

GENERATION OF SPATIALLY CONSISTENT RAINFALL DATA

REPORT TO THE MINISTRY OF AGRICULTURE, FISHERIES AND FOOD,
BY:

DEPARTMENT OF CIVIL AND ENVIRONMENTAL ENGINEERING
IMPERIAL COLLEGE OF SCIENCE, TECHNOLOGY AND MEDICINE
LONDON SW7 2BU

AND

DEPARTMENT OF STATISTICAL SCIENCE
UNIVERSITY COLLEGE LONDON
GOWER STREET
LONDON WC1E 6BT

VOLUME I: MAIN REPORT

Principal Investigators:

Professor H.S. Wheeler Professor V.S. Isham
Dr. C. Onof Dr. R.E. Chandler

Researchers:

Dr. P.J. Northrop Dr. P. Guiblin
Mr. S.M. Bate

Consultants:

Professor Sir David Cox (Department of Statistics, University of
Oxford)

Dr. D. Koutsoyiannis (Faculty of Civil Engineering, National
Technical University of Athens)

February 2000

Contents

1	Introduction	6
1.1	The context of the research	6
1.1.1	Introduction	6
1.1.2	The importance of spatial rainfall for flood estimation	7
1.2	Background to rainfall model selection	7
1.3	Modelling strategy	8
1.3.1	Spatial-temporal models	8
1.3.2	Generalized Linear Models (GLMs)	9
1.3.3	Hybrid Modelling Approaches	9
1.4	Data	10
1.5	Structure of the report	10
2	Spatial-temporal modelling using radar data I: modelling of event interiors	11
2.1	Model development	11
2.2	Model description – event interiors	12
2.3	Distributional assumptions	14
2.4	Model properties	15
2.5	Further properties	16
2.6	Event interior parameter identification	17
2.7	Rain event selection	17
2.8	Parameter identification	17
2.8.1	Selection of fitting properties	18
2.8.2	Velocity estimation	19
2.8.3	Estimation of λ, μ_X and μ_C	20
2.8.4	Estimation of temporal parameters	20

2.8.5	Estimation of spatial parameters	20
2.8.6	Summary of parameter identification procedure	22
2.9	Model performance	22
2.10	Parameter estimate summary	26
3	Spatial-temporal modelling using radar data II: modelling of event sequences	31
3.1	Rain Event Arrivals	31
3.1.1	Introduction	31
3.1.2	Observation of Rain Events	31
3.1.3	Modelling Event Durations and Inter-Event Intervals	33
3.1.4	Distributions of Durations of Events and Dry Periods	36
3.1.5	Fitted Distributions for Event Durations and Dry Periods	39
3.1.6	Model Simplification	45
3.2	Modelling the event sequence	48
3.2.1	Leading and trailing edges	48
3.2.2	Selection of parameter sets	48
3.2.3	Simulation and data storage	50
3.3	Continuous Simulation Performance	51
3.3.1	Calculation of summary statistics	51
3.3.2	Results	54
3.3.3	Extreme-value performance	62
3.4	Summary of continuous space-time modelling	67
4	Generalized Linear Modelling of Daily Data	68
4.1	Model development	69
4.1.1	Interactions	70
4.1.2	Model fitting	70
4.1.3	Modelling nonlinear dependencies	71
4.1.4	Model checking	72
4.1.5	Modelling spatial structure	77
4.2	Model performance	82
4.2.1	West of Ireland	82
4.2.2	The Brue	106

4.2.3	The Blackwater	113
4.3	Summary of GLM modelling	119
5	Applications with limited data	121
5.1	Spatial-temporal model implementation with raingauge data	121
5.1.1	Modelling strategy	121
5.1.2	Model parameters and rainfall statistics	122
5.1.3	Case study: Catchment and data	123
5.1.4	Basic statistical properties of rainfall	124
5.1.5	Model fitting	128
5.2	Continuous simulation performance with raingauge data	131
5.2.1	Simulation model	131
5.2.2	Statistical checking criteria	133
5.2.3	Results	133
5.2.4	Extreme-value performance	137
5.2.5	Conclusions	138
5.3	Spatial-temporal rainfall disaggregation	141
5.3.1	Problem formulation	141
5.3.2	Estimation of the stochastic structure at the hourly level	142
5.3.3	Modelling approach	143
5.3.4	Models involved	144
5.3.5	The simplified multivariate rainfall model	145
5.3.6	The transformation model	146
5.3.7	Specific difficulties	148
5.3.8	Application to the Brue catchment	149
5.3.9	Conclusions	151
5.4	Temporal disaggregation	153
5.4.1	Disaggregation of fixed daily totals	156
5.4.2	Direct generation of fine-resolution data	156
5.5	Summary of work on limited data applications	157
5.5.1	Spatial-temporal modelling using raingauge data	157
5.5.2	Spatial-temporal rainfall disaggregation	157

6	Summary and Recommendations	159
6.1	Summary of methods and results	159
6.1.1	Spatial-temporal modelling using radar data	159
6.1.2	Spatial-temporal modelling using raingauge data	160
6.1.3	Generalized Linear Models (GLMs)	161
6.1.4	Hybrid approaches	162
6.2	Recommendations	162
6.3	Strategic Priorities	163
6.3.1	Medium term applications	164
6.3.2	Short term applications	164
	Bibliography	165
	APPENDICES (VOLUME II)	171
A	Documentation for spatial-temporal model	171
A.1	Summary of model specification	171
A.2	Model fitting: algorithms and software	171
A.3	Simulation: overview of software	171
A.4	Data analyses and parameter estimates	172
A.5	Model Performance	229
B	Documentation for Generalized Linear Models	293
B.1	Generalised Linear Modelling — summary of models	293
B.1.1	Case study 1: West of Ireland	293
B.1.2	Case study 2: the Brue	326
B.1.3	Case study 3: the Blackwater	330
B.2	Generalized Linear Modelling — overview of software	339
B.2.1	Obtaining the software	339
B.2.2	Generating the executables	340
B.2.3	Using the package	342
B.2.4	Summary of outputs	352
B.2.5	Hints on use	353
C	Notes for limited data applications	355

C.1	Computational tools for GDSTM applied to raingauge data	355
C.2	Computational tools for multivariate disaggregation	355
D	Single-site temporal disaggregation using rejection sampling methods	357
D.1	Introduction	357
D.2	The Bartlett-Lewis Rectangular Pulses Model	358
D.3	A Generalised Linear Model for daily rainfall totals	359
D.4	Rejection sampling	360
D.5	The distribution of daily rainfall under the BLRPM	361
D.5.1	Rainfall occurrence	362
D.5.2	Positive daily totals	363
D.6	Rejection sampling with gamma proposal and target	363
D.7	Strategies to increase efficiency	364
D.7.1	Conditioning on positive rainfall	364
D.7.2	Satisfying $f(y)/g(y) > u$	365
D.8	Sequential rejection sampling	366
D.9	Concluding remarks	367
D.10	Algebraic addendum	368
D.10.1	Cells	368
D.10.2	Existing storms	368
D.10.3	New storms	370

Chapter 1

Introduction

1.1 The context of the research

1.1.1 Introduction

Rainfall is the main input to most hydrological systems, and hence a wide range of hydrological analyses, for flood, water resource, water quality or ecological studies, require quantification of rainfall inputs. This may be possible using empirical observations, but there is often a need to extend available data in terms of record length, temporal resolution and/or spatial coverage. Hence rainfall models are required, and for almost all applications, it is rainfall over an area, i.e. the spatial distribution of rainfall, that must be characterized.

In the last decade, radar data for the UK have become routinely available as a means of recording spatial rainfall. Although radar measurement has a number of limitations with respect to performance (see, e.g. Collier (1989)) and long records of continuous data are not yet available, nevertheless it represents an important source of information which allows, for the first time, the continuous spatial distribution of rainfall to be studied. In parallel, new research into spatial rainfall modelling has produced a range of tools with potential hydrological application. However, most methods of hydrological design and simulation are relatively primitive. Areal reduction factors, assuming uniform spatial distribution, are widely applied (e.g. NERC (1975)), and modelling of raingauge networks has received relatively little research attention or application. Radar data have been applied in the context of real time flood estimation using simple short-term translation of observed images, but not for the evaluation of more complex modelling approaches. There is a need, therefore, to combine the strengths of new data sources and new modelling methods to produce a new generation of rainfall modelling tools to support hydrological practice. Given the evident complexity of spatial rainfall fields, this is an extremely challenging task.

In the context of UK flood design, recent research at the Institute of Hydrology (Calver, Lamb & Morris 1999), supported by MAFF, has been exploring the use of continuous simulation models for flood estimation. As discussed in the Flood Estimation Handbook (Institute of Hydrology 1999), this approach has potential benefits in overcoming the problems of representing the joint distribution of precipitation and antecedent conditions required when event simulation is used, but a number of new problems arise. One of these is the requirement for long time-series of rainfall

inputs, and hence for appropriate rainfall models. The research reported here was stimulated by this need and the recognition that the potential of new developments in spatial rainfall data and modelling needed to be explored. The report represents recent research at Imperial College and University College London, supported in part by The Ministry of Agriculture, Fisheries and Food, and in part by the Natural Environment Research Council.

1.1.2 The importance of spatial rainfall for flood estimation

While many studies have highlighted the importance of the temporal distribution of rainfall on flood hydrograph properties, a general understanding of the importance of spatial rainfall for flood management is not yet available; this will vary, for example, with the spatial scale of the catchment (which will determine the spatial and temporal scale of the rainfall input), the catchment properties and rainfall type. A recent review of the effects of spatial and temporal properties of rainfall upon the stream flow hydrograph is given by Singh (1997). Several studies have demonstrated that effects of storm movement can be important. For example, Ngirane-Katashaya & Wheeler (1985) showed that storm direction and speed of movement can be significant for rapidly responding catchments; downstream movement can enhance flood peaks and upstream movement can reduce them, with resonance effects generated as the storm speed approaches that of runoff routing. The effects of the spatial distribution of rainfall will depend on the nature and spatial distribution of catchment properties. Naden (1992) found that for the Thames basin, the spatial variation of rainfall could have a marked effect on channel network response, but that the slow response of chalk catchments damped out effects of rainfall variability. In complete contrast, Michaud & Sorooshian (1994), investigating flood runoff from convective thunderstorm rainfall in the arid South West of the USA, have shown that high spatial resolution of rainfall is essential to simulate flood peaks (coarse resolution data led to underestimation of flood peaks by 50-60%). It can be concluded that the spatial and temporal variability of rainfall can be extremely important in influencing flood hydrograph shape and volume, but that the importance will vary greatly as a function of catchment and rainfall properties.

1.2 Background to rainfall model selection

There are several different approaches to the modelling and simulation of precipitation. Cox & Isham (1994) classified precipitation models into three broad categories. ‘Empirical statistical models’ are, as the name implies, essentially based on empirical analysis of raingauge data and do not attempt to model explicitly rainfall structure or processes. They generally consider a single time scale, for example, daily or hourly (Foufoula-Georgiou & Lettenmaier 1987), and have most widely been used to represent single site rainfall. While it is in principle straightforward to aggregate to larger time intervals, such models cannot be directly used to disaggregate to smaller intervals. Multi-site developments include the Generalized Linear Models (GLMs) of Stern & Coe (1984) and Chandler and Wheeler (1998a,b), which can readily represent spatial non-stationarity and temporal trend in daily rainfall.

In contrast, ‘dynamic meteorological models’ encapsulate the physical processes of mass, momentum and energy transport in the atmosphere explicitly in a continuum-based representation, requiring detailed data input of all atmospheric conditions and atmosphere-surface interactions

and extensive computational resources (Mason 1986). Such methods are widely applied in weather forecasting and Global Climate Models (GCMs).

The third category is that of ‘intermediate stochastic models’, which are modelled in continuous space and time and hence can be aggregated to any required spatial or temporal scale. A simplified conceptual representation of rainfall processes is defined, using simple probabilistic assumptions which lead to highly parameter efficient models. The model parameters thus represent observable rainfall features and can be used to simulate physical precipitation processes. The approach is based on single-site models developed by Rodriguez-Iturbe *et al.* (1987a, 1988), in which storm arrivals are modelled using a Poisson process and associated with each storm arrival is a random number of cells, of random duration and intensity, dependent on the model process. Much work has been carried out in further development and validation of point models over the last decade (Onof & Wheater 1993, Onof & Wheater 1994, Kakou 1997, Samuel 1999, Velghe, Troch, de Troch & Van de Velde 1994, Khaliq & Cunnane 1996, Verhoest, Troch & de Troch 1997, Gyasi-Agyei & Willgoose 1997, Calenda & Napolitano 1999). The framework for extension to spatial-temporal modelling was developed by Cox & Isham (1988) and Cox & Isham (1994), with further recent development by Northrop (1998).

An additional category which has emerged over the last decade and a half is that of scaling and multi-scaling models (Foufoula-Georgiou 1998). These focus on the preservation of certain properties across spatial scales and are able to reproduce the scaling nature of the rainfall with a very small number of parameters (Tessier, Lovejoy & Schertzer 1993, Gupta & Waymire 1993). These are not related to the observed storms and cells of rainfall, although they can be linked to climatological features (Perica & Foufoula-Georgiou 1996).

For continuous simulation modelling in the context of flood design, computational constraints preclude the use of dynamic meteorological models. Scaling and multi-scaling models have difficulty with the accurate reproduction of the wet/dry field (Gupta & Waymire 1993) and the full space-time modelling of rainfall using random cascades is an ongoing area of research as far as theoretical development is concerned (Over & Gupta 1996). Such modelling has been mostly developed for short-term forecasting (Marshall 1996) and not long term simulation. The approach in this study is therefore based on exploration of the potential and relative merits of empirical statistical models and intermediate stochastic models, represented by the Generalised Linear Models of Chandler and Wheater and the Poisson-process based spatial-temporal models of Northrop, respectively.

1.3 Modelling strategy

1.3.1 Spatial-temporal models

The modelling of rainfall in continuous space and time gives complete flexibility of application, and the Poisson models discussed above are also attractive in being parsimonious in the number of their parameters, and the fact that the model parameters have physical significance. A major thrust of this study has therefore been to investigate the potential applicability of this family of models for continuous simulation. One limitation is that, at their present stage of development, these models represent rainfall as stochastically stationary in space. Hence systematic effects, such as topographic variation in rainfall, cannot be represented. The consequences of this are obviously dependent on the spatial scale and location of application, and remain to be defined for

rainfall-runoff application. The models are also stationary in time. Seasonality can be included, for example by allowing different parameters for each month, and in principle long-term variability (e.g. climate change) could be similarly included. However, data support for this is not available at present. A final point concerning these models is that they require adequate spatial data to allow identification of model parameters representing spatial structure. This is most obviously provided by radar measurements, although there are associated problems of data quality and available record lengths. The extent to which raingauge data can be used to identify local spatial structure has also been investigated in this study and is reported below.

1.3.2 Generalized Linear Models (GLMs)

The GLM approach represents point rainfall at a number of locations by what is essentially an extension of a multiple regression approach (spatial correlation — that is, association over and above deterministic regional effects — is introduced through a model of the noise). In this way, any important explanatory variables can be included (for example elevation, rainshadow effects, distance from the sea, as well as previous rainfall amounts to allow for temporal dependence). The model is thus extremely flexible, and can incorporate spatial non-stationarity. Once the underlying independent controls on spatial non-stationarity have been defined, it can be used to simulate rainfall at any location within the modelled field, not just those gauges used in calibration. In addition, its efficient model identification structure allows for possible long-term trend or periodicity to be rigorously investigated, and included in simulations as appropriate. However, given the complexity of the spatial-temporal rainfall process, such models have only been applied at the daily time-scale. Although there are important exceptions, for most flood estimation problems at the scale of UK catchments, sub-daily rainfall is required for rainfall-runoff modelling. Ideally, therefore, this approach should be combined with some other method if finer temporal-scale behaviour is required.

1.3.3 Hybrid Modelling Approaches

It was originally proposed for this study that a hybrid modelling approach could be used to retain the benefits of the GLM in representing spatial and temporal non-stationarity, and of the spatial-temporal model in representing fine-scale temporal and spatial behaviour. This would be done through use of the GLM to define probabilities of rainfall at locations of interest on a given day in a sequential simulation, and rejection sampling to identify those simulations of the spatial-temporal model consistent with the (non-stationary) daily simulation. This approach was investigated first for the single-site problem, but after theoretical and simulation studies, the method was rejected as infeasible for the spatial-temporal problem. Nonetheless, a brief account of the work is presented in chapter 5 of this report.

An alternative hybrid approach has been explored, using the concept of spatial-temporal disaggregation. The general problem considered is as follows: given daily rainfall data (observed or simulated) from a number of locations, and one or more sub-daily data sets (again, observed or simulated, but in the latter case constrained to be consistent with the daily values), can an appropriate spatially-distributed sub-daily set of values be generated? Methodology due to Koutsoyiannis for addressing this type of problem (Koutsoyiannis 1994, Koutsoyiannis & Manetas 1996, Koutsoyiannis 2000) has been adapted and applied here. For the requirements of a full simulation model, it is necessary to provide spatially-distributed daily rainfall (which can be obtained from

the GLM) and one finer temporal-scale data series. In this case single site temporal disaggregation is required, and two approaches are considered.

1.4 Data

The primary source of continuous spatial data used in this study has been a radar data set from the Warden Hill radar in south-west England. Data are available for the period October 1993 to March 1997 at a resolution of $2 \times 2 \text{ km}^2$ over a radius of 76km and $5 \times 5 \text{ km}^2$ over a 210 km radius. Some gaps in the record occur; obvious sources of error have been removed. However no attempt has been made to re-calibrate the data series. Due to artefacts, completely dry images are very rarely observed. The impact of this is discussed in detail in the report.

Raingauge data have been used for three areas. The Brue catchment in south-west England lies within the Warden Hill radar field, and has a dense network of 49 0.2mm tipping bucket raingauges within an area of approximately 120 km^2 , installed as part of the HYREX experiment. Further details may be found in Wheater *et al.* (1999). Data are available from September 1993 to September 1998. A second data set comes from an area of some $40 \times 50 \text{ km}^2$ in the south Galway region of western Ireland. A network of 23 daily raingauges with records spanning the period 1941–1996 is augmented by 2 monthly records from the late 19th century to 1994. The third data set used is a network of 44 daily raingauges from an area of $30 \times 40 \text{ km}^2$, in and around the catchment of the river Blackwater, a tributary of the Thames located between Guildford and Basingstoke. Further details of each of these networks are given in section 4.2.

1.5 Structure of the report

Following the strategy above, the report is structured as follows: In chapter 2, the structure of the spatial-temporal model is described in detail, together with its fitting to individual events using the Warden Hill radar data. In chapter 3, the sequence of rainfall events is analysed and a continuous simulation model developed and tested, again using radar data. Chapter 4 introduces the Generalized Linear Modelling methodology for daily data, illustrated using raingauge data from three areas. These represent a range of climate conditions and raingauge densities.

Applications with limited data are presented in chapter 5. This includes implementation of the spatial-temporal model with raingauge data, the development and testing of a new method for spatial-temporal disaggregation based on daily point rainfall data, and the supporting temporal disaggregation required to generate the sub-daily structure. Chapter 6 concludes the main report with Summary and Recommendations.

More detailed information is given in a volume of Appendices. Appendix A includes specification of the spatial-temporal model and the algorithms for model fitting, the derived parameter values (for event interiors and arrivals) and full results. Appendix B presents documentation for the GLM, including a User's Guide. Appendix C describes the spatial-temporal disaggregation procedures, and Appendix D gives details of the work done on single-site hybrid model development.

Chapter 2

Spatial-temporal modelling using radar data I: modelling of event interiors

We aim to approximate the complex spatial-temporal rainfall process using a simple stochastic mechanism. In particular we incorporate some physical knowledge about spatial-temporal rainfall structure into a parsimonious stochastic model, parameterized in terms of physically meaningful quantities. This approach can be described as *stochastic-mechanistic*. The model is a spatial analogue of the point process based models that have been used to represent the temporal process of rainfall at a single site (Rodriguez-Iturbe *et al* 1987, 1988) and a generalisation of the simpler spatial-temporal models of Cox & Isham (1988). The model is envisaged as having greater physical realism than these earlier models, albeit at the expense of a certain amount of mathematical tractability. Its formulation has been guided by inspection and analysis of rainfall radar data (Northrop 1996). A full account of the derivation of model properties can be found in Northrop (1998).

2.1 Model development

We develop a model based on a point process and localised areas of relatively intense rainfall called *rain cells*. It is assumed that rain cells are located according to some random mechanism in space and time and that individual rain cells have random durations, spatial extents and intensities. Observational studies (Petterssen 1956, Austin & Houze 1972) have shown that there is a tendency for new rain cells to form in the vicinity of existing cells, so that rain cells tend to cluster within larger scale structures that we will call *storms*. This hierarchical structure can be reflected by specifying a point process in which rain cell locations are clustered in space and time. Storms themselves tend to cluster in similar manner to form *rain events*. These rainfall elements can be characterised in terms of their typical spatial extents and durations (see table 2.1). The model has been fitted to rainfall radar data covering an area of approximately 10,000 km² with rainfall intensities averaged over pixels of area 4 km². In particular we concentrate on a *radar window* consisting of a 52 × 52 array of 2 km × 2 km pixels. The temporal separation of the radar images

rainfall element	spatial extent / km ²	duration
rain cell	10–50	up to 40 minutes
storm	100–1000	a few hours
rain event	> 1000	several hours

Table 2.1: Typical spatial extents and durations of rainfall elements

is 5 minutes and we use images produced by the lowest radar beam, i.e. the level of the atmosphere nearest the ground. Within a radar image we can reasonably expect to detect rain cells and storms while a rain event will typically cover the entire radar window at some point during its lifetime. Therefore, we fit our model to the spatial and temporal *interior* of a rain event, i.e., without including the development and dissipation of rain events in time or their periphery in space and hence require only one level of clustering in our model – rain cells clustering to form storms at this stage of the report. In principle we could introduce another layer of clustering, storms within rain events, if data were available over an area large enough to contain a number of rain events. In chapter 3 we consider the problem of modelling a continuous temporal succession of events.

The model is constructed in continuous space and time but can be fitted to both rainfall radar data and raingauge data by evaluating the theoretical properties of the discretized model over the appropriate spatial or temporal scales respectively. Once we have estimated the parameters of the model we are able to simulate realizations of rainfall fields in continuous space and time. A single parameterization can produce many different realizations for use in Monte Carlo simulation studies of catchment response.

2.2 Model description – event interiors

We assume that, within a rain event, storm centres occur in a homogeneous Poisson process of rate λ in two-dimensional space (in practice, this means some large region containing the area of interest) and time. Following each storm centre, cell *origins* arrive in a temporal Bartlett-Lewis-type cluster i.e. in a Poisson process of rate β starting with a cell located in time at the storm centre (see Cox & Isham (1980)). The process of cell origins terminates after a time which is exponentially distributed with mean $\mu_L = 1/\gamma$. This time is referred to as the storm duration. Thus, the number of cells per storm, C , has a geometric distribution with mean $\mu_C = 1 + \beta/\gamma$. The temporal process of cell arrivals is identical to one of the single-site, i.e. purely temporal, models proposed by Rodriguez-Iturbe, Cox & Isham (1987).

Each cell within the storm is displaced from the storm centre by a vector which is drawn from a bivariate Gaussian distribution with mean $\mathbf{0}$ and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{pmatrix}.$$

Displacement vectors are drawn independently for each cell. The components of Σ vary randomly from storm to storm so that distinct storms can have different sizes and shapes. For example, when $\sigma_x = \sigma_y$, storms for which ρ is close to +1 or -1 will tend to have a banded structure whereas storms for which ρ is near 0 will tend to be circular.




Figure 2.1: Schematic diagram showing the spatial structure of storms within the spatial-temporal model for event interiors. Elliptical rain cells are scattered randomly about a storm centre, according to a bivariate normal distribution. The cells and storm centre move with the same velocity. This schematic can be regarded as a snapshot of a storm at a single point in time.

Each cell is elliptical, with a semi-major axis of random length A_c , depositing rain at a constant intensity X on all points in space covered by its defining ellipse during its duration D . We assume that D has an exponential distribution with mean $1/\eta$. The elliptical cell assumption is desirable since spatial autocorrelation plots of radar data often have elliptical contours, even at small spatial lags. For simplicity we assume that each rain cell has the same eccentricity E and orientation Θ as the storm within which it is born, so that, in particular, E and Θ are fixed given Σ . The assumption that the rain cell's intensity X is constant over its area and duration is also made for simplicity. Although the results do not appear to be sensitive to this assumption, it is possible to consider other shapes for the cell intensity function or to insert irregularity in the form of a high-frequency 'jitter' (Rodriguez-Iturbe et al. 1987). A schematic diagram, showing the displacement of elliptical cells about a storm centre, is shown in figure 2.1.

The total rainfall intensity at a spatial location \mathbf{u} at time t , $Y(\mathbf{u}, t)$, is the sum of the contributions from all cells active at (\mathbf{u}, t) . The variables A_c , D and X are assumed to be mutually independent between cells and independent of Σ . All cells within a storm and the storm centre itself move with the same random velocity $\mathbf{V} = (V_x, V_y)$. Thus, all cells are born within the existing structure of the storm and the storm shape is not distorted by the motion of the storm. Cell clusters belonging to distinct storms are independent. We shall refer to this model as the Gaussian displacements spatial-temporal model (GDSTM).

A variant of this model which has also been studied is the random ellipse spatial-temporal model (RESTM), in which the spatial displacement of rain cells about the storm centre has a uniform distribution over an ellipse with a major axis of random length A . In the special case (of the RESTM and of the GDSTM) when $\beta = 0$, so that each storm consists of a single rain cell, the model is that of Cox & Isham (1988) with the generalisation that the rain cells are elliptical rather than circular.

2.3 Distributional assumptions

We have assumed that rain cell and storm durations are exponentially distributed primarily for mathematical convenience. [In fact the mean length of time μ_T that storms produce rainfall also depends on the mean cell duration μ_D and the mean number of cells per storm μ_C . Rodriguez-Iturbe et al. (1987), pg 281, give an expression for μ_T . Generally, $\mu_T > \mu_L$ unless μ_D and/or μ_C are small.] The distributional assumptions that we can make for cell and storm areas and for cell intensity are more flexible.

Earlier work (Northrop 1996) indicated that the increased flexibility afforded by assuming a gamma distribution (rather than an exponential distribution) for the cell semi-major axis A_c is worthwhile. In particular, under the exponential distribution a large number of small cells are produced giving images simulated from the model a ‘spotty’ appearance. The key point in alleviating this phenomenon is to allow the distribution of A_c to have a mode away from zero. If $A_c \sim \Gamma(\alpha_1, \alpha_2)$ the mean cell area μ_A is given by

$$\mu_A = \frac{\pi\sqrt{1-e^2} \alpha_1(1+\alpha_1)}{\alpha_2^2},$$

where e is the eccentricity of the cell. Further investigation reveals that we can parameterize the assumed gamma distribution for A_c in terms of a single parameter while losing little in terms of goodness of fit. We choose to fix α_2 (at a value determined from an initial study of a selection of rain events) and estimate α_1 .

Any definition of the area covered by a storm will be somewhat arbitrary because, under the bivariate normal displacement distribution, rain cells can be located arbitrarily far from the storm centre. A sensible measure of the spatial extent of a storm is based on the area within which we would expect a certain proportion of the rain cell centres within a storm to lie (see Northrop (1996)). Let σ^2 denote the variance of cell displacements in a direction parallel to the storm and cell major axes. We assume, primarily for mathematical convenience, that $1/\sigma^2$ has a gamma distribution with shape parameter $\psi/2$ and scale parameter $\varsigma/2$. For a given value of the cell eccentricity e the mean storm area μ_S is defined to be

$$\frac{\pi\varsigma\sqrt{1-e^2}}{\psi-2}.$$

This expression corresponds to the average area within which we would expect 40% of rain cell centres to lie and thus can be thought of as an indication of the extent of the relatively intense part of a storm. In a manner similar to that used for the cells above we parameterize the distribution of $1/\sigma^2$ in terms of a single parameter.

We assume that the cell intensity X has an exponential distribution with mean μ_X . However,

in terms of the properties we will use for parameter estimation, it is only the relationship between $E(X)$ and $E(X^2)$ that matters. In this case $E(X^2) = 2\mu_X^2$.

2.4 Model properties

We wish to fit the model to temporal sequences of radar images. Each radar image consists of a regular array of values, each value giving the average rainfall intensity over a pixel of dimensions h km by h km. Typically $h = 2$ or 5 and the temporal separation of images is 5 minutes. In order to fit the model to these data, we need to derive properties of the spatially averaged rainfall process under the model. If $Y_{i,j}^{(h)}(t)$ is the average rainfall intensity over pixel (i, j) , of dimensions h km by h km at time t , then

$$Y_{i,j}^{(h)}(t) = \frac{1}{h^2} \int_{(i-1)h}^{ih} \int_{(j-1)h}^{jh} Y(\mathbf{u}, t) \, du_x \, du_y,$$

where $\mathbf{u} = (u_x, u_y)$. The mean of the continuous rainfall process is given by the product of the cell arrival rate ($\lambda\mu_C$), and the expected area (μ_A), duration (μ_D) and intensity (μ_X) of a rain cell. Therefore, the mean of the spatially averaged process is also given by

$$E[Y^{(h)}] = \lambda\mu_C\mu_A\mu_D\mu_X.$$

The covariance $c(\mathbf{u}, t) = \text{cov}[Y(\mathbf{0}, 0), Y(\mathbf{u}, t)]$ between points displaced by t hours in time and $\mathbf{u} = (u_x, u_y)$ km in space is given by

$$c(\mathbf{u}, t) = \frac{\lambda\mu_C}{\eta} E(X^2) e^{-\eta|t|} E_{A_c, \mathbf{V}, E, \Theta} [A_c^2 \mathcal{C}(\delta/A_c)] + \frac{\lambda\beta\mu_C\mu_X^2\psi}{4\pi\eta(\gamma + \eta)\zeta} e^{-\gamma|t|} E_{\mathbf{V}, E, \Theta} [g(\delta)],$$

where

$$\begin{aligned} \Delta_1 &= (u_x - v_x t) \cos \Theta + (u_y - v_y t) \sin \Theta, \\ \Delta_2 &= [(u_y - v_y t) \cos \Theta - (u_x - v_x t) \sin \Theta] / \sqrt{1 - E^2} \end{aligned}$$

and $\delta^2 = \Delta_1^2 + \Delta_2^2$. $\mathcal{C}(x)$ is the area of intersection of two discs of common unit radius whose centres are a distance x apart and

$$g(\delta) = E_{A_{c1}, A_{c2}} \left\{ \int_{|\mathbf{w}_1| \leq A_{c1}} \int_{|\mathbf{w}_2 - \mathbf{\Delta}| \leq A_{c2}} \frac{1}{\left(1 + \frac{(\mathbf{w}_1 - \mathbf{w}_2)^T (\mathbf{w}_1 - \mathbf{w}_2)}{2\zeta}\right)^{\frac{\psi}{2} + 1}} d\mathbf{w}_1 d\mathbf{w}_2 \right\}. \quad (2.1)$$

where A_{c1} and A_{c2} are the semi-major axes of two distinct cells within the same storm and $\mathbf{\Delta} = (\Delta_1, \Delta_2)$ (Northrop 1998). It can be shown that

$$\mathcal{C}(x) = \begin{cases} 2 \cos^{-1}\left(\frac{x}{2}\right) - x \sqrt{1 - \frac{1}{4}x^2} & x \leq 2 \\ 0 & x \geq 2. \end{cases}$$

In the interests of computational efficiency we make use of approximations in evaluating $\mathcal{C}(x)$ and $g(\delta)$ (see Northrop (1996)). We use a simple piecewise linear approximation to $\mathcal{C}(x)$,

$$\mathcal{C}_l(x) = (\pi - lx)^+ \quad (l > 0),$$

where $z^+ = \max(z, 0)$. We set $l = 1.84$, the value that minimises the integrated squared error

$$\text{ISE} = \int_0^2 [\mathcal{C}(x) - \mathcal{C}_l(x)]^2 dx,$$

in approximating $\mathcal{C}(x)$ by $\mathcal{C}_l(x)$.

If we assume that $A_c \sim \Gamma(\alpha_1, \alpha_2)$, then

$$\begin{aligned} \mathbb{E}_{A_c} \left\{ A_c^2 \mathcal{C}_l(\delta/A_c) \right\} &= \frac{\alpha_1}{\alpha_2^2} \left\{ \pi(\alpha_1 + 1) \left[1 - \Gamma\left(\alpha_1 + 2, \frac{l\alpha_2}{\pi}\delta\right) \right] \right. \\ &\quad \left. - l\delta\alpha_2 \left[1 - \Gamma\left(\alpha_1 + 1, \frac{l\alpha_2}{\pi}\delta\right) \right] \right\}. \end{aligned}$$

We approximate $g(\delta)$ using a Taylor series expansion of the integrand in (2.1) under the assumption that rain cells are small relative to the storm within which they are born.

Evaluation of the second order properties of the spatially averaged rainfall process requires the numerical evaluation of a double integral. Specifically,

$$\begin{aligned} \text{var}[Y^{(h)}] &= \frac{1}{h^4} \int_{-h}^h \int_{-h}^h (h - |u_1|)(h - |u_2|) c(\mathbf{u}, 0) du_1 du_2 \\ c^{(h)}(\mathbf{k}, t) &= \text{cov}[Y_{i,j}^{(h)}(0), Y_{i+k_1, j+k_2}^{(h)}(t)], \\ &= \frac{1}{h^4} \int_{-h}^h \int_{-h}^h (h - |u_1|)(h - |u_2|) c(\mathbf{u} + h\mathbf{k}, t) du_1 du_2, \end{aligned}$$

where $\mathbf{k} = (k_x, k_y)$. The autocorrelation function $\text{corr}[Y_{i,j}^{(h)}(0), Y_{i+k_1, j+k_2}^{(h)}(t)]$, denoted by $\rho^{(h)}(\mathbf{k}, t)$, of the model process aggregated over h km by h km pixels is defined by

$$\rho^{(h)}(\mathbf{k}, t) = \frac{c^{(h)}(\mathbf{k}, t)}{\text{var}[Y^{(h)}]}.$$

2.5 Further properties

Other properties of the model are of potential interest for use in parameter estimation and assessing goodness of fit, particularly those that relate to the wet/dry pattern of rainfall, but most seem inaccessible to analytic study. For example, the aggregation of wet/dry properties over space creates analytical difficulties. It is possible to derive an expression for the probability that an arbitrarily chosen point in space-time is dry, although we can only extend this to consider an arbitrarily chosen pixel when $\beta = 0$. Properties of rainfall fields to which thresholds have been applied (i.e. rainfall intensities below a given level are set to zero) may also be of interest. However, simulation from a fitted model can readily be used to assess the model's performance with respect to its ability to reproduce such properties.

2.6 Event interior parameter identification

The model is stationary in time and homogeneous in space so, in order to fit models to identify model parameters from sequences of radar data, we need to select periods of time within which the radar data are consistent with these assumptions. We must, therefore, restrict attention to the interior of rain events in both space and time, i.e. without including the development and dissipation of rain events in time or their periphery in space. In addition we make the simplifying assumption in the model that $V(=v)$, $E(=e)$ and $\Theta(=\theta)$ are fixed and common to all storms and cells within the rain event, an assumption which is reasonable over the relatively small temporal and spatial scales involved. We fit the GDSTM to sequences (52×52 arrays at 5 minute intervals) of $2 \text{ km} \times 2 \text{ km}$ resolution radar data recorded by the Wardon Hill radar station in south-west England.

2.7 Rain event selection

In order to fit the model to the interior of a given rain event, ideally we require the event to extend over the radar window. To obtain reliable estimates of the temporal features of the model we need this to happen for a time period of at least an hour. We use time series plots of the mean and variance of the rainfall over the radar area and the proportion of wet pixels to identify time periods over which these criteria may be met. Radar images from the chosen time period are then used to assess the assumptions of spatial homogeneity, constant velocity and constant cell shape. Obviously there will be rain events that do not meet all of the above criteria. For example an event may pass along one side of the radar window rather than moving directly over it or be of insufficient spatial extent to cover the radar window. Thus, we may consider two populations of rain events - those which are amenable to model fitting and those which are not. In section 3.1 we model the arrival process of rain events of these two types, and find that events of the second type contribute only a very small proportion of the total rainfall recorded by the radar.

2.8 Parameter identification

The complex dependencies produced by the structure of the model mean that a maximum likelihood approach is not feasible (it is difficult to obtain a likelihood in useful form). Indeed, it is not clear that such an approach would be appropriate. The rain cell intensity structure (constant cell intensity across space and time) means that under the model the rainfall field is a step function and the likelihood function places undue emphasis on such localised deterministic features.

A more subjective approach is to select properties of the data which are regarded as being important and find values for the parameters that produce as close a fit to these properties as possible. This is in the spirit of the Generalised Method of Moments of Hansen (1982) in which an *objective function* of a weighted sum of squared differences between observed values of selected properties and their model values is minimised numerically with respect to the parameters of the model. The properties must be chosen carefully to enable the parameters of the model to be estimated reliably. Properties not used for fitting can then be used to assess the adequacy of fit of the model.

2.8.1 Selection of fitting properties

Clearly the estimates obtained for the parameters of the model are dependent on the properties selected for calibration purposes. It is possible that quite different parameter sets may give very similar fits to a given set of properties. In some cases multiple minima may exist in the objective function. Problems of parameter identification can be alleviated by estimating the velocity \mathbf{v} directly from the data using a cross-correlation method before estimating the other model parameters. This also has the effect of reducing the computing time required to estimate the remaining parameters. Prior estimation of \mathbf{v} helps distinguish between, for example, a few large storms moving quickly and many small storms moving slowly. In addition we consider which properties of the data provide information on certain aspects of the model structure and design our fitting procedure accordingly. For example, the decay of spatial autocorrelation gives information on the typical sizes of rain cells and storms and their shape.

Before proceeding we summarize the notation used for the parameters of the GDSTM.

- λ – rate of storm arrivals (number of storms per km² per hour);
- μ_C – mean number of cells per storm;
- μ_X – mean cell intensity (mm/h);
- μ_D – mean cell duration (h);
- μ_L – mean storm duration (h);
- μ_A – mean cell area (km²);
- μ_S – mean storm area (km²);
- e – common cell and storm eccentricity;
- θ – common cell and storm orientation (degrees from east);
- v_x – west-east component of common cell and storm velocity (km/h);
- v_y – south-north component of common cell and storm velocity (km/h).

We use the following properties of the spatially averaged rainfall process for the purposes of parameter identification. We have reparameterized these expressions in terms of the 11 parameters given above.

The mean

$$E[Y^{(h)}] = \lambda\mu_C\mu_X\mu_A\mu_D,$$

the variance

$$\text{var}[Y^{(h)}] = a_1f_1(\mathbf{0}, \mu_A) + a_2f_2(\mathbf{0}, \mu_A, \mu_S).$$

the spatial autocorrelation function

$$\rho^{(h)}(\mathbf{k}, 0) = [a_1f_1(\mathbf{k}, \mu_A) + a_2f_2(\mathbf{k}, \mu_A, \mu_S)]/\text{var}[Y^{(h)}]$$

and the maximum spatial autocorrelation

$$\rho^{(h)}(\mathbf{v}t, t) = [a_1 e^{-|t|/\mu_D} f_1(\mathbf{0}, \mu_A) + a_2 e^{-|t|/\mu_L} f_2(\mathbf{0}, \mu_A, \mu_S)] / \text{var}[Y^{(h)}], \quad (2.2)$$

for positive time lags, where $f_1(\mathbf{w}, \mu_A) = E_{A_c} \{A_c^2 C_k(\delta/A_c)\}$, $f_2(\mathbf{w}, \mu_A, \mu_S) = g(\delta)$, $\mathbf{w} = (w_1, w_2)$,

$$\begin{aligned} a_1 &= \lambda E(X^2) \mu_C \mu_D \\ a_2 &= \frac{\lambda \mu_X^2 \mu_C (\mu_C - 1) \mu_D^2}{\mu_D + \mu_L} \end{aligned}$$

and

$$\delta^2 = (w_1 \cos \theta + w_2 \sin \theta)^2 + \frac{(w_2 \cos \theta - w_1 \sin \theta)^2}{1 - e^2}.$$

Intuitively, we expect

- the form of the spatial autocorrelation function $\rho^{(h)}(\mathbf{k}, 0)$ to depend strongly on the ‘spatial’ parameters μ_A, μ_S, e and θ ;
- that given \mathbf{v} , $\rho^{(h)}(\mathbf{v}t, t)$ should depend strongly on the ‘temporal’ parameters μ_D and μ_L ;
- the remaining parameters λ, μ_X and μ_C to influence the mean and variance of the rainfall field.

and this is borne out on inspection of the expressions given above. Note that

- e and θ determine the shape of contours of equal $\rho^{(h)}(\mathbf{k}, 0)$ and μ_A and μ_S determine the manner in which $\rho^{(h)}(\mathbf{k}, 0)$ decays with δ ;
- given \mathbf{v} the manner in which $\rho^{(h)}(\mathbf{v}t, t)$ decays with t depends only on μ_D and μ_L ;
- the parameters λ, μ_X and μ_C appear only in a_1 and a_2 .

It is possible that different spatial patterns of rainfall produce similar spatial autocovariance functions. The manner in which $\text{var}(Y^{(h)})$ decreases as the scale of spatial averaging h increases may distinguish between these patterns.

In the following sections we suppress the arguments of $\rho^{(h)}(\mathbf{k}, t)$ when it is clear to which autocorrelation we refer.

2.8.2 Velocity estimation

We note that $\rho^{(2)}(\mathbf{k}, t)$ is maximized for $\mathbf{k} = \mathbf{v}t$. If $\hat{\mathbf{k}}$ is an estimate of the position of the maximum of ρ , $\hat{\mathbf{v}} = \hat{\mathbf{k}}/t$ provides an estimate of \mathbf{v} . Let $\hat{\rho}$ denote the observed values of ρ estimated using the data within the interior of a rain event. We are only able to estimate ρ at a discrete set of points, namely $k_x = \dots, -2, 0, 2, \dots$ km and $k_y = \dots, -2, 0, 2, \dots$ km. Thus, the position $\tilde{\mathbf{k}}$ of the observed maximum of $\hat{\rho}$ may prove to give an imprecise estimator of \mathbf{v} .

However, since $\rho^{(2)}(\mathbf{k} + \mathbf{V}t, t)$ is an even function of \mathbf{k} , the centroid of ρ is at $\mathbf{v}t$. We use an estimator of \mathbf{v} based on the centroid of $\hat{\rho}$, noting that as t increases, the observed spatial-temporal correlation function, $\hat{\rho}^{(2)}(\mathbf{k} + \mathbf{V}t, t)$ tends to become less even particularly for \mathbf{k} far from $\mathbf{v}t$. However, $\hat{\rho}$ is approximately even for $t = 5$ minutes. Thus, we estimate \mathbf{v} using

$$\begin{aligned}\hat{v}_x &= \frac{12 \sum \sum k_x \hat{c}^{(h)}(\mathbf{k}, 1/12)}{\sum \sum \hat{c}^{(h)}(\mathbf{k}, t)}, \\ \hat{v}_y &= \frac{12 \sum \sum k_y \hat{c}^{(h)}(\mathbf{k}, 1/12)}{\sum \sum \hat{c}^{(h)}(\mathbf{k}, t)},\end{aligned}\tag{2.3}$$

where the summation is over values of (k_x, k_y) in the vicinity of $\tilde{\mathbf{k}}$.

