


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

RainSim: A spatial–temporal stochastic rainfall modelling system

A. Burton^{a,*}, C.G. Kilsby^a, H.J. Fowler^a, P.S.P. Cowpertwait^b, P.E. O'Connell^a^aWater Resource Systems Research Laboratory, School of Civil Engineering and Geosciences, Cassie Building, Claremont Road, Newcastle University, Newcastle upon Tyne NE1 7RU, UK^bInstitute of Information and Mathematical Sciences, Massey University, Albany Campus, Private Bag 102-904, Auckland, New Zealand

ARTICLE INFO

Article history:

Received 4 August 2007

Received in revised form 31 March 2008

Accepted 2 April 2008

Available online 3 June 2008

Keywords:

Rainfall

Precipitation

Simulator

Stochastic

Poisson process

Spatial

Temporal

Multi-site

Extremes

NSRP

Shuffled Complex Evolution

ABSTRACT

RainSim V3 is a robust and well tested stochastic rainfall field generator used successfully in a broad range of climates and end-user applications. Rainfall fields or multi-site time series can be sampled from a spatial–temporal Neyman–Scott rectangular pulses process: storm events occur as a temporal Poisson process; each triggers raincell generation using a stationary spatial Poisson process; raincells are clustered in time lagging the storm event; each raincell contributes rainfall uniformly across its circular extent and throughout its lifetime; raincell lag, duration, radius and intensity are random variables; orographic effects are accounted for by non-uniform scaling of the rainfall field. Robust and efficient numerical optimization schemes for model calibration are identified following the evaluation of five schemes with optional log-transformation of the parameters. The log-parameter Shuffled Complex Evolution (InSCE) algorithm with a convergence criterion is chosen for single site applications and an effort limited restarted InSCE algorithm is selected for spatial applications. The new objective function is described and shown to improve model calibration. Linear and quadratic expressions are identified which can reduce the bias between the fitted and simulated probabilities of both dry hours and dry days as used in calibration. Exact fitting of mean rainfall statistics is also implemented and demonstrated. An application to the Dommel catchment on the Netherlands/Belgian border illustrates the ability of the improved model to match observed statistics and extremes.

© 2008 Elsevier Ltd. All rights reserved.

Software availability

Name of software: RainSim V3

Developer: School of Civil Engineering and Geosciences, Newcastle University, NE1 7RU, UK

Contact: Aidan Burton, School of Civil Engineering and Geosciences, Newcastle University, NE1 7RU, UK, aidan.burton@ncl.ac.uk

Year first available: 2007

Hardware: PC with windows 2000 or XP

User interface: Command line with a simple visual interface

Size: 1 Mb

Availability: May be available for research purposes on application to the authors

1. Introduction

Daily and hourly stochastic rainfall models provide useful supporting roles in the analysis of risk and vulnerability within

hydrological and hydraulic systems. These roles include the generation of synthetic precipitation records where none are available; the extrapolation of short observed records; temporal downscaling of observed records; the downscaling of climate change scenarios in both space and time. Applications of synthetic rainfall data may then be made in such diverse fields as flood modelling and urban drainage (e.g. Moretti and Montanari, 2004; Brath et al., 2006; Dawson et al., 2006; Hall et al., 2006), pesticide fate modelling (e.g. Nolan et al., in press), landslide modelling (e.g. Bathurst et al., 2005), desertification vulnerability (e.g. Bathurst and Bovolo, 2004), water resource assessment (e.g. Fowler et al., 2005) and flood risk assessment (e.g. Kilsby et al., 2000).

Traditional approaches to stochastic rainfall modelling used Markov chains to simulate the occurrence of wet and dry days in the precipitation process (e.g. Gabriel and Neumann, 1962). These readily extend to multi-site models of both amounts and occurrence, and sophisticated extensions now exist, e.g. Wilks (1998) obtains spatially smooth transitions to zero rainfall using spatially correlated random state variables. Although such models acknowledge the event-based nature of the precipitation process, they are generally inadequate in the modelling of extremes and persistence (e.g. Gregory et al., 1992). A number of alternative types of model have since developed. These include models that, at least

* Corresponding author. Tel.: +44 (0)191 222 8836; fax: +44 (0)191 222 6669.
E-mail address: aidan.burton@ncl.ac.uk (A. Burton).

in part, represent the rainfall process in terms of its scaling properties (e.g. Pegram and Clothier, 2001; Seed et al., 1999; Jothityangkoon et al., 2000); conceptual point-process event-based precipitation models which model the occurrence of precipitation events as a Poisson process in continuous time (more detail later); atmospheric state simulators linked with downscaling methodologies based on multi-site autoregressive models (e.g. Bardossy and Plate, 1992; Stehlik and Bardossy, 2002); Markov models (e.g. Fowler et al., 2005; Mehrotra and Sharma, 2005); Generalized Linear Models (GLMs) (Chandler and Wheeler, 2002; Furrer and Katz, 2007); resampling approaches (Wilby et al., 2003); other procedures whose objective is to mimic observed rainfall statistics (e.g. Bardossy, 1998). More detailed reviews are provided by Wilks and Wilby (1999) and Srikanthan and McMahon (2001).

Poisson cluster models were originally developed in a spatial context by Neyman and Scott (1958) and were first applied to precipitation modelling by Le Cam (1961). One of their natural advantages is that they can be extended to simulate continuous spatial-temporal precipitation (e.g. Gupta and Waymire, 1979) which is increasingly important in support of distributed hydrological modelling.

Rodriguez-Iturbe et al. (1987a,b) first developed the Bartlett-Lewis rectangular pulses (BLRP) model and the Neyman-Scott rectangular pulses (NSRP) model. BLRP model developments have included the use of random sampling of parameters and the use of the gamma distribution for rainfall intensity (e.g. Velghe et al., 1994; Onof et al., 2000). The model has also been used in a methodology to disaggregate multi-site daily rainfall from a GLM to hourly aggregations (Segond et al., 2006). The GLM's simulation was conditioned on atmospheric circulation properties (Chandler and Wheeler, 2002) which led to the possibility of generating rainfall simulations for future climate scenarios.

Brief histories of the NSRP model can be found in Cowpertwait (1991) and Onof et al. (2000). NSRP or BLRP storms occur as Poisson processes with characteristic timescales and so theoretically cannot exhibit either persistence or apparent scaling behaviour over more than a limited range of temporal scales (e.g. see Marani, 2003). Olsson and Burlando (2002), however, provide an empirical examination of this issue consisting of a comprehensive evaluation of the apparent scaling behaviour of an NSRP model. Power spectra shapes were found to be well reproduced (though biased) as were the apparent scaling behaviours of statistical moments of various rainfall aggregations (20 min to 2 weeks). Model deficiencies were ascribed to exponential rather than hyperbolic rainfall intensities and poor modelling of the dry period probability. Comparisons between the NSRP and BLRP models are limited in the literature. However, random parameter versions of both models were developed, to address deficiencies in dry period probability statistics, and compared by Velghe et al. (1994). They found that the NSRP model parameterization was less affected by the arbitrary choice of fitting statistics than the BLRP model and preferred sampling the number of raincells using a geometric rather than a Poisson distribution. Wheeler et al. (2005) concluded that the NSRP and BLRP model differences were probably negligible. However, the availability of an analytical expression for the probability of a dry h -hour period (e.g. Cowpertwait, 1994) removes the need for the random parameter NSRP model. Also an analytical expression for the third order moment property improves the modelling of extreme events (Cowpertwait, 1998; Cowpertwait et al., 2002). Together these developments may also address the biases noted by Olsson and Burlando (2002).

Other recent extensions to the NSRP model include raincell specialization into stratiform and convective types (Cowpertwait, 1994); regionalization and seasonal smoothing of model parameters (Cowpertwait and O'Connell, 1997); a spectral maximum likelihood fitting methodology (Chandler, 1997; Montanari and

Brath, 2000); a two site NSRP process using bi-variate distributions for raincell properties (Favre et al., 2002); a spatial-temporal NSRP model (Cowpertwait, 1995); calibration using second order moments and the fluctuation lengths of the rainfall process (Calenda and Napolitana, 1999; Favre et al., 2004).

The Generalized Neyman-Scott Rectangular Pulses (GNSRP) model (Cowpertwait, 1994, 1995) was previously implemented into a modelling package called RainSim. This simulates rainfall time series either at a single location or distributed across a region of up to ~200 km in diameter and is used in the UK Water Industry as the commercially available STORMPAC software. RainSim V2 was developed to include third moment properties, important for the modelling of extreme rainfall, with the software releases only providing point simulation with a single raincell type. Recent applications using developmental versions in a broad variety of climatic and end-user contexts have demonstrated the practical utility of the RainSim approach and led to improvements in model design. This paper consolidates these developments into a new version of the software, RainSim V3. In particular this version includes recent developments in model calibration that address a number of practical modelling deficiencies and which provide a full spatial-temporal modelling capability.

2. Recent applications of the NSRP model

This section summarizes recent model applications using the NSRP model and, in particular, developmental versions deriving from RainSim V2. These applications have provided practical experience in rainfall modelling and have indicated the need for the model calibration developments described later in the paper.

2.1. Single site applications

Single site applications of the NSRP methodology have been demonstrated by a number of authors (e.g. Cowpertwait et al., 1996a,b; Calenda and Napolitana, 1999; Fowler et al., 2000; Burlando and Rosso, 2002; Olsson and Burlando, 2002). Similarly, RainSim V2 has been used for single site applications in a wide range of locations. The first application was for a raingauge in Switzerland (Kilsby et al., 2000) as the basis of an assessment of the impact of climate change on flood risk. Another example was a synthetic study of rainfall-runoff model errors using the Slapton Wood catchment in the UK (Ewen et al., 2006) which required a long synthetic rainfall time series. More recently an application in the FOOTPRINT¹ project (Blenkinsop et al., 2006; Nolan et al., in press) attempted to identify the characteristics of climatic variability that most affect the hydrological fate of pesticides and their degradation products. Four rainfall characteristics were considered: annual total rainfall amount, seasonality, probability of a dry day and behaviour of extremes. Four sites representing diverse European climates were selected and for each combination a range of NSRP parameterizations was identified to represent a spectrum (increase and decrease) of the above characteristics. A pesticide fate model was then used to evaluate the transport of the applied pesticides.

