guttorp, P., 1993. A stochastic approach for assessing the effect of changes in synoptic circulation patterns on gauge precipitation, *Wat. Resour. Res.*, **29**, 3303–3315.

Hutchinson, M.F., 1995. Stochastic space-time weather models from ground-based data, *Agric. For. Meteorol.*, **73**, 237–264.

Jones, P.D., Hulme, M. and Briffa, K.R., 1993. A comparison of Lamb circulation types with an objective classification scheme, *Int. J. Climatol.*, **13**, 655–663.

Katz, R.W. and Parlange, M.B., 1993. Effects of an index of atmospheric circulation on stochastic properties of precipitation, *Wat. Resour. Res.*, **29**, 2335–2344.

Katz, R.W. and Parlange, M.B., 1996. Mixtures of stochastic processes: application to statistical downscaling, *Clim. Res.*, **7**, 185–193.

Lall, U. and Sharma, A., 1996. A nearest-neighbor bootstrap for resampling hydrologic time series, *Wat. Resour. Res.*, **32**, 679–693.

Landwehr, J.M., Matalas, N.C. and Wallis, J.R., 1979. Probability weighted moments compared with some traditional techniques in estimating Gumbel parameters and quantiles, *Wat. Resour. Res.*, **15**, 1055–1064.

Linsley, R.K., Kohler, M.A. and Paulhus, J.L.H., 1988. *Hydrology for Engineers.* McGraw-Hill, London, 492 pp.

Rajagopalan, B. and Lall, U., 1995. A nearest-neighbor bootstrap for resampling daily precipitation and other weather variables. Working paper WP–95-HWR-UL/013, Utah State University, Logan, Utah.

Richardson, C.W., 1977. *A model of stochastic structure of daily precipitation over an area.* Hydrology paper No. 91, Colorado State University, Fort Collins, Colorado, 46 pp.

Richardson, C.W., 1981. Stochastic simulation of daily precipitation, temperature, and solar radiation, *Wat. Resour. Res.*, **17**, 182–190.

Woolhiser, D.A., 1992. Modelling daily precipitation—Progress and problems. In: *Statistics in the Environmental and Earth Sciences*, (A.T. Walden and P. Guttorp, eds.), 71–89, Edward Arnold, London.

Zorita, E., Hughes, J.P., Lettenmaier, D.P. and von Storch, H., 1995. Stochastic characterization of regional circulation patterns for climate model diagnosis and estimation of local precipitation, *J. Climate*, **8**, 1023–1042.

Zwiers, F.W. and Ross, W.H., 1991. An alternative approach to the extreme value analysis of rainfall data, *Atmosphere-Ocean*, **29**, 437–461.