We could devise more complex estimators of \mathbf{v} (for example using estimates of the centroid for $t = 5, 10, 15, \dots$ minutes) but we have found that the estimator in (2.3) works well in practice.

2.8.3 Estimation of λ, μ_X and μ_C

Given the other parameters we find the values of λ, μ_X and μ_C for which the observed and fitted values of $E[Y^{(2)}]$, $\text{var}[Y^{(h_1)}]$ and $\text{var}[Y^{(h_2)}]$ are equal. Closed form expressions for $\hat{\lambda}, \hat{\mu}_X$ and $\hat{\mu}_C$ result. We take $h_1=2$ and $h_2 = 8$.

2.8.4 Estimation of temporal parameters

Ideally we would like to use $\rho(\hat{\mathbf{v}}t, t)$ to estimate μ_D and μ_L . The best we can do is to use $\rho(\mathbf{k}_t, t)$ with \mathbf{k}_t chosen to be as close to $\hat{\mathbf{v}}t$ as possible. In this case the dependence of ρ on the spatial parameters is small. Given the other parameters we find the values of μ_D and μ_L for which the observed and fitted values of $\rho(\mathbf{k}_t, t)$ and $\rho(\mathbf{k}_{2t}, 2t)$ are equal. Closed form expressions for $\hat{\mu}_D$ and $\hat{\mu}_L$ result in this case. We find that $t = 10$ minutes works well in practice.

2.8.5 Estimation of spatial parameters

We note that for given values of e, θ, μ_A and μ_S we can estimate the remaining parameters simultaneously using the arguments of the previous two sections.

Under the model the spatial autocorrelation function $\hat{\rho}^{(2)}(\mathbf{k}, 0)$ is very nearly elliptically isotropic (the slight distortion caused by spatial averaging is negligible). The *elliptical space lag* z is given by

$$z^2 = (k_x \cos \theta + k_y \sin \theta)^2 + \frac{(k_y \cos \theta - k_x \sin \theta)^2}{1 - e^2}.$$

This space lag accounts for the fact that, unless $e = 0$, the decay of the spatial autocorrelation depends not only on distance but also direction.

For good estimates of e and θ we expect $\hat{\rho}$ to decay with z with little variation in $\hat{\rho}$ for a given value of z . For poor estimates we will see more variation in $\hat{\rho}$ for a given z . Figure 2.2 illustrates these points for the rain event of 6 February 1994, 1300-1400. Firstly, we use the estimates \hat{e} and $\hat{\theta}$ of e and θ obtained by fitting the GDSTM to these data (see example in section 2.9 and in

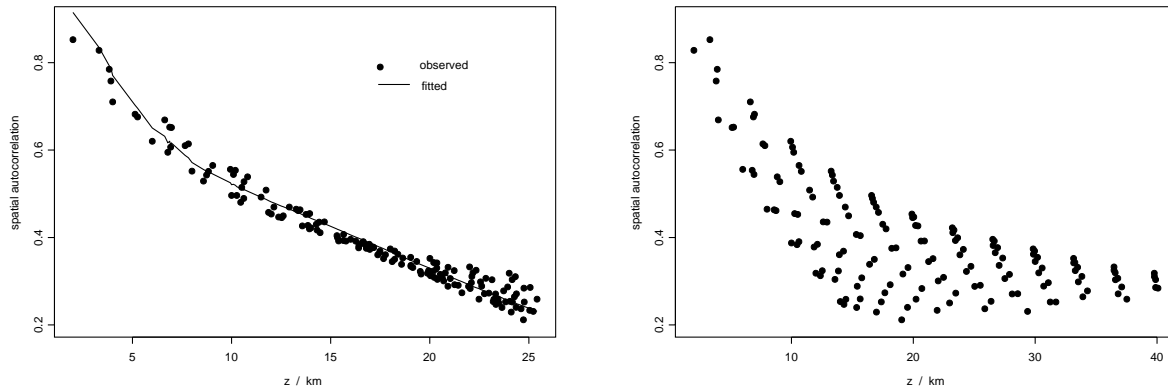


Figure 2.2: Spatial autocorrelation vs. elliptical space lag z . Left: z defined using \hat{e} and $\hat{\theta}$. Right: cell orientation = $\hat{\theta} - 90^\circ$.

particular the spatial autocorrelation plots in figure 2.3) and then we set the cell orientation at right angles to $\hat{\theta}$. Note that in the right-hand plot (where we have used a poor value of θ) there is much greater variation in $\hat{\rho}$ for a given θ than in the left-hand plot.

For given values of e , θ , a_1 and a_2 we could estimate μ_A and μ_S using a (non-linear) least squares regression of $\hat{\rho}$ on z . However, it is well known (Fisher 1915) that for finite random samples drawn from a bivariate normal population,

- $\hat{\rho}$ is biased for ρ with the bias of $\hat{\rho}$ and $\text{var}(\hat{\rho})$ depending on ρ and the size of the sample from which $\hat{\rho}$ is calculated;
- the sampling distribution of $\hat{\rho}$ may be far from normal, particularly when the underlying variables are non-normal (as is clearly the case for rainfall data).

Therefore, although there will be some limited dependence in our estimates, we make use of a transformation of ρ , namely Fisher's logarithmic transformation

$$F = \frac{1}{2} \ln \left(\frac{1 + \rho}{1 - \rho} \right). \quad (2.4)$$

Guyen (1951) demonstrates that the assumption of normality is a remarkably good approximation for non-normal samples of a reasonable size and gives approximate expressions for $E(F)$ and $\text{var}(F)$ (under the assumption that ρ and other higher order joint moments are known). We use these expressions (plugging in estimates of the quantities in $E(F)$ and $\text{var}(F)$ where necessary) to estimate e , θ , μ_A and μ_S in a weighted least squares regression of F on z . Specifically, we minimize

$$\sum_{i=1}^n \frac{[\hat{F}_i - E(F_i)]^2}{\text{var}(F_i)}, \quad (2.5)$$

with respect to e , θ , μ_A and μ_S , where we include all $\hat{\rho}$ corresponding to values of z above a chosen threshold.

model property	fitted value
cell arrival rate (λ ; $\text{km}^{-2} \text{h}^{-1}$)	0.00034
expected cell duration (μ_D ; mins)	39
expected cell area (μ_A ; km^2)	23
expected cell intensity (μ_X ; mm/h)	0.54
expected number of cells per storm (μ_C)	650
expected storm duration (μ_L ; h)	2.7
expected storm area (μ_S ; km^2)	310
x component of cell/storm velocity (v_x ; km/h)	36
y component of cell/storm velocity (v_y ; km/h)	34
cell/storm eccentricity (e)	0.83
cell/storm orientation (θ ; degrees from east)	94

Table 2.2: Parameter estimates: GDSTM fitted to the rain event of 6 February 1994, 1300–1400

2.8.6 Summary of parameter identification procedure

The general strategy is to estimate the spatial parameters e, θ, μ_A and μ_S using observed spatial autocorrelation structure, updating the estimates of the other parameters in the background.

We require initial estimates of e, θ, μ_A and μ_S to start the procedure. We calculate initial estimates $\hat{e}, \hat{\theta}$ of e and θ directly from $\hat{\rho}(\mathbf{k}, 0)$ in the following manner. The sample of pairs (k_x, k_y) for which $\hat{\rho}(\mathbf{k}, 0)$ is greater than a chosen threshold will span an elliptical region in the (k_x, k_y) plane. We estimate the variance-covariance matrix for this sample and calculate the eigenvalues ϕ^-, ϕ^+ and associated eigenvectors of this matrix. The orientation of the eigenvector associated with the largest eigenvalue ϕ^+ and $\sqrt{1 - \phi^-/\phi^+}$ are used as estimates of θ and e respectively. This procedure provides estimates that are close to the final estimates from the minimization of (2.5).

We perform a grid search over μ_A and μ_S to identify promising initial estimates of μ_A and μ_S given \hat{e} and $\hat{\theta}$. The minimisation of equation (2.5) is then carried out by an iterative numerical procedure using a modified Newton scheme. We obtain the same estimated parameter set regardless of the initial estimates of μ_A and μ_S in the vast majority of cases.

2.9 Model performance

For illustration we present results from the fitting of the GDSTM to a one hour sequence (13 images) of radar images recorded by the Wardon Hill radar station on the 6th February 1994.

Table 2.2 gives the estimated values of the parameters of the model. The cell and storm dimensions are in broad agreement with those given in table 2.1. Table 2.3 shows the observed and fitted values of a variety of properties over a range of levels of spatial aggregation. Note that the fitted and observed values of properties used in the estimation procedure will be close (or equal) by default. The fitted values for the rainfall *coverage* - the proportion p_h of h km by h km pixels in an image that are wet - are estimated using a sequence of images simulated from the fitted model. The close agreement between many properties not used for parameter estimation indicates that the model is performing well. In particular the model reproduces properties over a range of levels of

property		Level of spatial aggregation (km)			
		$h = 2$	$h = 4$	$h = 8$	$h = 16$
mean		1.78 [†]			
(mm/h)		(1.78)			
variance		3.65 [†]	3.18	2.59 [†]	1.92
(mm ² /h ²)		(3.65)	(3.22)	(2.59)	(1.93)
$\hat{\rho}^{(h)}(0, 0, t)$	t=1/12	0.66	0.74	0.84	0.93
		(0.62)	(0.71)	(0.84)	(0.91)
	t=1/6	0.45	0.51	0.61	0.77
		(0.45)	(0.51)	(0.63)	(0.76)
	t=1/4	0.32	0.36	0.43	0.59
		(0.33)	(0.37)	(0.45)	(0.58)
$\hat{\rho}^{(h)}(x, 0, 0)$	x=2	0.83 [†]	0.76	0.65	0.36
		(0.82)	(0.72)	(0.63)	(0.41)
	x=8	0.45	0.22	-0.23	-0.24
		(0.45)	(0.22)	(0.01)	(0)
$\hat{\rho}^{(h)}(0, y, 0)$	y=2	0.86 [†]	0.81	0.76	0.66
		(0.92)	(0.84)	(0.78)	(0.67)
	y=8	0.56	0.45	0.33	-0.03
		(0.57)	(0.47)	(0.20)	(0.01)
$\hat{\rho}^{(h)}(x, 0, 1)$	x=-2	0.56	0.53	0.47	0.22
		(0.52)	(0.51)	(0.46)	(0.27)
	x=2	0.72	0.80	0.78	0.51
		(0.77)	(0.84)	(0.77)	(0.53)
$\hat{\rho}^{(h)}(0, y, 1)$	y=-2	0.59	0.60	0.60	0.56
		(0.57)	(0.60)	(0.63)	(0.54)
	y=2	0.69	0.79	0.82	0.70
		(0.66)	(0.76)	(0.80)	(0.71)
\hat{p}_h		0.92	0.97	1	1
		(0.93)	(0.97)	(0.99)	(1)

Table 2.3: Rainfall recorded by the Wardon Hill radar. 6 February 1994, 1300–1400. Observed and (in parentheses) fitted values under the GDSTM. [†] denotes a property used for parameter estimation.

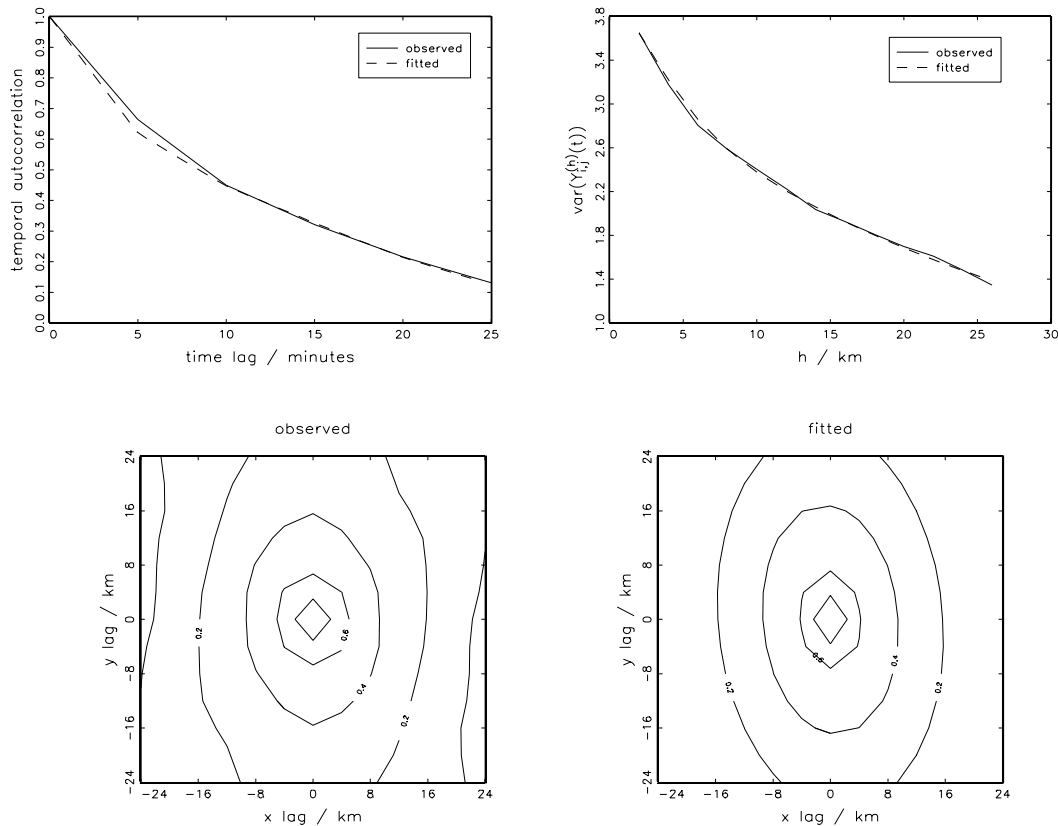


Figure 2.3: Assessment of the fit of the GDSTM to the rain event of 6 February 1994, 1300–1400. Top left: temporal autocorrelation, top right: scaling of variance, bottom: spatial autocorrelation.

spatial averaging.

Figure 2.3 illustrates the goodness of fit of the model in terms of the temporal and spatial autocorrelation functions (at the 2 km scale of spatial aggregation) and the space-time variance across scales of spatial aggregation h . There is good agreement between the observed and fitted values of these functions, the latter plot (in conjunction with table 2.3) demonstrating that the model is able to reproduce properties of the empirical data over a range of pixel sizes.

As a further assessment of model fit we investigate the effect of imposing thresholds on the empirical data and on data simulated from the fitted model. To impose a threshold of 1 mm/h on an image we simply set all rainfall intensities below 1 mm/h to zero. In particular we consider the proportion of pixels with a rainfall intensity over a given threshold, the *coverage*, and the mean rainfall intensity of images that have been thresholded. Figure 2.4 shows how the coverage and the mean intensity of empirical and simulated images is related to the threshold imposed. Each curve has been estimated by averaging the results from a 1 hour sequence (i.e. 13 images) of data. The curves derived from the empirical data are generally consistent with the curves derived from the simulated data although the empirical data tends to have a greater coverage for large thresholds than the model. A further investigation examines the rainfall totals accumulated at each pixel over




Figure 2.4: Threshold analysis of the intensity field. Left: coverage vs. threshold. Right: mean intensity of thresholded image vs. threshold. Solid line is the empirical data curve. Dotted lines are 10 independent simulations from the fitted model.




Figure 2.5: Threshold analysis of the depth field. Left: coverage vs. threshold. Right: mean depth of thresholded image vs. threshold. Solid line is the empirical data curve. Dotted lines are 10 independent simulations from the fitted model.

the 1 hour sequence of data. We define the rainfall *depth* at a given pixel to be the sum of all the rainfall intensities observed at the pixel over a time period of interest. Figure 2.5 shows how the coverage and the mean depth of the empirical and simulated images is related to the threshold imposed. These plots exhibit broadly similar features to figure 2.4.

Figure 2.6 gives a visual illustration of the ability of the model to reproduce the internal structure of rain events. We have included images from outside the time period used for parameter estimation, to illustrate the movement of the rain event across the area (see section 3.2). The images have been simulated from the fitted model in a manner intended to mimic the movement of the rain event. This is achieved by defining a rain event area within which the model process is simulated and moving this area with the same velocity as that of the storms within it. The radar images and the realisation simulated from the fitted model exhibit broadly similar features, although the assumption of constant rainfall intensity over the area of a rain cell means that data simulated from the model is not as smooth as the empirical data. The peripheral drizzle apparent in the empirical




Figure 2.6: Visual assessment of the GDSTM fitted to the rain event of 6 February 1994, 1300–1400. The temporal separation of the images is 1 hour. Top: Observed data. Bottom: Simulation.

Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
20	31	20	18	11	5	6	7	15	21	34	19

Table 2.4: Variation in the number of events fitted over the year

data is not present in the simulated data. We may have to address this feature if it becomes clear that the drizzle surrounding rain events is important hydrologically. It remains to be seen whether output from rainfall-runoff models is sensitive to this feature of the simulated data. Close agreement between the two sequences in terms of exactly *where* areas of relatively intense rainfall occur is not expected, just as close agreement between two realisations from the same stochastic model would not be expected.

2.10 Parameter estimate summary

The GDSTM has been fitted to data from the interior of all suitable rain events from the Wardon Hill data record. The data record spans a time period from September 1993 to March 1997 inclusive. For the purposes of simulation we will need additional information for each event, namely the manner in which the rain event enters the radar window and its duration which we define to be the length of time during which the radar coverage is above 25% (see section 3.1). Missing radar images may mean that this information is not available for some rain events. In such cases we are forced to exclude these events. Thus the model has been fitted to a sample of 207 events; table 2.4 shows how the number of these events per month varies over the year. It can be seen that there are fewer such events in the summer months than in the winter. Figures 2.7 and 2.8 summarise the direction and

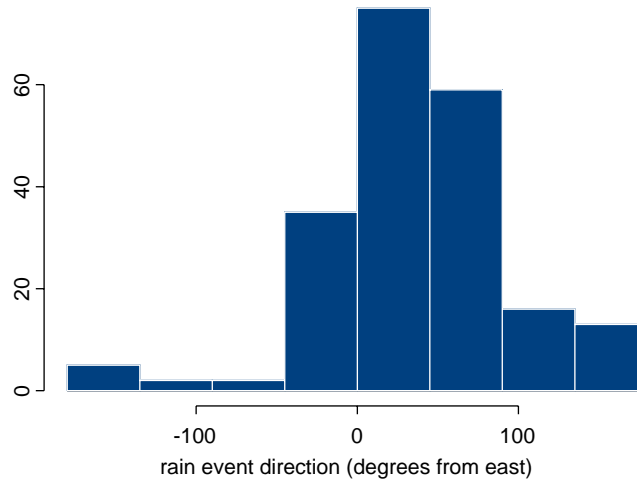


Figure 2.7: Direction of movement of rain events.

	λ_t (h^{-1})	μ_D (mins)	μ_L (h)	μ_A (km^2)	μ_S (km^2)	μ_C	μ_X (mm/h)	e	speed km/h
mean	12	16	1.5	21	340	1400	0.74	0.78	47
median	7.2	14	1.1	18	270	770	0.41	0.84	42

Table 2.5: Summary statistics for model parameters

speed of the rain events to which the GDSTM has been fitted. (Following mathematical convention, the rain event direction is defined as the direction in which the events are moving.) It can be seen that rain events tend to come from west to south-westerly directions and that rain events moving faster than 100 km/h are rare. Figure 2.9 shows a relatively even spread of cell orientations with perhaps a slight preference for the major axes of cells to lie in a west-east direction.

Table 2.5 gives the mean and median of the distributions of the marginal parameter estimates for the year as a whole. These values agree well with those in table 2.1. Note that we have used the variable λ_t , the temporal rate of storm arrivals in the radar window, rather than λ for ease of interpretation. The fact that the median is less than the mean in all cases except for e indicates that these distributions are positively skewed. Figure 2.10 illustrates how the mean of each of the 11 model parameters varies over the year. There is less of a clear seasonal pattern than we might expect. However, we should note that we have selected rain events of a particular type and so we might expect more similarity between events in, January and July, say, than if we had selected rain events at random from these months. There is clearer evidence of seasonality in the temporal rain event arrival process studied in section 3.1. In addition we should note that the estimated means for the summer months are based on a small number of rain events.

Tables A.1 to A.12 in appendix A.4 list the time periods of data to which the model has been fitted and the resulting parameter estimates.

The histograms in figure A.1 of appendix A.4 show the marginal distributions of each of the

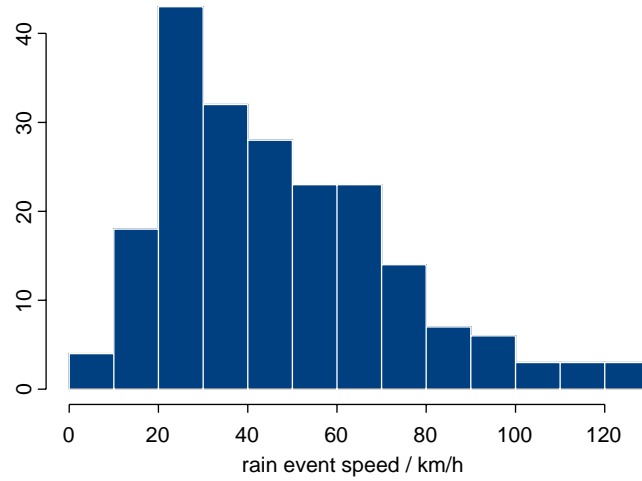


Figure 2.8: Rain event speed.

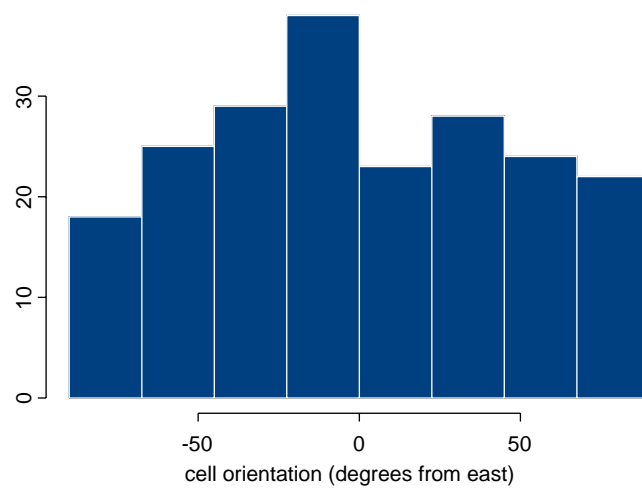


Figure 2.9: Cell orientation.

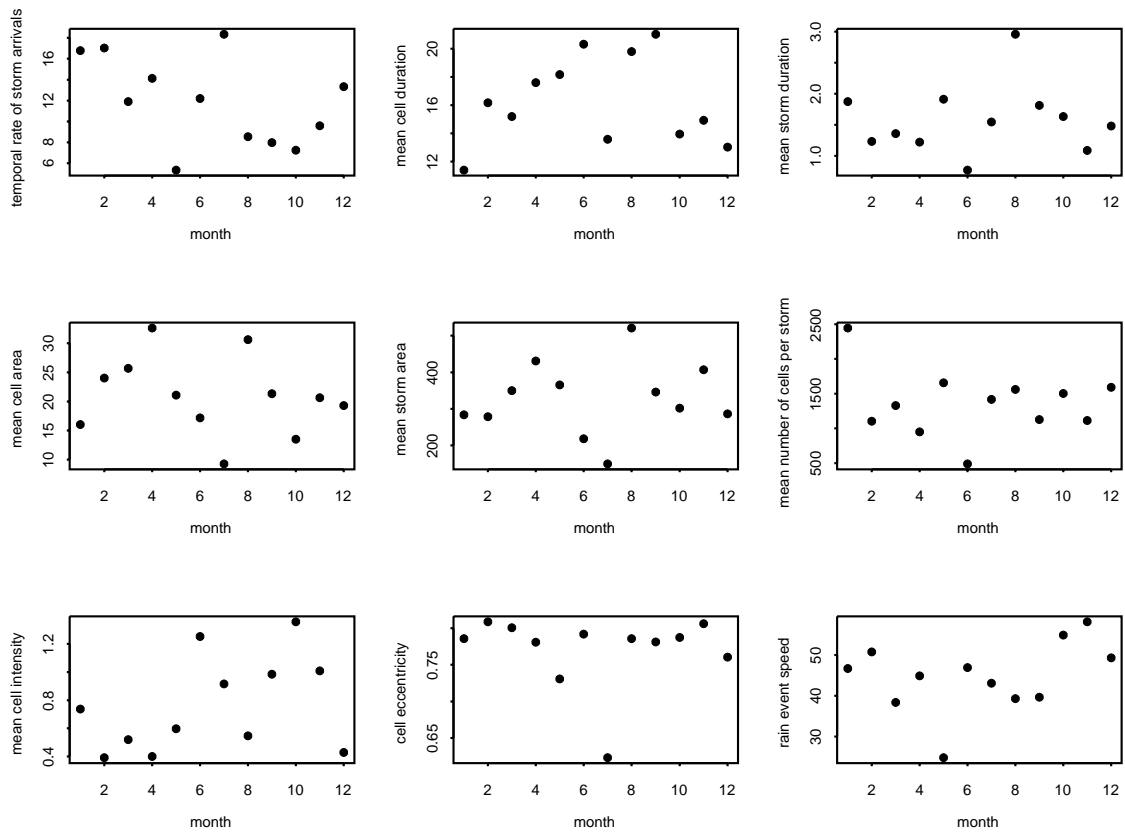


Figure 2.10: Monthly variation in the mean of parameter estimates.




Figure 2.11: Association between the parameter estimates for mean storm area and mean cell area

11 model parameters pooled over all rain events. Figures A.2, A.3, A.4 and A.5 show equivalent histograms pooled over events from winter, spring, summer and autumn respectively.

We do not expect parameter estimates to be mutually independent. For example, we might expect μ_A and μ_S to be positively correlated and indeed a plot of $\hat{\mu}_S$ against $\hat{\mu}_A$ for our sample of 207 rain events reveals a positive association (see figure 2.11; note that the four outliers with atypically large mean storm areas relative to their respective mean cell areas could be due to storms with widespread showers or to artefacts in the data).

Rather than attempt to specify a multivariate distribution for 11 parameters of the model, for the purposes of simulation, we propose to sample parameter sets from a parameter library (see section 3.2.2 for more details). This will allow for inter-parameter dependence.

Chapter 3

Spatial-temporal modelling using radar data II: modelling of event sequences

3.1 Rain Event Arrivals

3.1.1 Introduction

So far, we have focused attention on the interior of rain events in both space and time. The rain events to which the GDSTM has been fitted have a spatial extent sufficient to cover the *radar window* (a square area of approximate dimensions 100 km by 100 km) for a time period of at least an hour. In the case of the Wardon Hill radar, for which $2\text{ km} \times 2\text{ km}$ resolution data in the radar window have been used for model fitting, the radar also provides data at a coarser resolution ($5\text{ km} \times 5\text{ km}$ pixels) over a larger, circular region of radius 210 km. From these data we observe that the vast majority of rain events in the Wardon Hill data record appear to be part of larger scale weather systems moving across the region. We now consider these events and their movement over the window, and model this process in two parts: the timing of the sequence of arrivals and departures of rain events and the shapes of these events in so far as they are relevant to the radar.

3.1.2 Observation of Rain Events

We first need to define what we mean by the arrival or departure of a rain event from the radar window. Observation of sequences of radar images leads us to distinguish two types of rain event:

1. those that extend over the whole radar window for at least an hour (suitable for model fitting);
2. those that produce a significant amount of rainfall for only a short period of time. These events often appear to have similar structure and size to the previous events but only to intersect the edge or corner of the radar window rather than covering it entirely. Alternatively, they may be events with very localised rainfall.

At other times there may appear to be very light and/or scattered rainfall. Unfortunately, it does not seem possible to distinguish true light rainfall from radar ‘clutter’ and other errors. In particular, we observe particular individual, or local configurations of, pixels that remain wet over long sequences when other surrounding pixels are recording no rain or changing from wet to dry and conversely.

Visual inspection of the radar images has lead us to base our definitions of rain events on the *coverage*, the proportion of pixels in the radar window recording positive rainfall (we term such pixels *wet*). We can calculate the coverage for each image over the radar window, at 5 minute intervals, thus obtaining a time series of the coverage. The time series for January 1996 is presented on figure 3.1; a gap in the series indicates missing data.




Figure 3.1: Time series of the coverage for Wardon Hill radar window ($2 \times 2\text{km}^2$ data) — January 1996.

Events to which the model is fitted are selected partly on the basis that they are widespread and the coverage gives a good indication of when such events are present in the radar window. In addition, measures such as the mean rainfall intensity over the image can be greatly affected by a single value in the data (e.g. a pixel falsely recording a very large intensity), whereas the time series of coverage will generally be smoother and more robust to outliers.

Inspection of the Wardon Hill data record suggests that imposing a threshold of 25% separates periods corresponding to dry or very light rainfall with background noise, from the two event types described above. In addition, the extended events generally have a maximum coverage above 50% whereas that for the shorter or very localised events is generally below 50%. These thresholds will therefore be used in the next section to define the arrival and departure times of two types of rain

event. For illustration, a time series of the coverage for the 12th of January 1996 is shown in figure 3.2, and a 15-minute image sequence corresponding to the time period 9.30–13.30 is displayed in figure 3.3. The coverage starts from 13% at 9.30am, crosses the 25% threshold between 9.45 and 10.00 and increases to a local maximum value of 44% at about 10.30am. At least 25% of the radar window is covered by rainfall continuously for about two and a half hours. At around 12.45 there is a down crossing of the 25% threshold and the coverage continues to decrease to 14% by 13.30. The following excursion above the 25% level lasts more than four hours and the coverage reaches a maximum of about 70%.




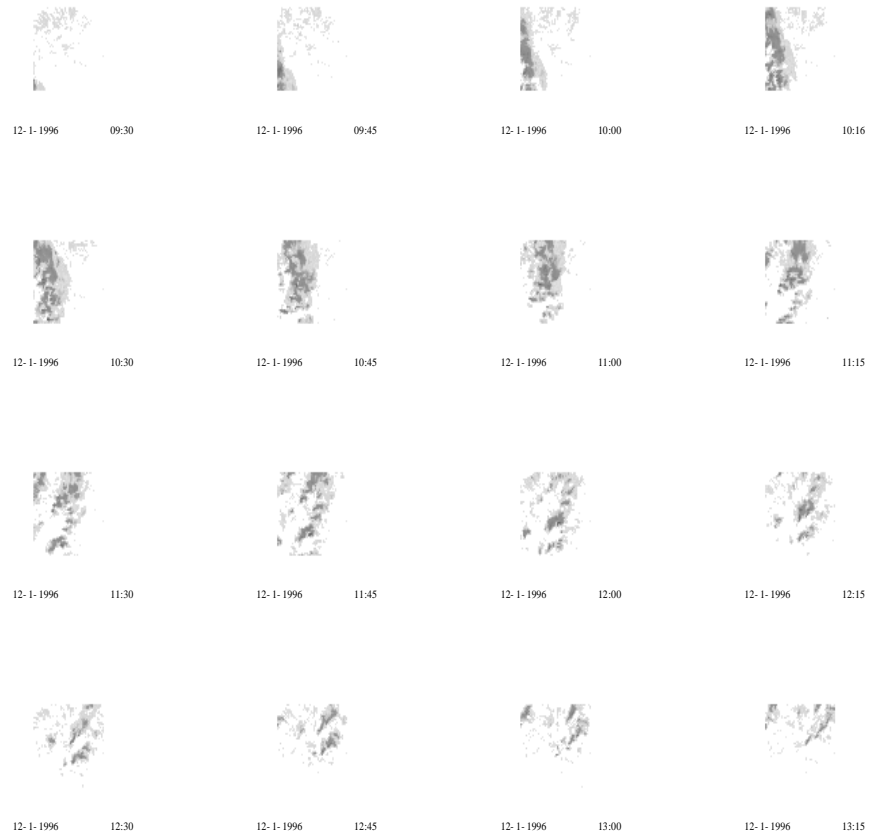

Figure 3.2: Time series of the coverage for Wardon Hill radar window ($2 \times 2 \text{ km}^2$ data) — 12 January 1996.

3.1.3 Modelling Event Durations and Inter-Event Intervals

The following definitions for the durations of two types of rain events over the radar window and the intervals between them are given in terms of a smoothed version of the time series of spatial coverages. We found that a 5-point smoother applied to the series was effective in smoothing out the effect of small oscillations around the 25% threshold (i.e. each value in the smoothed series of the coverage is an average over a 20 minute period).

A rain event ‘arrives’ at the window when there is an upcrossing of the 0.25 threshold in the smoothed coverage series and ‘departs’ from the window at the next downcrossing of this threshold, with its *duration* or *lifetime* being the interval between these two time points. We distinguish two types of rain events: those that attain a coverage of at least 50% at some point during their lifetimes (*type 1*) and those that do not (*type 2*). The type 1 rain events so defined include essentially all

Rainfall intensity
(mm/hr):



Wardon Hill

Radar



Figure 3.3: Radar images from Wardon Hill radar window, from 09:30 to 13:30 on 12 January 1996

the events to which the spatial-temporal model was fitted as described in section 2. The interval between a downcrossing of the threshold and the next successive upcrossing will be described as a *dry interval* although, as noted above, there may be small amounts of genuine rainfall mixed in with the radar errors during this interval. Treating periods where the coverage is below 25% of the window as dry, and ignoring all positive pixel values in these periods, means that some genuine events of light rainfall are unavoidably excluded. However, we can see from figure 3.4, that even if all positive values in the so-called ‘dry’ periods represented genuine rainfall, their contribution to the total rainfall intensity would be generally very small. Nevertheless, in summer months, the ‘dry’ periods would contribute amounts of rain comparable to the type 2 events, and in July, the type 2 events and the ‘dry’ periods would each contribute more rain than the type 1 events. Whilst treating the ‘dry’ periods as literally *dry* has only very minor effects on the total rainfall statistics, it remains to be seen (this will be the subject of future research) whether such events play an important role as regards rainfall-runoff properties. Note that the values plotted in the figure are monthly averages over the smoothed rainfall record, which have been scaled up to allow for missing data.




Figure 3.4: Average monthly amounts of rain attributed to different event types.

We model the stochastic process of events and dry periods using a 3-state semi-Markov process (we identify dry intervals with state 0, and type i events with state i , $i = 1, 2$). Thus, we assume the succession of events and dry periods follows a Markov chain, while the durations of events and dry periods are mutually independent. Note that the definitions of the type 1 and type 2 events imply that these rain events are always separated by a period when the coverage falls below the 25% threshold, that is by a ‘dry’ interval. Hence the transition matrix of the Markov chain has the

form

$$\begin{pmatrix} 0 & p & 1-p \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

We model the distributions of the durations in the three states by three Weibull distributions. The 3-parameter Weibull distribution is chosen because it has a flexible shape, allowing the representation of considerable skewness and a positive lower bound for the variable if appropriate. The model is applied independently to each month of the year to allow for seasonality in the rainfall record.

The appropriateness of these assumptions has been investigated for the Wardon Hill data, and the results are described below. In particular, we need to verify that the event and dry period durations can be assumed to be mutually independent and that the sequence of states visited is Markovian. The unknown parameter p is estimated by the proportion of dry periods that are followed by type 1 events while the parameters for the three Weibull distributions are estimated by maximum likelihood. Missing radar images result in censored observations of durations of events and dry periods, but these are easily incorporated in fitting the Weibull distributions.

3.1.4 Distributions of Durations of Events and Dry Periods

Empirical Marginal distributions

A complete set of histograms and summary statistics for the durations of events and dry periods on a month by month basis is given in the Appendix (figures A.6 to A.17), using data available from the Wardon Hill radar; four years of data for September to March and three years for the remaining months. The results for January are also given here in figure 3.5 for illustration. The basic statistics of the empirical data for all months are summarised in tables 3.1 and 3.2. It should be noted that the censored observations have been included without any allowance for censoring in these histograms and summary statistics, but the values of these observations are shown on the figures.

We note the following points:

There are only 5 type 1 events for July over the three years data available, and only 12 each for June and August.

The proportion of censored event durations is generally less than 10% of the observations (the exceptions are in May, September and November) while the proportion of censored dry periods is usually above 10% (except January and February). This is not surprising as, for most months, the dry periods tend to be longer than the rain events. The proportions of censored observations make it important to include these data in fitting the duration distributions.

For type 1 events, the monthly minima of the durations are systematically substantially greater than zero. This is to be expected because, over the duration of such a rain event, the coverage has to increase from 25% to 50% and then decrease back below 25%. Given the typical range of sizes and speeds of rain events, this process is likely to take at least an hour or two in almost all cases. We therefore fit three-parameter (right-shifted) Weibull distributions to the durations, where the third parameter specifies the minimum value of the support of the distribution.

The monthly maxima of the dry intervals can reach very high values (May, 13 days; July 7

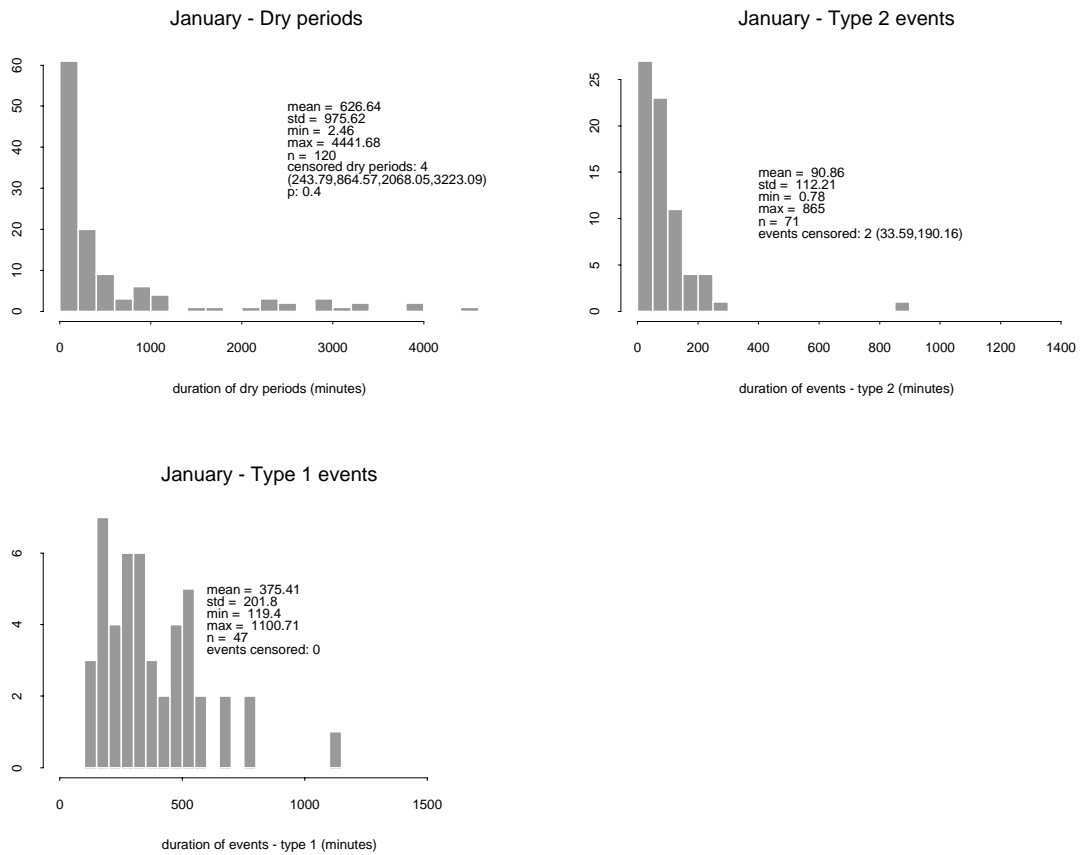


Figure 3.5: Histograms of durations of events and of ‘dry’ periods — January

days). For events of type 1, the largest observed maxima correspond to the months of April (21 hours), May (27 hours) and November (26 hours). For events of type 2, the largest maxima are in March (11 hours) and November (21 hours). Generally, the duration distributions appear very skew, so that isolated maxima in the upper tails may make the fitting procedure unstable when there are small numbers of observations.

Independence assumptions

In order to justify modelling the process of rain events as a semi-Markov process, we need to verify that it is reasonable to assume that successive event or dry-period durations are mutually independent. To this end, for each month a set of nine scatterplots has been produced; see the Appendix, figures A.18 to A.29.