2.2. Spatial applications

Applications of the Spatial-Temporal NSRP (STNSRP) methodology are rarer in the literature. Cowpertwait (1995) proposed the form of such a model and demonstrated it with the skewness statistic for a catchment in Northern Italy (Cowpertwait et al., 2002) and for part of the Thames basin (Cowpertwait, 2006). Moretti and

¹ <http://www.eu-footprint.org/home.html>

Montanari (2004) and Brath et al. (2006) provide two examples of spatial applications to Italian catchments for estimating flood frequency curves using rainfall–runoff modelling. In the latter case, the consequences of land use change were evaluated.

RainSim V2 has been applied in a spatial context in various applications. The model was applied to help assess the vulnerability of groundwater aquifers in the Palestinian West Bank in the DFID SUSMAQ² project. This was the first application of the model to a highly arid catchment (5600 km², completely dry for 4 months of the year), rainfall was simulated at points on a 2 km resolution spatial grid (rather than at irregularly located raingauges) and a methodology to represent interannual variability was also implemented.

More recently, three spatial–temporal model applications have been carried out for climatically different regions in the AQUA-TERRA³ project (Burton and Fowler, 2005). These were for (a) the Centa, a sub catchment of the Brenta in Northern Italy with an area of 23 km², a region characterised by an Alpine climate with higher precipitation in the summer than the winter, (b) the Dommel, a low-lying catchment on the Dutch–Belgian border with an area of 1350 km², which does not have a pronounced seasonal rainfall cycle and (c) the Gallego, a high relief tributary to the Ebro catchment which lies on the Spanish side of the Pyrenees and has an area of 4009 km². These applications were used to generate daily rainfall data for climate change impact studies.

2.3. Simulations conditional on weather types

A derivative of the RainSim V2 model involved direct conditioning of the single site NSRP process using daily time series of Lamb weather types (LWTs) (Fowler et al., 2000). The typical parameter calibration by calendar month was discarded in favour of calibrating for each of three weather ‘states’, groups of LWTs, and a two season partition of the year, ‘Summer’ and ‘Winter’, to reintroduce seasonality (six parameter sets in total). This approach was extended to spatial modelling of the Yorkshire region and a methodology to generate future climate spatial rainfall scenarios was demonstrated (Fowler et al., 2005). Further details of the piecewise stationary storm generation process used in this approach are given in Appendix A.

2.4. Modelling rainfall under a changing climate

RainSim V2 has been used frequently as a stochastic downscaling methodology to provide climate change projections for hydrological applications. Such schemes explicitly address the problem that Global Climate Model (GCM) simulations are at an inappropriate resolution in both time and space to be of direct use for hydrological impact studies.

In the FRAMEWORK⁴ project (Burton and O’Connell, 2000; Kilsby et al., 2000), extending the work of Kilsby et al. (1998), two downscaling relationships were identified between NCEP reanalysis data (Kistler et al., 2001) and 860 UK raingauges that allowed mean daily rainfall (PR) and the probability of a dry day (PDD) in a calendar month to be estimated from atmospheric circulation variables. The annual cycles of these atmospheric circulation variables were then evaluated for the GCM control (1961–1990) and future (2070–2099) scenarios. A perturbation approach was adopted, whereby observed values of PR and PDD were modified by the percentage change in the monthly mean indicated by the GCM for the future scenario, ensuring that the coefficient of variation

remained constant by also changing variance statistics. This approach was extended in the WRINCLE⁵ project by Kilsby and Burton (2001) to produce a climate change atlas of rainfall and potential evapotranspiration for Europe at a 0.5° resolution. This has been used in climate change impact assessments of desertification vulnerability in southern Italy and south–west Portugal (Bathurst and Bovolo, 2004) and landsliding in northern Italy (Bathurst et al., 2005).

The development of Regional Climate Models (RCMs) has provided physically based climatic model outputs at resolutions approaching those relevant to hydrologists. However, RCM rainfall still needs to be bias-corrected before use in hydrological impact studies (e.g. Wood et al., 2004; Fowler and Kilsby, 2007). Deficiencies, in particular, exist in the representation of high extremes (Fowler et al., 2007) and of dry periods (e.g. Blenkinsop and Fowler, 2007). Consequently, although RCM data can be used directly as input to impact studies after bias correction, there is still an important requirement for rainfall modelling based on RCM outputs to provide long representative rainfall time series for hydrological risk assessment. Currently, the perturbation approach is the most widely used, where change factors are calculated between the future and control scenarios of an RCM. These change factors are then applied to the observed rainfall statistics.

The most sophisticated use of the perturbation approach with a single site RainSim model is in EARWIG (the Environment Agency Rainfall and Weather Impacts Generator), a specialist climate scenario generator, designed for the UK (Kilsby et al., 2007). Control scenarios (1961–1990) are based on meteorological observations for a 5 km UK grid (as Perry and Hollis, 2005a,b) and future climates are based on outputs from HadRM3H, part of the UKCIP02 climate scenarios (Hulme et al., 2002), with four emission scenarios (low, medium-low, medium-high and high) and three future time slices (2020 s, 2050 s and 2080 s). The user can select a catchment of interest, a future scenario and time slice, and generate both control and future representative single site daily precipitation time series. A regression-based weather generator then conditionally simulates consistent weather time series such as temperature, wind speed and potential evapotranspiration.

3. Model description

This section provides an overview of the new software and the following section describes modelling developments implemented into this software. RainSim V3 operates in three modes: *Simulation*, *Fitting* and *Analysis*. In brief, *Simulation* generates synthetic rainfall time series based on a parameter set, *Fitting* uses numerical optimization to identify the parameter set such that the simulation best corresponds to a selected set of rainfall statistics and *Analysis* derives rainfall statistics at various time aggregations from either observed or simulated rainfall time series at a number of sites (e.g. mean daily rainfall, variance of hourly rainfall). Typically an application involves four steps: Analysis to characterise observed time series; Fitting to calibrate the model; Simulation; and finally Analysis again to check that the simulated time series is consistent with observations.

3.1. Rainfall simulation

The spatial–temporal and generalized aspects of the GNSRP model (Cowpertwait, 1995) are extensions of the stochastic point-process NSRP model (Cowpertwait, 1994) which can generate

² <http://www.ceg.ncl.ac.uk/research/water/projects/susmaq.htm>

³ <http://www.attempto-projects.de/aquaterra/>

⁴ <http://www.diiar.polimi.it/framework/>

⁵ <http://www.ncl.ac.uk/wrincle/>

synthetic rainfall time series for a raingauge. In the NSRP conceptualization, storms give rise to a cluster of raincells, the aggregated contributions of which provide a rainfall time series.

The stochastic NSRP model structure is illustrated in Fig. 1 and is constructed as follows:

- (a) storm origins occur as a uniform Poisson process with the occurrence rate represented by a parameter λ ;
- (b) each storm origin generates a Poisson random number, C with parameter ν , of raincells that each follows the storm origin after a time interval that is independent and exponentially distributed with parameter β ;
- (c) each raincell produces a uniform rainfall rate throughout its lifetime. The duration and the intensity, X , of each raincell are independent and are exponentially distributed with parameters η and ξ , respectively;
- (d) the rainfall intensity is equal to the sum of the intensities of all the active cells at that instant in time.

This process is continuous in time and so a time series is generated by discretizing the process into hourly or daily time steps. Different parameterizations for each calendar month provide an annual cycle of rainfall properties.

In the spatial-temporal version of the model (STNSRP) (Cowpertwait, 1995) the raincell generation process of the single site model, the first part of step (b), is replaced by a uniform Poisson process in space with density ρ to generate the centres of spatially circular raincells. Additionally, the radius of each raincell is exponentially distributed with parameter γ . During each cell's lifetime rainfall occurs with a uniform intensity across its spatial extent and throughout its duration. This process is spatially stationary and so a necessary final step is to account of orography by non-uniform scaling of the rainfall field. Time series sampled at each site m are scaled by a factor, ϕ_m , proportional to each sites mean rainfall. Sampling the simulated rainfall field at locations without observed records therefore requires interpolation of these factors. The parameters of the STNSRP models are summarized in Table 1.

Table 1
Parameters of the NSRP/STNSRP simulators

Parameters	Descriptions	Units
λ^{-1}	Mean waiting time between adjacent storm origins	(h)
β^{-1}	Mean waiting time for raincell origins after storm origin	(h)
η^{-1}	Mean duration of raincell	(h)
ν	Mean number of raincells per storm	(-)
ξ^{-1}	Mean intensity of a raincell	(mm/h)
γ^{-1}	Mean radius of raincells	(km)
ρ	Spatial density of raincell centres	(km ⁻²)
Φ	A vector of scale factors, ϕ_m , one for each raingauge, m	(-)

Five are used for single site, NSRP, applications and seven for spatial applications, STNSRP. All vary by calendar month.

3.2. Fitting the model

The model is calibrated separately for each calendar month in turn. A numerical optimization scheme is used to find the best choice of parameters to minimize an objective function, $D(\lambda, \beta, \dots, \xi)$, which describes the degree to which a simulation is expected to correspond to a selected set of observed rainfall statistics, with possibly varying aggregation periods, where the parameters are $\{\lambda, \beta, \nu, \eta, \xi\}$ for single site and $\{\lambda, \beta, \rho, \gamma, \Phi, \eta, \xi\}$ for spatial applications.

Analytical expressions are available for expected statistics of arbitrary period (e.g. 1 day or 2 h) accumulations of the STNSRP process at any site for the mean, variance, lag-autocovariance, lag-autocorrelation, dry period probability, probability of dry-dry (or wet-wet) transition probabilities and the third order central moment (e.g. Cowpertwait, 1995, 1998). Inter-site properties can be estimated as cross-covariances and correlations (Cowpertwait, 1995). The third order moment property (Cowpertwait, 1998) is particularly important for applications where extreme rainfall events are important, such as flood risk assessment. This is implemented in RainSim V3 as the skewness coefficient, Eq. (1), where $E(\cdot)$ indicates statistical expectation, Y_h is an h hour accumulation and $\sigma_{Y_h}^2$ its variance.

$$E\left[\frac{(Y_h - E(Y_h))^3}{\sigma_{Y_h}^3}\right] \quad (1)$$

3.3. Analysis of time series

The Analysis capabilities of the RainSim V3 software allow the user to quickly evaluate rainfall statistics from a set of rainfall time series, whether observed or simulated. The statistics are selected by the user and may be either single site statistics such as the aggregation moments or the PDD, or dual-site statistics such as the correlation or covariance between sites. Each selected statistic is evaluated separately for each month of the year for each time series. During an application, time series analysis is typically used both to characterise the observed data sets and to analyse the synthetic time series. Comparison of these two sets of statistics provides assurance that the synthetic data sets are indeed a good representation of observed rainfall data sets.

4. Improved calibration in RainSim V3

The model developments described in this section were motivated from a requirement to improve the practicality of the model calibration. A numerical optimization algorithm able to obtain robust fits for spatial applications was implemented as the existing scheme was not considered satisfactory. A new objective function was also implemented as the existing one was found to fit to low absolute magnitude values of observed statistics with disproportionate accuracy. Expressions were developed to address a bias in the analytical expressions for the probability of dry days and hours,

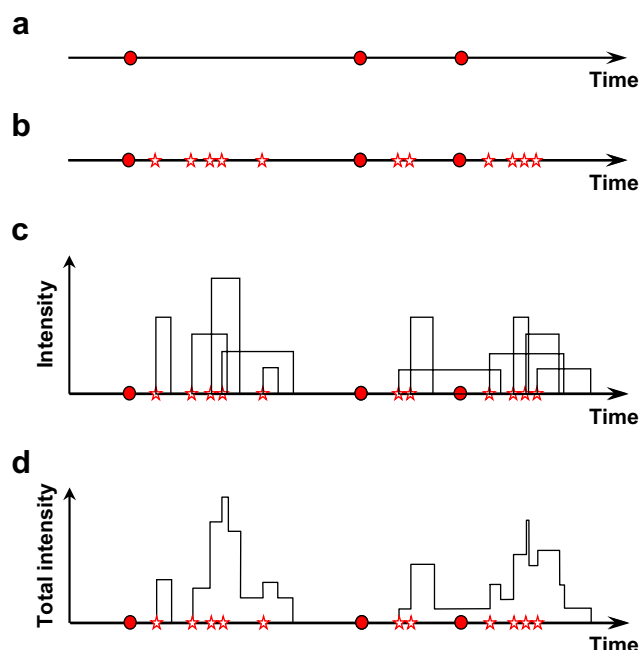


Fig. 1. Schematic of the Neyman-Scott Rectangular Pulses model.

allowing improved fitting of these statistics. Finally, a procedure to fit exactly to the mean rainfall intensity was introduced.

4.1. An improved optimizing algorithm

Single site applications of RainSim V2 used the Downhill Simplex method modified so that it was implemented iteratively (referred to as Spxl), restarting the algorithm from the best point following convergence. Each restart expands the Simplex and allows escape from local minima. After five iterations little improvement was found in the estimated optimum, though in practice 10 restarts were used. However, recent applications of the full spatial implementation of the STNSRP methodology found that observed spatial statistics were fitted poorly by the Iterative Simplex method. Therefore a study of alternative optimization schemes was instigated to find a methodology able to provide a robust and reasonably efficient fit for a range of applications.

A collection of 12 test applications of the STNSRP model were selected to provide a basis for the evaluation of the optimizing algorithms. These provided a range of climates and fitting difficulty levels. Each application also had a seasonal cycle so that a single application consisted of 12 optimization problems, one for each month of the year. The three classes of application consisted of the following.

1. Five synthetic *perfect* single site applications for which exact parameterizations existed and for which the parameters were based on those used for UK applications.
2. Five single site applications: three based on hourly statistics (Gatwick, UK, Ringway, UK, and Glize-Rijen (Dommel), Netherlands) and two on daily statistics (Sallent de Gallego, Spain, and Boxtel (Dommel), Netherlands).
3. Two spatial applications each calibrated using five sites, from the Dommel catchment in the Netherlands and the Gallego catchment in Spain.

The locations of these applications are shown in Fig. 2.




Fig. 2. The locations of test applications and case study sites.

To provide a reliable scheme five optimizing methodologies were selected from the variety available in the literature to provide a range of approaches and according to the availability of their implementation codes. In brief the selected algorithms were

- (1) The deterministic Downhill Simplex method (Sp_x) (Nelder and Mead, 1965; Press et al., 2002) without any restarts.
- (2) The Iterative Simplex method (Sp_xl) as used in RainSim V2.
- (3) The Simulated Annealing Simplex (SAS) method (Press et al., 2002), an algorithm in which the ideas of Simulated Annealing (Metropolis et al., 1953) and the Simplex method are combined to provide a directed stochastic search for the optimum.
- (4) The Shuffled Complex Evolution (SCE) algorithm (Duan et al., 1993), an optimizing scheme widely used in hydrology for calibrating hydrological models (indeed, this algorithm was involved in the fitting of an NSRP model by Montanari and Brath (2000)). A population of points is ranked and then partitioned (shuffled) into groups (complexes) in a manner similar to dealing cards. The complexes converge (evolve) independently and randomly by substituting new points in a manner related to the Simplex method and additionally by testing random points. Finally, the complexes are recombined to form the new population and the process repeated.
- (5) The Evolutionary Simplex Annealing (ESA) algorithm (Rozos et al., 2004), a scheme based on Simulated Annealing Simplex with additional transformations (e.g. multiple expansions), a modification allowing it to attempt directed climbs out of local minima and an adaptive annealing cooling schedule.

Each of these optimizing algorithms has a number of parameters that affect its performance. Therefore, prior to intercomparison, each algorithm was tuned to ensure the optimal choice of internal parameters.

Whilst the rainfall model parameters described in Table 1 take only positive values, these may vary by an order of magnitude from one application to another. For example storm arrival rate in a Mediterranean catchment may vary from once in 2 days, 0.02 h^{-1} , to one event in 3 years during that calendar month (once in 90 days), 0.0005 h^{-1} . Therefore a log-transformation of the parameter space was also considered. The utility of using such a transformation was investigated in parallel with the five selected algorithms.

The rainfall model parameters are also restricted to physically meaningful ranges, specified by the user or consistent with preset defaults. Therefore the optimization problem must be implemented as bounded and in a manner consistent with optimizer intercomparison. Implementing the bounds in a manner that affected the shape of the simplex (e.g. rejecting infeasible points and moving them to the boundary) was found to lead to the simplex being trapped incorrectly in vertices of the feasible region. A consistent scheme was therefore implemented to allow the optimizers to test infeasible points in the parameter space. A modified-objective function then returned the objective function value of the nearest feasible parameter set plus a high cost proportional to the distance of the test point from the feasible region. Each optimizer was allowed to consider the feasible parameter set as a possible optimum but proceeded as if the modified-objective function value was returned from the infeasible point in the parameter space. This procedure has the advantage of a continuous modified-objective function field with optima that are guaranteed to be feasible, and efficient retention of objective function values of parameter sets on the feasible boundary.

The first experiment determined which scheme was best able to reliably locate an optimum with reasonable effort, by evaluating the tuned optimizing schemes using the three classes of test data sets.

Each optimum was realized 11 times for optimizers with stochastic algorithms (algorithms 3–5) to sample the distribution of estimated optima, as these algorithms return different results when repeated for the same optimization problem. The number of realizations chosen here represents a compromise between computational effort and improved sampling. Each optimizer, l , was allowed to converge until a fixed number of iterations, mit , had occurred. A relative skill score, $sk_{ijkl}(mit)$, was then evaluated for each month, i , test data set, j , and realization, k :

$$sk_{ijkl}(mit) = \log_{10}\left(\max\left(10^{-6}, D_{ijkl}(mit)/D_{ij}^{ref}\right)\right) \quad (2)$$

where $D_{ijkl}(mit)$ was the best objective function value located and D_{ij}^{ref} a corresponding baseline value estimated using an un-tuned Downhill Simplex method with 100,000 iterations. Skill quartiles were evaluated for each set of 11 realizations, then averaged across all months and test data sets in each class to provide an aggregate score for both normal and log-parameter space for each class: perfect, single site and spatial.

Fig. 3 illustrates how the skill improves with increasing effort for each of the optimization schemes. The deterministic schemes (Spx and Spxl) are shown as simple curves whereas for the stochastic schemes the quartiles of skill are plotted. The upper quartile (75th percentile) curve thus provides a measure of robustness. Note that the results of the deterministic schemes may be sensitive to small changes in the application that affect the particular convergence path of the scheme. Conversely, the sensitivity to such small changes is effectively tested by the high level of noise in the stochastic schemes and included in their quantiles. The effort may also be interpreted in terms of computation time. At roughly 800 iterations per second, a typical 12-month calibration of 10,000 iterations will take about 2.5 min.

It was found that the Downhill Simplex method initially converges the most rapidly for the single site class of applications but may not converge to the optimum, e.g. see Fig. 3. Fig. 4 shows the results for the three classes for near optimum skill values. It can be seen that the SAS and InSAS methods (the *ln-* prefix indicating that the optimization is carried out in log-parameter space) are the only

ones to converge robustly in all cases. The SAS scheme converges within 20,000 iterations and the InSAS scheme within 50,000 iterations. The spatial case causes the most difficulty overall but the InSpxl method provides a relatively good performance. The InSCE method fails on only one count, that it cannot be considered robust for the spatial class. Despite this, the median result is of equal skill to the InSpxl method. However, with only 2000 iterations the InSCE method converges reliably for both single site classes and the median value is equal to or better than the InSpxl for the spatial class.

This final result suggests that a restarted InSCE scheme should perform more robustly than the other schemes. Consequently, the InSCE scheme was implemented with three restarts (InSCEx3), carrying over the previous best result in each case. The skill of this scheme was evaluated and the 75th percentile of skill is shown in Fig. 4. These plots confirm that this scheme provides the most robust and efficient method, even when limited to 5000 iterations.