To explain these scatterplots we define some notation as follows. Consider a sequence, $\{X_t, (Y_t, I_t)\}$,

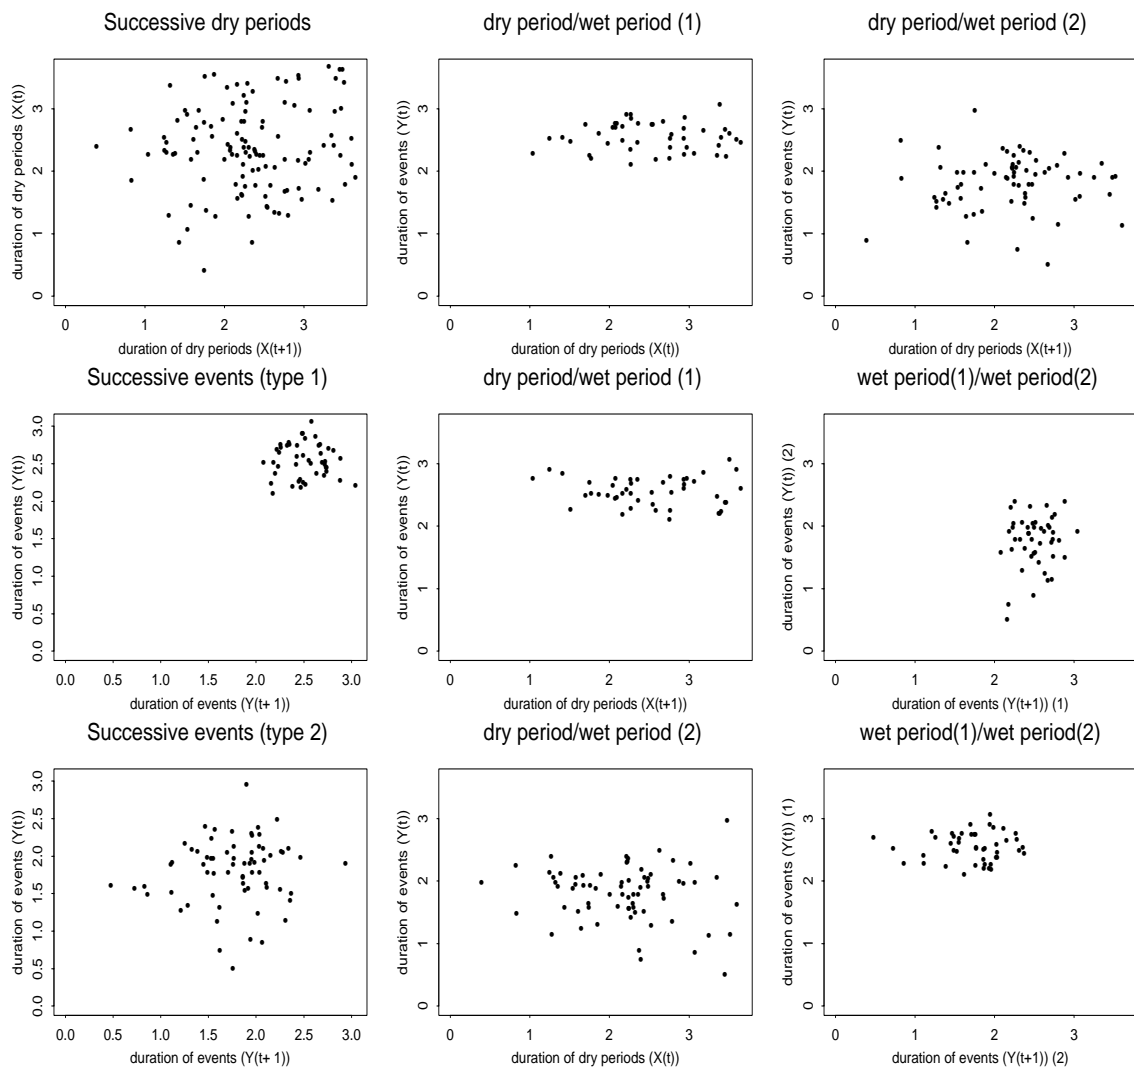


Figure 3.6: Scatterplots showing relationships between durations of successive events and dry periods (January). Axis units represent $\log_{10}(\text{duration in minutes})$

Month	Dry periods(mins)					Type 1 durations (mins)				
	n	mean	std	min	max	n	mean	std	min	max
Jan	120	626.64	975.62	2.46	4441.68	47	375.41	201.8	119.4	1100.71
Feb	131	689.44	944.47	2.69	5003.08	57	385.42	213.68	92.16	961.71
Mar	73	1006.99	1355.36	12.6	7315.94	27	490.76	263.82	111.54	1214.7
Apr	66	922.33	1320.17	16	7919.98	26	403.19	275.90	123.06	1267.74
May	69	1319.7	3336.31	12.64	19418.32	39	493.97	430.81	35.87	1609.78
Jun	33	1495.26	1868.4	18.6	6807.92	12	280.88	181.13	77.68	672.25
Jul	40	1371.18	2128.47	6.6	10848.52	5	303.34	135.54	163.58	503.87
Aug	44	1307.11	1880.19	6.96	8366.22	12	398.75	343.45	117.37	1135.51
Sep	67	1031.52	1484.81	6.75	6445	25	363.57	185.9	53.82	844.72
Oct	89	828.98	1423.61	6.77	9885	34	417.24	279.26	118.01	1066.31
Nov	90	872.81	1704.75	6.73	11241.03	36	472.42	334.35	133.23	1575.9
Dec	112	861.41	1381.15	11.32	8519.52	48	397.06	254.09	87.56	1110.92

Table 3.1: Basic Statistics of the empirical distributions for durations of dry periods and of events of type 1

of intervals corresponding to alternating durations of dry periods, X_t , and rain events, Y_t , where $I_t = i$ if the rain event is of type i ($i = 1, 2$). The scatter plots are chosen to show any association between the following pairs of intervals:

- the lengths of successive dry periods (X_t against X_{t+1});
- the lengths of successive events of the same type i (Y_t against Y_{t+1} , for t such that $I_t = I_{t+1} = i$) for $i = 1, 2$;
- the lengths of successive wet and dry periods in either order (that is, of Y_t against X_t , and of Y_t against X_{t+1} , for $I_t = i$ and $i = 1, 2$);
- the lengths of successive events of types 1 and 2 in either order (that is, of Y_t against Y_{t+1} when $I_t = 1, I_{t+1} = 2$, and when $I_t = 2, I_{t+1} = 1$).

The scatterplots described above show no evidence of dependence between successive event and inter-event durations. Those for January are given in figure 3.6 to illustrate the results obtained. The data plotted in this figure represent $\log_{10}(\text{duration in minutes})$.

3.1.5 Fitted Distributions for Event Durations and Dry Periods

If X follows a 3-parameter Weibull distribution, the probability density function of X has the following form (Johnson & Kotz 1970, chapter 20):

$$p_X(x) = (c/\alpha)[(x - \epsilon)/\alpha]^{c-1} \exp[-[(x - \epsilon)/\alpha]^c]$$

for $x \geq \epsilon$, where c is the shape parameter, α is the scale parameter and ϵ is the shift of origin.

Month	Type 2 durations (mins)					\hat{p}
	n	mean	std	min	max	
Jan	71	90.86	112.21	0.78	865	0.4
Feb	72	84.12	83.64	4.14	440.22	0.44
Mar	43	94.34	110.62	6.01	683.7	0.39
Apr	39	80.13	80.27	8.63	384.38	0.4
May	29	120.82	125.42	5.73	438.51	0.57
Jun	19	113.23	93.32	17.7	362.74	0.39
Jul	35	116.36	120.69	5.25	462.93	0.12
Aug	29	129.49	108.64	5.55	446.86	0.29
Sep	39	115.81	96.76	6.17	429.58	0.39
Oct	53	116.59	107.82	12.63	568.36	0.39
Nov	52	108.93	184.14	7.01	1283.73	0.4
Dec	63	89.61	75.28	7.5	403.42	0.43

Table 3.2: Basic Statistics of the empirical distributions for durations of events of type 2 and the estimated Markov transition probability \hat{p}

The corresponding survival function is given by

$$P(X \geq x) = \exp(-((x - \epsilon)/\alpha)^c)$$

while the first two moments are $E(X) = \alpha\Gamma(1+c^{-1})+\epsilon$, and $\text{Var}(X) = \alpha^2\Gamma(1+2c^{-1}) - (\Gamma((1+c^{-1}))^2)$.

Weibull distributions, as above, have been fitted by maximum likelihood (Johnson & Kotz 1970, pp.255–256) to the durations of events and dry periods for each month. Figures A.30 to A.38 in the Appendix show probability plots for each month as a means of assessing the adequacy of the Weibull distributions to model these durations (censored observations are not included in these plots). The results are generally encouraging; see figure 3.7 which shows the fitted distributions for January. The assumption of Weibull distributions for the durations seems broadly justified, although there are some obvious departures in the tails of the distributions. We also note that, in May, durations of type 1 events appear to cluster in compact groups (see Appendix, Fig A.34). We have found no obvious explanation of this effect.

The fitted parameters are given in table 3.3. In addition to examining their seasonal variation, we can also determine whether simpler distributional forms are appropriate. For example, there might be no need for a shift of origin ($\epsilon = 0$) for type 2 rain events and dry periods, while $c = 1$ corresponds to an exponential distribution. From inspection of the parameter values it appears that the shape parameters for the distributions of dry periods and type 2 events are fairly constant over the year, taking values between 0.5 and 0.7 for the dry periods, and around 1.0 for the type 2 events. However, the shape parameters for the type 1 events show more seasonal variation. Given the time resolution of the raw data (5 minute intervals), and the small number of data points involved in the fitting, shifts of origin (ϵ) of under 20 minutes do not seem significant, suggesting that it should be adequate to assume that $\epsilon = 0$ in fitting the durations of both the dry periods and the type 2 rain events.

To gauge the accuracy of these estimates and to investigate hypotheses about the parameter values in more detail, confidence intervals can be calculated. Formulae for approximate variances

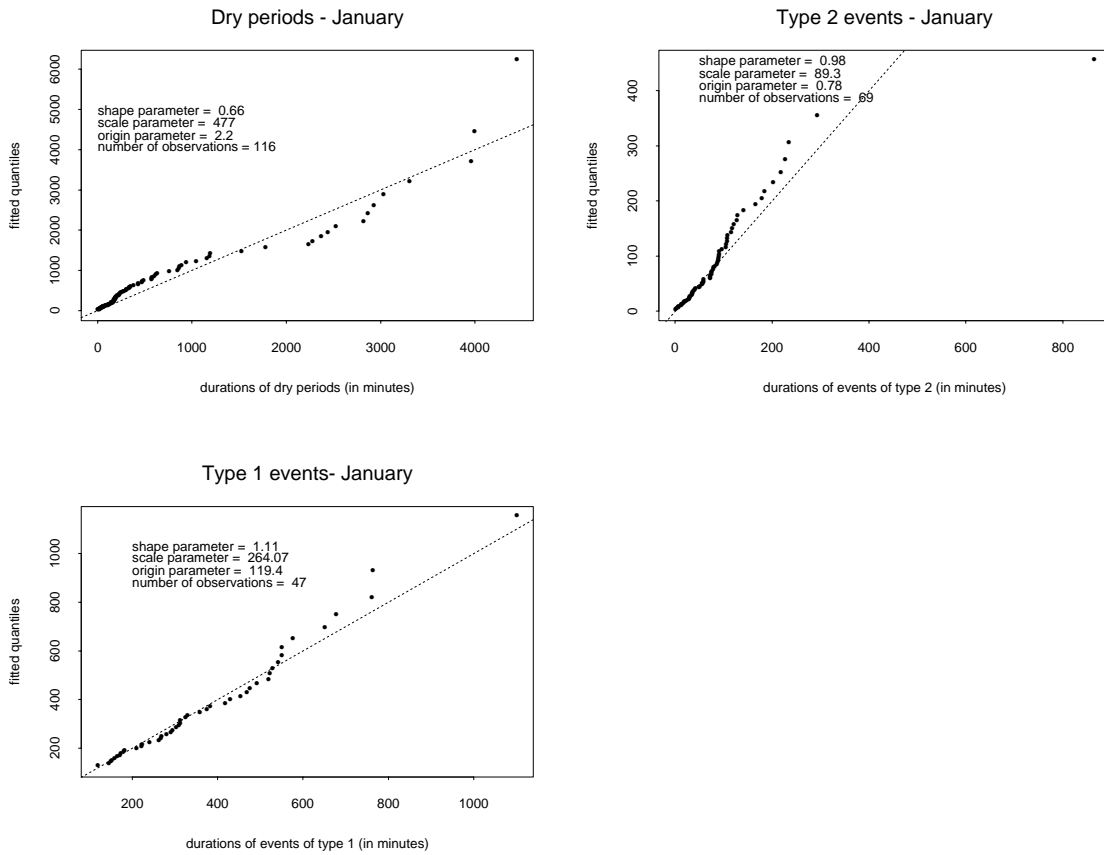


Figure 3.7: Fitting a Weibull Distribution, January (3 parameters per month and per type)

month	Dry periods (mins)			Type 1 durations (mins)			Type 2 durations (mins)		
	shape, c	scale, α	shift, ϵ	shape, c	scale, α	shift, ϵ	shape, c	scale, α	shift, ϵ
1	0.66	477.00	2.20	1.11	264.07	119.40	0.98	89.30	0.78
2	0.71	547.93	2.69	1.20	308.56	85.00	1.03	81.23	4.05
3	0.62	716.92	12.60	1.56	440.00	95.00	1.01	91.79	6.00
4	0.68	582.10	16.00	1.05	292.98	123.10	0.79	63.41	8.63
5	0.46	528.68	12.64	0.77	395.26	35.87	1.13	139.67	5.73
6	0.51	1021.78	18.60	0.66	175.13	77.68	1.25	111.37	17.60
7	0.44	399.11	6.60	0.41	86.57	163.58	0.82	90.00	5.20
8	0.55	894.86	6.96	0.62	217.50	117.30	0.85	118.27	5.55
9	0.70	807.59	6.75	1.98	390.50	17.00	0.98	105.38	6.17
10	0.68	548.53	6.77	1.05	306.50	116.20	0.93	97.25	12.63
11	0.54	379.60	6.73	0.87	320.70	133.23	0.92	107.00	6.95
12	0.64	496.29	10.57	1.00	313.79	87.56	1.30	97.55	7.50

Table 3.3: Estimated parameters of the Weibull distributions

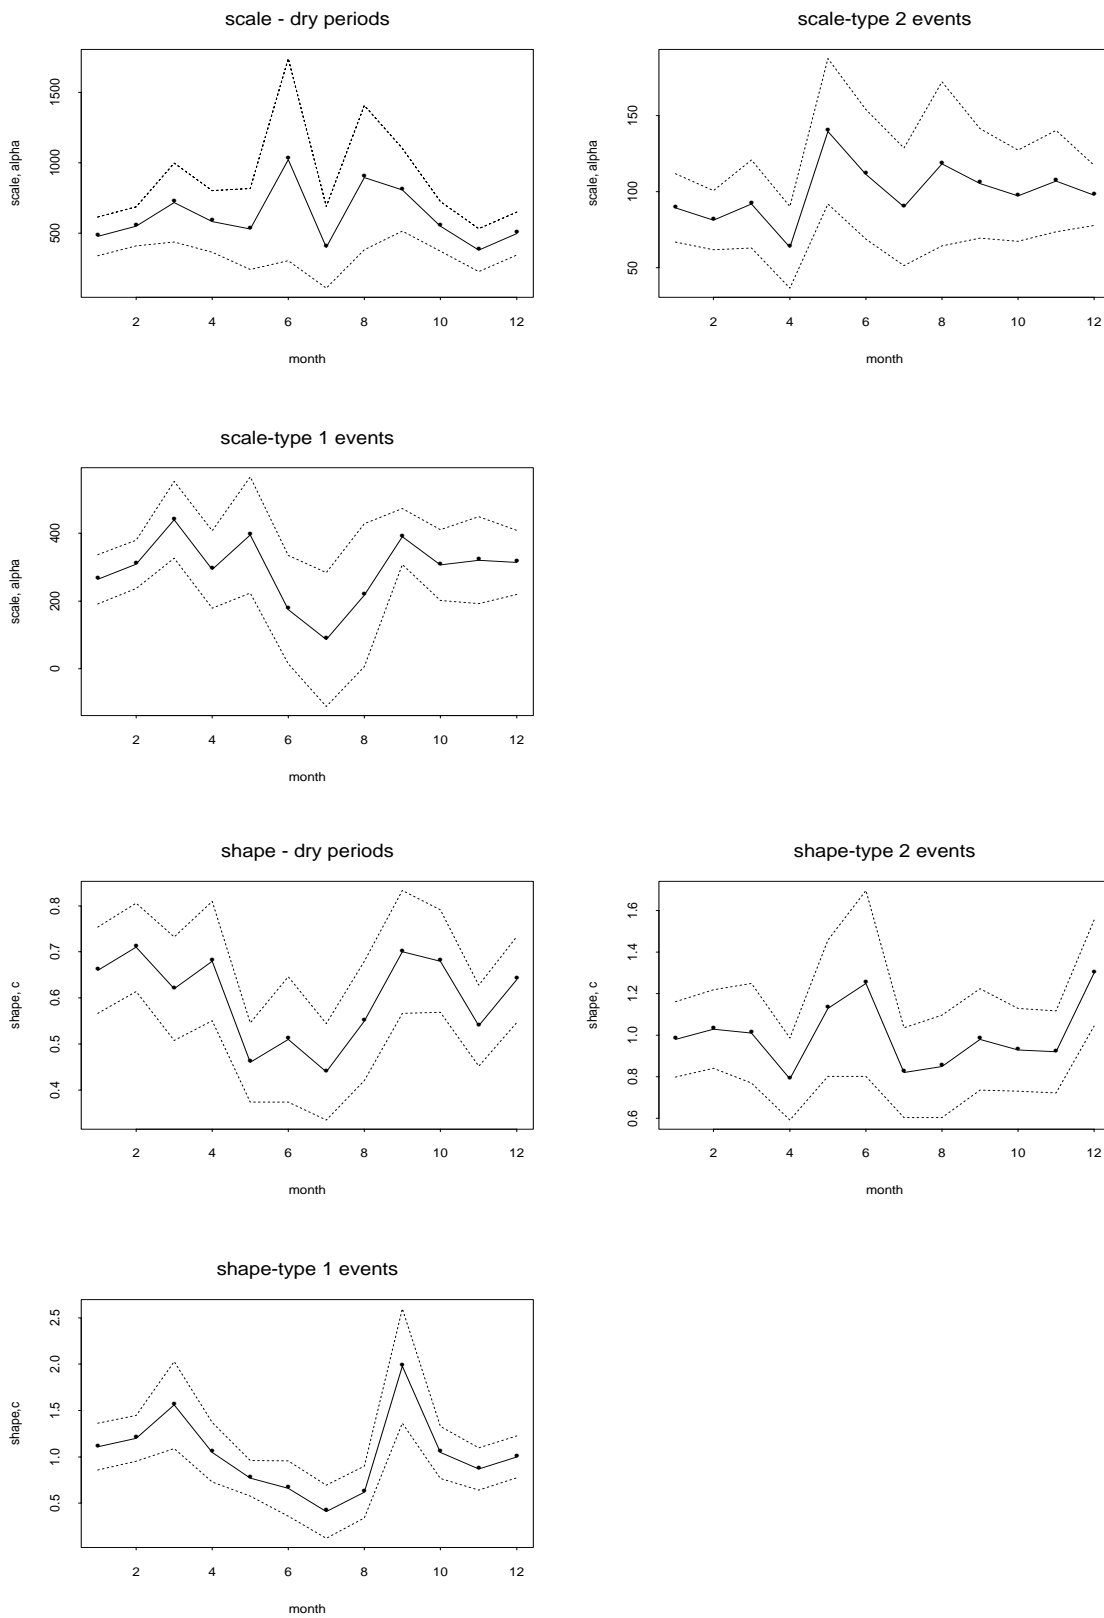


Figure 3.8: Weibull Parameters and Confidence Intervals. Top: scale parameter, α (in minutes). Bottom: shape parameter, c .

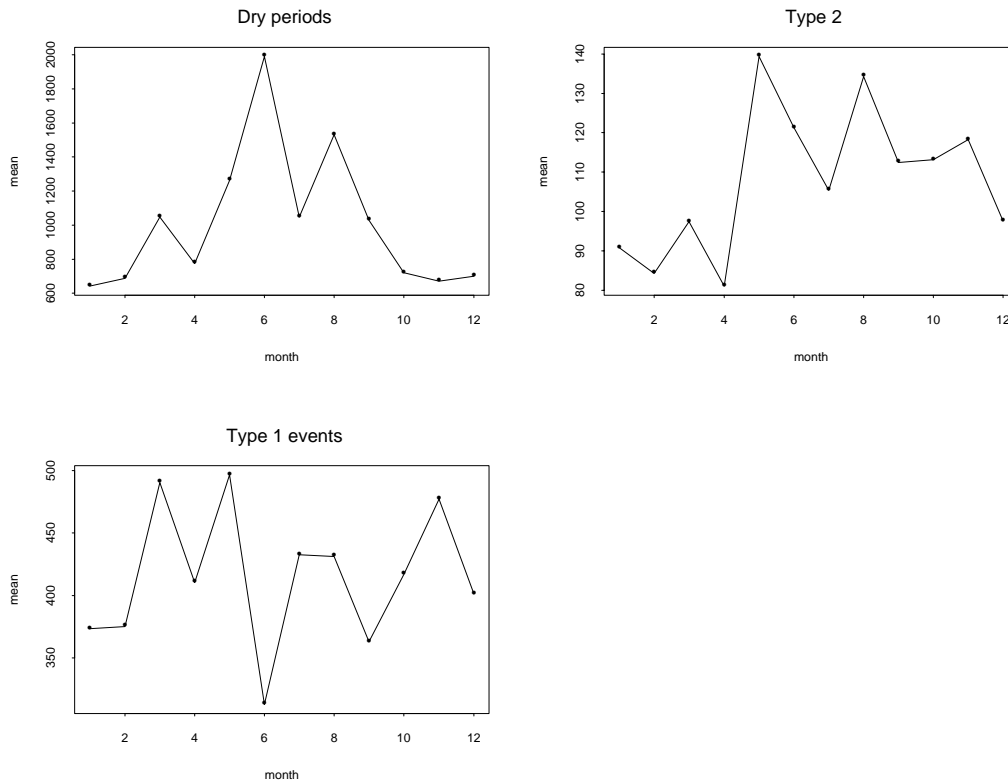


Figure 3.9: Parameters of the fitted Weibull Distributions — mean (in minutes)

of the estimates of c and α based on n independent observations are given by Johnson and Kotz (1970, p 256):

$$\text{Var}(\hat{c}) = 0.608\hat{c}^2/n, \quad \text{Var}(\hat{\alpha}) = 1.087(\hat{\alpha}/\hat{c})^2/n$$

leading to approximate 95% confidence intervals for these estimates which are shown (dotted lines) together with the estimates in figure 3.8. Note that, as discussed in connection with figure 3.6, there is no evidence against the assumption of independent observations.

From figure 3.8, it can be seen that an assumption of constant values (over months) for the scale and the shape parameters of each of the three distributions is implausible. The data are not consistent with the assumption of an exponential distribution ($c = 1$) for the durations of either the dry periods or the durations of type 1 events, although this assumption might be acceptable for the type 2 events.

For the durations of dry periods, and for the type 2 events, the estimates of the shift parameter ϵ are generally close to zero and always less than 20 minutes. In addition, there is no physical reason to suggest a positive lower bound for these durations, so assuming $\epsilon = 0$ in each case seems reasonable. However, we have already noted that, on physical grounds, type 1 events are likely to have a minimum duration of around $1\frac{1}{2}$ –2 hours, and this is confirmed by the majority of estimates in table 3.3. In fact, in most cases here the estimate is itself the minimum observed duration.

In an attempt to see whether there would be a more obvious seasonal pattern were the Weibull distributions reparameterised, we have plotted in figures 3.9, 3.10 and 3.11, the mean, standard

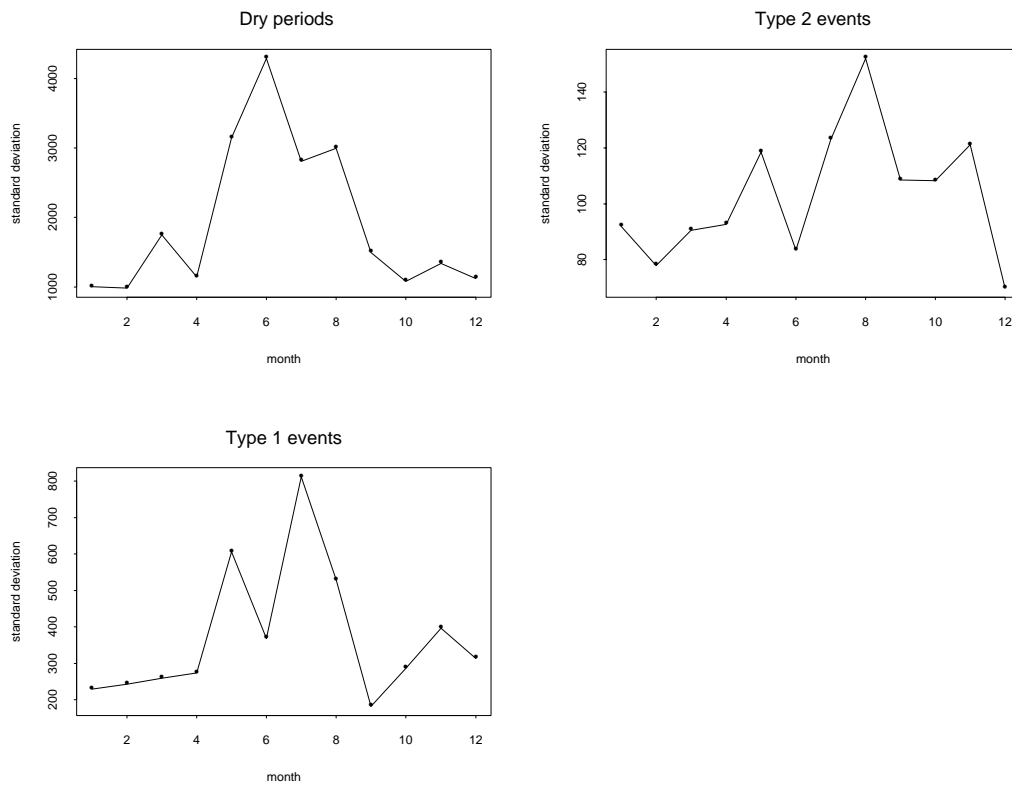


Figure 3.10: Parameters of the fitted Weibull Distributions — standard deviation (in minutes)

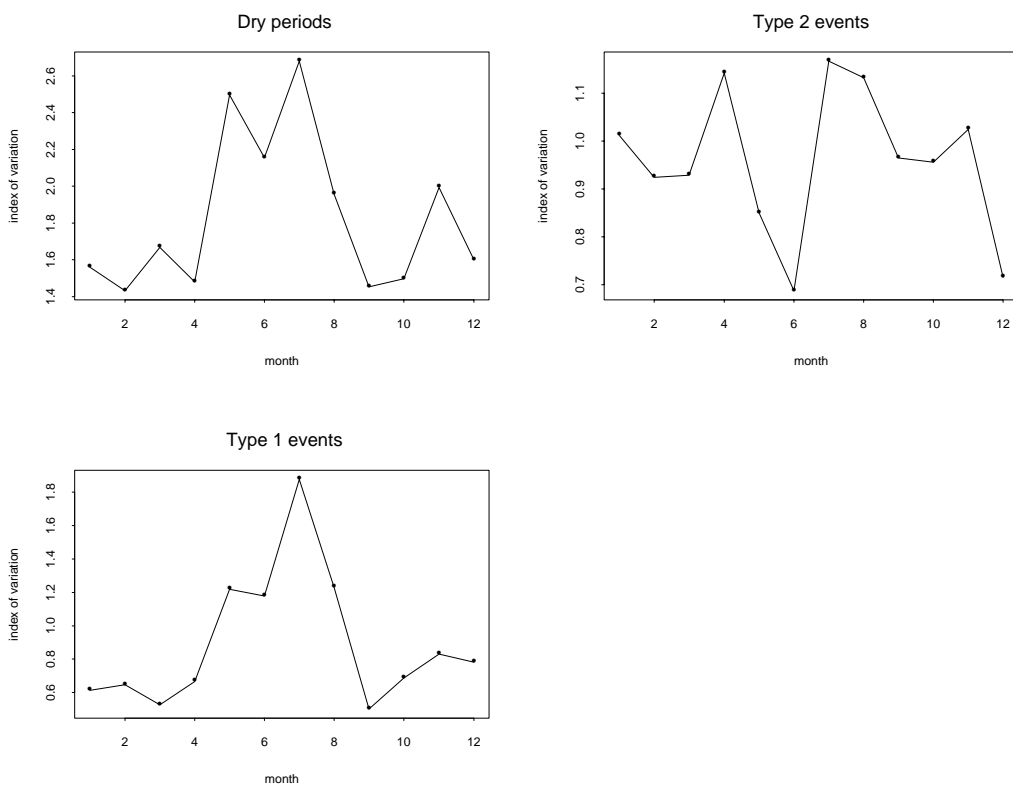


Figure 3.11: Parameters of the fitted Weibull Distributions — index of variation

deviation and an index of dispersion (here the coefficient of dispersion i.e. standard deviation/mean) of the corresponding fitted Weibull distributions. There is a clear and intuitively reasonable seasonal pattern in the mean durations of the dry periods and type 2 events. There is no particular seasonal pattern in the means for the type 1 events. This presumably reflects the definition of such events, which could be expected to occur with different frequencies over different months but to correspond to a distinct and identifiable weather type. The standard deviations of all three types of durations are generally greater in the summer. While, at least for dry periods and type 2 events, one might be tempted to postulate that the standard deviation increases in proportion to the mean, it may simply be an effect of the relative scarcity of rain events in summer months.

These fitted properties can also be used as a means of assessing the goodness of fit of the Weibull distributions, by comparison with the corresponding properties of the empirical distributions (see figures 3.12, 3.13 and 3.14). The fitted and empirical means seem in generally good agreement; the exception is in July, for type 1 events, but we have previously noted the very small number of such events in the data record. As is to be expected, the plots for the standard deviation and (therefore) the index of dispersion show rather more variation, with July being an obvious outlier throughout.

3.1.6 Model Simplification

In view of the results of the preceding section, we can reduce the number of parameters by setting $\epsilon = 0$ for the durations of dry periods and type 2 events. The results of fitting the reduced models have been compared with the previous ones. Fitted parameter estimates are given in table 3.4. Probability plots for each fitted Weibull distribution against the empirical distribution of the durations are shown here for January in figure 3.15, and in the Appendix, figures A.39 to A.44 for all 12 months. The model fits seem generally acceptable. The reduction in parameters has not led to any obvious reduction in the adequacy of fit.

month	Dry periods (mins)			Type 1 durations (mins)			Type 2 durations (mins)		
	shape, c	scale, α	shift, ϵ	shape, c	scale, α	shift, ϵ	shape, c	scale, α	shift, ϵ
1	0.68	463.62	0	1.11	264.07	119.40	1.03	92.27	0
2	0.74	564.19	0	1.20	308.56	85.00	1.15	88.89	0
3	0.7	791.38	0	1.56	440.00	95.00	1.07	93.21	0
4	0.73	749.21	0	1.05	292.98	123.10	1.09	82.96	0
5	0.55	689.93	0	0.77	395.26	35.87	0.98	119.94	0
6	0.66	1135.86	0	0.66	175.13	77.68	1.4	125.47	0
7	0.62	938.99	0	0.41	86.57	163.58	0.98	115.27	0
8	0.65	964.83	0	0.62	217.50	117.30	1.15	135.81	0
9	0.70	805.45	0	1.98	390.50	17.00	1.23	124.21	0
10	0.72	652.63	0	1.05	306.50	116.20	1.22	125.23	0
11	0.59	528.35	0	0.87	320.70	133.23	0.95	105.46	0
12	0.73	690.18	0	1.00	313.79	87.56	1.27	97.47	0

Table 3.4: Estimated parameters of the Weibull distributions (simplified models).

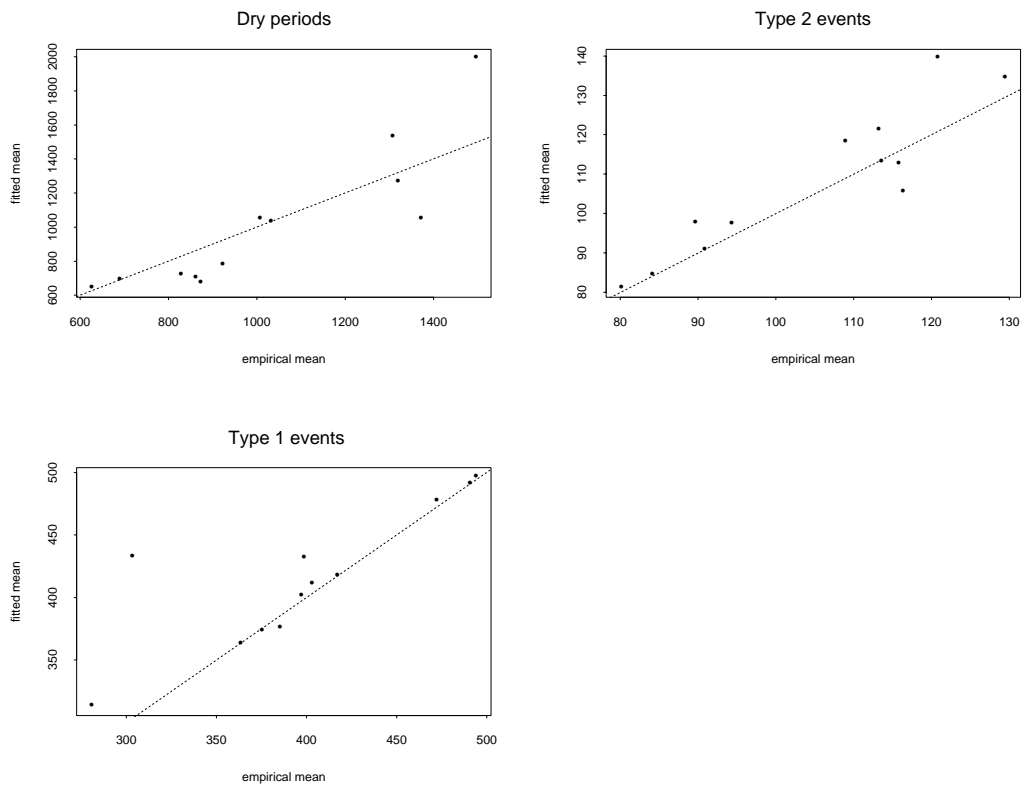


Figure 3.12: Mean of the fitted Weibull distribution against empirical mean (in minutes)

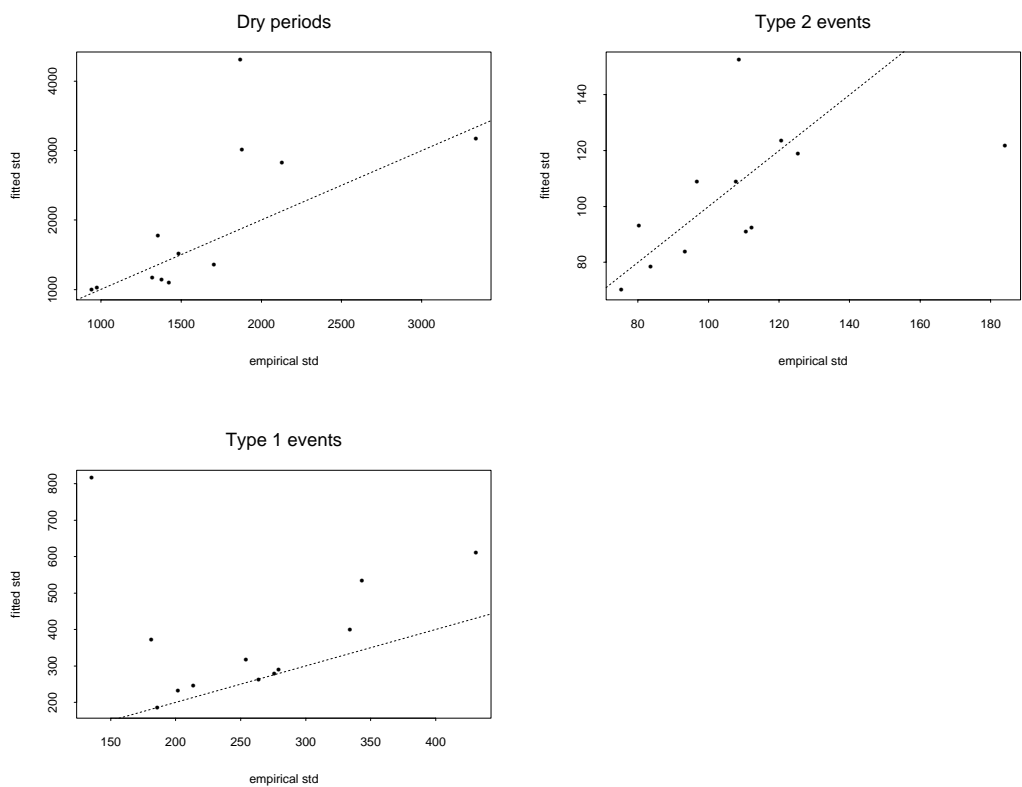


Figure 3.13: Standard deviation of the fitted Weibull distribution against empirical standard deviation (in minutes)

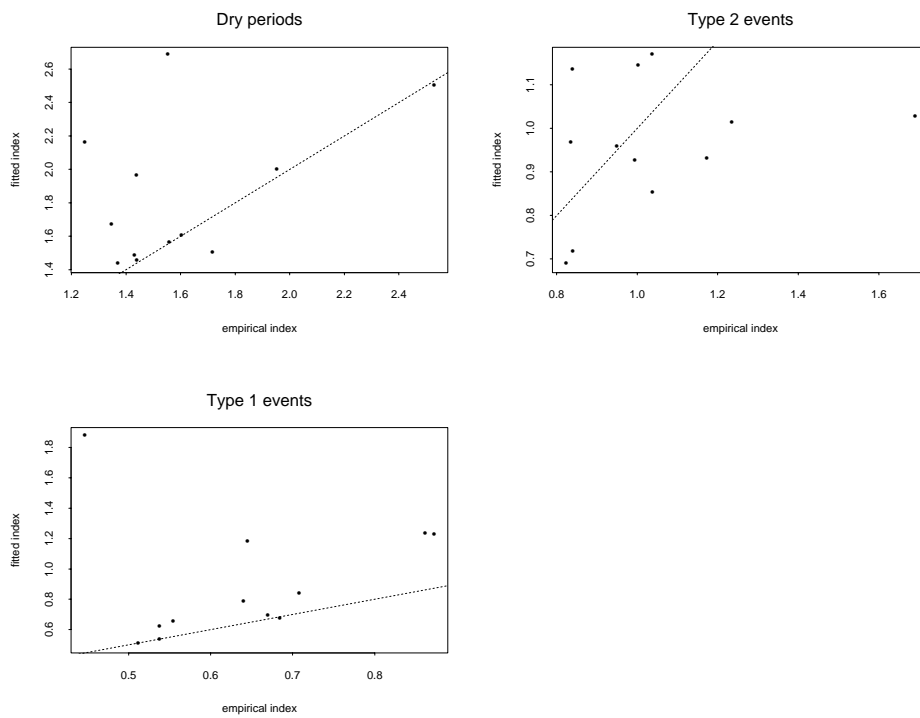


Figure 3.14: Index of variation of the fitted Weibull distribution against empirical index of variation

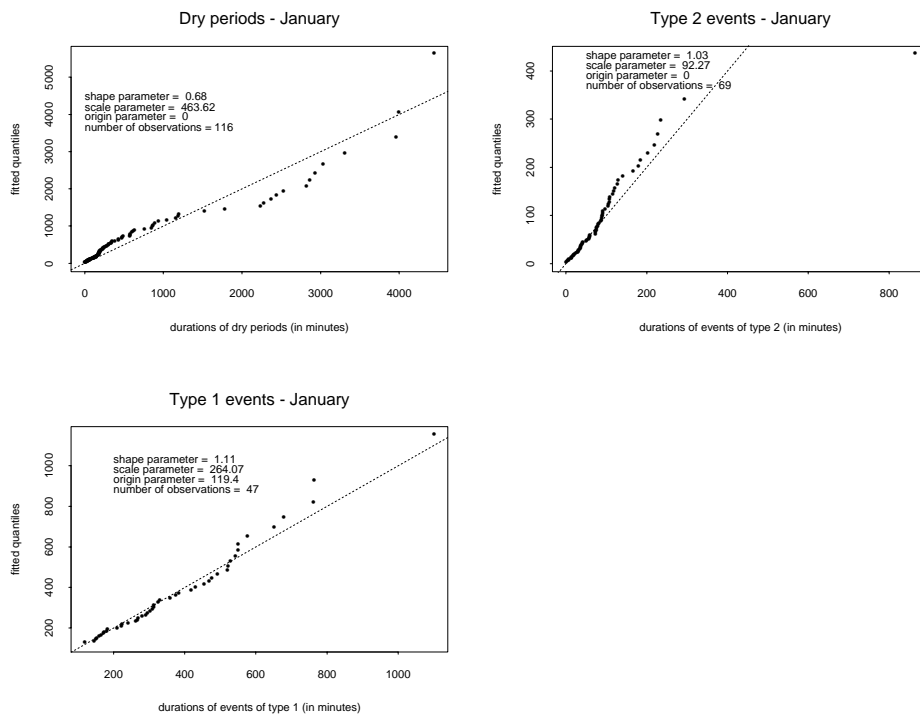


Figure 3.15: Fitting a Weibull distribution, January (simplified model: 2 parameters per month for dry periods and events of type 2, 3 parameters per month for events of type 1)

3.2 Modelling the event sequence

We are now almost ready to combine our stochastic model for the interiors of rain events with the above model for the sequence of rain events passing over the radar window. As noted earlier, the vast majority of rain events in the Wardon Hill data record appear to be part of larger scale weather systems moving across the radar window. In our model, the arrival of a rain event coincides with the time when the radar coverage rises above 25% and its departure with the time when the radar coverage drops below 25% again. We will therefore generate rainfall data over the window by simulating from the GDSTM within a quadrilateral rain event area that moves across the radar window. The shape and size of the rain event area is determined by the interval between the arrival and departure times of the event, its velocity and an analysis of the orientation of leading and trailing edges of rain events in the data record which is described below.

3.2.1 Leading and trailing edges

In order to specify the shape of the rain event areas we need to estimate the orientation of leading and trailing edges for the rain events to which the model has been fitted. In an ideal situation we would have distinct wet and dry regions and a clearly defined straight edge to the incoming/outgoing rain event. Many rain events exhibit this behaviour to a reasonable extent, in that the approximate angle at which the rain event enters/leaves the radar window can be estimated by eye on inspection of radar images. We devise an automated method of performing this function as follows.

We classify each pixel of the image as ‘wet’ (positive rainfall intensity) or ‘dry’ (zero rainfall intensity) so that we have a sample of (x, y) locations of pixel centres for each of these two populations. We use Fisher’s Linear Discriminant Function (LDF) (Fisher 1936) to find a straight line that splits the image into 2 regions - which can be termed ‘wet’ and ‘dry’ (the ‘wet’ region will contain mainly ‘wet’ pixels, while ‘dry’ pixels will predominate in the ‘dry’ region). Fisher’s LDF chooses this dividing line to maximize the ratio of the between-region variance to the within-region variance of the pixels. In fact we are only interested in the gradient of the estimated straight line.

We estimate the orientations θ_L and θ_T of the leading and trailing edges of the rain event using Fisher’s LDF as described above. We select one radar image for the purpose of estimating θ_L and another for θ_T . These images are selected from the time periods during which the rain event is in the process of entering and leaving the radar window respectively. Let $r(t)$ denote the radar coverage at time t and (t_s, t_e) denote the time period used for parameter estimation. We estimate θ_T using the image before the first image after t_e with a coverage less than $r(t_e)/2$. We use a similar selection rule to select an image before t_s for use in estimating θ_L . This procedure is preferred over an alternative strategy in which we use the first and last images of a rain event with coverages greater than 25%. For some rain events it is apparent that the latter approach is overly influenced by peripheral drizzle surrounding the main part of the rain event and may produce misleading results.

3.2.2 Selection of parameter sets

We construct a library of estimated parameter sets for each month of the year. For each rain event, we add to this library the estimated duration of the event, together with estimates of the

orientations of its leading and trailing edges. Suppose that we wish to select a parameter set with which to simulate a rain event of duration d_{sim} , where d_{sim} has been sampled from one of the month-specific Weibull distributions estimated in section 3.1. A simple strategy is to sample a parameter set from the parameter library for the month in question at random, taking no account of the value of d_{sim} . However, we might expect d_{sim} to be related to some of the parameters of the model. For example, physical considerations lead us to expect rain event speed to be inversely related to rain event duration - fast moving events tending to produce relatively short durations and vice versa. Although we do not observe a strong relationship between these two variables we find that we do not tend to observe fast moving events with long durations (see figure 3.16).




Figure 3.16: Association between speed and duration of rain events

There is no immediately apparent relationship between rain event duration and any of the other model parameters.