An alternative to halting optimization schemes by means of an iteration limit is the use of a convergence criterion such as $r_{conv} < \tau_c$ where r_{conv} is given by Eq. (3), τ_c is a threshold and D_{min} and D_{max} are the minimum and maximum function values in the pool of points after an optimization iteration. Testing the most promising schemes generally found no improvement, however, for single site applications the convergence criterion $r_{conv} < 10^{-4}$ and an iteration limit of 20,000 made the InSCE scheme more efficient and as robust as the InSCEx3 scheme.

$$r_{conv} = 2 \frac{|D_{max} - D_{min}|}{|D_{max} + D_{min}|} \quad (3)$$

Although these results are potentially sensitive to the chosen case studies, the range of both simple and more problematic applications selected from European catchments provides confidence in the robustness of this result. These two schemes are now implemented in the RainSim V3 software package.

4.2. A new objective function

RainSim V2 used the objective function, D_2 , given by Eq. (4) (e.g. Cowpertwait, 1995; Favre et al., 2004) where \mathcal{Q} is a set of statistics,




Fig. 3. A plot showing how optimizer skill varies with computational effort for single site applications with log-transformed parameters. The range of skill for the stochastic optimizers is shown using three curves indicating the 25th, 50th and 75th percentiles of skill.




Fig. 4. Plots showing how optimization skill varies by computational effort, optimizer, application type and whether the optimization is carried out in normal- or log-parameter space. In each plot the skill axis is exaggerated to show near optimum values. The worst cases of the stochastic optimization methods are emphasised by showing only the upper and median quartiles of skill. Plots (a) and (b) are for perfect single site applications, (c) and (d) are for single site applications and plots (e) and (f) are for spatial applications. Plots (b), (d) and (f) refer to optimization of log-transformed parameters.

g, used to characterise the rainfall process. These can have varying aggregation periods and each concerns a single site (24 h variance at site m , say) or a pair of sites (for cross-correlation statistics). The observed sample estimate of a statistic is \bar{g} , the corresponding expected mean value of each statistic arising from the Neyman–Scott process is expressed analytically in terms of the model’s parameters by $\hat{g}(\lambda, \beta, \dots, \xi)$ and w_g is a weight set by the user to regulate the accuracy with which each statistic is fitted. The summed terms may be viewed as the error of a given fitted statistic standardized by the observed value squared. This function is optimum when fitting errors are in proportion to each statistic’s magnitude which leads to the fitting of observed statistics with values that have low absolute

magnitude more accurately than statistics with values that have high absolute magnitude. This is generally beneficial, typically similar proportional fitting errors will be obtained for statistics with different observed magnitudes. However, noisy observed statistics may have small magnitude values which will be fitted with relative accuracy, alternatively the accuracy with which PDD or correlation statistics are fitted should not be proportional to magnitude (arguably an observed PDD value of 0.9 should be fitted more accurately than a value of 0.5). A striking effect consequently occurs when a statistic with a spatially uniform *analytical* value approximates a range of spatially varying *observed* values, whereupon the smallest magnitude observation dominates as it is fitted

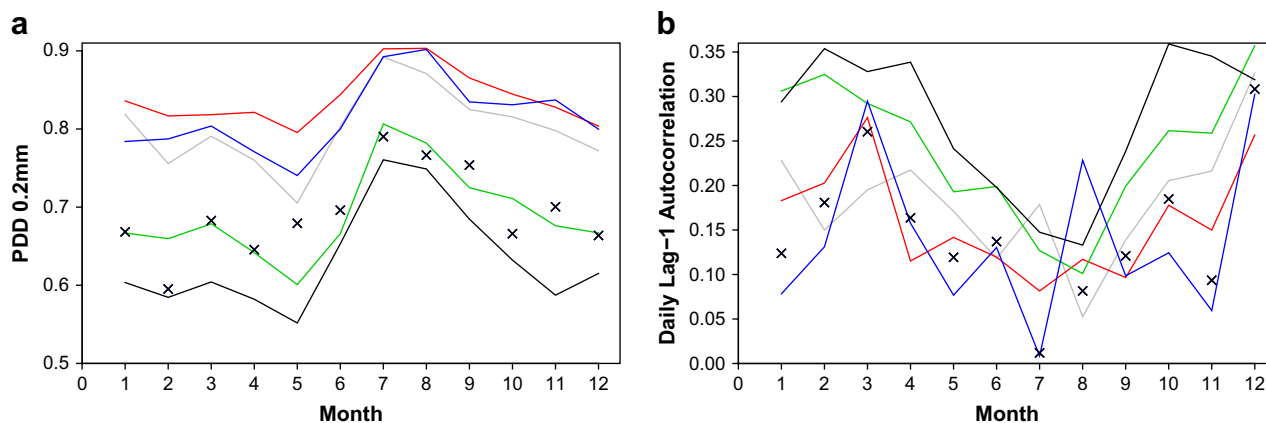


Fig. 5. Observed (lines) and fitted (crosses) monthly values of (a) 0.2 mm dry day probability and (b) daily lag-1 autocorrelation, for sites in the Gallego catchment. Each site is represented by a different colour. The spatial fit to both statistics is uniform so the fitted values overlie each other.

more accurately. Such stationary statistics arise from the conceptual construction of the STNSRP model and include, for example, the coefficient of variation, PDD and skewness coefficient. Two examples are shown in Fig. 5, from an application to the Gallego catchment, where the spatially uniform fit to both PDD and lag-1 autocorrelation is biased low in all months of the year. In particular, a low magnitude observation of the noisy lag-1 autocorrelation statistic at a site in July dominates the model fit for this month.

$$D_2(\lambda, \beta, \dots, \xi) = \sum_{g \in \Omega} w_g \left(1 - \frac{\hat{g}(\lambda, \beta, \dots, \xi)}{\bar{g}} \right)^2 \quad (4)$$

$$D(\lambda, \beta, \dots, \xi) = \sum_{g \in \Omega} \frac{w_g^2}{g_s^2} (\bar{g} - \hat{g}(\lambda, \beta, \dots, \xi))^2 \quad (5)$$

To avoid these effects in RainSim V3, a new objective function was adopted, Eq. (5), in which a scaling term, g_s , was introduced. This is set to one, for a probability dry or correlation statistic, or to the annual mean of \bar{g} , otherwise. Thus, each statistic at a site is effectively standardized by its average annual value, except for dry probability and correlation statistics which are not standardized. Consequently, each statistic is fitted with an accuracy roughly in proportion to a typical value (except for dry probability and correlation statistics which are fitted relative to one). Further, in Eq. (4), fitting errors will be distributed roughly in inverse proportion to the square root of the weights. To make the

fitting errors behave in a more user friendly manner (i.e. in inverse proportion to the weight) the weights have been squared in Eq. (5).

Fig. 6 shows results from the Gallego using the new objective function. The fitted values of PDD and lag-1 autocorrelation can now be seen to be located nearer to the mean value than in Fig. 5. In particular the value fitted for lag-1 correlation in July is significantly improved.

4.3. Bias correction for fitted probability of dry hours and days

Trace amounts of simulated or observed rainfall may occur which when recorded (typically with a precision of 0.1 mm) may be rounded down to zero and considered dry. Small rainfall amounts may also be classified as dry if they are below a specified threshold. For example in climate change analysis a dry day threshold of 1.0 mm is typical (e.g. Conway and Jones, 1998; Haylock and Nicholls, 2000; Haylock, 2004; Kilsby et al., 2007, Fig. 1). However, the analytical expression used in model fitting defines a dry period as one containing no raincells (e.g. Cowpertwait, 1994). This leads to a bias in, what we will denote as the old method, STNSRP applications whereby the simulated values of PDD statistics are greater than those fitted (e.g. see Cowpertwait, 1998). This bias can be as large as 20% for a 1 mm dry day threshold which shows the importance of explicitly recording the dry thresholds used in such statistics.




Fig. 6. Gallego application using the new objective function. Observed (lines) and fitted (crosses) monthly statistics of (a) 0.2 mm dry day probability and (b) daily lag-1 autocorrelation. Each site is represented by a different colour. The spatial fit to both statistics is uniform so the fitted values overlie each other.

We will define PDD_{τ} as the probability that a day's rainfall is strictly less than a threshold value τ (mm). By considering a time series record $\{x_t\}$ with a precision of δ , typically 0.1 mm, corresponding to continuously valued rainfall accumulations $\{X_t\}$ it can be seen that:

$$x_t < \tau \Rightarrow \text{Dry}; \quad \text{is the same as} \quad X_t < \tau - \delta/2 \Rightarrow \text{Dry} \quad (6)$$

and

$$x_t \leq \tau \Rightarrow \text{Dry}; \quad \text{is the same as} \quad X_t \leq \tau + \delta/2 \Rightarrow \text{Dry} \quad (7)$$

Therefore the omission of equality from the definition is significant and should be noted. An exception will be made for notational convenience, however, simply writing PDD_0 for the probability of no rainfall occurring.

A single site application of RainSim was fitted to each of 115 daily raingauge records from the period 1961–1995, representing the full range of climatic zones in the UK. This produced a parameter set and the corresponding analytically estimated fitted statistics, including daily PDD_0^{fit} , for each site. Rainfall time series were generated and the PDD statistics of these were estimated with thresholds of both 1.0 mm, $PDD_{1.0}^{\text{sim}}$, and 0.2 mm, $PDD_{0.2}^{\text{sim}}$. This provided 1380 data points of fitted and simulated PDD values for each threshold.

The data sample was split for validation purposes. A quadratic function, Eq. (8), was fitted achieving an R^2 of 92.2% with a validation R^2 of 91.3%. This curve was truncated according to the source data's approximate range at (0.15, 0.2821) and (0.75, 0.8045) and continuous linear expressions used to extend it to the full possible range of PDD_0^{fit} , see Fig. 7(a).

$$PDD_{1.0}^{\text{sim}} = 0.05999 + 1.603 PDD_0^{\text{fit}} - 0.8138 (PDD_0^{\text{fit}})^2 \quad (8)$$

A similar result was sought for the 0.2 mm threshold. Eq. (9) was found to have an R^2 of 98.2% and a validation R^2 of 98.0%. This curve was truncated at (0.2, 0.2405) and (0.75, 0.7617) and continuous linear expressions used to extrapolate the curve, see Fig. 7(b).

$$PDD_{0.2}^{\text{sim}} = 0.007402 + 1.224 PDD_0^{\text{fit}} - 0.2908 (PDD_0^{\text{fit}})^2 \quad (9)$$

The improvements introduced by using the quadratic models described by Eqs. (8) and (9) were evaluated separately and in detail through single site applications to four sites. For each site, fits were made to five observed daily statistics including the PDD statistic with either a 1.0 mm or a 0.2 mm threshold. The observed,

fitted and simulated statistics are shown for the $PDD_{1.0}$ threshold case for Eskdalemuir in Fig. 8. The new method significantly improves on the precision with which $PDD_{1.0}$ is estimated. The residual is due to a random error in the quadratic fit to the simulated PDD rather than from stochastic variation in the simulation; supported by the consistent sign of the residual at Eskdalemuir and the other sites for at least 9 calendar months. The new method also improves the fitting and simulation of the variance and the skewness coefficient. A similar analysis using $PDD_{0.2}$ (Eq. (9)) demonstrated a similar improvement in the simulated value, generally improved the variance and skewness coefficient fits and again found a consistent residual bias (though less pronounced than in the 1.0 mm case).

Similar results have also been developed for the probability of a dry hour (PDH) statistic, Eqs. (10) and (11). For this aggregation period 0.1 mm and 0.2 mm thresholds are in common usage. Linear fits were found sufficient to match the simulated statistics well, achieving calibration and validation R^2 values of 99.3% and 99.1%, respectively, for $PDH_{0.1}$ and 97.8% and 97.5%, respectively, for $PDH_{0.2}$. Applications on two UK sites demonstrated improvements where these results were used.

$$PDH_{0.1}^{\text{sim}} = 0.114703 + 0.884491 PDH_0^{\text{fit}} \quad (10)$$

$$PDH_{0.2}^{\text{sim}} = 0.239678 + 0.758837 PDH_0^{\text{fit}} \quad (11)$$

4.4. Exact fitting of mean rainfall statistics

Whilst recognising that observed rainfall statistics are not exact, for applications requiring long simulations of comparisons between scenarios it is desirable to obtain exact fits to the mean daily rainfall statistics. First, all statistics, including the mean, are fitted. This results in a parameterization with expected simulated means, $\hat{\mu}_h^m$, close to the observed, μ_h^m . We make these exact by a uniform perturbation of all scale factors or of the random variable modelling raincell intensity. The analytic expression for the mean rainfall of the STNSRP process in a period h can be written in terms of the model parameters, see Table 1 (e.g. Cowpertwait, 1995)

$$\hat{\mu}_h^m(\cdot) = h\lambda\phi_m \frac{E(C)E(X)}{\eta} \quad (12)$$

where the statistical expectations are simple functions of model parameters. The scale factor is the only site varying parameter and

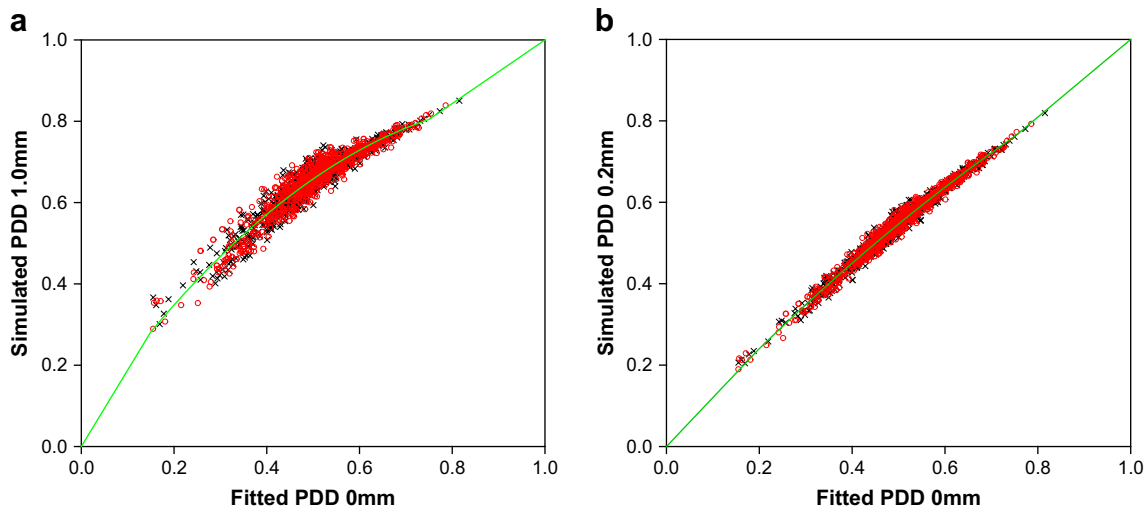


Fig. 7. Simulated (a) $PDD_{1.0}$ and (b) $PDD_{0.2}$ against fitted PDD_0 for the 115 UK raingauges. Crosses show calibration data; circles show validation data; the fitted functions are shown as curves.




Fig. 8. Observed, fitted and simulated statistics of the two fitting techniques for Eskdalemuir using a dry day probability with a 1.0 mm threshold.

during fitting is chosen to be proportional, say with constant k , to the observed mean rainfall:

$$\phi_m = k\bar{\mu}_h^m \quad (13)$$

We require the perturbed mean, $\hat{\mu}_h^m$ (with the prime indicating a perturbed term) to match the observed mean. A perturbation factor, s_m , may therefore be defined for each site:

$$s_m = \frac{\bar{\mu}_h^m}{\hat{\mu}_h^m} = \frac{\hat{\mu}_h^m(\cdot)}{\bar{\mu}_h^m(\cdot)} = \frac{h\lambda\phi'_m E(C)E(X')}{h\lambda\phi_m E(C)E(X)} \frac{\eta}{\eta} = \frac{k'E(X')}{kE(X)} \quad (14)$$

Since the final expression for s_m in Eq. (14) is independent of it's site all of the s_m are equal, so the suffix may be dropped. RainSim V3 uses an exponentially distributed raincell intensity

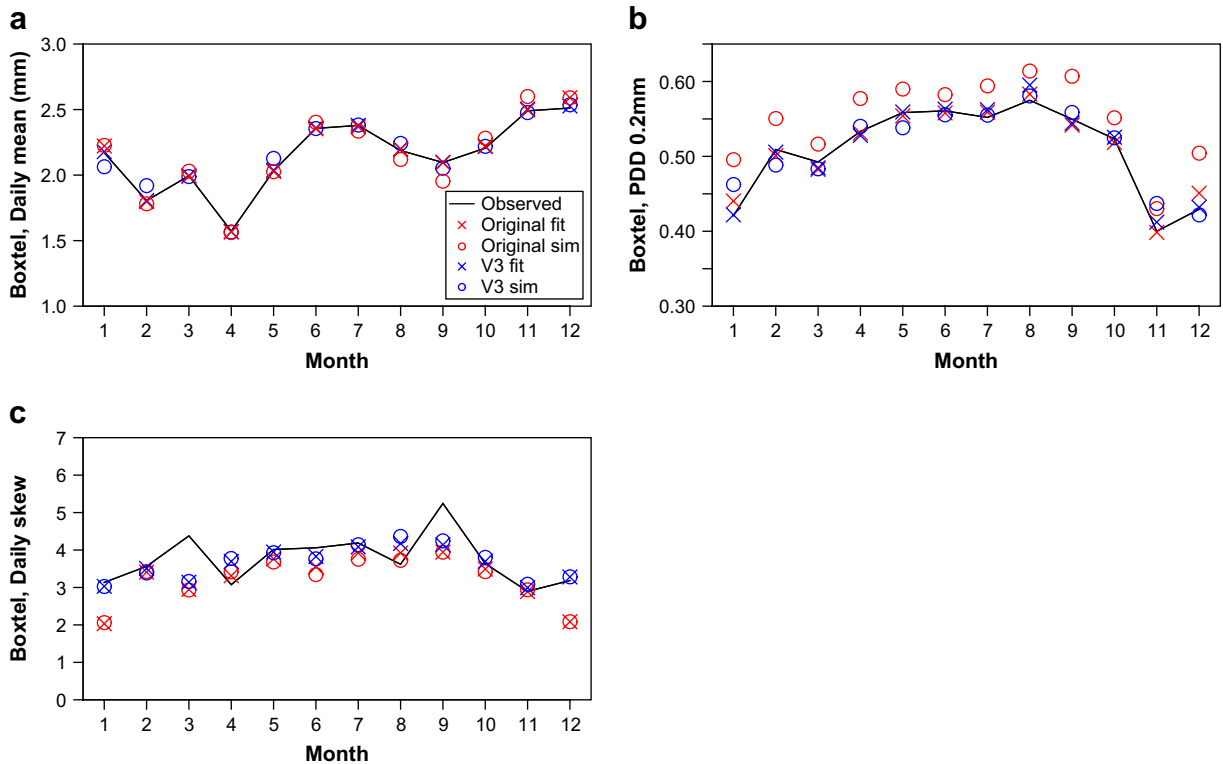


Fig. 9. Selected daily observed, fitted (fit) and simulated (sim) statistics for the five site Dommel catchment comparing the Original and the V3 fit. Only results for the Boxtel site are shown: (a) mean; (b) PDD_{0.2}; (c) skewness coefficient.