In order to respect the dependence between rain event speed and duration (and any less obvious dependence between d_{sim} and the other parameters) we restrict our choice of parameter set to those with a duration d_{obs} close to the simulated duration d_{sim} . Additionally, we should ensure that we select different types of parameter sets in the same proportions that they were observed. For example, if we had observed 1 very intense event in the 4 Januaries from 1993 to 1997, we should ensure that the expected number of such events in a simulation of 4 Januaries is 1. These considerations lead us to the following parameter set selection scheme:

1. partition the time interval $(0, \infty)$ into m sub-intervals, where m is the number of parameter sets in the library for the month in question, based on the percentiles of the relevant Weibull distribution for rain event duration;
2. order the parameter sets with respect to their durations (shortest duration first);
3. generate a duration d_{sim} from the relevant Weibull distribution and determine the percentile i of the Weibull distribution within which d_{sim} lies.

4. select the parameter set with the i th largest duration.

The scheme achieves the two objectives detailed above – the duration associated with the selected parameter set will tend to be close to d_{sim} and the unconditional (on d_{sim}) probability of selection is $1/m$ for each parameter set, where m is the number of parameter sets (for the month in question) in the parameter library. It is the scheme used to produce the simulated data analysed in section 3.3.

The simulated rain event duration d_{sim} for a given parameter set is restricted to a single percentile of the relevant Weibull distribution. If this is deemed too restrictive we may wish to increase the range of values d_{sim} can take for a given parameter set. One possibility is to partition $(0, \infty)$ into n (n is a factor of m) sub-intervals so that each percentile of the Weibull distribution is associated with a collection of parameters sets, from which we select one at random. If n is not a factor of m a more complex rule is required to ensure that the probability of selection is $1/m$ for all parameter sets. A better approach is to replace 4. above by

$$\text{select parameter set } \begin{cases} i - 1 & \text{with probability } 1 - p/2 \\ i & \text{with probability } p \\ i + 1 & \text{with probability } 1 - p/2 \end{cases}$$

for $i = 2, \dots, n - 1$ and some suitably chosen p , and a similar rule for $i = 1$ and $i = n$. That is, we may select the parameter set either side of the set initially selected (if such a parameter set exists). This scheme can be extended in an obvious way to allow the selection of parameter sets further away from the parameter set initially selected.

Finally, we note that very few type 1 events were observed in certain months of the year (see table 2.4). In particular the parameter libraries of fitted parameter sets for June and July contain only 5 and 6 parameter sets respectively. It is reasonable to expect only small differences between adjacent months in the year. Therefore, when simulating rainfall for February, for example, we might select parameter sets from a combined library containing all the parameter sets from January, February and March. Such a scheme would have the effect of smoothing seasonal variation in parameter values. A similar problem pertains to the treatment of rain event durations.

3.2.3 Simulation and data storage

The simulation program described and listed in appendix A.3 simulates the temporal process of dry periods and type 1 and type 2 rainfall described in section 3.1. Within each period of type 1 rainfall we generate rainfall data by simulating from the spatial-temporal model described in section 2.1 within a quadrilateral rain event area that moves across the radar window. The selection scheme given in section 3.2.2 is used to select a parameter set for this purpose, the velocity of the selected parameter set determining the velocity of the quadrilateral rain event area. At this stage we do not simulate data within periods of type 2 rainfall. By definition these events do not cover the entire radar window and further investigation is required to characterise their movement, size and shape. The absence of type 2 events from the simulation is anticipated to have relatively little impact as events of the second type contribute only a very small proportion of the total rainfall recorded by the radar (see section 3.2).

The data produced by the simulation program are stored in a similar format to the Wardon Hill radar data themselves. Simulated rainfall intensities are stored to the nearest $1/32$ mm/hr in

all cases. The Wardon Hill radar data are stored to the nearest 1/32 mm/hr for small intensities rising to the nearest 2 mm/hr for very large rainfall intensities. Thus, the data produced by the simulation program can easily be adjusted so that it is in the same form as the Wardon Hill data. For further details on data storage see appendix A.3.

3.3 Continuous Simulation Performance

We now assess the performance of the spatial-temporal model by comparing basic summary statistics for simulated and empirical data, to see whether their statistical properties are similar. The statistics calculated include the mean, standard deviation, spatial autocorrelation function and temporal autocorrelation function. The proportion of wet pixels through both space and time is also calculated.

3.3.1 Calculation of summary statistics

The performance of the spatial-temporal model is assessed at various spatial and temporal scales. The spatial scales considered are $4km^2$, $16km^2$, $64km^2$, $256km^2$ and $10,000km^2$, and the temporal scales considered are 15 minutes, 1 hour and 24 hours. We focus mainly on the hourly time scale and the $64km^2$ spatial scale. Therefore, each time scale is considered at the $64km^2$ spatial scale and each spatial scale at the hourly time scale. The actual combinations of spatial and temporal scales considered are:

- Hourly & $64km^2$
- Hourly & $4km^2$
- Hourly & $16km^2$
- Hourly & $256km^2$
- Hourly & $10,000km^2$
- 15 mins & $64km^2$
- Daily & $64km^2$

For both the simulated and the empirical data the rainfall intensity values are stored as 5 minute images which consist of 52×52 $4km^2$ square pixels. Therefore, to obtain measurements at larger spatial scales we simply average over the relevant number of $4km^2$ pixels. Similarly, to obtain measurements over longer temporal scales we average over the relevant number of 5 minute images. By averaging the images over both space and time we obtain resulting images at the required spatial and temporal scale which are then used to calculate the summary statistics.

Before taking averages of the original 5 minute, $4km^2$, images for the empirical data, all of these images with a coverage of less than 25% have been set to zero (i.e. all pixels within the image have been set to a rainfall intensity of zero). As explained earlier (section 3.1), such images are badly contaminated by radar ‘noise’ making their use in assessing the performance of the model

problematic. This procedure was also carried out on the original 5 minute, 4km², simulated images to allow for a more informed comparison.

For the remainder of this section, when we refer to pixels and images we are referring to those which have already been aggregated to the appropriate space and time resolution.

All of the summary statistics are given on a monthly basis. As there are 3 years and 7 months (September 1993 – March 1997) of empirical data, we obtain 3 or 4 estimates of the summary statistics for each month of the year. A weighted average of these 3 or 4 estimates is calculated to give the final summary statistics. Generally, the weights are determined by the number of images in each month (to allow for the missing data). For the simulations, we have 4 complete years of data, so the monthly summary statistics given are each averages over 4 values.

The statistics calculated fall into two main categories: unconditional statistics and conditional statistics. The unconditional statistics are calculated over all pixels, whereas the conditional statistics are only calculated over pixels which are ‘wet’ (i.e. pixels where the mean rainfall intensity is nonzero). While the unconditional statistics provide us with an overall picture, the conditional statistics allow us to focus specifically on rain events. These statistics provide us with valuable and complementary information. The various statistics calculated are as follows:

Unconditional Statistics:

- (1) Overall Mean - This is simply the mean rainfall intensity, in millimetres per hour, taken over all pixels through both space and time.
- (2) Overall Standard Deviation - This is the standard deviation of rainfall intensity taken over all pixels through both space and time.
- (3) Within Image Standard Deviation - For each image the spatial within image variance of rainfall intensity is calculated over all pixels. Having calculated this for all images, the square root of the average of these within image variances is taken, to give an average within image standard deviation.
- (4) Overall Proportion Wet - This statistic is simply the proportion of pixels through both space and time that are wet.
- (5) Spatial Autocorrelation Function (SACF) - Each aggregated image is made up of a square grid of square pixels. Any given pixel p within the image will generally have 8 neighbouring pixels, as shown in the diagram below.

(-1,1)	(0,1)	(1,1)
(-1,0)	p	(1,0)
(-1,-1)	(0,-1)	(1,-1)

The labels for the neighbouring pixels have been obtained by defining the pixel p to be at the origin in a Cartesian (x, y) co-ordinate grid. From this diagram we can see that there are eight SACF lags to calculate if we just consider neighbouring pixels. To calculate the image SACF at lag $(1,1)$, for example, we use all pairs of pixels within the

image having the same relative orientation as pixels p and $(1,1)$. A SACF at lag $(1,1)$ is calculated for each image which contains some wet pixels and an average of these values is taken. The SACF for all other lags is calculated similarly. It is obvious from the above description that the SACF at lag $(1,1)$ is identical to that at lag $(-1,-1)$. Similarly, the SACF at lag $(-1,1)$ equals that at lag $(1,-1)$, the SACF at lag $(1,0)$ equals that at lag $(-1,0)$ and the SACF at lag $(0,1)$ equals that at lag $(0,-1)$. Therefore, we need only calculate the SACF at 4 of the 8 lags: $(1,0)$, $(0,1)$, $(1,1)$, and $(-1,1)$.

- (6) Temporal Autocorrelation Function (TACF) - The TACF is calculated at time lags 1, 2 and 3. To calculate the TACF at lag 1 we proceed as follows. A particular pixel location is chosen in space. Fixing on this spatial location we then step through time identifying all pairs of values which are separated by one time lag. This procedure is carried out for each spatial pixel location. From these pairs of values we can then obtain the TACF at lag 1. The TACF at lags 2 and 3 are calculated similarly.

Conditional Statistics:

- (1) Overall Mean - This is the mean rainfall intensity of all wet pixels taken through both space and time.
- (2) Overall Standard Deviation - This is the standard deviation of rainfall intensity for all wet pixels taken through both space and time.
- (3) Within Image Standard Deviation - For all images which contain two or more wet pixels, the spatial within image variance of rainfall intensity is calculated over the wet pixels only. Once this within image variance has been calculated for all partially wet images, the square root of the average of these variances is taken to give an average within image standard deviation.
- (4) Proportion Wet Conditional On Wet Images - This statistic is slightly different from the other conditional statistics in the sense that it is not calculated over wet pixels only. Instead it is conditional on partially wet images. For all images which have one or more wet pixels, the proportion of pixels which are wet within the image is calculated. Once this has been calculated for all partially wet images an average of these proportions is taken.
- (5) Spatial Autocorrelation Function - This statistic is calculated in a similar fashion to its unconditional counterpart; however, when calculating each image SACF we only compare pairs of values if they are both non-zero.
- (6) Temporal Autocorrelation Function - This statistic is calculated in a similar fashion to its unconditional counterpart; however, in this instance we only compare pairs of values for a particular pixel location if they are both non-zero.

When calculating the spatial autocorrelation functions above it should be noted that due to the geometry of an image there are fewer pairwise pixel comparisons for the diagonal lags than for the other lags. To compensate for this we have used the number of pairwise pixel comparisons as the divisor in the calculation of the autocovariance function, as opposed to the more traditional divisor which is the total number of pixels in the image. A similar procedure was also adopted for the calculation of the temporal autocorrelation functions where there are more pairwise pixel comparisons at lag 1 than there are at each of lags 2 and 3. This is especially important for the

empirical data where there are periods of missing images. The advantage of our procedure is that we are able to make a more informed comparison between the ACF at various lags; however, we no longer guarantee that the ACF will be restricted to the range $-1 \leq ACF \leq +1$. Some discussion regarding the choice of estimator for an autocorrelation function may be found in Priestley (1981, section 5.3.3) and in Kendall & Ord (1990, chapter 6).

3.3.2 Results

Four simulations, each of 4 years in length, were produced. Each of the 4 simulations was kept separate when calculating the summary statistics. Therefore, for a given spatial and temporal combination, we have 4 sets of summary statistics for the simulated data and 1 set of summary statistics for the empirical data. To simplify the comparison between the simulated and the empirical data, an average set of simulation summary statistics was produced over the four simulations.

Within this section we focus attention primarily upon the results for the hourly and 64km^2 scale which are represented graphically in figures 3.17 to 3.19. These results will be discussed, and comparisons made with the results obtained at the other spatial and temporal scales. Tables and plots giving the results at all spatial and temporal scales can be found in Appendix A.5.

From the plot of the unconditional overall means we can see that the simulated data are able to capture the seasonal effect which is apparent in the empirical data, with more rain being produced in the winter months. For the majority of the months, the means for the average simulation are close to those for the empirical data, although there are a few months where the fitted model overestimates the mean rainfall intensity. If we consider the conditional means we see that for some months, such as June and November, these are almost identical for the average simulation and the empirical data. However, for other months (e.g. August), the fitted model again overestimates the observed value. We note that the model of storm interiors matches the means exactly. It is therefore likely that observed discrepancies are an artefact of the event sampling procedure.

The plots for the unconditional and conditional overall standard deviations are generally quite encouraging. Notice the similarity in the plots for the unconditional overall standard deviation and overall mean. The plots for the unconditional and conditional within image standard deviations provide a very similar picture to that of the overall standard deviations.

For the overall proportion wet, the values for the average simulation and the empirical data are in close agreement. It does appear though that for the majority of the months the overall proportion wet is slightly greater in the empirical data. This suggests that it rains slightly less frequently in the simulations but, when it does rain, the intensity is greater. The results for the conditional proportion wet statistic also look generally good. At this stage it is worth comparing the plot for this statistic at the hourly and 64km^2 scale against the same plot at the hourly and 4km^2 scale (results for this scale are presented graphically in figures 3.20 to 3.22). At the 4km^2 scale the average simulation line is generally above that for the empirical data, whereas at the 64km^2 scale both lines have shifted upwards but now the empirical data line is generally above that for the average simulation. This indicates that the empirical data have more large scale structures than the simulations at the 64km^2 scale, but that these structures are less homogeneous than those in the simulations. Referring to the model parameters graphed in figure 2.10 (page 29), structures on this scale would correspond to large rain cells in the model. The discrepancy between observed and simulated statistics here is probably due to the fact that the model represents rain cells as solid

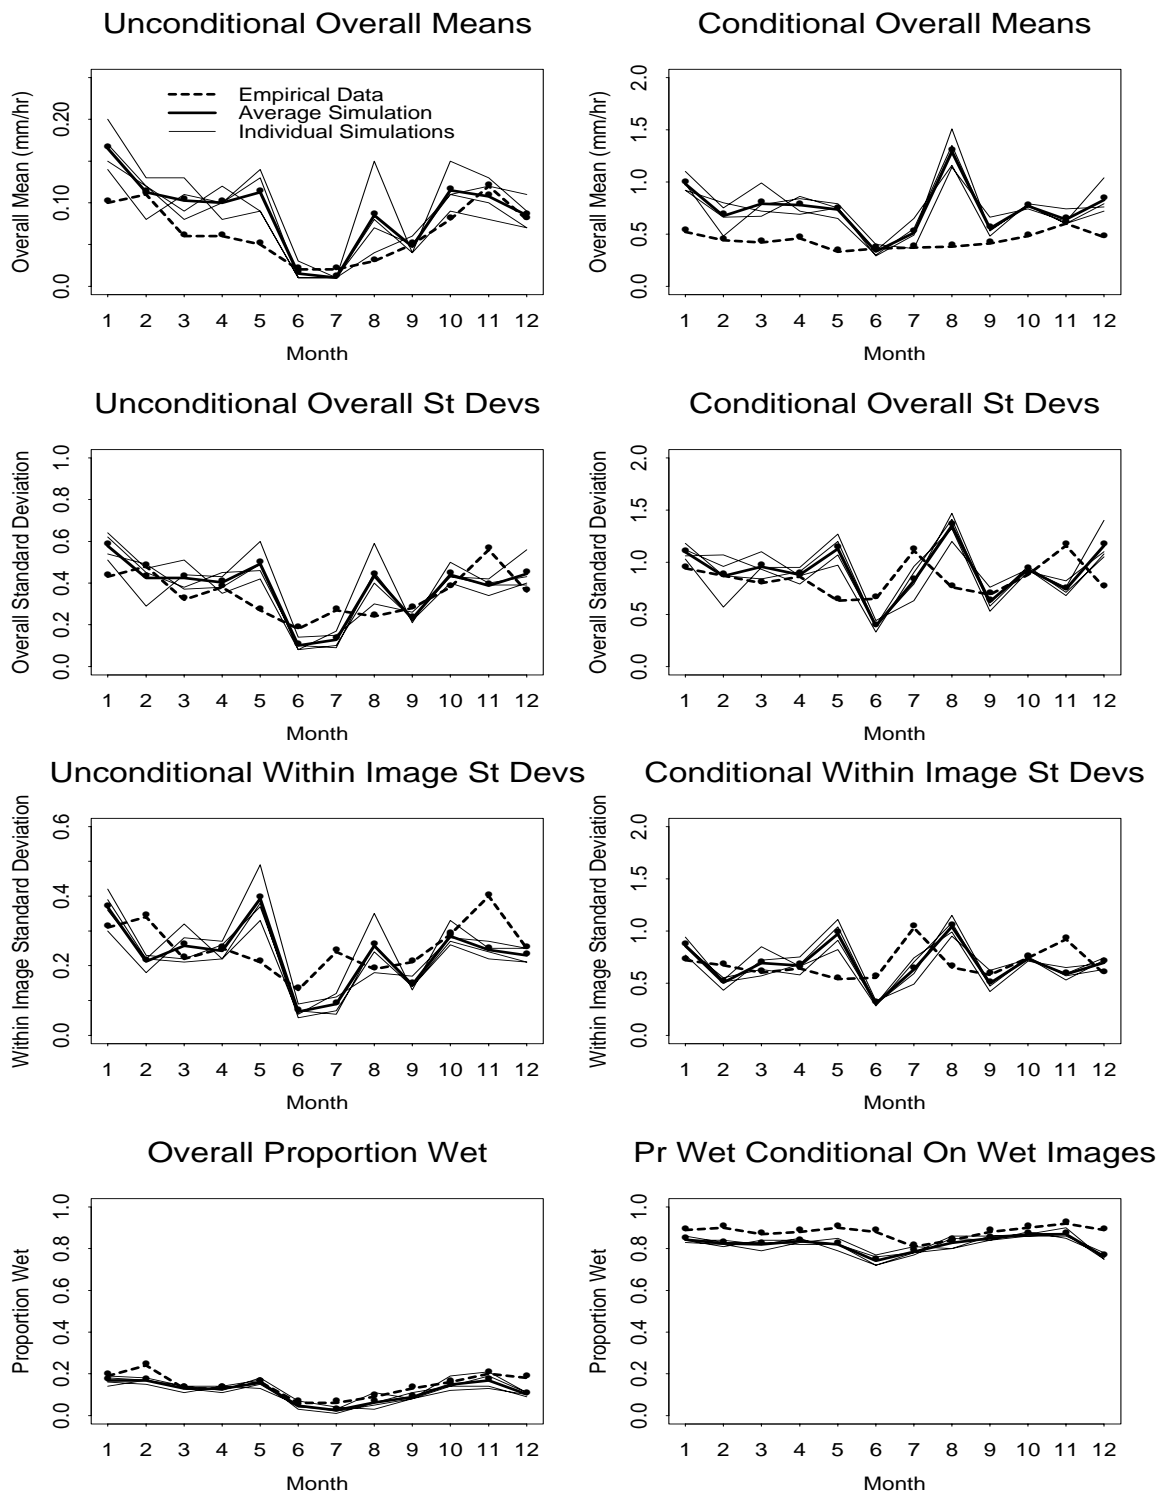


Figure 3.17: Plots of summary statistics for the simulated and empirical data (hourly & $64km^2$).

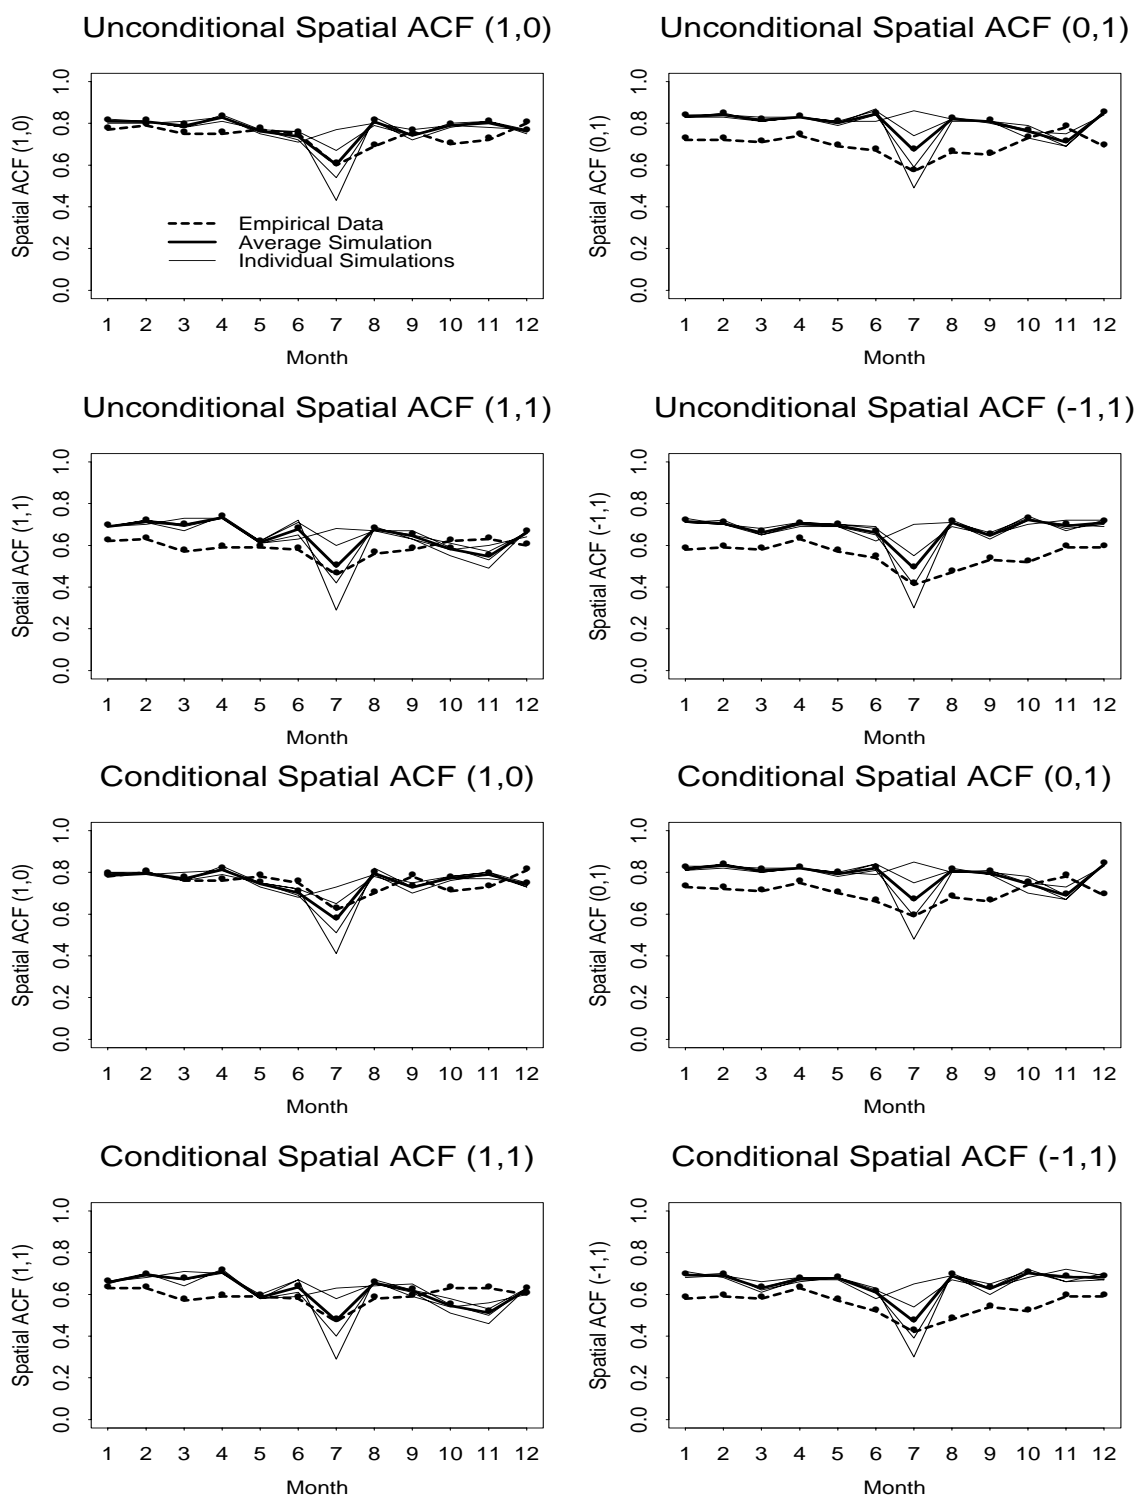


Figure 3.18: Plots of summary statistics for the simulated and empirical data (hourly & $64km^2$).

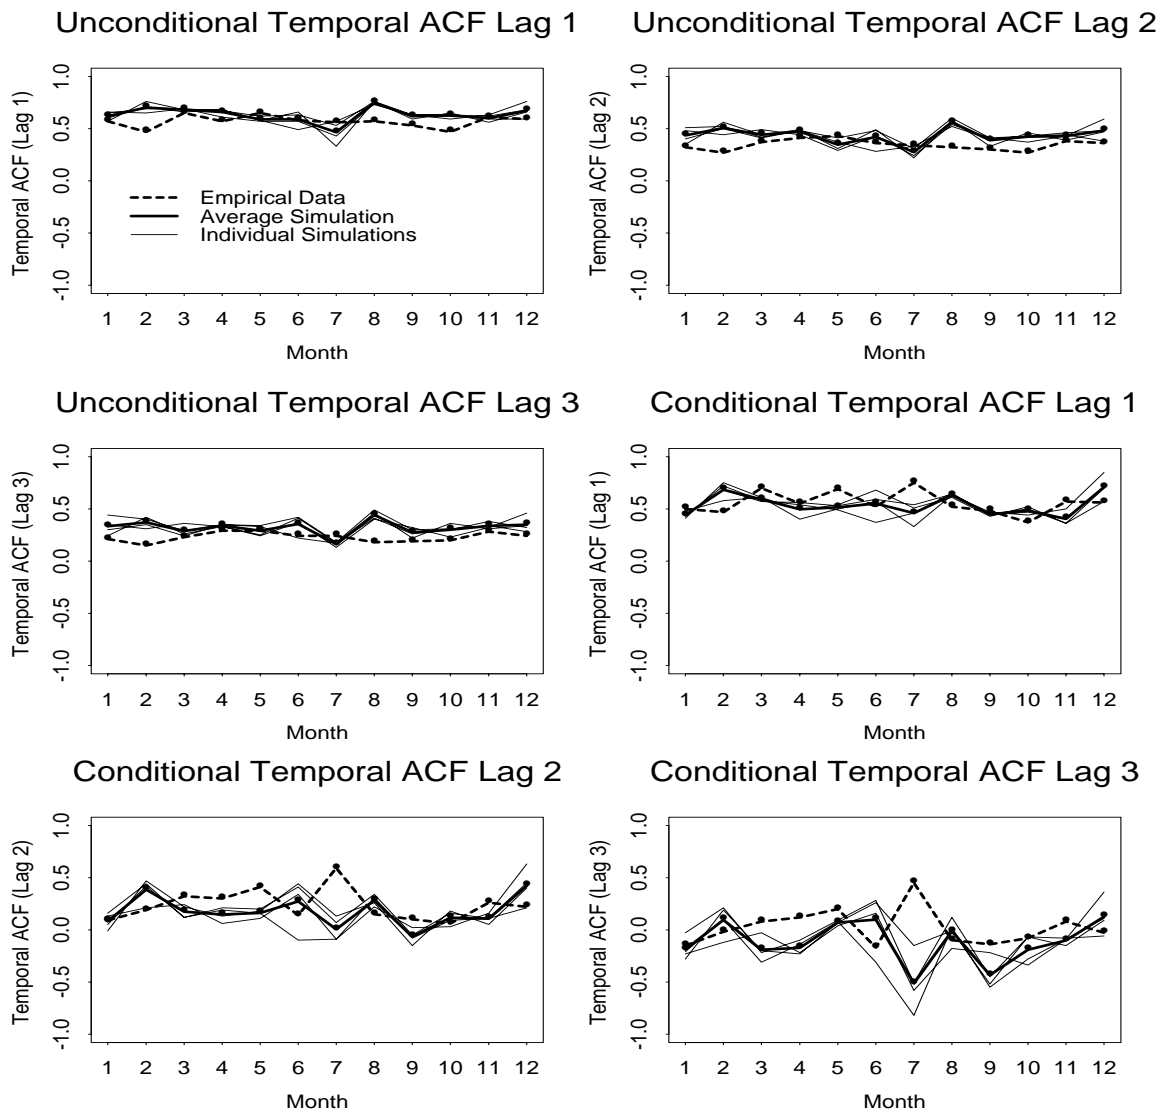


Figure 3.19: Plots of summary statistics for the simulated and empirical data (hourly & $64km^2$).

ellipses, whereas in reality cell boundaries are likely to be less well defined.

Turning now to the unconditional spatial autocorrelation function (SACF), the results are generally very good. For all 4 spatial lags, the values for the average simulation and the empirical data follow similar patterns and are in close agreement, although the values for the average simulation do tend to be slightly greater than those for the empirical data. It is worth noting that for both the average simulation and the empirical data, the SACF is greater at lags (1,0) and (0,1) than it is at lags (1,1) and (-1,1). This is to be expected given that in the SACF at lags (1,0) and (0,1) we are comparing pixels that are closer together than at the other 2 lags. Another noticeable feature of all 4 plots is that the value for July dips for both the average simulation and the empirical data, and that there is a large variation between the July values for the 4 individual simulations. This is almost certainly simply a reflection of the scarcity of data for type 1 events in July (just 5 were observed in the empirical data — see table 2.4). At this stage it is worth comparing the 4 plots for the unconditional SACF at the hourly and 64km² scale (figure 3.18) with the equivalent plots at the hourly and 4km² scale (figure 3.21). We can see that the SACF is lower at the larger spatial scale, as we would expect, for both the average simulation and the empirical data. Furthermore, the difference between the average simulation and the empirical data seems to increase at the larger spatial scale. Turning to the conditional SACF we see that the results are very similar to those for the unconditional values.

These discrepancies in the spatial autocorrelations may be explained as follows. The conditional SACF focuses on pairs of wet pixels, which are likely to be part of the same rain cell. In the model, such pixel values are perfectly correlated, because each cell has a constant intensity, whereas the empirical correlations of such points will be lower. It is noteworthy that the discrepancy of the spatial ACF plots in an east-west direction (spatial lag (1,0)) is minor compared with that in other directions, which is consistent with the predominant orientation of the rain cells (see figure 2.9).

The unconditional temporal autocorrelation function (TACF) plots show good agreement between simulated and empirical data. For all three lags, the two sets of values are in close agreement. The TACF decreases with increasing lag, as expected, for both the simulated and empirical data. The equivalent plots and tables at the daily time scale (see Appendix A.5) show that the average simulation and the empirical data are still in close agreement but that there is practically no correlation evident at the daily time scale. If we now turn to the conditional TACF plots, we see that again the results are in quite good agreement, although the agreement decreases with increasing lag. There is also less correlation evident in the conditional plots than there is in the unconditional plots which reflects the fact that the latter is inflated by pairs of successive zero values. Finally, notice that the largest difference in the two sets of results occurs in July at lag 3, again not surprising given the small number of type 1 events present in the empirical data.

Overall, if we now compare the set of results obtained at the hourly and 64km² scale with the other sets of results (given in the Appendix A.5) we can observe the effect of spatial and temporal scale upon the various summary statistics. Tables 3.5 and 3.6 summarise these effects in broad terms.

The changes indicated in the tables are expected on intuitive grounds and are generally consistent in both the simulated and the empirical data. This suggests that the spatial-temporal model performs consistently and acceptably over the various spatial and temporal scales.

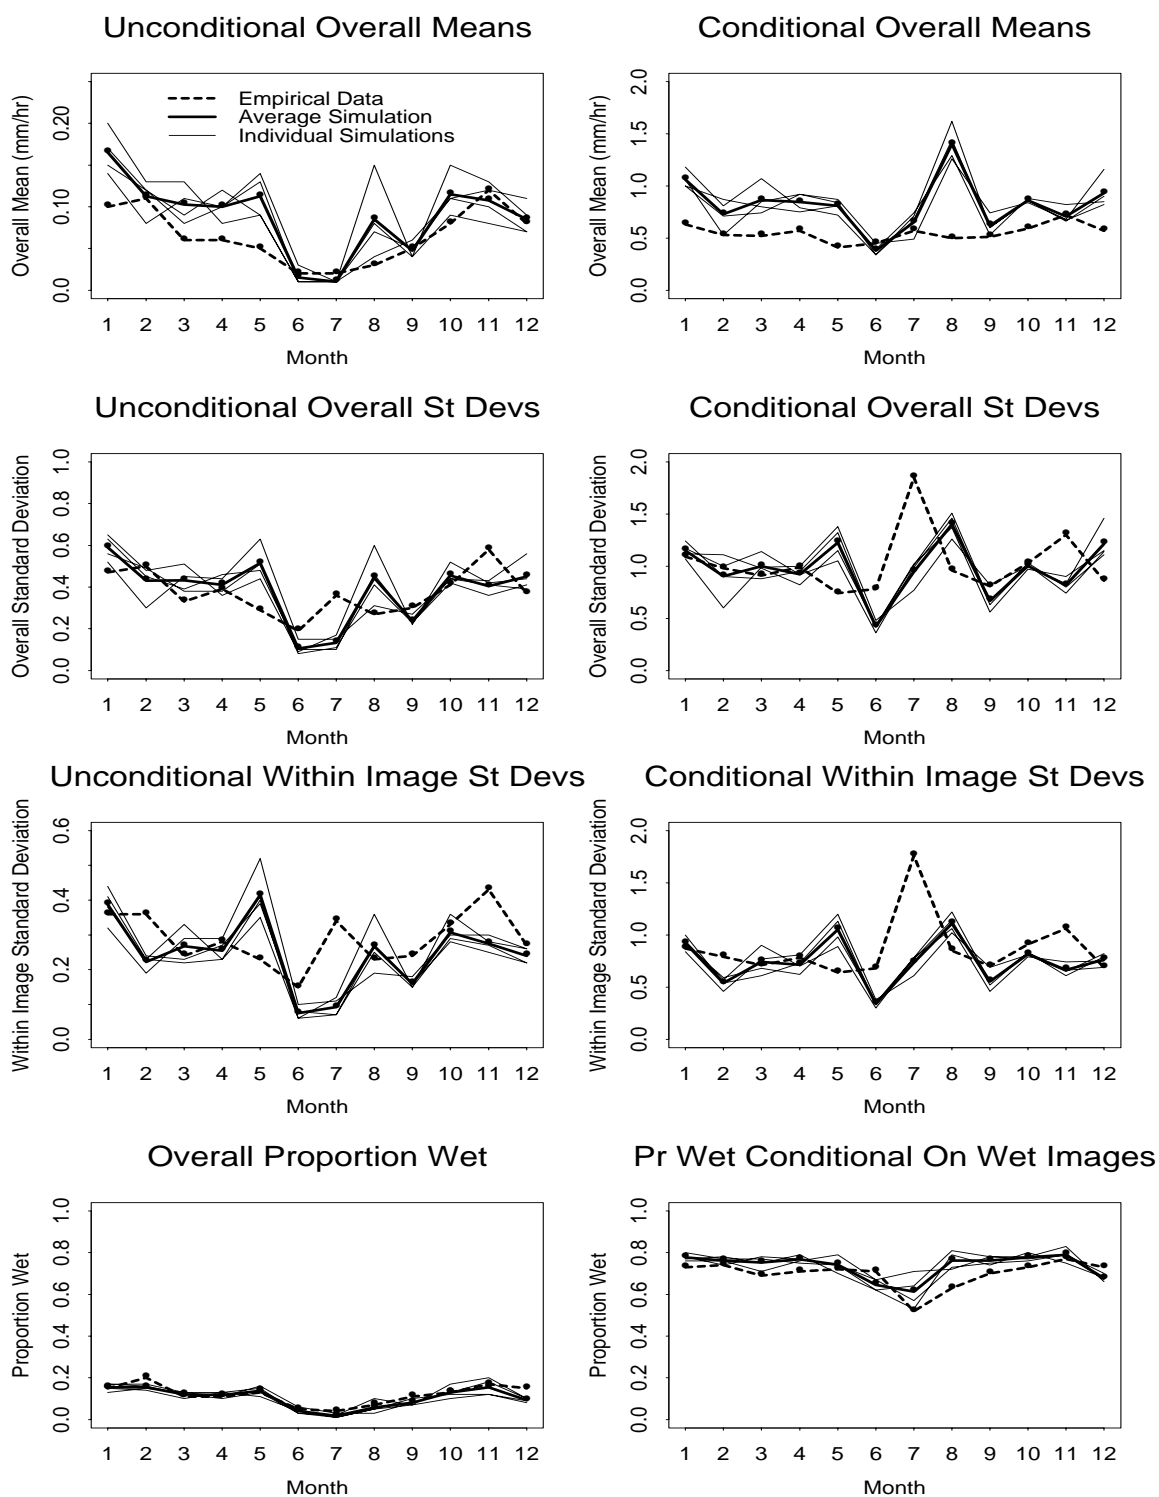


Figure 3.20: Plots of summary statistics for the simulated and empirical data (hourly & $4km^2$).

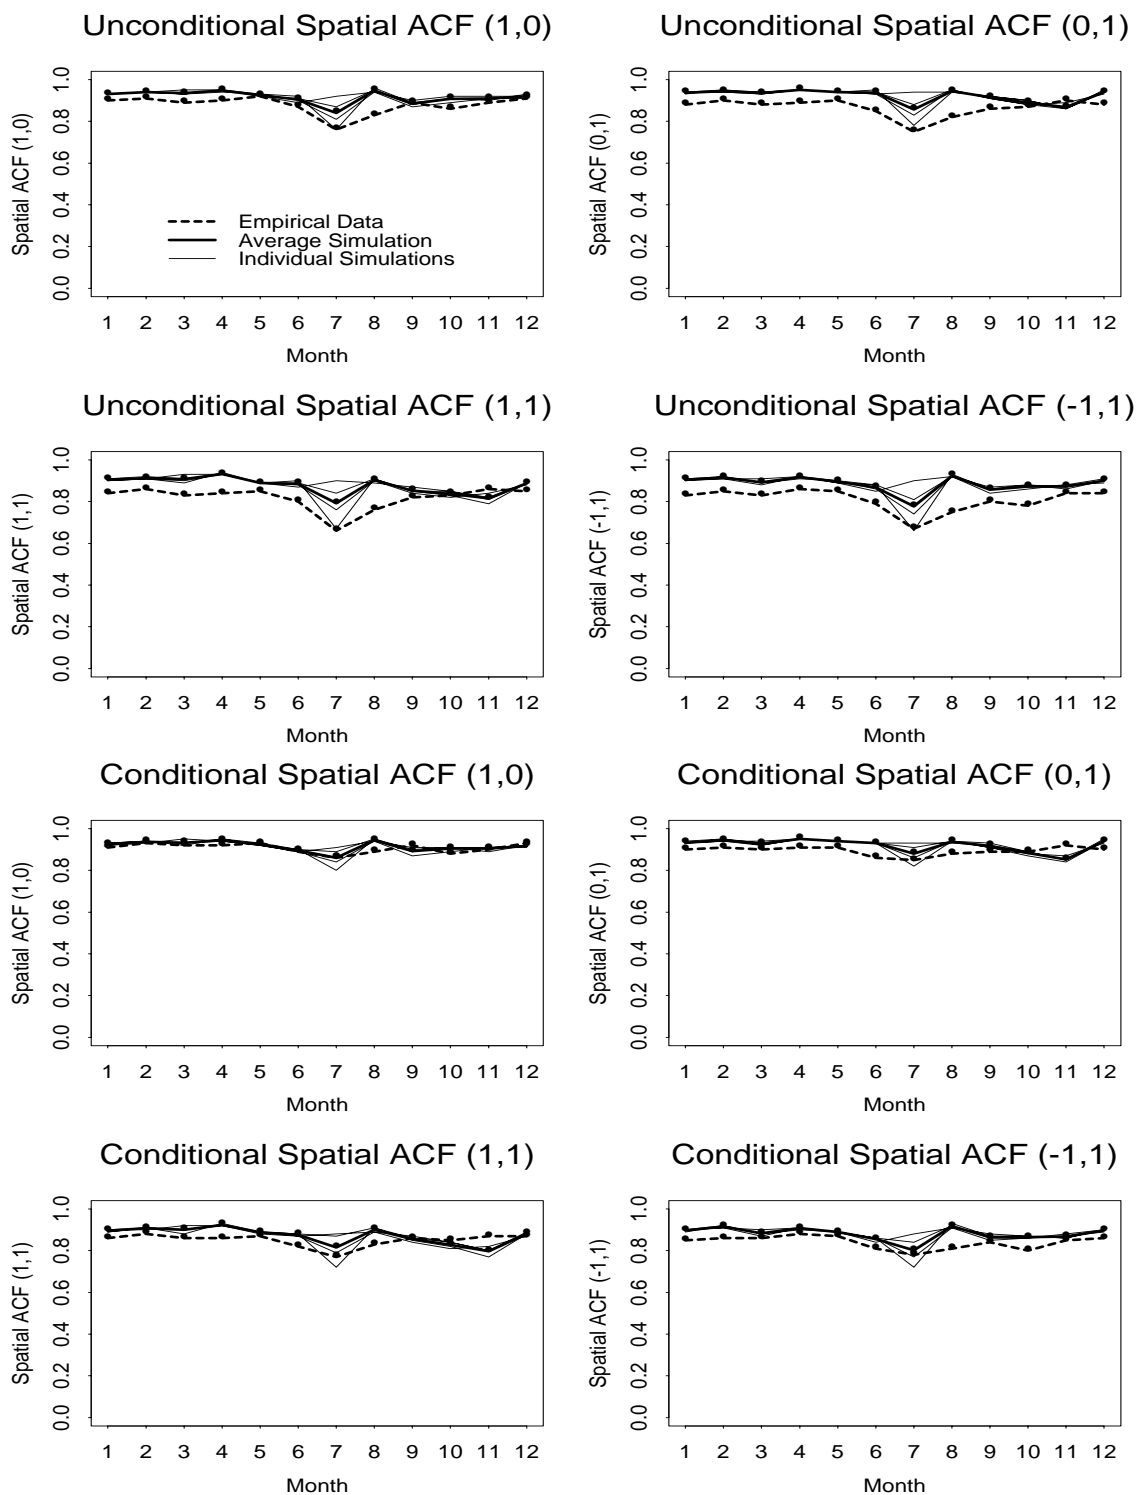


Figure 3.21: Plots of summary statistics for the simulated and empirical data (hourly & $4km^2$).