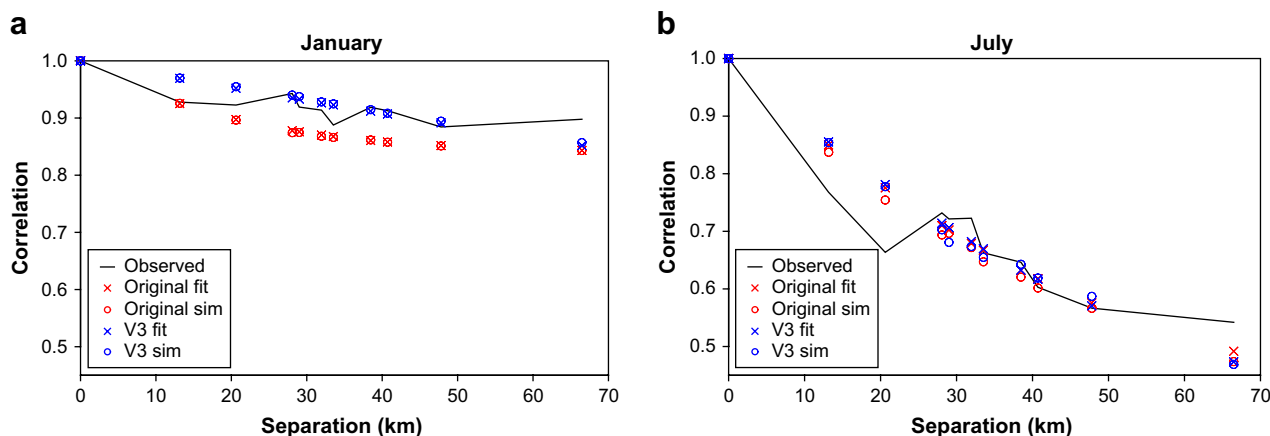


Fig. 10. Comparison of the spatial cross-correlation with distance relationships for the Original and the V3 fits, for January and July. Observed, fitted (fit) and simulated (sim) cross-correlations are shown, each value corresponding to a pair of sites from the five site Dommel catchment.

with parameter, ξ , and so $E(X) = 1/\xi$. An exact fit to the mean may therefore be obtained by either perturbing the scale factors as in Eq. (15) or the intensity parameter as in Eq. (16).

$$\phi'_m = s\phi_m \tag{15}$$

$$\xi' = \xi/s \tag{16}$$

5. Illustrative application for the Dommel

To provide an illustration of the improvements described in the previous section, a comparative study was carried out using five raingauges for the Dommel catchment in the Netherlands (see Fig. 2). First, observed daily rain gauge data from the five sites were analysed to find the daily mean, variance, lag-1 autocorrelation, skewness coefficient and PDD (with a 0.2 mm threshold).

The three fittings evaluated are denoted Original, V3 and V3ex. The Original fit was based on using the Iterative Simplex method with 10 iterations. Five such fits were made in total, manually adjusting statistics' weights and the parameter bounds at each step in order to obtain a suitable parameterization. Additionally, $PDD_{0.2}$ was approximated as PDD_0 . The V3 fit used the new optimization routine, the correction of the bias in the analytical expression for $PDD_{0.2}$ and the reformulation of the objective function (using weights adjusted for the new formulation). The V3ex fit was as for V3 but in addition using the exact mean fitting. For all three fits, a 1000 year spatial-temporal simulation was carried out, daily time series sampled at the rain gauge locations and the statistics of the multi-site time series evaluated.

Fig. 9 shows selected observed, fitted and simulated statistics for both the Original and the V3 fit. For both fits, Fig. 10 shows the cross-correlation between daily time series of all pairs of sites plotted against separation distance for 2 months from different seasons exhibiting the most extreme spatial correlation scales. For the observed and simulated cases cross-correlation is calculated from pairs of time series and for the fitted case the value corresponds to that estimated analytically from the optimum parameter set. The V3 fit was achieved with a single application of the new optimization routine and makes a reasonable overall match to all of the observed statistics used in the fitting. There is also an improvement in the quality of the fit, compared with the Original fit, particularly in the January and December skewness coefficient and spatial correlation statistics (e.g. see Figs. 9 and 10). This demonstrates the considerable

practical benefit of the new optimization algorithm, which significantly reduces the need for user intervention and generates better simulations for the case study catchments. The ability of the STNSRP model to simulate the range of spatial correlation scales over the annual cycle is also demonstrated. The benefit of the new PDD bias correction is shown by the significant improvement in the V3 fit compared with the Original fit.

Table 2 provides summary error statistics for the single site statistics for the three fitting procedures. The improvement from the Original to the V3 fit is clear in all statistics except PDD (but a much larger fitting bias has been corrected here). This result derives from a combination of the improved fitting algorithm and the revised objective function, so their contributions cannot be clearly separated. However, a reduction in fitting bias achieved by the objective function appears likely as the summary statistics for variance, autocorrelation and skewness coefficient are negative for the Original case and improved for the V3 case. This is consistent with the Original objective function fitting smaller magnitude statistics more precisely (negative bias) as these three statistics have the greatest ratio of maximum to minimum value of all of the statistics used. Exact fitting of the mean is demonstrated in the V3ex application. This is seen to be at the expense of a worse fit to the variance, however, the quality of the fits to the other statistics are not affected.

Comparisons of the observed and simulated daily annual maximum rainfall are shown in Fig. 11 by means of Gumbel plots of the best and worst results of the five simulated sites. Each plot shows 45 years of observed data and the inter-quartile range of the extremes at each corresponding return period for the 22 member ensemble of 45-year series extracted from the simulation. The range of simulated extremes is seen to match the observed data well, in all cases matching the median extreme value, the slope, the curvature and the greatest extreme.

Table 2

Fitting error averaged over all months and sites for each of the three fitting procedures for the Dommel application

	Mean	Var	PDD	Accorr	Skew
Original	0.010196	-0.692	0.005008	-0.0033	-0.327
V3	0.005903	-0.338	0.006263	0.0016	0.028
V3ex	0	-0.440	0.006263	0.0016	0.028

The statistics are the daily: mean (mm), variance (mm^2), $PDD_{0.2}$, lag-1 autocorrelation and skewness coefficient.




Fig. 11. A comparison of observed and simulated daily annual maximum for the (a) best and (b) worst of the five sites simulated for the Dommel. The inter-quartile ranges of the simulated results are shown at each return period.

6. Discussion and conclusions

This paper describes model developments leading to RainSim V3, a practical implementation of the Spatial–Temporal Neyman–Scott Rectangular Pulses (STNSRP) stochastic rainfall generator. The simulation of rainfall as a continuous spatial–temporal process, strongly supports the developing field of distributed hydrological modelling as the process may be sampled at arbitrary spatial locations and integrated to provide time series with arbitrary time steps. This software is routinely used to provide spatial rainfall fitted to daily statistics (e.g. Burton and Fowler, 2005) and is found to match annual extremes well, an ability improved by the use of the third order moment (Cowpertwait, 1998). Whilst spatial hourly simulations have not been extensively evaluated, such simulations are possible and have been successfully demonstrated for a similar model (Cowpertwait et al., 2002; Cowpertwait, 2006). However, single site hourly applications have been well tested (e.g. Kilsby et al., 2000). The downscaling of single site daily rainfall for future climate scenarios is now common, typically using a perturbation of observed rainfall statistics. This approach is used in EARWIG, a specialist climate scenario generator, which facilitates the assessment of climate change impacts on hydrological systems by providing single site simulations of rainfall and consistent weather time series for UK catchments (Kilsby et al., 2007) and which will be used in the UKCIP08⁶ scenarios. RainSim V3 is robust, computationally efficient and well tested. It has a modular design, allowing increased ease of maintainability and rapid testing of alternative model formulations.

The RainSim modelling package has been applied in a wide variety of contexts. These have included single site simulations, multi-site raingauge networks and gridded fields. Whilst the application locations have mainly been European, they span the full range of climates from the relatively wet Atlantic coast to the mountainous regions of the Pyrenees and the Alps, to arid regions such as central Spain and the Middle East. The model has been used in the context of modelling sensitivity studies, model error analysis, flood risk assessment, pesticide fate modelling, water resource planning, urban drainage, landslide modelling and desertification risk.

To provide a new and robust optimizing algorithm for RainSim V3 we compared five methods. The Shuffled Complex Evolution algorithm (SCE) (Duan et al., 1993) using log-transformed parameters with three restarts (InSCEx3) limited to 5000 iterations was

identified as the best scheme for spatial applications. However, for single site applications the InSCE scheme with a convergence criterion (limited to 20,000 iterations) performed best. A new objective function was also implemented and showed significant qualitative improvement in the fitting of previously biased statistics in a test application on the Gallego catchment and quantitative improvement for the Dommel catchment.

Quadratic polynomials in the fitted probability of a dry day (PDD) were found sufficient to provide excellent predictions of simulated PDD with either a 1 mm or 0.2 mm threshold. These considerably reduced the fitting bias at four locations in the UK and for a spatial application to the Dommel catchment in the Netherlands. The UK applications also showed consequential improved fitting of the variance and skewness coefficient. Simpler linear expressions were found to provide excellent predictions of the simulated probability of a dry hour with thresholds of either 0.2 mm or 0.1 mm. These expressions are implemented in the RainSim V3 software.

A procedure to obtain exact fitting of mean statistics by a perturbation to the raincell intensity parameter has also been implemented in RainSim V3. This approach naturally extends to the generalized form of the NSRP, with different cell types (Cowpertwait, 1994). It is important to recognise, however, that mean rainfall statistics contain observational and sampling errors and so cannot be considered exact observations of mean daily rainfall.

Poisson cluster models such as RainSim provide time and space resolutions suitable for the hydrological modelling of catchments (of up to 5000 km²) and hydraulic modelling of large urban areas as they generate rainfall with a defined structure at time scales ranging from hourly to yearly and generate a modest quantity of data. Whilst the simple geometric structure of these models (e.g. circular raincells) may appear unrealistic they are sampled at discrete locations and aggregated to typically daily or hourly time steps for which such simplifications are assumed appropriate. At the annual level such models may underestimate the variance and at larger spatial scales their spatial stationarity may limit their utility. However, the RainSim approach provides better simulations of rainfall than the ‘scaling’ approaches which typically do not model seasonal cycles or orographic effects. RainSim can also be conditioned on atmospheric circulation patterns (e.g. Fowler et al., 2000, 2005) and used to model daily rainfall over regions of up to synoptic scale or on climate model outputs to model perturbed rainfall sequences under climate change. RainSim V3 therefore provides a robust, well tested and broadly applicable practical implementation of a Spatial–Temporal NSRP.