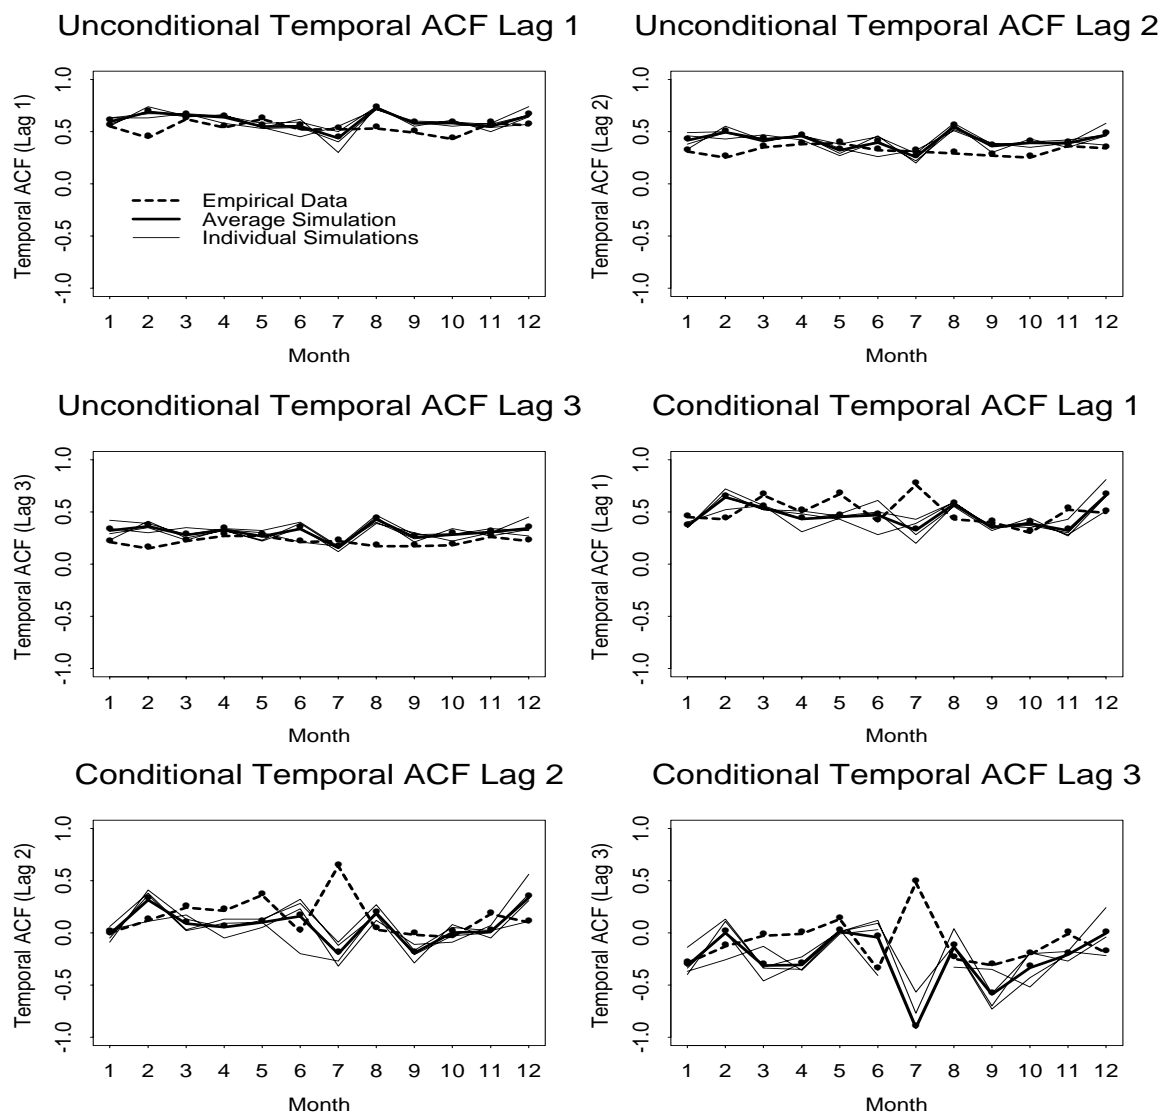


Figure 3.22: Plots of summary statistics for the simulated and empirical data (hourly & $4km^2$).

Change	Effect
Increase Spatial Scale	Unconditional Mean - Unchanged Conditional Mean - Decrease Unconditional and Conditional St Devs - Decrease Unconditional and Conditional Prop Wet - Increase Unconditional and Conditional SACF - Decrease Unconditional and Conditional TACF - Increase

Table 3.5: Effect of increasing spatial scale upon summary statistics for rainfall intensity

Change	Effect
Increase Temporal Scale	Unconditional Mean - Unchanged Conditional Mean - Decrease Unconditional and Conditional St Devs - Decrease Unconditional and Conditional Prop Wet - Increase Unconditional and Conditional SACF - Increase Unconditional and Conditional TACF - Decrease

Table 3.6: Effect of increasing temporal scale upon summary statistics for rainfall intensity

3.3.3 Extreme-value performance

The evaluation of the model's ability to reproduce the extreme-value behaviour of rainfall should be carried out for different spatial and temporal scales. We shall consider the usual hourly and daily scales while areas of $2 \times 2 \text{ km}^2$ and $16 \times 16 \text{ km}^2$ will be examined to see if the model is able to reproduce annual extremes at different spatial scales. (It might, for example, be expected that convective rainfall has a greater effect on the smaller-scale extremes.)

There are less than 4 years of radar data with which model simulations can be compared, so that it is necessary to pool data from different locations over the analysed region to be able to carry out a meaningful extreme-value analysis. Although dependent data sets are being used in fitting the extreme-value distributions, the same procedures are used for observed and simulated data, and it was felt that further refinement was not justified in the light of the data limitations. In selecting locations for this pooling, we have chosen regions within the observed and simulated spatial fields from which raingauge data are also available, enabling additional comparisons to be made. Overall, data from five regions are considered for the pooling, corresponding to raingauges at Godminster Farm, Yeovilton, Downside Abbey, Sidmouth and Boscombe Down (see figure 3.23). To compare a 4-year realisation of the continuous-simulation model with the radar data therefore, both empirical and simulated data are aggregated to the appropriate temporal and spatial scale for the analysis, and simulated and observed annual extremes are pooled across all five regions to obtain samples of size 20 for extreme-value analyses. Each of the four simulations considered in the previous has been studied in this way.

For each of the data sets obtained by pooling over regions, an extreme-value analysis is per-

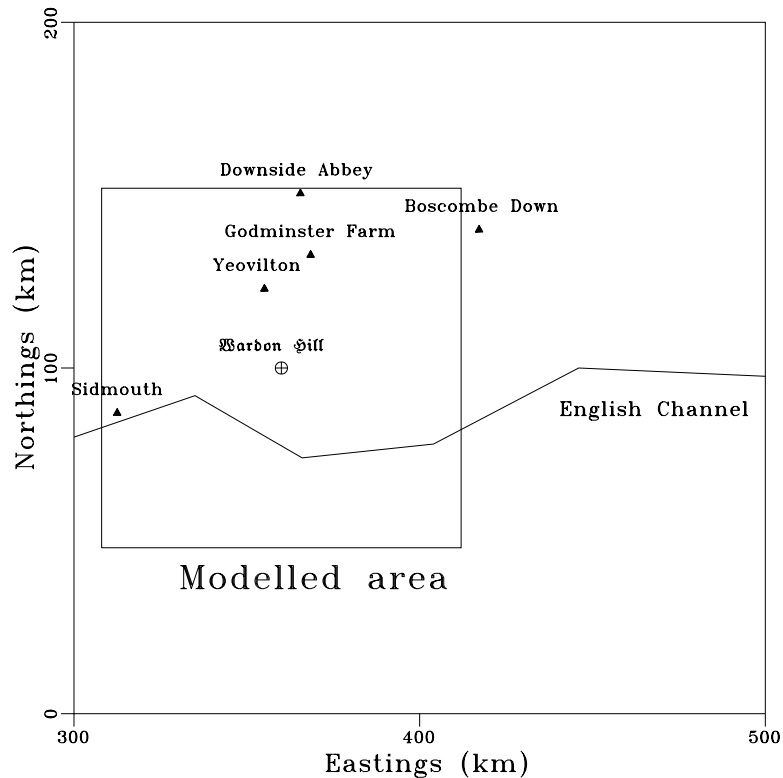


Figure 3.23: Sites at which the GDSTM extreme value performance is evaluated

formed on the annual maxima at hourly and daily timescales. Using extreme-value theory, we expect the distribution of these maxima to converge towards one of three extreme-value distributions, EV-I (Gumbel), EV-II (Fréchet) or EV-III. In practice, since EV-III has an upper bound, it is not useful for rainfall maxima. The EV-II distribution is more skewed than EV-I which has a fixed skewness coefficient of 1.1396. It is sometimes preferred to the EV-I for rainfall, but on the basis of the available data (where only 20 annual maxima are available), it is not possible to decide between an EV-I and an EV-II at a 95% confidence level. Moreover, if a Generalised Extreme Value distribution (which incorporates the EV-I, EV-II and EV-III as special cases — see, for example, Joe, 1997, p.170) is fitted to the data here, the shape parameter is not such as to warrant the use of an EV-II — see Kotegoda (1980), for example. We shall therefore focus upon fitting an EV-I distribution to all the data sets. This is done using the method of moments, and plots are produced showing the extreme rainfalls for different return periods, represented on the horizontal axis using the reduced variate:

$$x = -\ln \left[-\ln \left(1 - \frac{1}{T} \right) \right]$$

where T is the return period.

In addition to carrying out conventional extreme-value analyses, Flood Studies Report extrema (NERC 1975) are estimated. In this method, the expected rainfall $R(D, T, A)$ for any duration D and return period T and area A is estimated on the basis of the 2 day rainfall with a 5 year return period, i.e. $R(2 \text{ day}, 5 \text{ year})$, as well as the ratio $R(1 \text{ hour}, 5 \text{ year})/R(2 \text{ day}, 5 \text{ year})$ and the Areal Reduction Factor (ARF) for given A and D .

Finally, estimates of extreme-value rainfall based upon 18 years of data from the Boscombe Down gauge (1980-1997) together with the above ARFs from the Flood Studies Report are used to provide a comparison based upon hourly rainfall which involves no regionalisation (although the Boscombe Down gauge is located slightly outside the radar window used in fitting the GDSTM — see figure 3.23).

There are two issues to bear in mind when comparing results obtained from radar data (which include the empirical radar data results along with those for the GDSTM) with those obtained using raingauge data. The first is the fact that radar-derived results reflect spatial averages whereas raingauge results are for points in space: this is corrected for by the use of appropriate ARFs, as detailed above. The second issue is that of radar/raingauge calibration. In general, radar tends to systematically underestimate rainfall amounts, the problem becoming more acute with increasing distance from the radar station. As a result, we would expect radar-derived extremes to be smaller than those derived from gauges even after applying an ARF to the gauge data.

Results for hourly data: The estimated Gumbel fits for the four simulations are found to be on either side of that for the pooled historical radar data both for the $2 \times 2 km^2$ and $16 \times 16 km^2$ resolutions (see figures 3.24 and 3.25). This indicates the model's ability to produce extreme values which are approximately distributed as the observed data.

The Flood Studies Report estimates diverge from the radar estimates for return periods which are of an order of magnitude larger than the size of the radar data set. This is not surprising given that they are estimated on the basis of long term data sets of which the four years of radar data may not be a representative sample.

Results for daily data: At this time-scale, the radar data appear much more representative of the long term raingauge behaviour as summarised in the Flood Studies Report statistics than at the hourly time-scale, as seen in figures 3.26 and 3.27.

The extreme values inferred from the Gumbel fit to the simulations are very close to those yielded by a Gumbel fit to the radar data at the resolution of $16 \times 16 km^2$. At the finer resolution of $2 \times 2 km^2$, the model underestimates the estimated radar data extreme values for large return periods: however, all of the simulations yield results which are well within the radar-derived confidence intervals. Notice also that at the finer resolution, the simulations are comparable with the results obtained from an EV-I fit to the 18-year record at Boscombe Down.

The overall conclusion here is that the GDSTM appears able to reproduce the extreme value statistics of the radar fields well, as far as can be determined from the limited data available. The different results obtained from the various observational series highlight the uncertainty associated with extreme value analyses on the basis of small data sets, and the potential discrepancies between radar and raingauge data. In particular, notice the difference between the two sets of results from Boscombe Down obtained using the Flood Studies Report methodology and by fitting an EV-I. The wide confidence intervals associated with the EV-I fit to the radar data also indicate the level of uncertainty here. Taking this into consideration, any discrepancies between the results from GDSTM simulations and from the radar data are minor.

Hourly Extreme Values




Figure 3.24: Hourly 2 × 2km² GDSTM extreme value performance

Hourly Extreme Values




Figure 3.25: Hourly 16 × 16km² GDSTM extreme value performance




Figure 3.26: Daily 2 × 2km² GDSTM extreme value performance




Figure 3.27: Daily 16 × 16km² GDSTM extreme value performance

3.4 Summary of continuous space-time modelling

In the work described in the previous chapter, a homogeneous model for the spatial and temporal interior of single rain events was developed and fitted to data from events that cover the radar window for a substantial period of time. Almost all rainfall over the window is generated by such events moving across the window from outside. In this chapter, two categories of event have been distinguished by analysis of the event sequence: in the first, the event covers the whole window for a substantial period of time (type 1 events), and in the second the event only partially covers the window, typically ‘clipping’ a corner (type 2 events). In addition, there may be light, spatially and temporally intermittent rainfall but this cannot be distinguished from noise in the data. The continuous simulation model developed here generates a random sequence of type 1 and type 2 events of appropriate durations separated by ‘dry’ intervals, according to a semi-Markov model. The durations of the events as they pass over the window, and of the dry periods in between them, are modelled by members of the Weibull family of probability distributions. In practice, type 2 events have been found to generate relatively small amounts of rainfall over the window, and a good fit to data has been achieved by setting the intensities of such events to zero.

Thus, the simulation proceeds by sampling a set of parameters for each event from a library of parameter sets obtained from historical events. These parameters are augmented by the velocity of the fitted event and the angles of its leading and trailing edges as it passes over the window. An event with these parameters is then simulated in such a way that it passes over the window for the required duration. Possibilities for assessing the performance of the continuous simulation model are restricted by the limited amount of empirical data available. Nevertheless, the model performed extremely well in reproducing basic statistical properties, including the proportion of wet pixels, over a range of time and spatial scales. Clearly, more data are needed before the extreme values can be fully assessed. However, the model reproduces very well the hourly extremes of the empirical data. While the daily extremes are underestimated for high return periods, nevertheless these are well within estimated 95% confidence intervals. We note that the assumption of spatial stationarity implicit in the model does not appear to have posed significant problems for this data set.

Chapter 4

Generalized Linear Modelling of Daily Data

The stochastic spatial-temporal models described in the previous chapters have great potential for the simulation of long rainfall sequences, so as to reproduce both spatial and temporal dependencies using a simplified, but realistic, representation of the structure of rainfall processes. Their main drawback is that they are stochastically stationary in both space and time. Seasonal variability in rainfall has been captured via the use of different parameter sets for each month of the year; however, at the present stage of model development, there is no scope for the incorporation of long-term changes in climate (temporal nonstationarity) or of systematic orographic effects (spatial nonstationarity).

The need for radar data to calibrate these models imposes a further restriction on their use. Even in regions where radar data are available, record lengths are limited as the technology is relatively new. In view of concerns that the climate of the UK has been unusual during the 1990s (see, for example, DOE (1996), DETR (1999)), it may be unwise to rely too heavily upon short runs of recent data to provide reliable model parameterisations which will hold into the future. This problem can, to some extent, be resolved by using long subdaily raingauge records (if available) to parameterise the model: ways in which this can be done are presented in section 5.1 of this report. However, long records of sub-daily rainfall may themselves be relatively sparse. A possible solution would be to parameterise the models using a combination of long single-site records and short spatial data sequences; however, further development is required to investigate this.

To address these areas of weakness in the Poisson-based model simulation strategy, a parallel line of research has investigated ways in which long records of daily rainfall data (which are relatively abundant within the UK) can be used to quantify spatial and temporal nonstationarity over an area. The methodology is that of Generalised Linear Models (McCullagh & Nelder 1989), which were introduced into the study of rainfall patterns by Coe & Stern (1982); Stern & Coe (1984). Such models for daily data may be of interest in their own right, or may be used in conjunction with the Poisson cluster models of the previous chapters. The simplest scenario occurs when a daily rainfall modelling exercise reveals no significant temporal or spatial nonstationarity: in this case, the Poisson cluster models can be used with confidence if suitable data are available. For example, over small areas with little topographic variability, it is likely that any spatial nonstationarity is minor and can effectively be ignored. However, for more general application, some way of incorporating

nonstationarity into Poisson-type models will become necessary. Some preliminary work in this area is reported below, in chapter 5.

We now give an overview of the Generalised Linear Modelling strategy. Subsequently, in section 4.2, three case studies are presented which illustrate the methodology.

4.1 Model development

The basic idea behind Generalised Linear Models (GLMs) is to predict a probability distribution for some quantity of interest using observations of various other related quantities. In our case the quantity of interest is the daily rainfall amount at a site; possible predictors include previous days' rainfall amounts, the month of the year and variables representing topographic effects.

Formally, a GLM for a $n \times 1$ vector of random variables $\mathbf{Y} = (Y_1, \dots, Y_n)'$, each dependent on p predictors (the values of which can be assembled into a $n \times p$ matrix \mathbf{X} whose (i, j) th element is the value of the j th predictor for Y_i), takes the form

$$g(\boldsymbol{\mu}) = \mathbf{X}\boldsymbol{\beta} \quad (4.1)$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)'$ is the vector mean of \mathbf{Y} , $g(\cdot)$ is a monotonic function (the *link function*) and $\boldsymbol{\beta}$ is a $p \times 1$ vector of coefficients (by $g(\boldsymbol{\mu})$ we mean the $n \times 1$ vector, $\boldsymbol{\eta}$ say, whose i th element is given by $g(\mu_i)$). Model (4.1) is a natural extension of the simple linear regression model. A constant term in the model can be defined by including a column of 1s in the matrix \mathbf{X} .

When, as here, the Y s arise as one or more time series and we wish to include previous values of the series as predictors, we are implicitly studying the conditional distributions of each Y given the past, and the usual GLM methodology carries over straightforwardly — see, for example, Fahrmeir & Tutz (1994), chapter 6. In implementation, we broadly follow Coe & Stern (1982); Stern & Coe (1984). They adopted a two-stage approach, as follows:

Stage 1: model the pattern of wet and dry days at a site using logistic regression. If we denote by p_i the probability of rain for the i th case in the dataset, conditional on a predictor vector \mathbf{x}_i , then the logistic regression model is given by

$$\ln \left(\frac{p_i}{1 - p_i} \right) = \mathbf{x}_i \boldsymbol{\beta} . \quad (4.2)$$

Stage 2: fit gamma distributions to the amount of rain on wet days. The rainfall amount for the i th wet day in the database is taken, conditional on a predictor vector $\boldsymbol{\xi}_i$, to have a gamma distribution with mean μ_i where

$$\ln \mu_i = \boldsymbol{\xi}_i \boldsymbol{\gamma} \quad (4.3)$$

for some coefficient vector $\boldsymbol{\gamma}$.

These two models will be referred to as ‘occurrence’ and ‘amounts’ models respectively.

4.1.1 Interactions

A common feature of the rainfall process, which can be incorporated into the GLM framework, is that predictors interact with each other. As an example of this, consider the temporal dependence structure within a rainfall time series. Intuitively, in the UK we would expect dependence between successive days' rainfall amounts to be weaker in the summer than in the winter. This is because summer rainfall tends to be dominated by short-lived convective events, rather than the longer-lasting frontal systems which predominate in winter. Any coefficients in a GLM, which correspond to predictors representing previous days' rainfall in (4.2) and (4.3), should therefore probably vary with the season. A simple strategy for incorporating this effect is to represent these coefficients themselves as linear combinations of other predictors. From the point of view of model specification, this is equivalent to adding an extra predictor to the model whose value is the product of the interacting predictors: hence interactions can be incorporated straightforwardly within the overall framework.

The presence or absence of interactions within a GLM can tell us a great deal about the mechanisms which are driving the rainfall process. For example, if significant interactions are found between a long-term trend and predictors representing seasonality, one of the effects of the trend may be to induce wetter winters and drier summers.

4.1.2 Model fitting

Fitting a GLM involves choosing an appropriate set of predictors (\mathbf{x} in (4.2) and $\boldsymbol{\xi}$ in (4.3)) and estimating the corresponding parameter vectors $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. In each case, if the responses are conditionally independent given the predictors, maximum likelihood estimates of parameters can be obtained using iterative weighted least squares (McCullagh & Nelder 1989), and standard techniques such as likelihood ratio tests (see, for example, Cox and Hinkley, 1974, section 9.3) can be used to assess the significance of individual predictors. For example, if a single extra predictor is added to a model and results in a log-likelihood increase of more than 1.92 (3.22), it is formally considered to be significant at the 95% (99%) level. However, in general it is necessary to fit models simultaneously to data from several sites, and simultaneous responses at different sites are not conditionally independent given the predictors because of spatial dependence between sites. Chandler & Wheeler (1998*a*) reviewed available methods for dealing with such spatial dependence when fitting models: they argued that $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ may still be estimated as though sites were conditionally independent, so long as individual sites have long records. Such an approach yields consistent parameter estimates (so that $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ will be well estimated); however, standard methods for assessing the uncertainty of such parameter estimates (e.g. confidence intervals and likelihood ratio tests) will tend to under-represent the true uncertainty.

Notwithstanding these objections, the 'likelihood'-based approach to parameter estimation, ignoring spatial dependence, does provide useful comparisons, and an informal appraisal of nominal standard errors and likelihood ratios can provide a useful guide to model selection. This is illustrated by the case studies considered in section 4.2.

4.1.3 Modelling nonlinear dependencies

In rainfall modelling applications, it is common to encounter situations where the response (rainfall occurrence or amount) is associated with a particular predictor, but where the relationship is best thought of as between the response and some nonlinear transformation of the predictor. Examples include the investigation of possible long-term cycles in the climate of an area (where the fundamental predictor for any day's rainfall is the year in which it occurs, but a cyclical pattern implies that the relationship is really with a sine wave derived from the year), and the realistic modelling of orographic variability (typically, the underlying predictors might be site eastings and northings, but any variability in rainfall amounts is unlikely to be well represented by putting these into equations (4.2) and (4.3) directly).

Such nonlinear transformations may be divided into two categories, as follows:

Category 1: in this case, there is an obvious parametric form for the transformation. The example of fitting a cyclical trend function falls into this category. In such a case the component of \mathbf{x} or $\boldsymbol{\xi}$, to be included in the model (4.2) or (4.3), takes the form

$$f(t, \boldsymbol{\theta}) \quad (4.4)$$

for some known function $f(\cdot)$, where t is the value of the underlying predictor and $\boldsymbol{\theta}$ is a vector of parameters in the transformation. Usually $\boldsymbol{\theta}$ is unknown and must be estimated: this can be done simultaneously with all the other parameters, using an extension of the usual iterative weighted least squares algorithm as described by Green (1984). The resulting estimates would be maximum likelihood estimates if all sites were conditionally independent. Direct application of the algorithm may not always produce convergence to the maximum likelihood estimate owing to computational instability: however, a stable algorithm can be obtained by making some small modifications, as described in section 2.3 of Wei (1997). This modified algorithm is the one which has been used here, and which is implemented in the software described in Appendix B.

Category 2: in this case, there is no obvious way in which a nonlinear transformation may be parametrised. This would be the case, for example, if we were interested in trying to model regional variability over a large area. In cases such as this, our approach has been to assume that the transformation of interest is a smooth one (i.e. continuous and differentiable) and is square integrable over the range of values of the underlying predictor. It is difficult to imagine any realistic practical applications where these assumptions do not hold. In this case the transformation can be expressed, to any desired degree of accuracy over the range of values of the underlying predictor, using an orthogonal series representation — see, for example, Priestley (1981, section 4.2.2). Specifically, suppose we wish to represent the unknown transformation $f(\cdot)$ over the interval (a, b) . Let $\{\psi_j : j = 0, 1, 2, \dots\}$ be a collection of functions satisfying

$$\int_a^b \psi_j(t)\psi_k(t)dt = \begin{cases} 1 & j = k \\ 0 & \text{otherwise.} \end{cases}$$

Then, for $t \in (a, b)$, $f(t)$ can be expressed as the infinite sum

$$f(t) = \sum_{j=0}^{\infty} A_j \psi_j(t) . \quad (4.5)$$

for some set of coefficients $\{A_j : j = 0, 1, 2, \dots\}$.

In practice, providing the ψ s are chosen intelligently, most of the coefficients A_j will be very small and can be neglected, so that $f(\cdot)$ can be represented to a very good degree of approximation using a small finite collection of ψ s. The point is that equation (4.5) reduces a highly nonlinear dependency into linear dependence upon a set of known functions: if we put a subset of the ψ s directly into the GLMs (4.2) and (4.3) as predictors, the coefficients $\{A_j\}$ will be subsumed into the coefficient vectors β and γ , and the problem is reduced into essentially linear form.

Orthogonality of the basis functions (ψ s) is not required for this approach to work. However, if they are orthogonal and the data points are scattered approximately uniformly over the range (a, b) , then they will produce predictors which are approximately uncorrelated. As a consequence, the model will be robust against mis-specification of any of the individual ψ s (see, for example, Chandler (1998b)).

The disadvantage of orthogonal series representation is that it may be quite parameter-intensive, as several ψ s may be needed to obtain an adequate representation of an effect. This problem can be minimised by careful selection of basis functions. For example, if a transformation is likely to be essentially monotonic, it might be represented efficiently using a polynomial basis such as Legendre polynomials (Abramowitz & Stegun 1965). Effects which are more oscillatory may be represented more parsimoniously using Fourier series.

The main use of orthogonal series in this work has been to represent regional variability as a bivariate function of site eastings and northings. A straightforward extension of the usual orthogonal series arguments shows that, if $\{\psi_j : j = 0, 1, 2, \dots\}$ forms an orthogonal basis for eastings effects and $\{\phi_k : k = 0, 1, 2, \dots\}$ forms a basis for northings effects, then the collection $\{\psi_j \phi_k : j, k = 0, 1, 2, \dots\}$ forms a basis for regional effects. Furthermore, this basis is orthogonal. But within the GLM framework, this collection consists simply of interactions between the ψ s and ϕ s (see section 4.1.1 above), and so representation of regional variability is straightforward. It is also worth mentioning that all the orthogonal basis sets used in this work have $\psi_0 = \phi_0 = 1$, which is incorporated in our models in any case as a constant term and does not need to be considered further.

There is one potential pitfall when using orthogonal series to model regional effects with few sites. If the total number of ψ s, ϕ s and their interactions approaches the number of sites, there is a danger of severely overfitting the model to match exactly the observed pattern of rainfall at the sites. As a general rule, the total number of site effects in the model (including interactions) should be kept below the number of sites available.

4.1.4 Model checking

Having fitted a GLM to records of daily data, there are several simple but informative checks which may be used to assess its adequacy. Questions which might be asked at this stage include:

- Will the results be reproducible?
- Can the model be simplified?
- Does the model need to be extended?

The question ‘Is the model correct?’ is hardly relevant here: atmospheric processes involve the interaction of many complex physical mechanisms for which data are unlikely to be available. Even if the data were available, it is unlikely that a perfect model could be found — especially using the rather simple formulations (4.2) and (4.3). Of far more importance is that the model is fit for purpose, and that results will be reproducible. See Chandler (1998*a*) for further discussion.

Formal methods for dealing with the questions above rely on methods such as hypothesis testing and calculation of confidence intervals. In the case where GLMs are fitted simultaneously to data from several sites as though they were independent, such methods are not available (recall the discussion in section 4.1.2 above). However, even if they were available it is not clear that they would be useful. This is because such methods implicitly assume that an underlying ‘true model’ exists, and that results will be reproducible if this model can be found. When modelling a complex process, any hypothesis testing procedure which tests a simple model against an extended model will tend, given enough data, to reject the null hypothesis that the simple model represents ‘truth’. In daily rainfall modelling problems, the amounts of data involved are typically large (running into hundreds of thousands of observations), whence formal methods may result in complex models which reflect nothing more than the fact that the rainfall process cannot be modelled perfectly.

This is not to say, however, that formal procedures are of no use in this situation. The danger is of over-interpreting the results of such procedures. Approximate confidence intervals, for example, may usefully be employed to give an informal assessment of uncertainty in parameter estimates, but should not be used in isolation to determine whether or not a predictor should be included in a model.

Other methods for assessing model adequacy are essentially informal, and fall broadly into three categories: assessment of the predictive ability of a model, checks that the forecast probability distributions are correct and checks that there are no discernible unexplained patterns in the available data.

Assessment of predictive ability: A natural way to decide whether or not a predictor should be included in a model is to examine its impact on the model’s predictive ability. For continuous responses such as rainfall amounts, predictive ability can usefully be assessed using a measure analogous to the R^2 of standard regression. This measures the proportion of variability in the data which is explained by the model, and is given by

$$\tilde{R}^2 = 1 - \frac{\text{Mean squared prediction error}}{\text{Variance of original observations}} . \quad (4.6)$$

It is necessary to use Mean Squared Error in the numerator, rather than the error variance, to allow for the fact that in a GLM, the mean of the errors may not be zero (this is a consequence of observations being individually weighted in the fitting procedure).

\tilde{R}^2 may also provide a useful summary measure for some categories of discrete response, but may be difficult to interpret for a binary response such as rainfall occurrence. This is a subject which has received some attention in the literature: for a recent article arguing in favour of its use, see Ash & Shwartz (1999). An alternative way of assessing predictive performance for a rainfall occurrence model is to imagine that we forecast rain for the i th case in the database if the predicted probability, p_i , is greater than 0.5, and forecast no rain otherwise. A table can then be constructed showing the numbers of correct and incorrect predictions for each group of forecasts (rain/no rain), and the overall percentage of correct predictions

may be interpreted as a summary measure of overall predictive performance. Of course, if the true probability of rain is close to 0.5, low predictive ability is unavoidable. Therefore, it is useful to include expected performance levels under the model in such tables of predictive performance, as an aid to interpretation. The computation of expected proportions of correct predictions is straightforward: details may be found in Chandler & Wheater (1998*a*).

Checks on forecast probability distributions: the GLM framework deals with uncertainty in a response variable by specifying a probability distribution conditional on the values of predictors. The family of distributions chosen (for example, the gamma family for rainfall amounts here), and the way in which the predictors are related to the mean of the forecast distribution (via the log odds transformation for the occurrence model (4.2) and the log transform in the amounts model (4.3)), are accounted for in the parameter estimation procedure. One implication of this is that, if the chosen family of distributions is unrealistic, parameter estimates may be biased. Correct specification of the forecast distributions is also important if the fitted models are subsequently to be used in simulations, particularly if extreme events are of interest.

For continuous responses such as rainfall amounts, the easiest way to check the form of the forecast distribution is via quantile-quantile plots of suitably-defined residuals. A residual measure can be chosen such that, if the forecast distributions are correct, all residuals have the same distribution. Examples for the gamma distribution are the Pearson residual:

$$r_i^{(P)} = \frac{Y_i - \mu_i}{\mu_i} \quad (4.7)$$

and the Anscombe residual:

$$r_i^{(A)} = \left(\frac{Y_i}{\mu_i} \right)^{1/3}, \quad (4.8)$$

where Y_i is the observed value for the i th case in the database, and μ_i is the fitted value. The Pearson residual is the proportional error in predicting Y_i : if the shape parameter of the gamma distribution is ν then $r_i^{(P)}$ is expected to have mean zero and standard deviation $1/\sqrt{\nu}$; if the gamma assumption is correct, then $r_i^{(P)} + 1$ has a gamma distribution with scale and shape parameters both equal to ν . Since this distribution is the same for all cases in the dataset, a Q-Q plot showing observed versus expected quantiles of the distribution of $r^{(P)} + 1$ should appear as a straight line if the gamma assumption is correct.

The reason for considering Anscombe residuals in addition to Pearson residuals is that their distribution is extremely close to Gaussian — see, for example, Hougaard (1982) — and this will provide an elegant means of incorporating spatial dependence into a rainfall amounts model (see section 4.1.5 below). The definition at (4.8) yields residuals having the same mean and variance for each case in the dataset: these depend on ν and can be determined straightforwardly using numerical methods, as described in Chandler & Wheater (1998*b*). A normal probability plot of Anscombe residuals from a rainfall amounts model therefore provides an alternative method of checking distributional assumptions here.

Checks on the forecast probabilities for a rainfall occurrence model can be made using an extension of the ideas discussed above for assessing predictive performance for such a model. The basic principle is that, if we collect together all of the days when the forecast probability of rain is close to some preassigned value p^* , then the overall proportion of these days upon

which rainfall was actually observed should be close to p^* . An overview of the ideas is given by Dawid (1986). For practical implementation, we collect together groups of days for which forecast probabilities are in the intervals $(0.0, 0.1)$, $(0.1, 0.2)$, \dots , $(0.9, 1.0)$ and compute observed and expected proportions of rainy days within each of these groups. Unless there is agreement across the whole range of forecast probabilities, there is something wrong with the probability structure of the model.

Checks for unexplained patterns: A further check is necessary to ensure that all effects, for which information is available, have been accounted for in an appropriate way. Traditionally in regression modelling, such checks are carried out by plotting suitably-defined residuals against predicted values, and against the predictors themselves. Such plots may be produced both for predictors which appear in the model, and for potential predictors which may need to be accounted for. Any apparent structure in these plots indicates a problem with the model.

A typical feature of daily rainfall datasets is their large size, which makes residual plots difficult to interpret (there are so many data points that the plot generally appears as a solid mass and any structure is obliterated). For this reason, rather than plotting individual residuals we focus on summary statistics for residual measures over subgroups of observations. For example, to check that seasonality is well reproduced we can compute the mean and root mean squared error of suitably-defined residuals for each month of the year, and plot these: any pattern in the plot, or values which are ‘significantly’ different from zero, indicate seasonal structure which has not been captured.

To aid visual interpretation of such plots, it is helpful to include approximate confidence bands, indicating the range within which mean residuals are expected to lie under the model. Again, such confidence bands should be interpreted informally with large datasets. If all residuals are expected to have mean μ_ϵ and variance σ_ϵ^2 under the model, and a mean residual (\bar{r} , say) is computed over a large subset of M cases, then 95% limits for this mean are at $\mu_\epsilon \pm 1.96\text{s.e.}(\bar{r})$, where $\text{s.e.}(\bar{r})$ is the standard error of the mean residual under the model. If all the M residuals are independent (e.g. if the model is correct and all residuals are from the same site), then this standard error is $\sigma_\epsilon M^{-1/2}$. Otherwise, if some of the M residuals represent different sites on the same day, this standard error will need to be inflated to account for spatial dependence between the sites. Ideally, a separate inflation factor should be computed for every subset of residuals considered. However, this is likely to be computationally expensive and, in view of the fact that the confidence limits are to be interpreted informally, probably over-sophisticated. The approach adopted here has been to calculate a single inflation factor, for the mean residual obtained from all of the data, and to apply this to standard errors obtained for all subsets considered. This is likely to provide a reasonable approximation to the correct standard errors, especially when (as is usually observed to be the case) all inter-site correlations are very similar.

The details of the calculation are as follows: the overall mean residual is defined as

$$\bar{r} = \frac{1}{N} \sum_{t=1}^T \sum_{s=1}^S \chi_{ts} r_{ts} , \quad (4.9)$$

where T is the number of days in the database; S the number of sites; χ_{ts} is an indicator taking the value 1 if a residual is available for site s on day t , 0 otherwise; r_{ts} is the residual at site s on day t ; and N is the total number of observations ($= \sum_{t=1}^T \sum_{s=1}^S \chi_{ts}$). If all cases

were independent, the standard error of \bar{r} would be $\sigma_\varepsilon N^{-1/2}$. When there is dependence between residuals at different sites on the same day (but independence between days), denote by $c_{s_1 s_2}$ the correlation between residuals at sites s_1 and s_2 on the same day (which is assumed constant over all days, and can be estimated straightforwardly). Then we have

$$\begin{aligned} \text{Var}(\bar{r}) &= \frac{1}{N^2} \sum_{t=1}^T \sum_{s_1=1}^S \sum_{s_2=1}^S \chi_{ts_1} \chi_{ts_2} \text{Cov}(r_{ts_1}, r_{ts_2}) \\ &= \frac{\sigma_\varepsilon^2}{N^2} \sum_{s_1=1}^S \sum_{s_2=1}^S \sum_{t=1}^T \chi_{ts_1} \chi_{ts_2} c_{s_1 s_2} \\ &= \frac{\sigma_\varepsilon^2}{N^2} \sum_{s_1=1}^S \sum_{s_2=1}^S n_{s_1 s_2} c_{s_1 s_2} , \end{aligned} \tag{4.10}$$

where $n_{s_1 s_2}$ is the number of days for which sites s_1 and s_2 both have data. Now comparing the square root of (4.10) with the ‘independence’ standard error, we see that the inflation factor required to correct for spatial dependence is

$$\sqrt{N^{-1} \sum_{s_1=1}^S \sum_{s_2=1}^S n_{s_1 s_2} c_{s_1 s_2}} . \tag{4.11}$$

The residual measures used here to check for unexplained patterns are: for the rainfall amounts model, the Pearson residuals defined at (4.7); and for the occurrence model, studentised (or Pearson) residuals defined by

$$r_i^{(S)} = \frac{Y_i - p_i}{\sqrt{p_i(1 - p_i)}} . \tag{4.12}$$

where here Y_i takes the value 1 if the i th case in the database was a wet day and 0 otherwise, and p_i is the forecast probability of rain for that case. The individual $r_i^{(S)}$ s are of little use, because they can only take two values; however, means based on large groups of residuals will be informative. Some discussion of the use of this residual measure may be found in Cox and Snell (1989, section 2.7).

The one remaining aspect of model checking relates to reproducibility of results. In this work, reproducibility is examined by carrying out all of the checks outlined above on extra data which were not used in a model fitting exercise. This requires decisions to be made, at the beginning of an analysis, regarding how a dataset is to be split into ‘fitting’ and ‘validation’ samples. Some possible strategies are as follows:

- Use the early part of the data record (typically, one might consider using the first 75–80%) for model fitting, and the remainder for validation. This is quite a stringent test of a model: in particular, it is likely to fail if rainfall patterns during the fitting and validation periods differ significantly (unless the difference is due to continuous changes which can be parameterised on the basis of the fitting data). Another potential problem is that slight errors in parameter estimates relating to trend functions may result in poor performance in the validation set, leading to exaggerated doubts regarding model adequacy. For example, if a long-term cycle

were fitted to the first half of a record, and the cycle maximum were estimated to occur in 1987, then during the validation period the model would predict gradually falling rainfall amounts after 1987. If in fact the cycle maximum were to occur in 1991, there would be quite a serious divergence between model and data over the validation period. From an engineering design point of view, however, such a model may still be a useful tool so long as it represents realistic rainfall amounts at the time of the cycle maximum, regardless of what this time actually is.

- Remove whole years, at regularly spaced intervals through the data record, and use these for validation after fitting the model to remaining years. Although this would overcome the problems of the strategy above, it may be regarded as not a sufficiently demanding test: on the validation data, the model might be expected to perform quite well since when it was fitted, it was allowed to ‘see’ what happened on both sides of the validation years. It could be argued that such a test does not really give an indication as to whether or not results will be reproducible into the future. Notice that classical cross-validation — whereby each observation in turn is deleted from the dataset and predicted using a model fitted on all the remaining observations — is an extreme example of this strategy, and is certainly inappropriate for our purposes.
- Use a subset of sites for model fitting, and another subset for validation. Again, this perhaps does not represent a stringent enough test, since there may be a high degree of association between the fitting and validation datasets as a result of spatial dependence between sites.

The strategy chosen for selecting a validation dataset will inevitably depend upon the available data. Ideally perhaps, the first strategy should be adopted, since it yields insights into how models may be expected to perform into the future. However, if available record lengths are relatively short (e.g. less than 30 years), then it may be considered appropriate to maximise the use of the limited temporal information available. In this case, the third strategy would be preferable, particularly in cases where data are only available from a few widely-separated locations (in which case spatial dependence between the sites will be limited).

4.1.5 Modelling spatial structure

The discussion so far has related to the parametrisation and checking of GLMs which will enable modelling of the marginal sequences of daily rainfall at each site. From the point of view of assessing evidence for temporal and spatial nonstationarity, to determine the appropriateness of using the stationary Poisson cluster models described in section 2, this may be all that is required in a particular application. However, it may also be of interest to model the daily rainfall at all sites simultaneously — either for purposes of simulating sequences of daily rainfall, or so as to derive a full multivariate distribution for comparison with output from a Poisson cluster model.

There are various ways in which spatial dependence might be included in models of daily rainfall: a review is given in Chandler & Wheater (1998*a*). Here, we use separate structures for the amounts and occurrence models. The structures are described separately. as follows:

Spatial structure — amounts model

The most straightforward way of incorporating spatial dependence into the gamma model for rainfall amounts is via the Anscombe residuals defined at (4.8). As these are normally distributed to a very good degree of approximation, all of the dependence structure in rainfall amounts can be expressed using correlations between Anscombe residuals at each pair of sites. One might choose to model the correlations as a function of inter-site distance and direction: however, in all the work reported here, correlations between Anscombe residuals at all pairs of sites are so similar that they can be regarded as effectively constant. The reason for this is presumably that the three study areas considered below are all small relative to weather systems (which correspond to ‘rain events’ in the terminology of section 2), so that on any particular wet day, all sites tend to be exposed to a very similar rainfall regime.

Simulation of rainfall amounts at ‘wet’ sites for a particular day therefore proceeds: firstly, by sampling a vector of Anscombe residuals from a multivariate normal distribution with an appropriate mean and covariance structure; and secondly by inverting the transformation (4.8) at each site. Simulation of multivariate normal random vectors is straightforward and relies on the fact that if \mathbf{X} has a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, and if $\mathbf{Y} = \mathbf{A}\mathbf{X}$ for some matrix \mathbf{A} , then \mathbf{Y} has a multivariate normal distribution with mean $\mathbf{A}\boldsymbol{\mu}$ and covariance matrix $\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}'$ (see, for example, Krzanowski, 1988, section 7.2). In our case, we generate a vector of independent standard normal variates (with identity covariance matrix) and find a matrix \mathbf{A} such that $\mathbf{A}\mathbf{A}' = \mathbf{C}$, where \mathbf{C} is the desired covariance matrix. Such an \mathbf{A} can be found by calculating the Cholesky decomposition of \mathbf{C} (see, for example, section 2.9 of Press *et al.* (1992) for details and an algorithm) — this will always work providing \mathbf{C} is nonsingular. Problems with singular covariance matrices are unlikely to occur in practice.

There is one potential drawback with this method for dealing with spatial dependence in rainfall amounts: it takes no account of ‘dry’ sites (since the amounts model is only defined at sites where the rainfall amount is nonzero). Therefore it is not guaranteed to produce small amounts of rain near sites which are dry, although wet sites which are close to each other will tend to have similar rainfall amounts. The extent to which this might be a problem is not known.

One final point relates to the use of the model to interpolate missing data in the historical record. If data from some sites are missing, but others are observed, then Anscombe residuals can be computed from the observed sites and the conditional distribution of the missing residuals, which remains multivariate normal, can be calculated (Krzanowski 1988). Missing residuals can then be simulated from this conditional distribution, and back-transformed to interpolate the missing data values. This technique can be used to determine uncertainty in the historical record due to missing data, by simulating missing data many times to construct uncertainty envelopes for historical rainfall statistics.

Spatial structure — occurrence model

Incorporating spatial dependence into the binary rainfall occurrence model is rather more difficult than for amounts. The structure chosen has been more or less dictated by the following requirements:

- It should be amenable to simulation. In particular, it needs to be able to cope with the fact

that on each successive day, the marginal probabilities of rain at each of the sites will change.