⁶ <http://www.ukcip.org.uk/scenarios/ukcip08>

Acknowledgements

This work was supported by the European Union FP6 Integrated Project AquaTerra (Project no. 505428) under the thematic priority sustainable development, global change and ecosystems. Dr. Hayley Fowler was supported by a NERC Postdoctoral Fellowship award (2006–2009) NE/D009588/1. The authors appreciate the constructive comments of Alberto Montanari and two anonymous reviewers which have contributed to the improvement of this paper.

Appendix A. STNSRP simulation conditional on weather types

A derivative of the RainSim V2 model was able to generate spatial rainfall with the STNSRP process directly conditioned a daily time series of Lamb weather types (LWTs) (Fowler et al., 2000, 2005). To achieve this, the typical monthly parameterization was discarded in favour of one using six atmospheric states, each representing both a class of LWTs and a season. This appendix clarifies the NSRP modelling treatment of this conditioning and highlights the modelling issues arising from this approach.

A semi-Markov chain process, conditioned by season, generated a daily time series of atmospheric states (Fowler et al., 2000). Provided a parameter set has been obtained for each state, a particular parameter set is used according to the atmospheric state corresponding to that day. Following the origin of a preceding storm, at say $t = 0$, it is necessary to estimate when the next storm will occur in the STNSRP process. If the storm arrival rate is constant with parameter, λ_0 , then the time of the next storm, t_s , will simply arise from the Poisson process as a random variable with distribution function:

$$F_{t_s}(x) = 1 - e^{-\lambda_0 x} \quad (A1)$$

The parameterization is different in different periods according to the conditioning time series of daily atmospheric states (or in the standard model, calendar month), so the arrival rate may change to λ_1 at time t_1 . If the parameterization is on a calendar month basis in an aseasonal climate, a storm simulated in a subsequent month by Eq. (A1) could be simulated with the new parameters with little introduction of error. However, in a seasonal climate or for daily varying parameterizations as here, more care must be taken with the sampling of the time of storm occurrence. Otherwise dry periods may persist beyond their true extent and daily changes in parameterization will not be correctly modelled as inter-storm periods may be much greater than the daily time scale at which weather states are modelled. Storm arrival is therefore simulated as a piecewise stationary Poisson process.

If t_s is simulated as greater than time t_1 using Eq. (A1) then there is no storm in the current period and instead the storm arrival sampling is restarted at t_1 with the new value of λ_1 . The distribution of this process is given in Eq. (A2) for this case. Similarly if the second period finishes at time t_2 with a change to λ_2 , Eq. (A3) can be used.

$$F_{t_s}(x|x \geq t_1) = 1 - e^{-\lambda_1(x-t_1)} \quad (A2)$$

$$F_{t_s}(x|x \geq t_2) = 1 - e^{-\lambda_2(x-t_2)} \quad (A3)$$

For example in the case of several relatively dry periods followed by a wet one, this sampling should be repeated period by period until the correct storm origin can be located. Note that if $\lambda_1 = \lambda_0$ then $F_{t_s}(x|x \leq t_2)$, which can be derived from Eqs. (A1) and (A2), simply reduces to Eq. (A1) until time t_2 . That is, the storm arrival in two consecutive periods with the same arrival rate using this sampling procedure is the same as in a longer period with that rate, as is required. Once a storm origin is simulated, the remaining

properties of the storm are simulated using the parameterization relating to that period.

A consequence of this modelling methodology is that the final effects of a simulated storm may lag the origin by several days. For example, simulation of a dry period is likely to be biased wet by a preceding wet period. Therefore the identification of observed rainfall statistics and model parameterization for each atmospheric state is difficult and must be carefully carried out (see Fowler et al., 2000, 2005 for further details).

References

- Bardossy, A., 1998. Generating precipitation time series using simulated annealing. *Water Resources Research* 34 (7), 1737–1744.
- Bardossy, A., Plate, E.J., 1992. Space–time model for daily rainfall using atmospheric circulation patterns. *Water Resources Research* 28 (5), 1247–1259.
- Bathurst, J.C., Moretti, G., El-Hames, A., Moaven-Hashemi, A., Burton, A., 2005. Scenario modelling of basin-scale, shallow landslide sediment yield, Valsassina, Italian Southern Alps. *Natural Hazards and Earth System Sciences* 5, 189–202.
- Bathurst, J.C., Bovolo, C.I., 2004. Development of Guidelines for Sustainable Land Management in the Agri and Cobres Target Basins, Deliverable 28 of the EU funded MEDACTION project, 37 pp. Available from: <http://www.ncl.ac.uk/medaction>
- Blenkinsop, S., Fowler, H.J., Burton, A., Nolan, B.T., Surdyk, N., Dubus, I.G., 2006. Representative Climatic Records, Report DL9 of the FP6 EU-funded FOOTPRINT project, 59 pp. Available from: <http://www.eu-footprint.org/>
- Blenkinsop, S., Fowler, H.J., 2007. Changes in drought characteristics for Europe projected by the PRUDENCE regional climate models. *International Journal of Climatology* 27 (12), 1595–1610.
- Brath, A., Montanari, A., Moretti, G., 2006. Assessing the effect on flood frequency of land use change via hydrological simulation (with uncertainty). *Journal of Hydrology* 324 (1–4), 141–153.
- Burlando, P., Rosso, R., 2002. Effects of transient climate change on basin hydrology. 1. Precipitation scenarios for the Arno River, central Italy. *Hydrological Processes* 16 (6), 1151–1175.
- Burton, A., Fowler, H.J., 2005. Development of Daily Rainfall Models for Catchment Case-Studies, Production of Daily Time Series for Catchments, AquaTerra Deliverable H1.3. European Union FP6 Integrated Project AquaTerra (Project no. 505428), 32 pp.
- Burton, A., O'Connell, P.E., 2000. Final Report by Partner 4 of the FRAMEWORK Project, A Project of the EU Environment and Climate Research Programme, project number ENV4-CT97-0529, 85 pp. Available from: <http://www.diiar.polimi.it/framework/doc/Partner4.PDF>
- Calenda, G., Napolitana, F., 1999. Parameter estimation of Neyman–Scott processes for temporal point rainfall simulation. *Journal of Hydrology* 225 (1), 45–66.
- Chandler, R.E., 1997. A spectral method for estimating parameters in rainfall models. *Bernoulli* 3, 1–22.
- Chandler, R.E., Wheeler, H.S., 2002. Analysis of rainfall variability using generalised linear models: a case study from the west of Ireland. *Water Resources Research* 38 (10), 1192.
- Conway, D., Jones, P.D., 1998. The use of weather types and air flow indices for GCM downscaling. *Journal of Hydrology* 212–213, 348–361.
- Cowpertwait, P.S.P., 1991. Further developments of the Neyman–Scott clustered point process for modelling rainfall. *Water Resources Research* 27 (7), 1431–1438.
- Cowpertwait, P.S.P., 1994. A generalized point process model for rainfall. *Proceedings of the Royal Society of London* 447, 23–37.
- Cowpertwait, P.S.P., 1995. A generalized spatial–temporal model of rainfall based on a clustered point process. *Proceedings of the Royal Society of London* 450, 163–175.
- Cowpertwait, P.S.P., O'Connell, P.E., Metcalfe, A.V., Mawdsley, J.A., 1996a. Stochastic point process modelling of rainfall. I. Single site fitting and validation. *Journal of Hydrology* 175, 17–46.
- Cowpertwait, P.S.P., O'Connell, P.E., Metcalfe, A.V., Mawdsley, J.A., 1996b. Stochastic point process modelling of rainfall. II. Regionalisation and disaggregation. *Journal of Hydrology* 175, 47–65.
- Cowpertwait, P.S.P., O'Connell, P.E., 1997. A regionalised Neyman–Scott model of rainfall with convective and stratiform cells. *Hydrology and Earth System Sciences* 1, 71–80.
- Cowpertwait, P.S.P., 1998. A Poisson-cluster model of rainfall: high-order moments and extreme values. *Proceedings of the Royal Society of London* 454, 885–898.
- Cowpertwait, P.S.P., Kilsby, C.G., O'Connell, P.E., 2002. A spatial–temporal Neyman–Scott model of rainfall: empirical analysis of extremes. *Water Resources Research* 38 (8) Art. No. 1131.
- Cowpertwait, P.S.P., 2006. A spatial–temporal point process model of rainfall for the Thames catchment, UK. *Journal of Hydrology* 330 (3–4), 586–595.
- Dawson, R., Hall, J., Speight, L., Djordjevic, S., Savic, D., Leandro, J., 2006. Flood risk analysis to support integrated urban drainage. In: *Proceedings of the Fourth CIWEM Annual Conference on Emerging Environmental Issues and Future Challenges*. 12–14 September 2006, Newcastle upon Tyne. Aqua Enviro.