- It should be estimable from historical data.
- It should take into account the fact that any network of sites is merely sampling an underlying stochastic process which operates in continuous space (in this way we hope to maintain the capability of simulating at a different set of locations from those which were used in model fitting).

A variety of different techniques have been investigated, including:

1. Specify the dependence via the correlation structure of the wet/dry field, in the spirit of Lunn & Davies (1998) and Park, Park & Shin (1996). However, it is well known that correlations between binary variables (Y_i and Y_j say, with $P(Y_i = 1) = p_i \geq p_j = P(Y_j = 1)$, without loss of generality) are constrained by the probabilities p_i and p_j . In the example here, we must have $P(Y_i = Y_j = 1) \leq p_j$, so that $\text{Cov}(Y_i, Y_j) \leq p_j(1 - p_i)$ and $\text{Corr}(Y_i, Y_j) \leq \sqrt{p_j q_i / (q_j p_i)}$ where $q_i = 1 - p_i$. At each site, the probability of rainfall occurrence changes from day to day and it is therefore very difficult to specify a correlation-based dependence structure which is guaranteed to be consistent with the marginal probabilities at each site over any extended period of time.
2. Generate a continuous-valued process with appropriate dependence structure, and define sites with values above some threshold to be wet, otherwise dry. An example of this approach is given by Emrich & Piedmonte (1991). However, it suffers from the same drawbacks as the previous approach in that dependence is usually specified via correlation structure. Moreover, every time the marginal probabilities change, typically a high-dimensional set of equations has to be solved numerically, which is likely to be inefficient for the present application.
3. Specify the dependence via odds ratios between pairs of sites (see, for example, Cox and Wermuth, 1996, section 3.7). The motivation here is that odds ratios are not constrained by the probabilities when just 2 sites are involved. However, for more than 2 sites simultaneously it is not easy to determine whether any set of odds ratios are consistent with each other, particularly in simulation when the probabilities change from day to day.
4. Incorporate the dependence by including other sites' rainfall occurrence on the same day as extra predictors in the logistic regression model. The result may be regarded as an example of a Markov Random Field auto-logistic model (Besag 1974). This approach, however, leads to complications in fitting the models, and simulation costs are increased because of the need to use iterative simulation methods such as the Gibbs Sampler (Geman & Geman 1984). In addition, our experience has been that this type of model is liable to yield unrealistic results for this particular application.

The approach finally adopted here makes use of the observation that, at spatial scales up to at least several hundred square kilometres, spatial dependence in daily rainfall occurrence is mainly due to the fact that all sites tend to be influenced by the same weather systems on particular days. This process can be modelled by including a hidden weather state which categorizes each day as 'wet' or 'dry' over the entire area. In order to develop a model whose parametrisation is not

constrained by the marginal probabilities of rainfall occurrence, the effect of the hidden weather state is modelled on the log odds scale.

In order to define the model, some notation is required as follows: daily rainfall amounts are to be generated at each of a set of spatial locations $\mathbf{x}_1, \dots, \mathbf{x}_S$. The random variable Y_j takes the value 0 if the rainfall amount at site \mathbf{x}_j is zero, 1 otherwise. We have a logistic regression model (4.2) which gives $P(Y_j = 1) = p_j$.

Associated with the current day is a random variable X which represents the ‘weather state’ on that day. $X = 1$ with probability α (representing a ‘wet day’) and 0 otherwise (a ‘dry day’). X is not directly observable.

The Y s are modelled as conditionally independent given X . A plausible stochastic process giving rise to this type of behaviour would be a nonstationary Boolean-type model, whereby discs (of either fixed or random radius) are attached to events of a nonstationary spatial Poisson process. Any sites covered by at least 1 disc give $Y = 1$, all other sites have $Y = 0$. Such a process will exhibit local autocorrelation between locations separated by less than the diameter of a disc, but beyond this sites will be independent (by virtue of the underlying Poisson process). This mechanism is similar to that used in the Poisson cluster models of chapter 2, and the ‘discs’ here might be thought of as corresponding to rain cells in those models: the point is that rain cells tend to be a few kilometres in diameter, whence it may only be reasonable to assume conditional independence of the Y s here as long as all pairs of sites are fairly well separated.

The probability of rain at each of the sites on day t is modelled using

$$\ln \left(\frac{P(Y_j = 1|X = x)}{1 - P(Y_j = 1|X = x)} \right) = \ln \left(\frac{p_j}{1 - p_j} \right) + x \ln a + \ln b(\alpha, p_j, a) . \quad (4.13)$$

Here $\ln a$ is constant for all sites and days, and is free over the range $(-\infty, \infty)$. $\ln b(\cdot)$ is a function of α , p_j and a , chosen to ensure that the unconditional probability of rain at site j is p_j . We abbreviate it to b_j for convenience. By rearranging equation (4.13) to find the unconditional probability that $Y_j = 1$, and setting this equal to p_j , we find that b_j must be a root of the quadratic equation

$$ap_j b_j^2 + [1 + \alpha(a - 1) - p_j(a + 1)] b_j - (1 - p_j) = 0 . \quad (4.14)$$

To determine which of the two roots of this equation should be used, first note that, from (4.13), $a > 0$ and $b_j > 0$. Note also that both α and p_j lie in the range $[0, 1]$ since they are probabilities. Writing the left hand side of (4.14) as $Ab_j^2 + Bb_j + C$, we must therefore have $4AC \leq 0$. If $4AC$ is strictly less than 0 then exactly one positive real root to the equation is guaranteed, and is given by

$$\frac{-B + \sqrt{B^2 - 4AC}}{2A} . \quad (4.15)$$

When $4AC = 0$, a valid solution for b_j can still be found: this will arise when $p_j = 0$ or 1 (the case $a = 0$ need not concern us since then $\ln a$ would be undefined in (4.13)). When $p_j = 0$ we obtain $b_j = (1 - \alpha + \alpha a)^{-1}$, and when $p_j = 1$ we obtain $b_j = 1 - \alpha + a^{-1}$. Both quantities are strictly positive for the allowable range of values of α and a .

It is of interest to examine the behaviour of (4.13) in the limit as $a \rightarrow \infty$, since this will affect the model’s ability to reproduce features such as the proportion of days when all sites are wet. It can be shown that the limiting value of the probability of rain at site j depends on both p_j and on α , on both wet and dry days. The following table summarises the behaviour as $a \rightarrow \infty$:

	$x = 0$	$x = 1$
Case I: $p_j > \alpha$	$(p_j - \alpha)/(1 - \alpha)$	1
Case II: $p_j < \alpha$	0	p_j/α

Notice that no matter how large a is, we cannot make $P(Y_j = 1|X = 1)$ arbitrarily close to 1 if $p_j < \alpha$. Similar results hold if we let $a \rightarrow 0$ (effectively letting $X = 1$ correspond to a ‘dry’ day).

A further aspect of interest is the covariance structure of the model (4.13). The covariance between Y_i and Y_j ($i \neq j$) is defined as $P(Y_i = Y_j = 1) - p_i p_j$, and can be shown to be equal to

$$\text{Cov}(Y_i, Y_j) = \frac{(1 - \alpha)p_i p_j (1 - p_i)(1 - p_j)(1 - b_i)(1 - b_j)}{\alpha [1 - p_i(1 - b_i)][1 - p_j(1 - b_j)]}. \quad (4.16)$$

The corresponding correlation is

$$\text{Corr}(Y_i, Y_j) = \frac{(1 - \alpha)(1 - b_i)(1 - b_j)}{\alpha [1 - p_i(1 - b_i)][1 - p_j(1 - b_j)]} \sqrt{p_i p_j (1 - p_i)(1 - p_j)}. \quad (4.17)$$

We now consider how to estimate the necessary parameters in the model (4.13) from historical data. There are two parameters to estimate on any given day: α and a . An obvious simple choice for α is to set it equal to the mean of the p s. In this case the marginal predictions of the GLM at all sites are used to determine whether a day will be ‘wet’ or ‘dry’. This will provide an automatic means of incorporating features such as seasonality into the weather states, since these ought to be reflected in the p s. If subsequently the model does not perform as expected, it may be worth considering more sophisticated approaches for determining α (it was noted above, for example, that even on a ‘wet’ day, the probability of rain at all sites cannot be made arbitrarily close to 1 unless α is less than all of the marginal probabilities).

How to choose a is less obvious. It can be estimated using observed Y values at some or all of the n sites on each of the T days, and also the corresponding p s predicted by the GLM. Where the p s are known, the corresponding α can be calculated. The sequence of X s is not observed. It is possible, numerically, to obtain a maximum likelihood estimate of a ; however, this may be undesirable since the model structure, while plausible, is probably a fairly crude approximation of reality (recall the discussion above indicating a Boolean process which could be considered as underlying the model). An alternative estimation procedure is a method of moments. From a practical point of view, the importance of representing spatial structure in rainfall occurrence patterns stems mainly from the need to reproduce the distribution of spatial coverage (i.e. the proportion of sites which are wet). Therefore we propose to use some feature of the observed coverage distribution to estimate a . The chosen feature is the variance: other features, such as the overall proportion of days on which all sites are wet (or dry), have the undesirable property that they are highly dependent on the number of sites with available data on each day. The idea is to choose a so as equate the observed and theoretical variances of the coverage distribution. The theoretical variance of the coverage, for a given value of a , can be derived from the expression for covariances at (4.16). The observed variance can be calculated straightforwardly from the time series of coverages, weighting each day by the number of active sites. Equating the two can be achieved using straightforward numerical methods.

As with the amounts model previously, the spatial dependence structure used here can be used to interpolate missing data in the historical record. Suppose on a particular day we observe

$Y_1 = y_1, \dots, Y_k = y_k$ ($k < n$), and wish to fill in the remaining values Y_{k+1}, \dots, Y_n . Rearranging (4.13), we find that

$$P(Y_1 = y_1, \dots, Y_k = y_k | X = x) = \prod_{y_i=1} \frac{a^x b_i p_i}{1 - p_i (1 - a^x b_i)} \prod_{y_i=0} \left(1 - \frac{a^x b_i p_i}{1 - p_i (1 - a^x b_i)} \right).$$

Now we can use Bayes' Theorem to find the conditional probability distribution of X on the basis of the observed Y s:

$$P(X = 1 | Y_1 = y_1, \dots, Y_k = y_k) = \frac{P(Y_1 = y_1, \dots, Y_k = y_k | X = 1)}{\sum_{x=0}^1 P(Y_1 = y_1, \dots, Y_k = y_k | X = x) P(X = x)}. \quad (4.18)$$

Recalling that $P(X = 1) = \alpha = 1 - P(X = 0)$, we can therefore interpolate missing data by simulating $X = 1$ with probability given by (4.18), and then sampling the remaining missing sites independently from (4.13) as before.

4.2 Model performance

To illustrate the GLM approach to daily rainfall modelling, we now present the results of three case studies. The first is from the Galway Bay region of Western Ireland and is presented in some detail. The others are from the Brue catchment in South-West England, and the Blackwater (a tributary of the Thames located between Guildford and Basingstoke). For the latter studies, only an overview of the modelling will be given, as the basic procedure is the same. The catchment sizes are roughly 1200km², 120km² and 400km² respectively, although for the Brue and Blackwater, gauges are used which lie outside the catchment (substantially so in the case of the Brue). The Galway study uses 23 daily rain gauge records; the Brue has only 4 moderately long records, together with 49 tipping-bucket gauges within the catchment for which data are available between 1994 and 1997; and the Blackwater has 44 daily records.

4.2.1 West of Ireland

The area around Gort, to the South of Galway in Western Ireland (see figure 4.1) has historically been subject to extreme flood events. The area affected is a low-lying Karst system, fed by rivers draining the Slieve Aughty mountains to the east. Under extreme conditions (associated with extended wet periods) ephemeral lakes, known as turloughs, overflow and coalesce, causing widespread flooding involving inundation of property and damage to livestock and roads. In the past such widespread flooding occurred once in the 1920s and again in the 1950s; then in early 1990, 1991, 1994 and 1995.

A preliminary report after the 1991 event (Daly 1992) identified changing rainfall patterns as a possible cause of the increased flooding. Subsequently an extensive investigation, funded by the Irish Office of Public Works, was carried out in an attempt to suggest and evaluate possible flood alleviation measures; the work reported here is an extension of work carried out as part of that study, which is reported in OPW (1998). The objectives of the study were:

1. To investigate the extent to which the flooding may be attributed to abnormal rainfall (rather than other factors such as changes in land use).

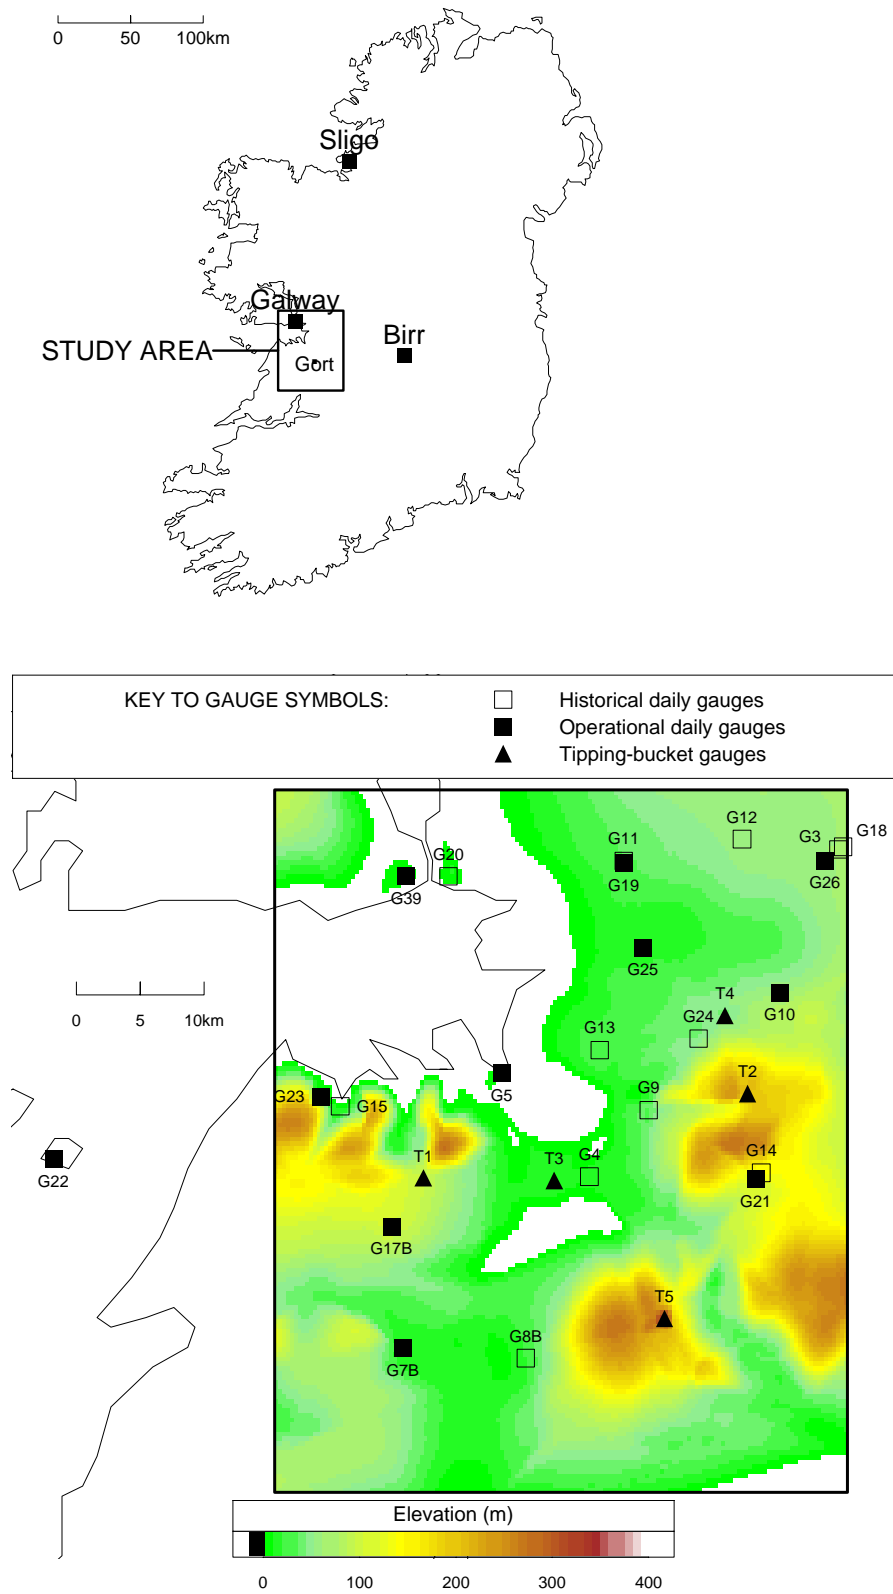


Figure 4.1: The Galway study area, and locations of raingauges. Top: location of the study area within Ireland, and positions of the gauges with long monthly records. Bottom: detail of the study area, showing positions of daily raingauges.

2. To examine the evidence for long-term temporal variability in the rainfall record, with a view to quantifying the nature and extent of any apparent changes.
3. To provide a tool for simulating future rainfall scenarios, which could be used to provide inputs to hydrological models of the area in order to assess the likely effectiveness of potential flood alleviation measures.

Two separate sources of data were used in the work reported here: daily rainfall data from a network of 23 gauges run by the Irish Meteorological Office, and monthly data from gauges at Birr and Sligo. Figure 4.1 shows the locations of all the gauges. The daily data span the period 1941–1996, although not all gauges have contemporary records. Typically, between 6 and 9 gauges have been operational at any time since the beginning of the 1960s. For these gauges, any non-zero rainfall value less than 0.1mm has been recorded as a ‘trace’ amount. The record from Birr runs from 1873 to 1994, and that at Sligo from 1890 to 1994. The periods of record for each gauge are shown in table 4.1.

Various exercises were carried out to ascertain the quality of the data used in this study. For details, see Chandler & Wheeler (1998*a*). To summarize: both daily and monthly records had previously been quality-controlled by the Irish Meteorological Service, and any value flagged as dubious was discarded. The study area was visited to inspect all gauges which are currently operational. In addition, simple exploratory statistical analyses were carried out to highlight unusual features of the data. The main conclusions were that some of the daily gauge records may be a little unreliable, and that over-detailed interpretations of any analyses should be avoided. A couple of particularly suspect gauges were discarded from any subsequent analysis.

The monthly records were tested, using a technique similar to the jackknife (Efron & Tibshirani 1993), to ensure that they could be regarded as representative of rainfall patterns within the study region. Again, Chandler & Wheeler (1998*a*) contains details of the procedure used. These records have not been incorporated formally into analyses reported here — they have been used only to suggest the nature of possible long-term trends in rainfall patterns.

To investigate long-term trends, time series plots of various summary statistics were examined, at monthly and annual timescales, for individual gauges and for the whole area. In general, records from individual gauges are too variable for any clear pattern to emerge, as are areal statistics at monthly timescales. However, the annual series of areal mean daily rainfalls indicates that during the 1960s, annual rainfall amounts tended to be rather lower than either before or since. This feature is most pronounced in the winter months (December–February): figure 4.2 shows the time series of mean winter daily rainfalls in the study area. To determine whether this apparent trend is part of a longer-term pattern, the long records from Birr and Sligo were examined. Figure 4.3 shows the mean winter rainfall, averaged over 5-year time periods, at Birr from 1875 to 1995. The pattern here is very similar to that in figure 4.2 for the period where the records overlap. The plot also indicates that the rainfall since the 1960s conforms to previous behaviour. There is some suggestion of a cyclical pattern (lows in the 1890s, 1960s and possibly the 1930s, and highs around 1920, 1990 and possibly 1950).

These results are in broad agreement with other studies of trends in Northern European climate. For example, the 1996 report of the UK Climate Change Impacts Review Group (DOE 1996) indicates that the decade from 1984–1995 was unusual relative to a ‘baseline’ climate defined over the period 1961–1990. Our conclusions agree with this, but also suggest that this choice of baseline

Gauge	Data type	Period of operation
G3	Daily	December 1941 – September 1991
G4	Daily	July 1943 – July 1971
G5	Daily	July 1943 – July 1996
G9	Daily	January 1954 – April 1972
G10	Daily	July 1943 – July 1996
G11	Daily	May 1944 – February 1951
G12	Daily	March 1951 – April 1989
G13	Daily	March 1952 – August 1976
G14	Daily	August 1952 – April 1982
G15	Daily	January 1960 – September 1975
G18	Daily	January 1971 – March 1979
G19	Daily	May 1975 – July 1996
G20	Daily	April 1978 – April 1990
G21	Daily	July 1982 – November 1996
G22	Daily	June 1982 – June 1995
G23	Daily	June 1984 – July 1996
G24	Daily	November 1984 – December 1994
G25	Daily	June 1985 – November 1996
G26	Daily	September 1993 – July 1996
G39	Daily	November 1965 – November 1996
G7B	Daily	June 1943 – July 1996
G8B	Daily	June 1943 – December 1990
G17B	Daily	January 1975 – November 1996
Birr	Monthly	January 1873 – December 1994
Sligo	Monthly	January 1890 – December 1994

Table 4.1: Periods of record for each rain gauge used in the Galway study. Refer to figure 4.1 for locations of the gauges.




Figure 4.2: Mean daily rainfall, during the December–February period, over the entire Galway study area for the years 1941–1995.

period is unrepresentative. Some other studies also report this period as being atypical — for example Pfister (1992) found that in Central Europe, winters during the period from 1965–1979 were 25% wetter than the long-term average of the previous 60 years.

A final piece of exploratory data analysis for the daily rainfall record uses analysis of variance (ANOVA) to indicate how predictable the daily rainfall sequence is. An ANOVA can be regarded as a regression model with categorical predictors (e.g. Dobson (1990), p.4). It is likely that gauge location and month of the year are both factors which affect daily rainfall amounts. Therefore a 2-way ANOVA has been carried out, which corresponds effectively to fitting a regression model with a separate parameter for every possible month/site combination. This model explains 2.86% of the variance in daily amounts, indicating that the rainfall sequence is dominated by noise at a daily timescale. However, at longer timescales the structure becomes clearer (for example, fitting the same ANOVA model to monthly data explains 24.0% of the variance).

For the modelling exercise, data for all years up to and including 1994 were used for model fitting and the remainder (1995–97) were kept for validation. The validation set is rather small: however, it was felt that the early 1990s appeared so atypical that they needed to be incorporated into the model fitting, in order adequately to assess the evidence for changing climate in the area. Models were fitted sequentially, starting with ‘obvious’ predictors and successively adding extra predictors and interactions. In the early stages of building a model, the value of extra predictors was assessed by examining ‘log-likelihoods’ for each model; in the later stages, predictive performance




Figure 4.3: 5-year mean winter rainfall at Birr, for the years 1875–1995.

and residual analyses, as described above in section 4.1.4, were used to guide the fitting. Broadly, the sequence in which predictors were added to a model was as follows:

1. Seasonality (initially represented by predictors $\sin(2\pi \times \text{month}/12)$ and $\cos(2\pi \times \text{month}/12)$).
2. Previous days' rainfalls. For each model, up to 5 previous days were considered.
3. Site effects, via orthogonal series and site altitude.
4. Interactions between the main effects already in the model.
5. Long-term trends, if residual analyses indicated that these were necessary.

In examining the evidence for changing rainfall patterns, three separate trend functions were considered, as follows:

- Constant linear trend. Although it is implausible to extrapolate such a function indefinitely outside the range of our data, it may well provide a good approximation to any monotonic trend over the period of record.
- No trend until time t_0 ; then linear trend afterwards. This is intended as a crude approximation to what would be observed if human-induced climate change were occurring (t_0 being the year in which the change started to occur).
- Cyclical trend, as suggested by the Birr and Sligo records (see figure 4.3).

Model	No. of parameters	Log-likelihood
1: No trend	26	-67695.2
2: Linear trend throughout record	33	-67546.8
3: Linear trend after t_0	34	-67545.1
4: Cyclical trend	35	-67537.4

Table 4.2: Log-likelihoods (ignoring spatial dependence) for rainfall occurrence models fitted to Galway daily rainfall data, 1941–1994.

Thus four models (corresponding to ‘no trend’, plus one for each of these three trend functions) were fitted for rainfall amounts, and four for rainfall occurrence.

For both occurrence and amounts models, there was compelling evidence for the inclusion of trend functions. In summarising the results here, we give log-likelihoods (ignoring spatial dependence) for all 4 models fitted, and compare the performance of the ‘no trend’ model with that of the best fitting of the other three. Full details of all 4 models fitted for both amounts and occurrence may be found in Appendix B.1.

West of Ireland — occurrence model

Table 4.2 gives the log-likelihoods attained by each of the fitted occurrence models. The three models incorporating trend functions perform similarly in likelihood terms, and are all clearly much better than the model with no trend function (the addition of 7, 8 and 9 parameters respectively yielding an increase of at least 150 in the log-likelihood). There is little evidence for the changepoint model over a simple linear trend; however, the cyclical trend model does appear to offer an improvement over both of these. The estimated cycle length is 75.46 years (with a nominal standard error of 6.46 years), and the estimated time of the cycle maximum is 1949.8 (nominal standard error 6.46). It is worth reiterating that the true standard errors are probably larger than this, as a result of spatial dependence which has not been accounted for here.

We now give an overview of the fitted models 1 and 4.

Fitted occurrence model — no trend: The predictors in the fitted ‘no trend’ model can loosely be summed up as follows:

- Constant term.
- Site effects, represented via:
 - Sine and cosine terms at the first Fourier frequency for both eastings and northings, together with their interactions (recall the discussion in section 4.1.2). The area over which the Fourier representation holds is illustrated in figure 4.4.
 - Site altitude.
- Seasonality, represented by an annual sine wave.
- Temporal dependence, represented by indicator variables for rainfall occurrence on each of 5 preceding days.

	Observed		% correct	
	Dry	Wet	Observed	Expected
Forecast dry	26126	12632	67.4	67.6
Forecast wet	17816	84084	82.5	82.6
OVERALL % CORRECT:			78.4	78.5

Table 4.3: Observed versus expected predictive performance for basic occurrence model. ‘Forecast dry’ row corresponds to days when the predicted probability of rain was less than 0.5; ‘forecast wet’ corresponds to other days.

- Temporal persistence, represented by an indicator for rainfall occurrence on all 5 preceding days.
- Interactions between:
 - Site effects and temporal dependence effects.
 - Site effects and temporal persistence effects.
 - Seasonal effects and temporal dependence effects.
 - Seasonal effects and temporal persistence effects.

Full details of the model may be found in Appendix B.1.

The interactions between seasonal and temporal dependence/persistence effects are as expected, and reflect the fact that temporal autocorrelation in rainfall sequences is stronger in winter than in summer. Interactions involving site effects are harder to visualise, but presumably reflect slight regional variation in exposure to different types of system. Interpretation of site effects is aided by extracting the corresponding terms in the fitted model, and contouring the function represented by these terms. Figure 4.4 shows the main site effects in the model, produced by extracting predictors corresponding to orthogonal series representations of eastings and northings. Note that the effect of site altitude, which appears in the model as a separate predictor, is therefore not included in the figure — the map shows systematic regional effects after accounting for altitude. The structure shown in the figure makes sense from a physical point of view — the main features are an area of increased rainfall occurrence centred on end of Galway Bay, and an area of decreased occurrence in the South-East of the area. Comparing with figure 4.1, this latter area is seen to correspond to mountains which lie to the East (i.e. in the rain shadow) of uplands in the South-West, and therefore we might expect reduced rainfall occurrence here.

To check the model, we first assess its predictive ability, as detailed above in section 4.1.4. Firstly, table 4.3 gives the observed versus expected performance of occurrence forecasts made on the basis of whether or not the predicted probability of rain was greater than 0.5 — for example, there are 26,126 cases in the database where the predicted probability of rain was less than 0.5 and where no rain was observed. This table shows good agreement between observed and expected performance.

The next check is that the forecast probabilities agree with observed rainday frequencies. Table 4.4 shows the results of this check. The columns represent ranges of forecast probabilities — the column headed “**1**” contains all days with forecast probabilities between 0 and 0.1, column “**2**” contains days with forecasts between 0.1 and 0.2, and so on. Again, the

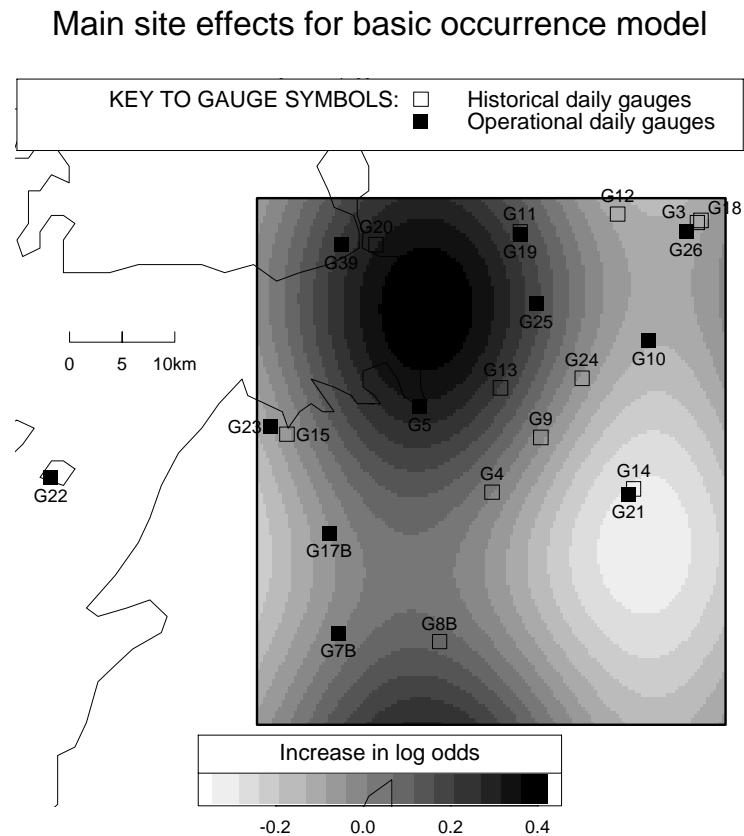


Figure 4.4: Main fitted site effects for basic rainfall occurrence model. The solid box is the area over which the Fourier representation of site effects is defined.

agreement between observed and expected rainday proportions is good, throughout the range of the distribution.

Having satisfied ourselves that the model is adequate as regards predictive performance, it is necessary to check for any apparent structure in residuals. Figure 4.5 shows the mean and mean squared error of studentised residuals from the model, for each month of the year and for each year in the fitting dataset. The studentised (or Pearson) residuals should all have a mean of zero and root mean squared error of 1. There is little seasonal structure in the monthly residual plots — the monthly means are, with the exception of December, not significantly different from zero at the 95% level, and December itself is not significant enough to excite attention considering 12 means are being tested here, and there is a 5% probability that any one of these will appear significant by chance. Moreover, the monthly root mean squares are all close to 1. The annual plots tell a different story, however, with a clear trend in the annual mean residuals and a decrease in the annual root mean square towards the end of the series. This is a strong indication that a trend function should be added to the model. Notice in particular that the series of annual mean residuals appears to decrease until the mid 1980s, and then to increase again — recall from the discussion above that the best-fitting model for rainfall occurrence contains a cyclical trend with period 75.46 years and maximum in 1949.8 (corresponding to a minimum in 1987.5).

A similar exercise was carried out to check that site effects were adequately captured by

Forecast decile	1	2	3	4	5	6	7	8	9	10
Observed	0.000	0.218	0.240	0.364	0.454	0.531	0.679	0.759	0.850	0.924
Expected	0.000	0.190	0.250	0.346	0.449	0.541	0.661	0.762	0.853	0.920
<i>N</i> days	0	367	17761	11965	8665	5853	5232	17483	48614	24718

Table 4.4: Observed versus expected proportions of days with rain, for data grouped according to forecast probability of rainfall occurrence (basic occurrence model).

the model. Out of 21 sites contributing to the fitting database, 5 had mean residuals which were significantly different from zero at the 95% level. These were sites G11 (mean 0.897 with associated standard error 0.0204), G14 (mean -0.0380, standard error 0.0099), G19 (mean -0.0531, standard error 0.0151), G21 (mean 0.0868, standard error 0.0153) and G17B (mean 0.0370, standard error 0.0122). Referring to figure 4.1, sites G11 and G19 are almost coincident, as are G14 and G21. The fact that the residuals at sites G11 and G21 are positive, whereas those at G19 and G14 are negative, suggests that these discrepancies are not due entirely to mis-specification of site effects, but may be related to another factor (such as different periods of operation for each site, for example).

The same sequence of checks was carried out on the residuals obtained by applying the fitted model to the validation dataset (years 1995–1997). The only point worth noting here is that the overall the mean Pearson residual for the validation dataset is significantly less than zero. This is a convincing demonstration that a trend function needs to be included in the model; in the light of this overall lack of fit, there is little point in reporting the results of any more detailed analyses on the validation data.

We now report the results from the best-fitting occurrence model incorporating a trend function.

Fitted occurrence model — cyclical trend: As reported above, the best trend function found to fit the historical rainfall occurrence pattern was a 75.46 year cycle with a minimum in 1949.8. There was significant evidence that this trend interacted with site effects and with seasonality. The seasonal interactions can be interpreted in terms of wetter winters and drier summers as the trend function increases. The interactions with site effects are best presented graphically, by extracting all terms corresponding to these site effects and plotting the value of the function defined when the value of the trend function is 1. This gives a map showing the effect of a unit increase in the trend function, as shown in figure 4.6. This is almost exactly the opposite of the main site effects shown in figure 4.4 (the fitted main site effects for the models incorporating a trend are all similar to those for the basic model). The interpretation is that as the trend function increases, regional variability over the area decreases and the occurrence regime becomes more homogeneous.

The same checks were carried out for this model as for the previous one. The results relating to predictive performance are very similar, and indicate a satisfactory fit. The seasonal structure of residuals, illustrated in figure 4.7, is much the same as before. However, note from figure 4.7 that there is now no trend in the mean annual residuals. Some of the annual means (notably those corresponding to 1954, 1958 and 1994) do fall outside the confidence

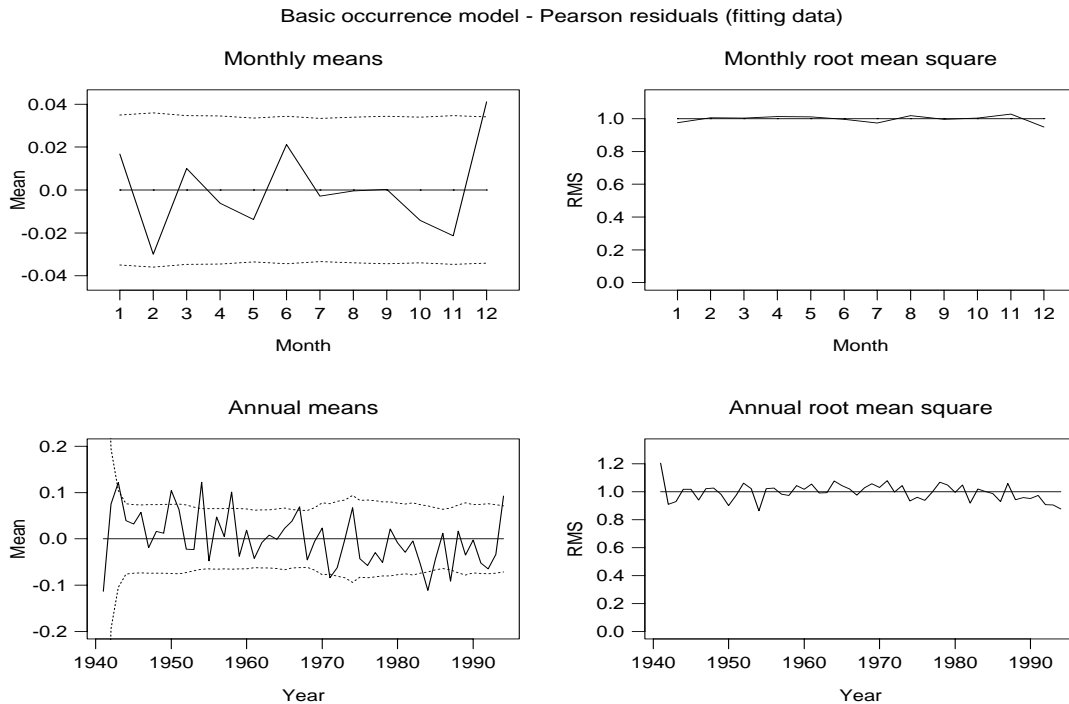


Figure 4.5: Seasonal and annual structure of Pearson residuals from basic occurrence model. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

limits, which suggests that there are still some inadequacies in the model, but this is a great improvement upon the results for the basic model presented in figure 4.5. The most obvious discrepancy between observed and expected residual behaviour for the new model is in the annual root mean square, which still shows a decrease over the last 10 or so years of the fitting dataset. The implications of this are not understood.

The remaining check is for site effects. Here, there are still some sites with mean residuals significantly different from zero. These are G14 (mean -0.0353 , associated standard error 0.0099), G21 (mean 0.0680 , standard error 0.0153), G25 (mean 0.0965 , standard error 0.0173), G26 (mean 0.1279 , standard error 0.0491) and G17B (mean 0.0686 , standard error 0.0122). Sites G14, G21 and G17B were problematic for the basic occurrence model as well: for G14 and G21, it was conjectured that part of the problem may have been to do with different periods of record and the absence of a trend function in the model. It appears from these results that there is something else causing problems for these sites, but it is difficult to see how the problem could be resolved. The problems with the basic model at sites G11 and G19 do indeed appear to have been related to their periods of record. However, sites G25 and G26 were well fitted by the basic model, and now cause problems. In the case of G26, this is probably because it only started recording in September 1993 (see table 4.1), so its results are primarily influenced by data from 1994 which, from figure 4.7, has significantly high residuals. The problem with G25 is not understood.

For the validation dataset, the overall mean Pearson residual for this model is negative (-0.0300) but not significantly different from zero. The good agreement between observed and expected predictive performance remains, as does that between observed and expected

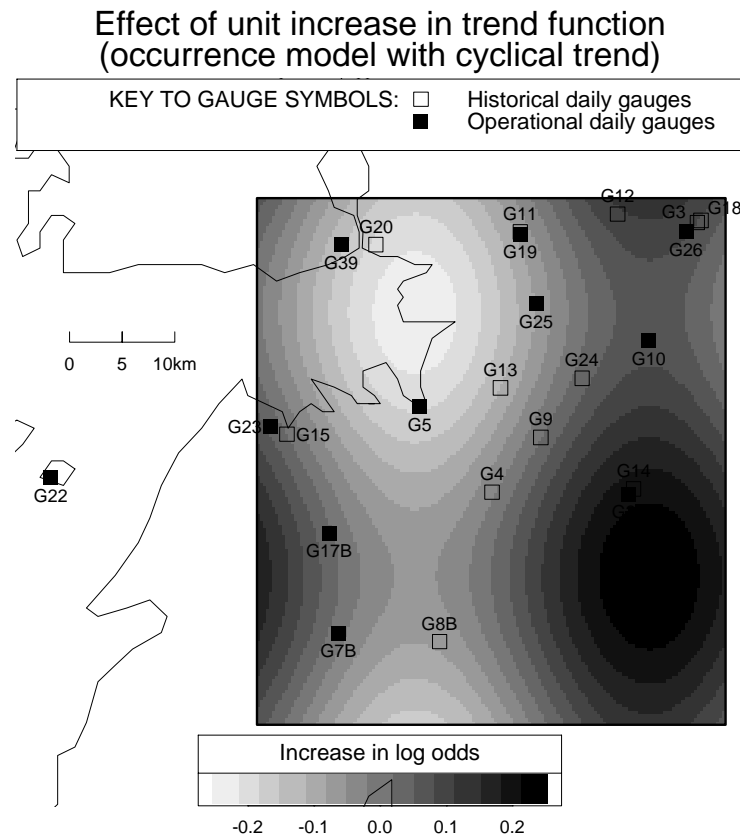


Figure 4.6: Regional effect of cyclical trend function upon log odds for rainfall occurrence, according to fitted model.

rainday frequencies. The one exception is for days when the forecast probability of rain is between 0.2 and 0.3 — here, the observed proportion of days with rain is 0.158 in contrast with the expected proportion of 0.247. Based as this is upon the results of 1130 days, this is a highly significant discrepancy, although its interpretation is unclear.

The seasonal and annual results from the validation dataset are much as expected. There is a tendency for negative residuals (in the monthly analysis, June and December both have significant negative mean residuals), and 3 of the 10 sites which were recording during the validation period (G5, G10 and G19) have significant negative mean residuals. In contrast, site G7B has a significant positive mean. Note that none of these four sites was flagged as problematic in the fitting dataset.

In summary: there is strong evidence for a trend in the pattern of rainfall occurrence over the area. Of the trend functions considered, the one most strongly supported by the data is cyclical. The trend appears to interact with site and seasonal effects. There remain some areas, particularly in the representation of site effects, where the model could possibly be improved. However, in view of concerns regarding the quality of some of the data used here, it is possible that some of the problems may be due to data errors. In any case, it is difficult to see how some of them could be resolved (particularly the lack of fit, in opposite directions, at sites G14 and G21 which are almost identically located).

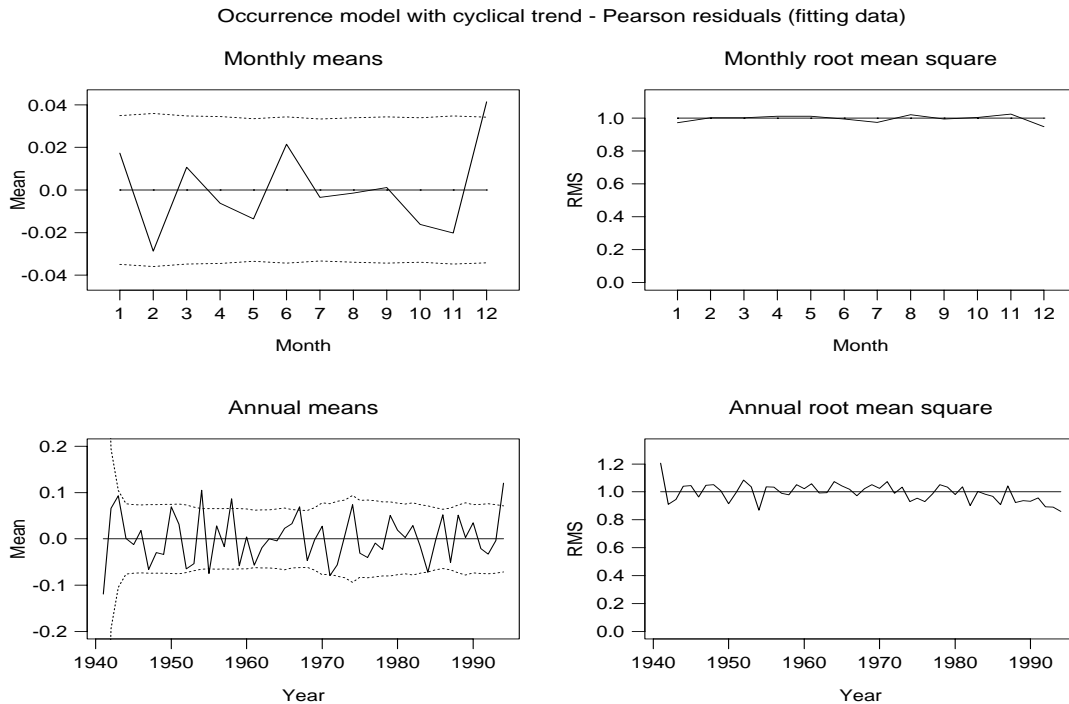


Figure 4.7: Seasonal and annual structure of Pearson residuals from occurrence model with cyclical trend. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

West of Ireland — amounts model

For the amounts model, the same fitting procedure was followed as for the occurrence model. Table 4.5 shows the log-likelihoods achieved by the models incorporating different trend functions, together with the number of parameters estimated for each model. Again, there is evidence for the inclusion of a trend function, but now the best supported function is linear after some time t_0 . The estimate of t_0 is 1972.6, with a standard error of 2.6 years.