- Duan, Q.Y., Gupta, V.K., Sorooshian, S., 1993. Shuffled complex evolution approach for effective and efficient global minimization. *Journal of Optimization Theory and Applications* 76 (3), 501–521.
- Ewen, J., O'Donnell, G., Burton, A., O'Connell, P.E., 2006. Errors and uncertainty in physically-based rainfall-runoff modelling of catchment change effects. *Journal of Hydrology* 330 (3–4), 641–650.
- Favre, A.C., Musy, A., Morgenthaler, S., 2002. Two-site modeling of rainfall based on the Neyman–Scott process. *Water Resources Research* 38 (12) Art. No. 1307.
- Favre, A.C., Musy, A., Morgenthaler, S., 2004. Unbiased parameter estimation of the Neyman–Scott model for rainfall simulation with related confidence interval. *Journal of Hydrology* 286 (1–4), 168–178.
- Fowler, H.J., Kilsby, C.G., O'Connell, P.E., 2000. A stochastic rainfall model for the assessment of regional water resource systems under changed climatic conditions. *Hydrology and Earth System Sciences* 4, 261–280.
- Fowler, H.J., Kilsby, C.G., O'Connell, P.E., Burton, A., 2005. A weather-type conditioned multi-site stochastic rainfall model for the generation of scenarios of climatic variability and change. *Journal of Hydrology* 308 (1–4), 50–66.
- Fowler, H.J., Kilsby, C.G., 2007. Using regional climate model data to simulate historical and future river flows in northwest England. *Climatic Change* 80 (3–4), 337–367.
- Fowler, H.J., Ekström, M., Blenkinsop, S., Smith, A.P., 2007. Estimating change in extreme European precipitation using a multi-model ensemble. *Journal of Geophysical Research – Atmospheres* 112, D18104, doi:10.1029/2007JD008619.
- Furrer, E.M., Katz, R.W., 2007. Generalized linear modeling approach to stochastic weather generators. *Climate Research* 34 (2), 129–144.
- Gabriel, K.R., Neumann, J., 1962. A Markov chain model for daily rainfall occurrence at Tel Aviv. *Quarterly Journal of the Royal Meteorological Society* 88, 90–95.
- Gregory, J.M., Wigley, T.M.L., Jones, P.D., 1992. Determining and interpreting the order of a 2-state Markov-chain – application to models of daily precipitation. *Water Resources Research* 28 (5), 1443–1446.
- Gupta, V.K., Waymire, E., 1979. A stochastic kinematic study of subsynoptic space-time rainfall. *Water Resources Research* 15, 637–644.
- Hall, J.W., Dawson, R.J., Speight, L., Djordjevic, S., Savic, D., Leandro, J., 2006. Attribution of flood risk in urban areas. In: Gourbesville, P., Cunge, J., Guinot, V., Liong, S.-Y. (Eds.) *Nice, 4–8 September Hydroinformatics 2006: Proceedings of the Seventh International Conference on Hydroinformatics*, vol. I. Research Publishing Services, Chennai, India, pp. 280–287.
- Haylock, M., 2004. STARDEX Core Indices, STARDEX project, 1 pp. Available from: <http://www.cru.uea.ac.uk/projects/stardex/>
- Haylock, M., Nicholls, N., 2000. Trends in extreme rainfall indices for an updated high quality data set for Australia, 1910–1998. *International Journal of Climatology* 20, 1533–1541.
- Hulme, M., Jenkins, G.J., Lu, X., Turnpenny, J.R., Mitchell, T.D., Jones, R.G., Lowe, J., Murphy, J.M., Hassell, D., Boorman, P., McDonald, R., Hill, S., 2002. *Climate Change Scenarios for the United Kingdom: The UKCIP02 Scientific Report*. Tyndall Centre for Climate Change Research. School of Environmental Sciences, University of East Anglia, Norwich, UK, 120 pp.
- Jothityangkoon, C., Sivapalan, M., Viney, N.R., 2000. Tests of a space-time model of daily rainfall in southwestern Australia based on nonhomogeneous random cascades. *Water Resources Research* 36 (1), 267–284.
- Kilsby, C.G., Cowpertwait, P.S.P., O'Connell, P.E., Jones, P.D., 1998. Predicting rainfall statistics in England and Wales using atmospheric circulation variables. *International Journal of Climatology* 18 (5), 523–539.
- Kilsby, C.G., Burton, A., Birkinshaw, S.J., Hashemi, A.M., O'Connell, P.E., 2000. Extreme rainfall and flood frequency distribution modelling for present and future climates. In: *Proceedings of the British Hydrological Society Seventh National Hydrology Symposium*, pp. 3.51–3.56.
- Kilsby, C.G., Burton, A., 2001. Rainfall modelling (Chapter 4). In: Kilsby, C.G. (Ed.), *Final Report of the WRINCLE Project*, Project Number ENV4-CT97-0452 of the EU Environment and Climate Research Programme, pp. 47–53. Available from: <http://www.ncl.ac.uk/wrincl/>
- Kilsby, C.G., Jones, P.D., Burton, A., Ford, A.C., Fowler, H.J., Harpham, C., James, P., Smith, A., Wilby, R.L., 2007. A daily weather generator for use in climate change studies. *Environmental Modelling and Software* 22, 1705–1719.
- Kistler, R., Kalnay, E., Collins, W., Saha, S., White, G., Woollen, J., Chelliah, M., Ebisuzaki, W., Kanamitsu, M., Kousky, V., van den Dool, H., Jenne, R., Fiorino, M., 2001. The NCEP–NCAR 50-year reanalysis: monthly means CD-ROM and documentation. *Bulletin of the American Meteorological Society* 82, 247–268.
- Le Cam, L.M., 1961. A stochastic description of precipitation. In: Neyman, J. (Ed.), *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 3. University of California, Berkeley, California, pp. 165–186.
- Marani, M., 2003. On the correlation structure of continuous and discrete point rainfall. *Water Resources Research* 39 (5) Art. No. 1128.
- Mehrotra, R., Sharma, A., 2005. A nonparametric nonhomogeneous hidden Markov model for downscaling of multisite daily rainfall occurrences. *Journal of Geophysical Research* 110, D16108, doi:10.1029/2004JD005677.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, M., Teller, E., 1953. Equations of state calculations by fast computing machines. *Journal of Chemical Physics* 21, 1087–1092.
- Montanari, A., Brath, A., 2000. Maximum likelihood estimation for the seasonal Neyman–Scott rectangular pulses model for rainfall. In: Claps, P., Siccaldi, F. (Eds.), *Mediterranean Storms, Proceedings of the Plinius Conference '99*. Editoriale Bios, Cosenza, Italy, pp. 297–310. GNDCI Publication No. 2012.
- Moretti, G., Montanari, A., 2004. Estimation of the peak river flow for an ungauged mountain creek using a distributed rainfall-runoff model. In: Brath, A., Montanari, A., Toth, E. (Eds.), *Hydrological Risk: Recent Advances in Peak River Flow Modelling, Prediction and Real-Time Forecasting – Assessment of the Impacts of Land-Use and Climate Changes*. BIOS, Cosenza, Italy, pp. 113–128.
- Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. *Computer Journal* 7, 308–313.
- Neyman, J., Scott, E.L., 1958. Statistical approach to problems of cosmology. *Journal of the Royal Statistical Society, Series B* 20 (1), 1–43.
- Nolan, B.T., Dubus, I.G., Surdyk, N., Fowler, H.J., Burton, A., Hollis, J.M., Reichenberger, S., Jarvis, N.J. Identification of key climatic factors regulating the transport of pesticides in leaching and to tile drains. *Pest Management Science*, in press, doi:10.1002/ps.1587.
- Olsson, J., Burlando, P., 2002. Reproduction of temporal scaling by a rectangular pulses rainfall model. *Hydrological Processes* 16 (3), 611–630.
- Onof, C., Chandler, R.E., Kakou, A., Northrop, P., Wheeler, H.S., Isham, V., 2000. Rainfall modelling using Poisson-cluster processes: a review of developments. *Stochastic Environmental Research and Risk Assessment* 14 (6), 384–411.
- Pegram, G.G.S., Clothier, A.N., 2001. Downscaling rainfields in space and time, using the String of Beads model in time series mode. *Hydrology and Earth System Sciences* 5 (2), 175–186.
- Perry, M., Hollis, D., 2005a. The development of a new set of long-term climate averages for the UK. *International Journal of Climatology* 25 (8), 1023–1039.
- Perry, M., Hollis, D., 2005b. The generation of monthly gridded datasets for a range of climatic variables over the UK. *International Journal of Climatology* 25 (8), 1041–1054.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2002. *Numerical Recipes in C++*, second ed. Cambridge University Press, Cambridge, UK.
- Rodriguez-Iturbe, I., Cox, D.R., Isham, V., 1987a. Some models for rainfall based on stochastic point processes. *Proceedings of the Royal Society of London, Series A* 410, 269–288.
- Rodriguez-Iturbe, I., Febres de Power, B., Valdes, J.B., 1987b. Rectangular pulses point process models for rainfall: analysis of empirical data. *Journal of Geophysical Research* 92, 9645–9656.
- Rozos, E., Efstratiadis, A., Nalbantis, I., Koutsoyiannis, D., 2004. Calibration of a semi-distributed model for conjunctive simulation of surface and groundwater flows. *Hydrological Sciences Journal/Journal des Sciences Hydrologiques* 49 (5), 819–842.
- Seed, A.W., Srikanthan, R., Menabde, M., 1999. A space and time model for design storm rainfall. *Journal Geophysical Research Atmospheres* 104 (D24), 31623–31630.
- Segond, M.L., Onof, C., Wheeler, H.S., 2006. Spatial-temporal disaggregation of daily rainfall from a generalized linear model. *Journal of Hydrology* 331, 674–689.
- Srikanthan, R., McMahon, T.A., 2001. Stochastic generation of annual, monthly and daily climate data: a review. *Hydrology and Earth System Sciences* 5 (4), 653–670.
- Stehlik, J., Bardossy, A., 2002. Multivariate stochastic downscaling model for generating daily precipitation series based on atmospheric circulation. *Journal of Hydrology* 256 (1–2), 120–141.
- Velghe, T., Troch, P.A., De Troch, F.P., Van de Velde, J., 1994. Evaluation of cluster-based rectangular pulses point process models for rainfall. *Water Resources Research* 30 (10), 2847–2857.
- Wheeler, H.S., Chandler, R.E., Onof, C.J., Isham, V.S., Bellone, E., Yang, C., Lekkas, D., Lourmas, G., Segond, M.-L., 2005. Spatial-temporal rainfall modelling for flood risk estimation. *Stochastic Environmental Research and Risk Assessment* 19, 403–416.
- Wilby, R.L., Tomlinson, O.J., Dawson, C.W., 2003. Multi-site simulation of precipitation by conditional resampling. *Climate Research* 23 (3), 183–194.
- Wilks, D.S., 1998. Multisite generalization of a daily stochastic precipitation generation model. *Journal of Hydrology* 210, 178–191.
- Wilks, D.S., Wilby, R.L., 1999. The weather generation game: a review of stochastic weather models. *Progress in Physical Geography* 23 (3), 329–357.
- Wood, A.W., Leung, L.R., Sridhar, V., Lettenmaier, D.P., 2004. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic Change* 62, 189–216.