We now report the results for models 1 and 3 in table 4.5.

Fitted amounts model — no trend: The predictors in this model may be summarised as follows:

- Constant term.
- Site effects, represented in the same way as for the occurrence model above.
- Seasonality, represented by:
 - An annual sine wave.
 - An additional indicator variable for the month of November.
- Temporal dependence, represented by $\log(1 + \text{rainfall } x \text{ days ago})$ for $x = 1, 2, 3, 4, 5$ and by indicators for ‘trace’ values on each of the preceding 5 days (i.e. values which were recorded as “less than 0.1mm”).
- Temporal persistence, represented by an indicator for rainfall occurrence on all 5 preceding days.

Model	No. of parameters	Log-likelihood
1: No trend	34	-241099.9
2: Linear trend throughout record	37	-241057.9
3: Linear trend after t_0	38	-241029.5
4: Cyclical trend	37	-241046.7

Table 4.5: Log-likelihoods (ignoring spatial dependence) for rainfall amounts models fitted to Galway daily rainfall data, 1941–1994.

- Interactions between seasonality and temporal dependence effects.
- Interactions between seasonality and temporal persistence effects.

Full details of the model may be found in Appendix B.1.

As with the occurrence model, the interactions between seasonality and temporal dependence reflect the seasonal variability in temporal autocorrelation of rainfall sequences.

A map showing the fitted site effect structure (again, omitting the effect of altitude) is given in figure 4.8. The pattern is very different from that for rainfall occurrence (figure 4.4) and is dominated by a trough running North-South through the centre of the study area, with enhanced rainfall amounts to the East and West. The coincidence of these areas of enhancement with the high ground in the area (see figure 4.1) may indicate that dependence of rainfall amounts upon altitude is not linear on the log scale (since altitude also appears as a predictor in the model). However, it should be noted that little data are available within the centre of the Eastern area of enhancement in figure 4.8, and the apparent enhancement here may simply be an artefact of the Fourier representation used (which effectively forces the fitted function to be periodic, so that the profile along the Eastern edge of the area must be the same as that along the Western edge where data, from sites G15 and G23, are available).

The fitted model explains 5.9% of the variance in observed non-zero rainfall amounts. Bearing in mind that only 2.9% of the variance is attributable to month and site effects (according to the ANOVA results reported previously), this actually represents quite a good level of predictive ability. The improvement is due to the inclusion of previous days' rainfalls as predictors. To check the assumption that rainfall amounts are gamma distributed, figure 4.9 shows a normal probability plot of Anscombe residuals from the model. The plot is straight, indicating an adequate fit, from the 10th to the 99.9th percentile of the distribution. The cases in the bottom decile have all been examined: almost all of these correspond to trace values, for which the exact rainfall amount is unknown and which, in the fitting software, are replaced by their approximate conditional expectations under the fitted model as described by Chandler & Wheater (1998*b*). The net effect of this is to under-represent the number of very small values in the database, which explains the pattern seen at the bottom end of the plot. The curvature at the top end may represent an error in the gamma assumption in this upper tail, or it may be related to errors in the normal approximation to Anscombe residuals. However, since the curvature only becomes apparent around the 99.9% point of the distribution, it is perhaps not a cause for concern — it will only affect events whose occurrence probabilities are less than 1 in 1000.

Main site effects for basic amounts model

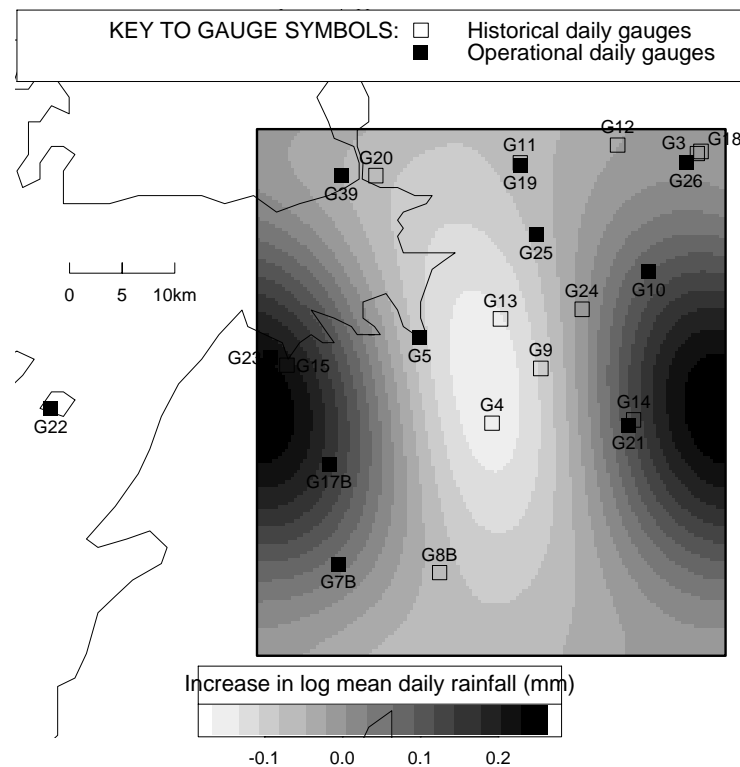


Figure 4.8: Main fitted site effects for basic rainfall amounts model.

The seasonal and annual structure of Pearson residuals from this model is shown in figure 4.10. There is no seasonal structure in the residual plots. Most of the annual mean residuals lie within the 95% confidence bands; however, there is a clear increasing trend through the 1970s and 1980s. From this plot, it might appear that either a cyclical trend function with a minimum around 1970, or a linear increase beginning around 1970, should be added to the model to capture this effect. This has already been confirmed by the log-likelihoods presented in table 4.5. The annual root mean squares of the residuals are all close to their expected value.

There is quite a lot of structure in the mean residuals from each site, for this model — 11 out of 21 sites have mean residuals which are significantly different from zero. Those with significant positive residuals are G4, G19, G23, G39 and G17B, and those with significant negative residuals are G5, G12, G13, G15, G18 and G8B. Referring to figure 4.1, there is little systematic pattern here, so the indication may be of some other effect (such as a trend) which has been omitted from the model.

For the validation dataset, this model overpredicts significantly (the mean Pearson residual is -0.063, with a standard error of 0.022). Despite this, its predictive performance in terms of root mean squared error does not suffer — in fact, \tilde{R}^2 as defined at equation (4.6) is 0.074 over the validation dataset, so that the model effectively explains 7.4% of the variability here, despite its bias. There is some seasonal pattern to the residuals for the validation set — a tendency for negative residuals between March and September and positive residuals




Figure 4.9: Normal probability plot of Anscombe residuals from basic rainfall amounts model.

elsewhere, although the only months in which the means are significantly different from zero are August and December. Of the 10 sites active during the validation period, only G7B and G17B had mean residuals which were significantly different from zero.

Fitted amounts model — linear trend after changepoint year: The extra predictors in this model, over and above the basic model just described, are the trend function itself, and its interactions with seasonal effects (no other significant interactions were found). As with the occurrence model, this indicates that an effect of the trend may be to induce wetter winters and drier summers.

This model explains 6.0% of the variance in daily rainfalls for the fitting data, which is a small improvement over the basic model. The normal probability plot of Anscombe residuals looks almost identical to that for the basic model in figure 4.9. The seasonal and annual structure of residuals is shown in figure 4.11. As previously, there is no seasonal structure here. The annual plots show that the inclusion of a trend has gone some way towards removing the structure apparent in figure 4.10, although there is now a suggestion of alternating blocks of low and high residuals, and one single year (1947) with a mean residual which is significantly different from zero. There is also a suggestion of a decline in the annual root mean squared residuals for this model during the 1990s.

The site-by-site analysis for this model remains slightly disappointing — residuals at sites G4, G19, G23 and G17B still have significant positive means, and those at G5, G12, G18, G21 and G8B have significant negative means (although at almost all the sites, the means are closer to zero than they were for the basic model). As mentioned above, there appears to be little systematic organisation of the locations of sites with positive and negative residuals, so it is difficult to see how the problem could be rectified.

For the validation data, this model actually performs worse than the basic model. The main problem is that it systematically overpredicts rainfall amounts (from the results of the basic

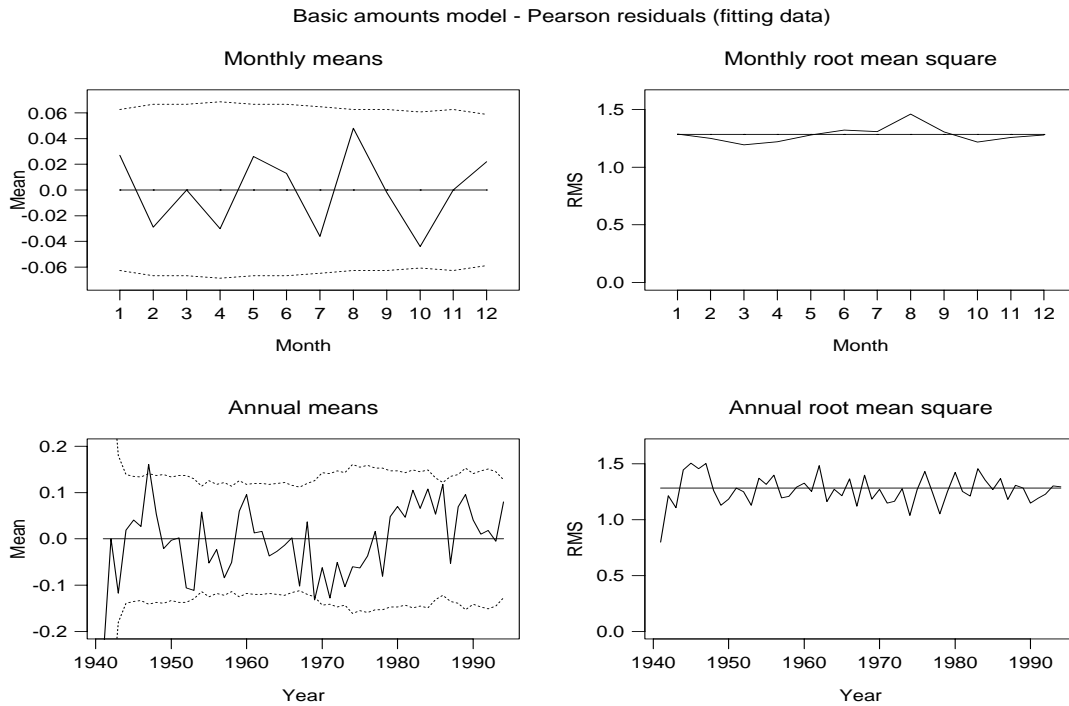


Figure 4.10: Seasonal and annual structure of Pearson residuals from basic amounts model. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

model, this is to be expected, since the basic model itself overpredicted, and the extended model incorporates an increasing trend in rainfall amounts). In terms of Mean Squared Error, the new model explains 7.0% of the variability in daily rainfall amounts for the validation data (compared with 7.4% for the basic model). Apart from this systematic overprediction, the basic structure of the residuals for the validation dataset is much the same as that for the basic model.

Simulation of the fitted models

So far, the checks on the fitted models have focused upon the adequacy of the forecast conditional probability distributions for each day, given preceding days' observations. The main area of concern appears to be in the representation of site effects, although there is also some indication from the validation exercise that any extrapolation of fitted trend functions outside the range of available data should be treated with caution.

For hydrological applications, it is necessary to focus additionally upon the unconditional performance of a daily model, since it is possible that small errors in the specification of conditional probability distributions may cumulate when the model is simulated over a long period of time. Moreover, the combined performance of the occurrence and amounts models taken together needs to be assessed. This can be achieved by simulating the fitted models together, and carrying out appropriate comparisons between simulated and observed sequences.

To simulate the fitted models, spatial dependence structures need to be estimated. As detailed

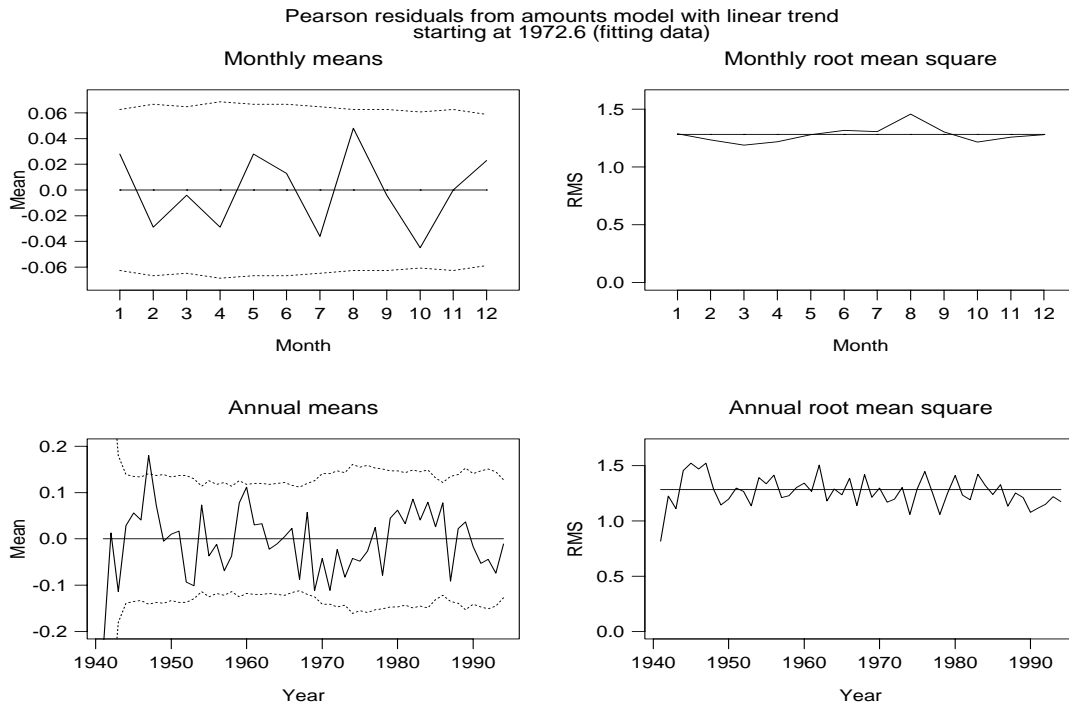


Figure 4.11: Seasonal and annual structure of Pearson residuals from amounts model incorporating linear trend starting at 1972.6. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

previously, for the amounts model spatial dependence is incorporated via the correlation structure of the Anscombe residuals, and a constant correlation between all pairs of sites has been found to be adequate. This is estimated by averaging the inter-site correlations for all site pairs (appropriately weighted to account for different numbers of contributing observations from each pair). For amounts model 3 (incorporating a linear trend after 1972.6), the estimated inter-site correlation is 0.7518.

Spatial dependence in the occurrence model is incorporated via the hidden weather state model defined at equation (4.13). The estimate of $\ln a$ in the occurrence model 4 (incorporating a cyclical trend) is 5.4064.

An initial exercise in checking the simulation performance involves assessing the ability of the model to reproduce winter rainfall amounts, since according to figure 4.2 these are closely associated with flooding events. To investigate this, a subset of sites (codes G10, G19, G21, G23, G25, G39 and G17B) was identified, all of which were operational and yielding reliable data throughout the 1990s. 1000 simulations of the fitted models were run, from December 1989 to February 1997: all available data from December 1989 were used to initialise each run, and then the models were allowed to simulate freely from January 1990 onwards. The winter rainfall, averaged over all 7 sites, for each year of each realisation was calculated so that 1000 sequences of simulated winter rainfall were obtained. The percentiles of the winter rainfall distributions thus obtained were plotted for each year: the results are shown in figure 4.12. For comparison, an estimate of the actual winter rainfalls, averaged over these sites, is required: since these sites occasionally have missing data, this has been obtained by simulating the model 1000 times in ‘interpolation mode’ (so that it merely generates missing values conditional upon all available data), and taking the median of the resulting winter rainfalls.

Simulated and observed winter rainfalls, 1990–1997




Figure 4.12: Simulated versus observed winter rainfalls for the Galway study area, 1990–1997. ‘Winter’ is defined as the months of January and February, together with December the previous year. The simulated distributions are obtained from 1000 realisations. The ‘observations’ series is the median of 1000 realisations of the model running in interpolation mode to fill in any missing data.

Figure 4.12 shows that the models do a reasonable job at reproducing the variability in winter rainfall distributions. This gives us some confidence that, despite the problems with site effects highlighted in the basic residual analyses, the models may still provide accurate simulations of catchment average rainfall over long time periods. From the figure, the 1990 and 1994 events both have exceedance probabilities between 0.01 and 0.05; that for the 1995 event falls just below the 0.01 mark. These probabilities cannot be quoted as conventional return periods due to the fact that we are working with nonstationary models: moreover, in view of uncertainty surrounding the adequacy of the models, they should perhaps not be interpreted too literally. However, for comparison a simple analysis assuming stationarity allocates a return period of around 950 years to the 1995 event (OPW 1998). We conclude that winter floods may be much less rare, in the presence of such trend functions as have been identified, than they would be in a stationary climate.

A further analysis examines summary statistics for simulated daily rainfall records. Comparison of such statistics with observations must be made rather carefully, allowing both for different record lengths in the observed data and for the fact that the data themselves are regarded as a single realisation of an underlying stochastic process. The comparison presented here has been carried out as follows:

- Three sites (numbers G39, G10 and G7B) were selected for comparison. The selection criteria were:
 - All the sites had at least 30 years of overlapping record (see table 4.1), whose quality

was generally deemed to be reasonable.

- The performance of the sites in the residual analyses for the fitted models did not highlight any serious problems.
 - The sites are widely spread across the catchment.
- A single run of the model, in interpolation mode, was made from January 1966 to December 1995 (the period of complete years for which data are available from all three sites), to fill in missing values. Data from all other sites were made available for this run, to maximise the accuracy of the interpolation. This run was regarded as a complete observational series.
 - Using observed data from the three selected sites for December 1965 only to initialise simulation, 1000 runs of the model were made from December 1965 to December 1995.
 - Annual means, averaged over the three sites, were calculated for the observed series and for each of the simulations. For each simulation, the mean squared difference (MSD) between simulated and observed annual means was calculated and this was used to identify the single simulation which followed the observations most closely in terms of overall trends. Daily statistics from this simulation were then compared with those of the observations.

The aim of this procedure is to obtain a simulation which can be regarded as comparable with observations on an annual timescale, in order to provide a viable starting point for comparison of the daily statistics. An alternative would have been to compute daily statistics for each of the 1000 simulated realisations, in order to derive a range of values for each statistic — however, this alternative is extremely labour-intensive. To give an idea of the variability between simulations, figure 4.13 presents plots of the observed annual rainfall series, together with series obtained from three simulations: the one used for comparison, the one with the largest MSD and the one with the median MSD. All three of the simulated annual series here look qualitatively similar, so the chosen simulation does not appear to be in any way unrepresentative.

Figures 4.14, 4.15, and 4.16 show a selection of observed and simulated summary statistics for the marginal time series at each site, for each month of the year. Although the comparison is with a single simulated sequence only, there is broad agreement between observation and simulation for most of these statistics at all sites. There is a tendency for the simulation to give a slightly higher mean and standard deviation on wet days than the observation, and the month of March particularly stands out in the standard deviation plots. However, the proportion of wet days in each month, and the autocorrelation at lags 1,2 and 3 at each site, are well reproduced. Given that this simulation represents only one realisation of the model, the performance indicated in these figures appears reasonable.

Figure 4.17 shows the performance relating to summary statistics for each pair of sites. The statistics for each pair are the proportion of days when both sites are wet, the overall lag zero cross-correlation, and the lag zero cross-correlation taken only on days when both sites experienced rain. It is seen that the simulation slightly under-represents the proportion of days when both sites are wet, for all three pairs of sites: this translates through into a slightly reduced inter-site cross-correlation. However, the inter-site correlation when both sites are wet is well reproduced. This indicates that the correlation structure of the amounts model is performing well, but that the dependence structure of the occurrence model may not produce enough days when all sites are simultaneously wet. This is either to do with the parameter estimation procedure for the model, or

Observed and simulated annual rainfalls, 3 sites, 1966-1995




Figure 4.13: Observed and simulated annual rainfall series, averaged over sites G10, G39 and G7B, 1966–1995. 1000 simulated series were generated and ranked according to the mean squared difference (MSD) between simulation and observation. The simulations shown are those with the smallest, largest and median MSD.

to do with the fact that the model structure does not allow the probability of rain to be arbitrarily close to 1 simultaneously at all sites (recall the discussion in section 4.1.5).

This completes the analysis for the Galway Bay study. The broad conclusions are as follows:

- The fitted models are interpretable, and provide insights into the mechanics of the rainfall process in the area. The extent of spatial and temporal nonstationarity can be quantified: from the analyses presented here, both are present to a significant degree.
- Simulations from the model reproduce various statistics of the rainfall reasonably, at daily and 3-monthly timescales. The main deficiency appears to be in reproducing the probability of simultaneous rainfall occurrence at more than 1 site.
- It is difficult to achieve good model performance at all sites individually, although the extent to which this reflects problems with data quality is not known. For practical purposes the problem may not be too serious, since processes such as runoff represent some sort of averaging over a catchment: overprediction at one site will tend to be compensated for by underprediction at another.
- Out-of-sample model validation exercises indicated that extrapolation of trend functions outside the range of available data should be treated with caution. One problem is the sensitivity of fitted trend functions to site effects: as a result of different periods of record from different sites, a small change in fitted site effects can have a marked impact upon the ‘optimal’ trend

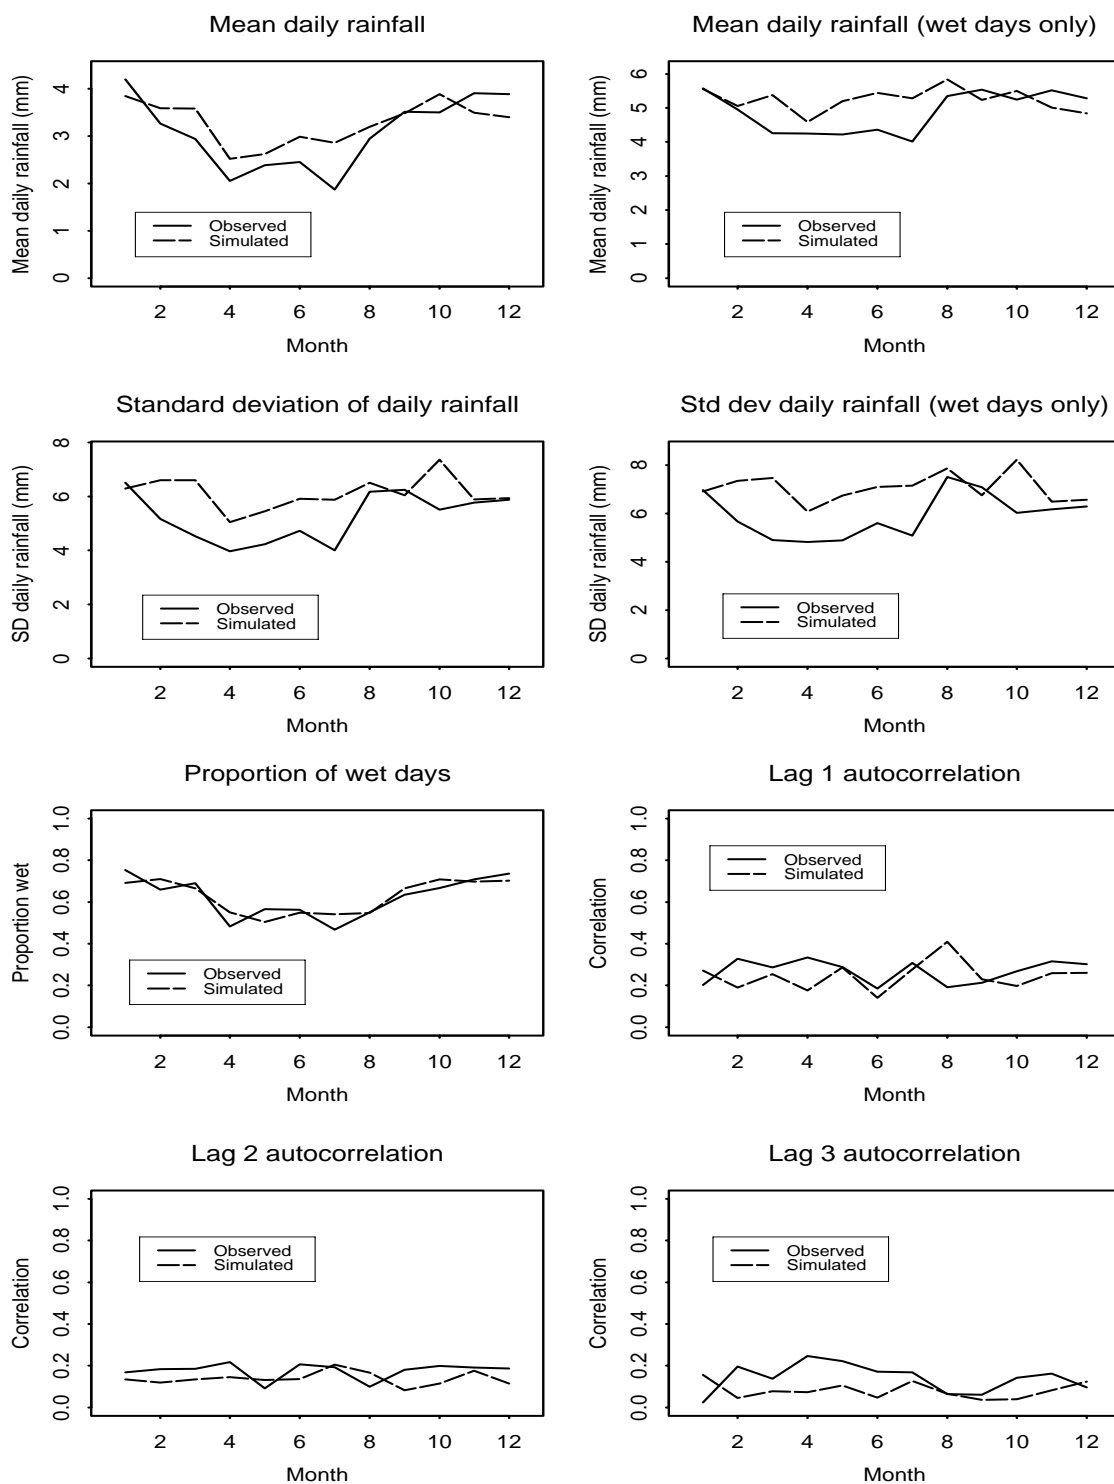


Figure 4.14: Observed and simulated statistics for daily rainfall at site G10, 1966-1995

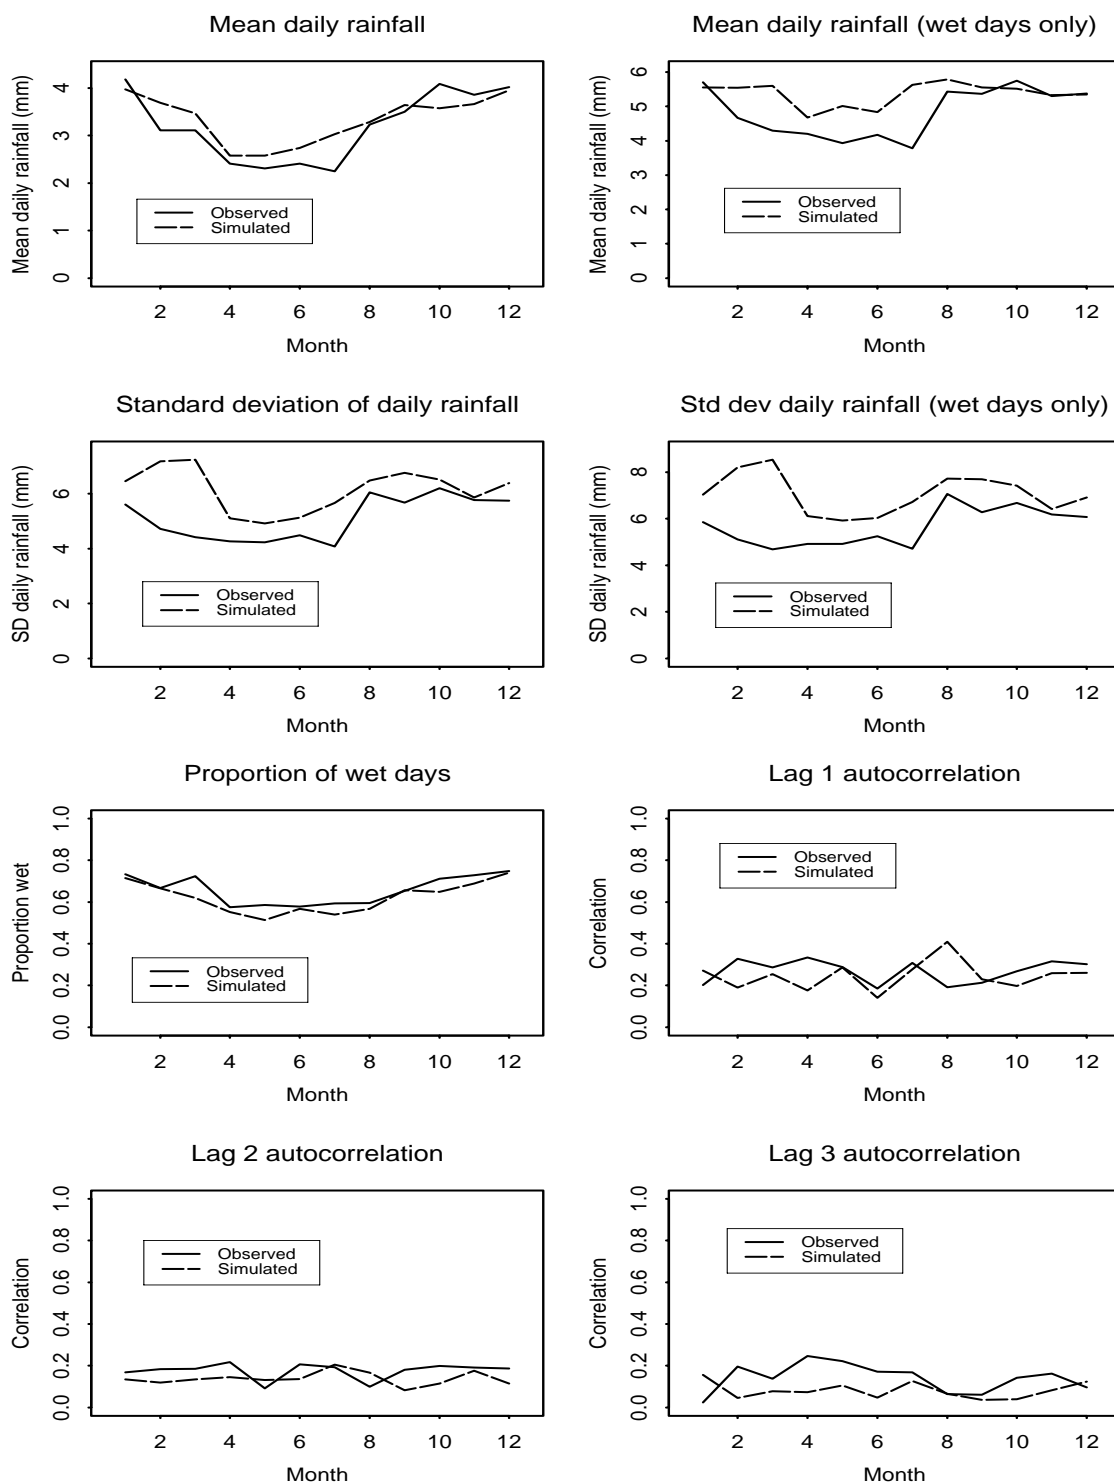


Figure 4.15: Observed and simulated statistics for daily rainfall at site G39, 1966-1995

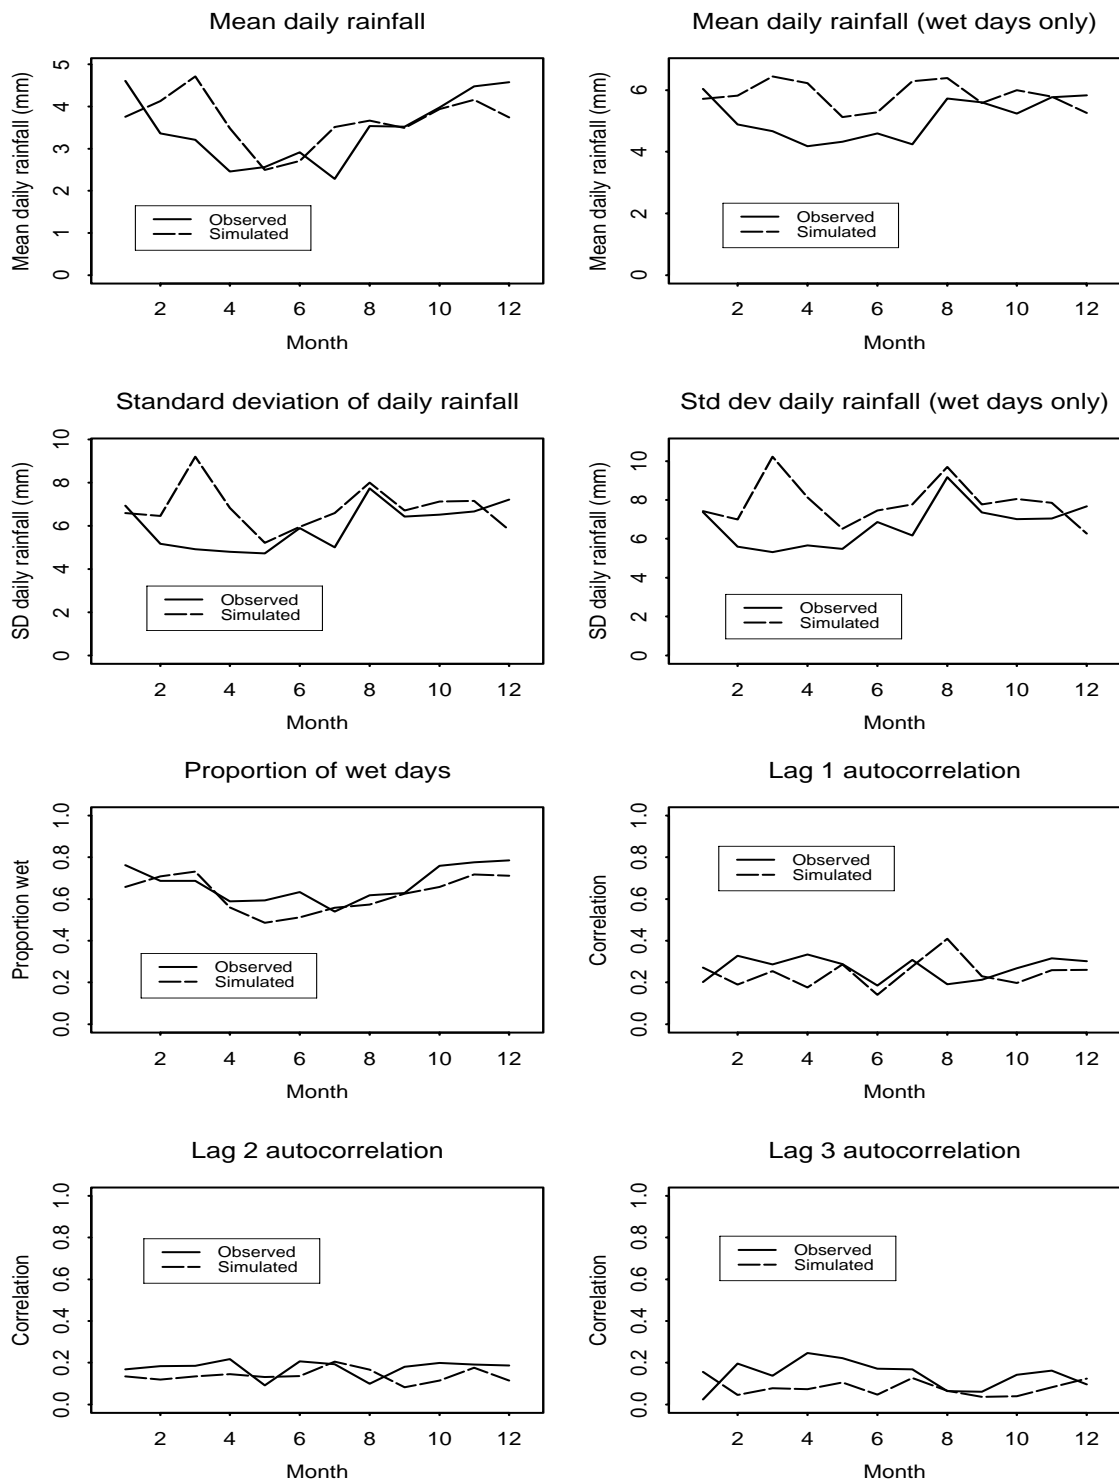


Figure 4.16: Observed and simulated statistics for daily rainfall at site G7B, 1966-1995

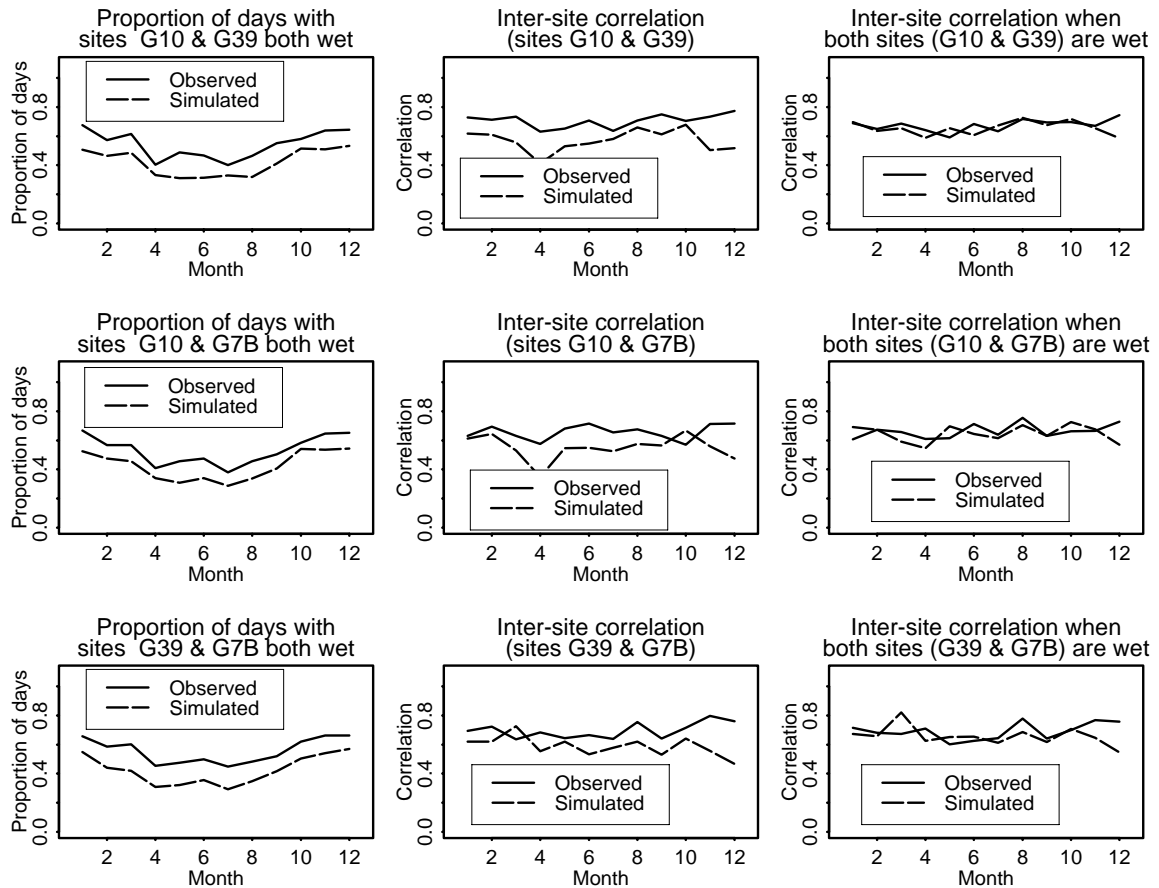


Figure 4.17: Observed and simulated statistics for daily rainfall at pairs of sites, 1966–1995.

function chosen. Another problem is that any parametric trend function fitted to the available data can at best be a plausible approximation to any underlying trend over this range. Recognizing these problems, a prudent analysis of the effects of changing rainfall patterns in the West of Ireland would include a variety of scenarios which appear well supported by the available data, and would recognise the inherent uncertainty associated with this type of problem.

4.2.2 The Brue

The second case study relates to the Brue catchment in South-West England, which is also considered elsewhere in this report. The data available are as follows:

- Data from 49 0.2mm tipping-bucket raingauges located in the catchment of the Brue itself. The earliest data from this network are from September 1993; data are available until August 1998. The data have been thoroughly checked and quality-controlled by the Institute of Hydrology and by the authors — for details of the procedures used, see Wheeler, Isham, Cox, Chandler, Kakou, Northrop, Oh, Onof & Rodriguez-Iturbe (1999). Dubious periods of data have been removed from each individual gauge's record.

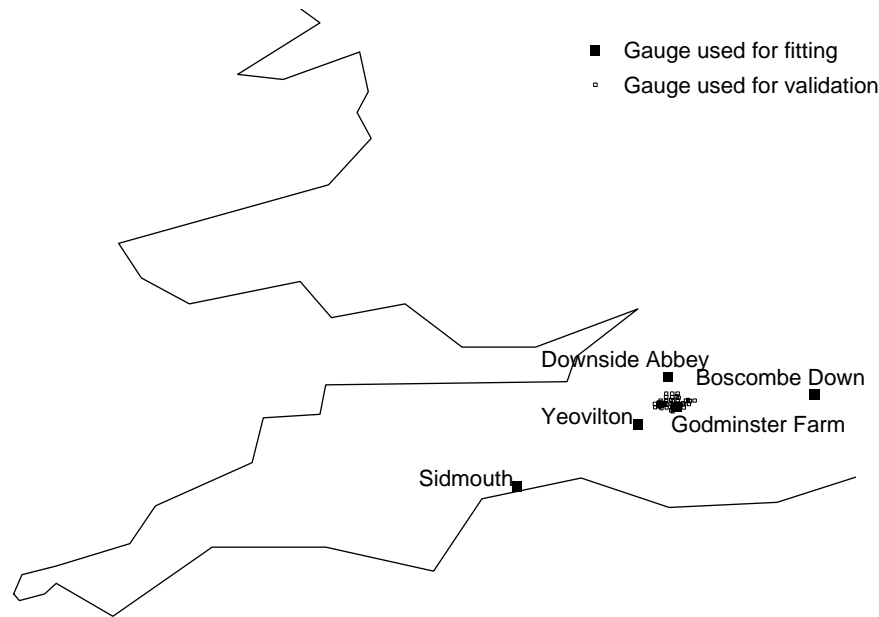


Figure 4.18: Locations of raingauges used in Generalised Linear Modelling exercise for the Brue catchment.

- Data from 3 daily gauges located at Sidmouth, Yeovilton and Downside Abbey. All of these gauges are outside the Brue catchment — in the case of Sidmouth, considerably so. The records run from 1980 to 1999.
- Data from a single recording raingauge at Boscombe Down, to the East of the catchment. These data have been supplied, and quality controlled, by the UK Meteorological Office. The record runs from 1980 to 1999 — however, years 1997 and 1998 contain unrealistically high values which are inconsistent with records from all of the other gauges, and have therefore been discarded for the analyses reported here.

A sketch map of South-West England and Wales, showing the locations of the gauges, is given in figure 4.18.

The data from the 49 gauges within the Brue catchment provide a unique opportunity to study the fine-scale structure of rainfall. However, as a result of their limited temporal extent they are unlikely to provide any useful information regarding trends in rainfall patterns. For this reason, we have chosen to use data from the four gauges outside the catchment, together with a single gauge inside the catchment, to fit models and then use the remaining data from within the catchment for validation purposes. In this way, the validation exercise will give a useful indication of how the GLM methodology might be expected to perform when very little data are available from within a catchment, and only sparse daily records are available nearby.

The gauge selected from the Brue network, for inclusion in the fitting dataset, is that at Godminster Farm: this was selected because there were few problems in its 5 years of record. Figure 4.18 shows the locations of the fitting and validation sites. The five gauges in the fitting set correspond to the locations used for extreme value analyses of the continuous space-time model, in

section 3.3 of this report.

The fitting dataset here is rather small compared to those in the other two case studies — around 26,000 observations. One consequence of this is that small effects, which might show up as statistically significant with a larger dataset, do not appear significant here. A further aspect is that with only 5 sites available for model fitting, we cannot hope to obtain complex representations of regional effects. As a result, the fitted models are simpler for this study than for either of the other two.

The occurrence and amounts models for this study were built up starting with ‘obvious’ predictors, such as seasonality and previous days’ rainfalls: extra predictors were then added and their usefulness assessed on the basis of increases in nominal log-likelihood, predictive performance and residual analysis. The final fitted models are now summarised — full details may be found in Appendix B.1.

Brue — rainfall occurrence model

The occurrence model for the Brue contains 13 predictors, as follows:

- A constant term.
- Site altitude.
- An annual cycle representing seasonality (sine and cosine terms).
- Indicators for rainfall occurrence up to 5 days previously.
- Indicators for rainfall occurrence on both the previous 2 days and on all 3 previous days.
- Interactions between the seasonal terms and the indicator for rainfall occurrence 1 day previously.

These are broadly similar to the predictors found for occurrence in the Galway study, with the exception of the detailed regional representation of site effects which cannot be carried out here due to lack of data.

There is nothing from the residual analysis of this model to suggest any obvious lack of fit anywhere. Tables 4.6 and 4.7 show the predicted performance, and figure 4.19 shows the seasonal and annual structure of model residuals. In the plot of annual means, there is possibly a suggestion of an increase towards the end of the record, but none of the means is significantly different from zero and it would be difficult to argue for the inclusion of a trend function on the basis of this plot.

We now report the results of the validation exercise using data from the dense raingauge network within the Brue catchment. Altitude information is not available for 4 of the 48 sites which were not used in fitting the model so, since site altitude is a predictor in the model, these sites were excluded from the validation.

In the validation data, the good predictive performance remains. The table of observed and expected proportions of wet days, grouped according to the forecast probability of rain, shows some apparent discrepancies (table 4.8). In fact, most of these are not statistically significant: the “independence” standard errors associated with each of the observed proportions need to be inflated, as

	Observed		% correct	
	Dry	Wet	Observed	Expected
Forecast dry	9035	3481	72.2	72.2
Forecast wet	4045	9617	70.4	70.4
OVERALL % CORRECT:			71.3	71.2

Table 4.6: Observed versus expected predictive performance for occurrence model for Brue catchment fitting data. ‘Forecast dry’ row corresponds to days when the predicted probability of rain was less than 0.5; ‘forecast wet’ corresponds to other days.

Forecast decile	1	2	3	4	5	6	7	8	9	10
Observed	0.000	0.172	0.243	0.355	0.440	0.560	0.661	0.749	0.809	0.000
Expected	0.000	0.182	0.240	0.346	0.450	0.547	0.661	0.752	0.810	0.000
<i>N</i> days	0	2632	5687	2373	1824	1622	4114	7001	925	0

Table 4.7: Observed versus expected proportions of days with rain, for data grouped according to forecast probability of rainfall occurrence (occurrence model for Brue catchment, fitting data).

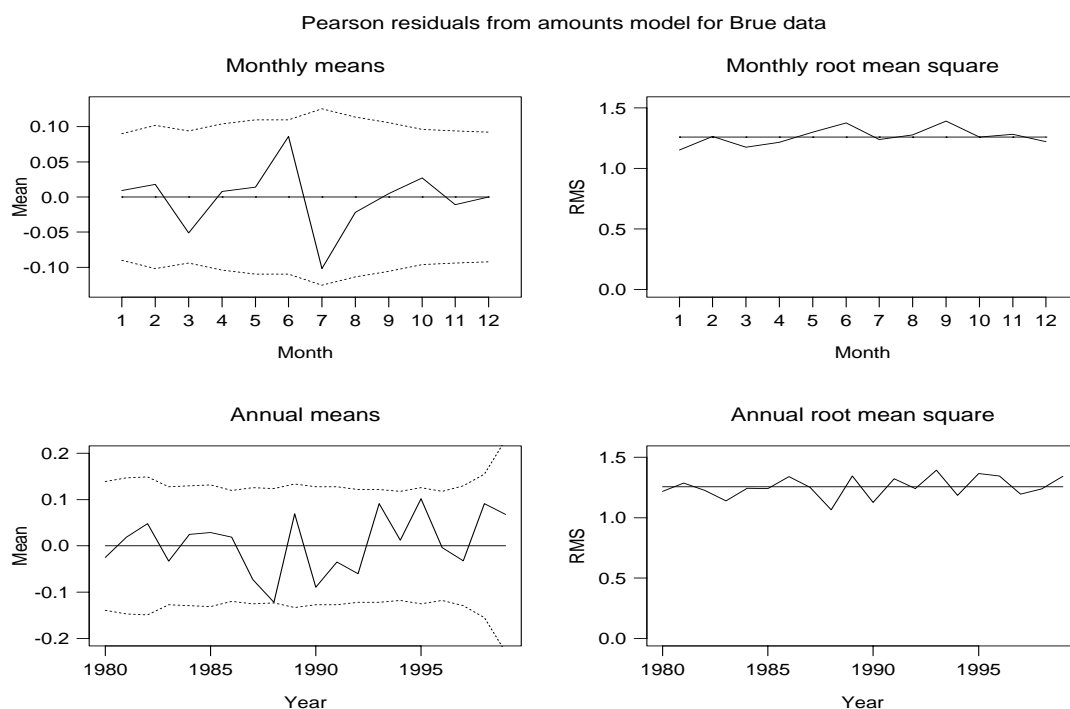


Figure 4.19: Seasonal and annual structure of Pearson residuals from occurrence model for Brue catchment. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

Forecast decile	1	2	3	4	5	6	7	8	9	10
Observed	0.000	0.135	0.212	0.412	0.523	0.597	0.653	0.767	0.854	0.000
Expected	0.000	0.184	0.240	0.345	0.449	0.540	0.659	0.756	0.806	0.000
<i>N</i> days	0	8305	11389	4287	4367	3540	7523	20039	2410	0

Table 4.8: Observed versus expected proportions of days with rain, for data grouped according to forecast probability of rainfall occurrence (occurrence model for Brue catchment, validation data).

described in section 4.1.4 previously, to allow for the strong spatial dependence within the validation dataset (a result of the large number of sites within a small area). The formula (4.11) in fact yields a value of 5.076 for the inflation factor with this dataset, and the only significant difference is in column 2 of the table. To illustrate the argument: column 4 in the table apparently shows the worst performance. Out of 4287 days in this group, 41.2% were observed to be wet instead of an expected 34.5%. If the true proportion in this group is 0.345 and all of the observations are independent, the associated standard error of the sample proportion will be $\sqrt{0.345 \times (1 - 0.345)/4287} = 0.0073$. Inflating this by a factor 5.076 yields a corrected standard error of 0.0369, so that 95% limits for the sample proportion are $0.345 \pm (1.96 \times 0.0369) = (0.273, 0.417)$. The observed proportion does fall within these limits.

The overall residual performance for the validation data is good: the mean Pearson residual is 0.009 with a standard error of 0.020, and the standard deviation of the Pearson residuals is 0.97. These agree well with the expected values of 0 and 1. There is, however, some seasonal structure to these residuals, with a tendency for negative values in the summer months and positive values in the winter. February and December both yield positive mean residuals which are significantly different from zero. Additionally, there appears to be annual structure to the residuals here — a monotonic decrease from 1994 to 1997. The annual mean residual for 1994 is 0.136, with an associated standard error of 0.042; that for 1997 is -0.163, with a standard error of 0.043. These are clearly highly significant, although there was no indication of this kind of structure in these years for the fitting data (figure 4.19).

The overall performance at individual sites in the validation dataset is remarkably good — out of 44 sites in the dataset, only 4 have mean residuals which are significantly different from zero. The means are positive in all 4 cases: in the worst case, the mean residual is 0.067, with an associated standard error of 0.026.

Brue — rainfall amounts model

The amounts model for the Brue contains 11 predictors, as follows:

- A constant term.
- Site altitude.
- An annual cycle representing seasonality (sine and cosine terms).
- $\text{Log}(1 + \text{rainfall } x \text{ days ago})$, for $x = 1, 2, 3, 4$.




Figure 4.20: Normal probability plot of Anscombe residuals from rainfall amounts model for Brue catchment.

- Indicator for rainfall occurrence on all 4 previous days.
- Interactions between the seasonal terms and $\log(1 + \text{previous day's rainfall})$.

The basic structure of this model is very similar to that of the occurrence model. The fitted model explains 3.0% of the variance in non-zero daily rainfall amounts. This is rather smaller than the corresponding figure for the Galway study. The reduction is due, at least in part, to the few sites used in the fitting — the variability due to differences between sites is simply not present in this dataset.

Figure 4.20 shows the normal probability plot of Anscombe residuals for this model. At the lower end, the shape is similar to that seen in the Galway study and, as before, indicates that very small rainfall amounts are not being recorded — this is not a matter for concern. There is a small amount of curvature in the upper tail, again starting around the 99.9% point of the distribution but this time in the opposite direction to that seen in the Galway study. As before, there should be few objections to a model which appears accurate for events as rare as 1 in 1000.

As with the occurrence model, there is no seasonal or annual structure to the residuals here (the plots are shown in figure 4.21). However, 4 of the 5 sites exhibit mean residuals which are significantly different from zero: Downside Abbey and Sidmouth have positive means, whereas those at Boscombe Down and Yeovilton are negative. Referring to figure 4.18, it is unlikely that this can be resolved in a constructive way on the basis of the available data.

For the validation sites, the model in fact overpredicts significantly on average (the average overprediction for daily rainfall is 0.5mm). This bias accounts for slightly more than half the reduction in \tilde{R}^2 , to 1.1% (from 3.0% for the fitting dataset) and, as expected, dominates the residual analyses broken down by month, year and site.

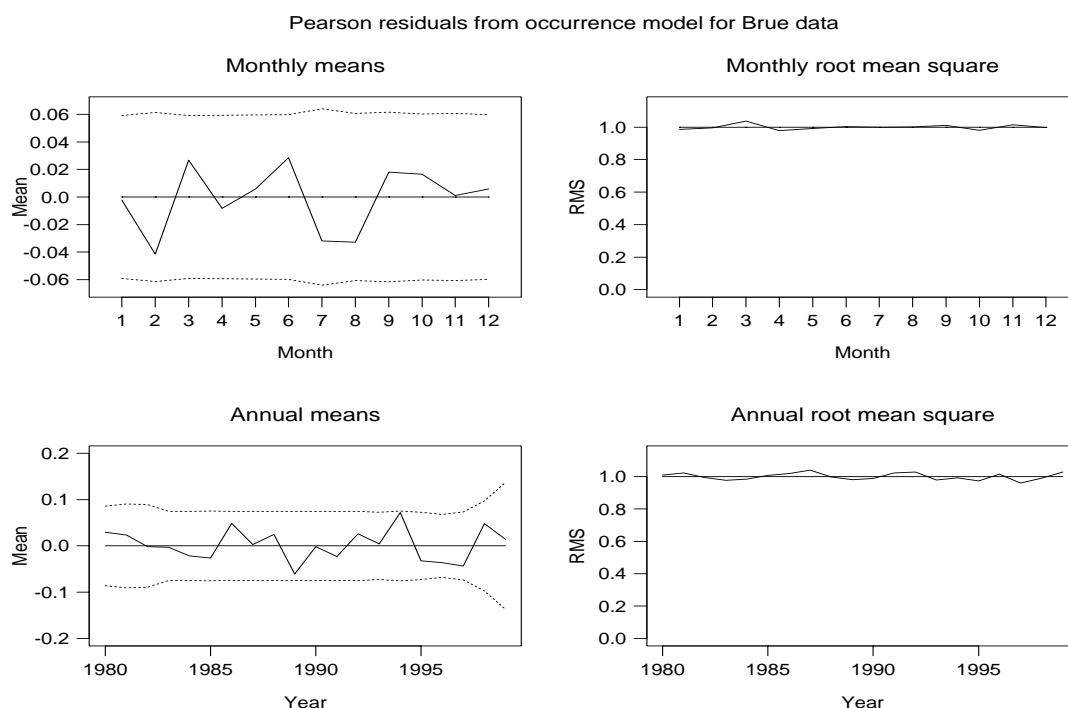


Figure 4.21: Seasonal and annual structure of Pearson residuals from amounts model for Brue catchment. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

Brue case study — summary

The results of this case study are of particular interest to a practitioner who is in the situation of having to model a catchment with extremely limited data. The validation exercises for both occurrence and amounts models indicated that it may be difficult to reproduce rainfall structure over a catchment, unless data are available from quite nearby. For the occurrence model, both seasonal and annual structure were present in the validation residuals, for which there was no evidence at all from the fitting sites. If anything, the fitting sites indicated an increase in residuals towards the end of the record, whereas the validation sites showed exactly the opposite. For the amounts model, the main problem was a systematic bias in the validation. It is worth noting that in the fitting data, the single gauge from within the catchment did yield a negative mean residual for this model, although it was not significantly different from zero. In retrospect, perhaps a prudent approach would have been to add an indicator variable to the model, taking the value 1 if a site lies within the catchment and 0 otherwise. The effect of this would be to adjust all of the fits from within the catchment upwards, which would have solved most of the problems with the validation for this model.

The other point to note from this case study is the lack of evidence for long-term trends. It should be noted, however, that the data are of limited temporal extent. Even the longer records only go back to 1980, and this may not be long enough to identify any long-term structure.

For this case study, the fitted models have not been simulated to compare observed and simulated rainfall summary statistics. However, the simulation exercise for the Galway study indicated that at locations where the models fit well (as assessed using residual analyses), properties of sim-

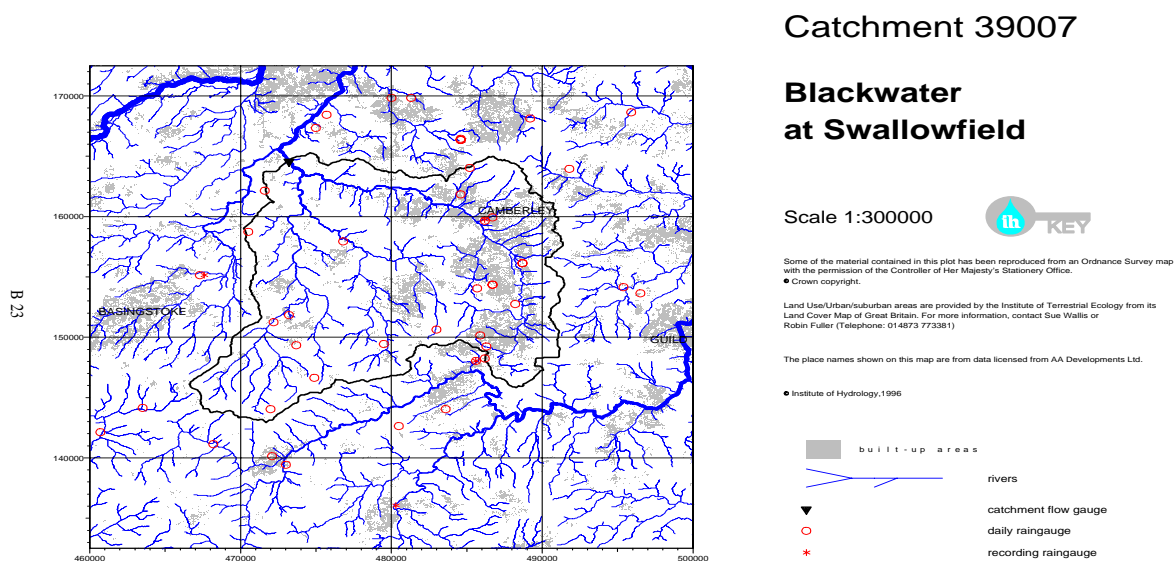


Figure 4.22: Map of Blackwater catchment, showing locations of raingauges.

ulated data agree well with those of observations. Such a conclusion seems intuitively reasonable, in view of the variety of residual checks which are carried out at the model fitting stage: we would therefore expect to obtain similar results from a simulation exercise applied to this case study.

4.2.3 The Blackwater

The final case study to illustrate the GLM methodology comes from the Blackwater, a tributary of the Thames located between Guildford and Basingstoke (figure 4.22). The catchment itself covers an area of some $20 \times 20\text{km}^2$. There is some topographic variation within the catchment: the North Downs run roughly East-West across its Southern boundary. Data are available from a network of 44 gauges in and around the catchment. 4 of these gauges provide subdaily data: the remainder are daily gauges. The lowest gauge is at 37m above sea level, and the highest is at 183m. The earliest records start in 1961; data are available to 1997. The dataset is large, containing some 350,000 observations.

The analysis reported in this section has been carried out by the Institute of Hydrology with the authors' assistance, by way of demonstrating the GLM fitting software. This has limited the time available to work on this dataset. As a result, there has not been time to perform any validation analyses, or to carry out many detailed checks of the fitted models. However, this example does indicate what can be achieved using the GLM methodology and software within a rather short space of time.

The basic approach to the modelling was the same as in the Galway study: start by including obvious predictors representing seasonality and temporal dependence, and then examine residual plots for guidance on how to parameterise regional effects and to look for evidence of trends. One

attractive feature of this dataset is the large number of sites available, which allows us to incorporate rather detailed representations of regional effects.

A summary of each of the final fitted models (occurrence and amounts) is now given. Full details may be found in Appendix B.1.

Blackwater — rainfall occurrence model

The final rainfall occurrence model for the Blackwater contains 38 predictors. The basic structure is as follows:

- A constant term.
- Site effects, incorporated via sine and cosine terms at the first 2 Fourier frequencies in an orthogonal series representation, over a rectangle from Eastings 460000 to 500000 and from Northings 138000 to 170000 (see figure 4.22 for co-ordinate system). Site altitude is not a statistically significant predictor.
- A cyclical trend function. The estimated cycle length is 51.8 years (nominal standard error 5.0 years) with a maximum at 1982.5 (nominal standard error 0.3).
- Seasonality, represented by whole-year and half-year cycles.
- Temporal dependence, represented by rainfall occurrence indicators for 1, 2 and 3 days previously.
- Temporal persistence, represented by an indicator for rainfall occurrence on each of the previous 3 days.
- Interactions between the trend and seasonal effects.
- Interactions between seasonal and temporal dependence effects.

These predictors broadly fit the pattern of the other two case studies: notice in particular that there is compelling evidence for a trend in rainfall occurrence and that this has significant interactions with seasonal effects.

A map showing the fitted site effects for this model is given in figure 4.23. The magnitude of the effects is generally quite small (compared, for example, with the results for the Galway Bay study). Where the map does show large effects (for example, significant enhancement to the South of gauges 32 and 33, and a corresponding reduction to the North of 25 and 26), there are few sites, and this may be an artefact of the fitting procedure, which puts the best possible smooth surface through the available gauge locations. There are no obvious physical features in the map of the area (figure 4.22) which would suggest any major regional variability.

The predictive performance of the fitted model is good. Tables 4.9 and 4.10 show the results. The only slight discrepancy between observed and expected performance is for days when the predicted probability of rain is between 0.1 and 0.2 (table 4.10). Here, the observed proportion of wet days is 0.214, which is significantly different from the expected proportion of 0.191 (the standard error for the observed proportion is around 0.0035). However, although this is formally

Fitted site effects for occurrence model, Blackwater catchment




Figure 4.23: Main fitted site effects for rainfall occurrence model, Blackwater catchment (also showing locations and numbers of raingauges).

significant in a statistical sense, the magnitude of the discrepancy is rather small and its practical consequences will be negligible.

Figure 4.24 shows the seasonal and annual structure of residuals from the fitted model. The seasonal structure of the rainfall has clearly been well captured by the model. 4 out of 37 mean annual residuals fall outside the 95% confidence bands: however, none of these is far outside and there does not seem to be any systematic overall structure to the plot, so this is not a major cause for concern. There is some suggestion that the annual root mean square was slightly greater than its expected value until around 1970: however, the effect is small, and does not necessarily represent a problem with the fitted trend function: it may equally well arise from a problem with some sites which were active before 1970.

The residual analysis for sites tells much the same story as that for the Galway study: around half the sites have mean residuals which are significantly different from zero. Those with significant negative residuals are numbers 9, 10, 13, 14, 19, 22, 29, 30, 32, 38, 39, 48 and 52; those with significant positive residuals are 16, 17, 20, 21, 33, 34, 37, 44, 49, 50 and 51. Referring to figure 4.23, most of these sites are towards the edge of the site network: there is an area towards the centre of the catchment (roughly corresponding to Eastings 470–485 and Northings 138–160) where there are few problems. Although there does appear to be structure regarding the locations of these sites,

	Observed		% correct	
	Dry	Wet	Observed	Expected
Forecast dry	115526	49564	70.0	70.0
Forecast wet	54968	132117	70.6	70.7
OVERALL % CORRECT:			70.3	70.4

Table 4.9: Observed versus expected predictive performance for occurrence model fitted to Blackwater catchment data. ‘Forecast dry’ row corresponds to days when the predicted probability of rain was less than 0.5; ‘forecast wet’ corresponds to other days.

Forecast decile	1	2	3	4	5	6	7	8	9	10
Observed	0.000	0.214	0.239	0.354	0.440	0.554	0.654	0.752	0.799	0.000
Expected	0.000	0.191	0.243	0.352	0.441	0.558	0.659	0.748	0.808	0.000
<i>N</i> days	0	11794	83602	42127	27567	14584	61468	103038	7995	0

Table 4.10: Observed versus expected proportions of days with rain, for data grouped according to forecast probability of rainfall occurrence (occurrence model for Blackwater catchment).

this does not appear carry through to their signs, so it would be difficult to rectify the problem in any systematic way. The indication may be of very local effects upon rainfall (relating to gauge placement, for example), since the fitted surface in figure 4.23 is already quite detailed, providing resolution of features down to a scale of around 5km.

Blackwater — rainfall amounts model

The predictors in the fitted rainfall amounts model for this catchment are more or less the same as for the occurrence model, except that temporal dependence is represented by $\log(1+\text{rainfall } x \text{ days ago})$ for $x = 1, 2, 3$ instead of indicator variables for rainfall occurrence on each of the previous 3 days. The amounts model also indicated a cyclical trend: the estimated cycle length was 64.8 years (nominal standard error 13.6 years) with a cycle maximum at 1955.1 (standard error 6.5). It is perhaps worth noting that the difference in the estimated cycle lengths for this model and for the occurrence model is not statistically significant (the observed difference is $64.8 - 51.8 = 13$, and the nominal standard error of this difference is approximately $\sqrt{5.0^2 + 13.6^2} = 14.5$). Moreover, the estimates of the trend parameters for this model would give a cycle *minimum* at 1987.5: the estimate of the maximum for the occurrence model was 1982.5. Taking into account the standard errors of all the estimates, it is therefore quite conceivable that the underlying trends for both amounts and occurrence in this area are the same, but that they work in opposite directions (in years of enhanced rainfall occurrence, amounts are reduced and vice versa).

Figure 4.25 shows a map of fitted site effects for this model. The structure is markedly different from that of the occurrence model for this catchment, and is dominated by a band of enhanced rainfall amounts running from North-East to South-West. As before, features such as that to the South of gauges 32 and 33, and that to the North of 48 and 49, should be interpreted with caution,

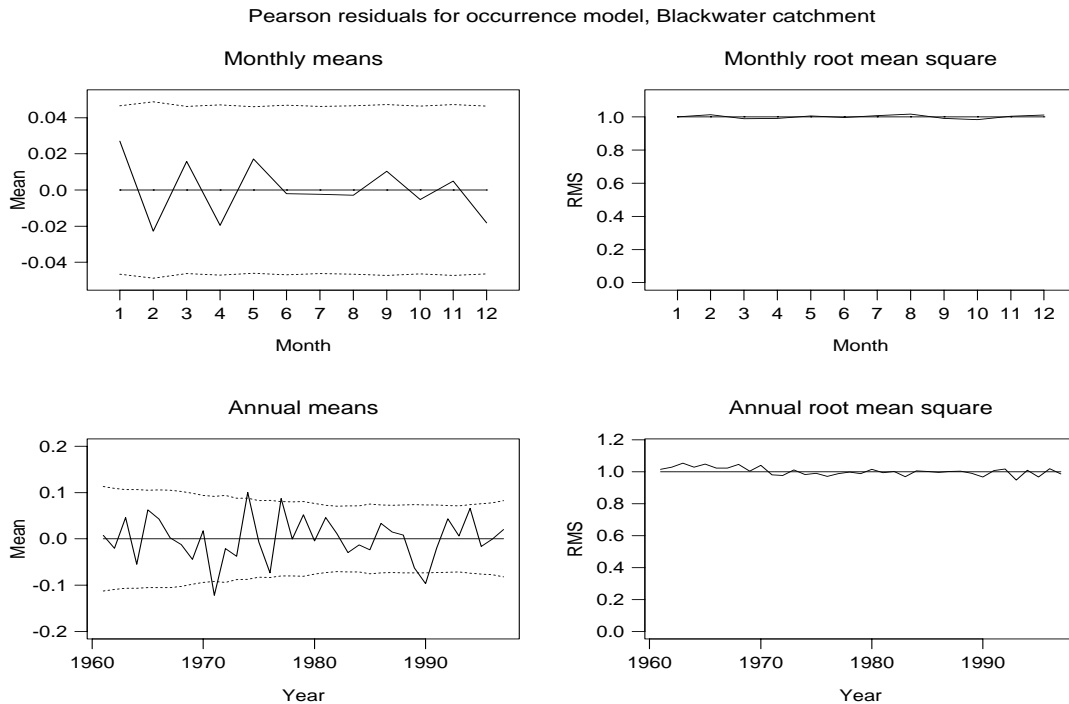


Figure 4.24: Seasonal and annual structure of Pearson residuals from occurrence model for Blackwater catchment. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

owing to scarcity of data.

The fitted model explains 3.3% of the variance in non-zero daily rainfall amounts. Again, this is rather smaller than the corresponding figure for the Galway study. In this case the reduction is probably a consequence of the more homogenous topography in this catchment. A normal probability plot of Anscombe residuals from the fitted model is presented as figure 4.26. The shape is similar to that for the Galway study, although the curvature in the upper tail is increased for this catchment and now starts around the 99.5% point of the residual distribution. This is still probably adequate for many purposes (representing events with an exceedance probability of 1 in 200).

Figure 4.27 shows the seasonal and annual structure of residuals from this model. These plots all indicate that the model has captured these aspects of rainfall variability in the area well. The only suggestion of a pattern is in the annual root mean square, which appears to be systematically below its expected value prior to 1970. This period was also identified as anomalous in the corresponding plot for the occurrence model — again, it may represent a feature of sites which were operating at that time.

The residual analysis by site highlights the usual problems: sites 16, 17, 20, 21, 28, 33, 34, 37, 46, 49, 50 and 51 all have significant negative mean residuals, and 9, 10, 13, 14, 19, 22, 29, 30, 32, 38, 39, 48 and 52 all have significant positive means. This pattern is almost identical to that for the occurrence model, with the signs reversed. Similar comments therefore apply. However, the striking coincidence between the two patterns of significant residuals suggests that part of the problem may relate to data quality. Suppose, for example, that the observer at site 37 tends to

Fitted site effects for amounts model, Blackwater catchment




Figure 4.25: Main fitted site effects for rainfall amounts model, Blackwater catchment (also showing locations and numbers of raingauges).

record rainfall amounts of less than 0.1mm of rain as zero, but that the observer at site 38 records all amounts of less than 0.2mm as zero. In this case, site 37 would record more wet days than site 38, but the mean rainfall on wet days would be lower. This would give exactly the kind of residual behaviour for the two models which has been observed at these sites.

Summary of Blackwater case study

This example indicates that it is possible to obtain a reasonable model for daily rainfall with a couple of days' work, using the GLM methodology. However, even with a lot of raingauges which enable resolution of site effects down to scales of around 5km, it is still not possible to obtain good fits at all sites. This indicates either that there is systematic variability in rainfall at a very local scale, or that differences between observers and siting of individual raingauges may be additional factors which affect the recorded rainfall to a degree which can be detected using this methodology.




Figure 4.26: Normal probability plot of Anscombe residuals from rainfall amounts model for Black-water catchment.

4.3 Summary of GLM modelling

The GLM approach has been studied as a powerful and flexible tool for studying temporal and spatial nonstationarity in rainfall, which is suitable for application to daily raingauge data. The broad conclusions from the case studies above are as follows:

- The methodology is capable of detecting patterns in noisy daily rainfall records. Although these patterns tend to be weak, in all the case studies considered they have been interpretable and realistic. Moreover, the overall structure of the models fitted is similar in all cases.
- Daily rainfall sequences have high noise levels, so that fitted models will tend to explain only a small fraction of the observed variability. However, at long time scales even weak signals can produce dramatic effects in hydrological terms (as in the Galway study).
- The assumed form of the probability distributions for daily rainfall can be expected to hold up to around the 99.9% point.
- The methodology is relatively quick and simple to use, given the publically-available software detailed in Appendix B.2. It would be fairly quick for a practitioner to fit a model, so as to judge the extent of temporal and spatial nonstationarity in an area.
- Fitted models are capable of reproducing many summary statistics for daily rainfall sequences well. There is, however, a slight tendency for the variability in daily rainfall sequences to be over-represented, and for spatial dependence in rainfall occurrence to be under-represented. There is scope for further work in this area to verify the extent to which these problems may arise in areas other than western Ireland.

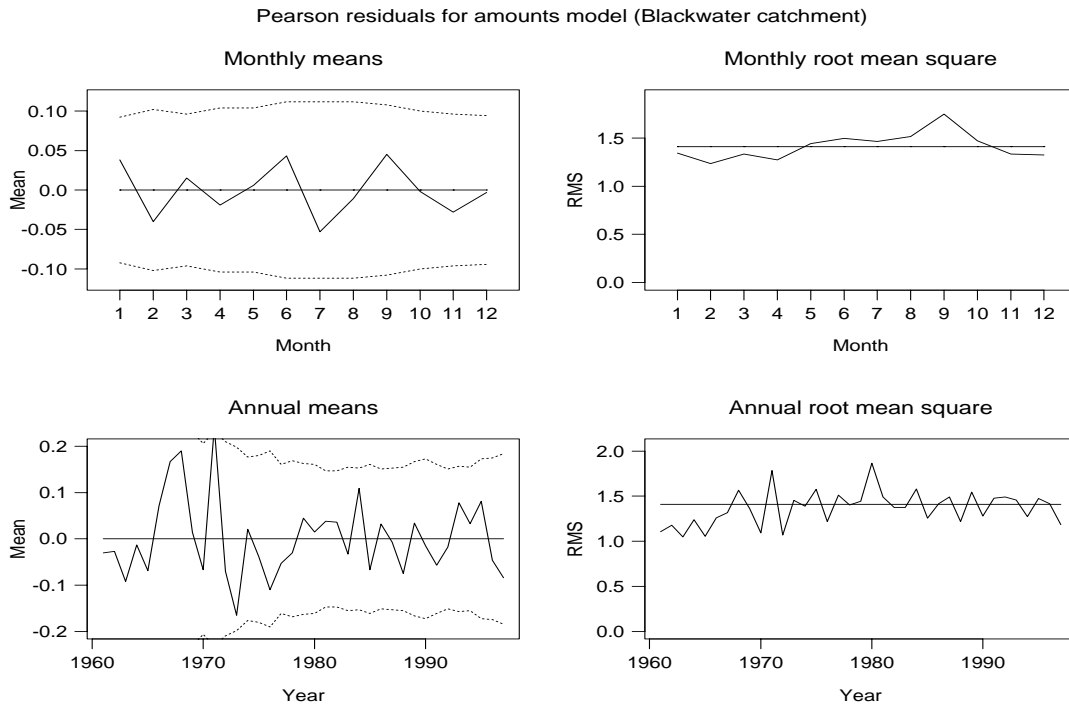


Figure 4.27: Seasonal and annual structure of Pearson residuals from amounts model for Blackwater catchment. Dotted lines on mean plots show approximate 95% confidence limits under the assumption that the model is correct, adjusted for spatial dependence between sites.

- It would be unwise to rely too heavily upon a single scenario regarding trends in climate, even if it appears well supported by this methodology over the range of time for which data are available. Validation exercises in these case studies have shown that such scenarios do not necessarily extrapolate well into the future. However, the methodology may enable an informed qualitative assessment of increasing or decreasing risk to be made, on the basis of available data.
- It appears difficult to capture exactly the rainfall patterns observed at individual sites. This is almost certainly a result of extremely local structure in observed rainfall fields. A natural, and plausible, way of overcoming this deficiency is to incorporate random effects into the model — see, for example, Fahrmeir & Tutz (1994).
- For catchments with limited data, but where data are available from outside the catchment, models should be structured in such a way as to exactly fit, as far as possible, the observed rainfall structure at gauges within the catchment. This can potentially be achieved through the use of appropriate indicator variables. In this case the role of the sites outside the catchment is to increase the precision of model parameter estimates, and to guide the choice of appropriate predictors for a model.

Chapter 5

Applications with limited data

In chapters 2 and 3, a continuous-simulation model was developed which can be calibrated from radar data and which is stochastically stationary in both space and time. Chapter 4 introduced an alternative strategy which uses daily rainfall data to quantify the extent of temporal and spatial nonstationarity in a catchment. In this chapter, we consider a set of problems associated with the use of limited data. First (sections 5.1 and 5.2), we return to the idea of using the GDSTM for continuous simulation, but consider the case when radar data are not available. In this case, the spatial-temporal rainfall model will have to be fitted using raingauge data only. Since we are looking at generating hourly sequences, we shall first consider the case when a set of hourly raingauge data is available. Secondly, we address the general problem of multivariate spatial-temporal disaggregation (section 5.3). This arises, for example, if daily raingauge data are available from several sites, but data at fine temporal resolution are limited (e.g. a single subdaily timeseries). We may wish simply to generate a consistent historical multivariate subdaily series, or we may wish to simulate subdaily series into the future in such a way as to respect the temporal and spatial nonstationarities embodied by a Generalized Linear Model for daily data. If GLMs are used for multivariate daily simulation, a consistent sub-daily time series at a single site is required for the spatial-temporal disaggregation; single-site temporal disaggregation is discussed in section 5.4,.

5.1 Spatial-temporal model implementation with raingauge data

5.1.1 Modelling strategy

The GDSTM presented in chapter 2 is a continuous space-time model and therefore is able to generate rainfall at a given point. The continuous time rainfall can then be used to obtain simulations of hourly rainfall. In calibrating the model from radar data, allowance must be made for the fact that the radar images can be regarded as instantaneous snapshots in time, but averages in space. Raingauge data, on the other hand, represent points in space but averages over time intervals. To fit the model to raingauge data therefore, properties of temporally but not spatially averaged data are required, in contrast with the fitting procedure for radar data in chapter 2. In this case, we have on the one hand a simplification since we do not need to integrate the point properties spatially, while on the other, they will need to be integrated over the hourly time-interval. The proportion dry at one point can be computed analytically for a given point in time, but the corresponding

expression for a time-interval is not easily usable for computations.

There are also differences in the spatial scales which are represented in the radar data and the set of raingauge data. Generally, the gauges will be too widely separated for the detection of such fine spatial features as are represented in radar data at $4km^2$ resolution. Typically, we might consider a raingauge network with a density of the order of one per $16km^2$ at the most. This will have implications for the model structure: the mean cell area as found in chapter 2 for Wardon Hill radar data is such that a typical cell would affect at most 2 of the raingauges. This is insufficient to allow for the estimation of the properties of cell areas. Consequently, the data provided by such a raingauge network must be viewed as allowing for the fit of a hierarchical spatial-temporal model from which the finer cellular structure is absent.

One obvious possibility is to fit the GDSTM to raingauge data, but regarding it as a model for storms clustered within rain events, rather than (as previously) for cells clustered within storms. By doing this, we are slightly changing the assumptions made about the arrival of rain events (in chapter 3, events arrived according to an alternating renewal process with Weibull durations, whereas if we regard the model as for storms clustered within rain events, rain events arrive according to a Poisson process) and about the distributions of storms inside events (in chapter 3, storms were uniformly distributed within an event whereas here they will be displaced from the center of an event according to a bivariate normal distribution). However, the practical consequences of these differences are likely to be minimal. In particular, since the gauged area will be of an order of magnitude which is small compared to that of the event, the assumption that storms are normally distributed around the centre of the event does not represent a significant difference in practice from the modelling strategy in chapter 2.

5.1.2 Model parameters and rainfall statistics

The GDSTM of chapter 2 is defined in terms of 11 independent parameters, which have the following physical or geometrical interpretations when applied to data from a raingauge network:

- Rate of event arrivals (number of events per km^2 per h): λ
- Mean storm duration (h): μ_D
- Mean event duration (h): μ_L
- Mean storm area (km^2): μ_A
- Mean event area (km^2): μ_S
- Mean number of storms per event: μ_C
- Mean storm intensity (mm/h): μ_X
- Component of storm/event velocity in the x direction (east) (km/h): V_x
- Component of storm/event velocity in the y direction (north) (km/h): V_y
- storm/event eccentricity: e
- storm/event orientation (in radians from east): θ

All other parameters are derived in terms of these 11 parameters using the equations of chapter 2.

In its initial formulation, the model assumes that the parameter values are constant during a “rainfall event” but vary among different events. However, in the application of this section we will assume that these parameters have constant values within each month, as the relatively small amounts of data available through raingauge measurements do not allow reliable parameter estimation for smaller time periods. The parameters can be estimated in terms of the first and second order rainfall statistics.

As we have seen above, rainfall statistics for raingauge data are only available for the temporally aggregated rainfall intensity process at the point of interest, that is, in terms of the discrete time process:

$$Y_i(x, y) = \int_{(i-1)h}^{ih} Y(x, y, t) dt \quad (5.1)$$

where $Y_i(x, y)$ denotes the mean rainfall intensity at the discrete time interval i with a fixed length h , and $Y(x, y, t)$ is the point instant rainfall intensity at point (x, y) at time t .

The mean rainfall intensity is independent of the time scale of aggregation h and is given by:

$$E[Y_i(x, y)] = \lambda \mu_D \mu_A \mu_C \mu_X \quad (5.2)$$

The second order properties depend on the scale of aggregation and are determined in terms of the point-instant covariance function:

$$c(u_x, u_y, \tau) := \text{Cov}[Y(x, y, t), Y(x + u_x, y + u_y, t + \tau)] \quad (5.3)$$

where u_x , u_y and τ are spatial and temporal displacements (lags) as in chapter 2. This is a complicated function of all the model parameters that can be evaluated only numerically. A Taylor series expansion for the expression is given by Northrop (1996).

Given the covariance function, the second order properties of the temporally aggregated process Y_i are given by ($k \geq 1$):

$$\text{Var}[Y_i(x, y)] = 2 \int_0^h (h - t) c(0, 0, t) dt \quad (5.4)$$

$$\text{Cov}[Y_i(x, y), Y_{i+k}(x + u_x, y + u_y)] = \int_{-h}^h (h - |t|) c(u_x, u_y, kh + t) dt \quad (5.5)$$

These statistics can only be evaluated by numerical integration (apart from special simple cases, e.g. when the velocity is 0). However, in practice we have found that the computational costs involved are small, and it is perfectly feasible to estimate model parameters by minimising the departures of theoretical statistics (as evaluated from the equations above) from historical rainfall series statistics.

5.1.3 Case study: Catchment and data

The Brue catchment located in South-West England is used as a case study area (figure 5.1). The catchment is equipped with 49 raingauges but only 8 of them were used to assess the ability of the methodology to perform with limited data. Five years of data were available, covering the period September 1993 to August 1998 (in fact the data of September 1993 were suspect in most of the raingauges and was excluded from the analyses).




Figure 5.1: Schematic of the case study area and the raingauges used. Each site is labelled with the gauge number used in this study, along with its altitude.

5.1.4 Basic statistical properties of rainfall

Statistical analyses of the rainfall time series, at both hourly and daily time scales, were performed for the eight raingauges. Their results for two representative months, namely January and July, are summarised graphically in a series of figures. Figure 5.2 shows marginal statistics (mean, variation, skewness and proportion dry) for each site; Figure 5.3 shows the autocorrelation functions. Figures 5.4 and 5.5 show cross-correlation coefficients among the different raingauges for time lags 0, 1 and 2: in Figure 5.4 they are plotted against the distance between gauges, whereas in Figure 5.5 they are plotted against the angle formed by each pair of gauges. Results for other months are very similar to those of January and July and in no case do they contradict the results and interpretation discussed below.

Stationarity of the rainfall field: Since the GDSTM rainfall is stationary in time, by fitting it to an entire month's rainfall record we are implicitly assuming that the rainfall process is stationary in time within each month. The analysis reported above indicates that the rainfall process can be regarded as spatially stationary in the case study area. As shown in Figure 5.2, marginal statistics do not vary significantly among different raingauges. The differences in the coefficients of skewness that appear in some cases are attributed to their high statistical variability and the short (5 years) record length available for their estimation, rather than to regional differences of the process. The sample autocorrelation functions, shown in Figure 5.3, provide additional evidence for the spatial stationarity of the process.

The conclusions here appear at odds with those determined from the GLM methodology in


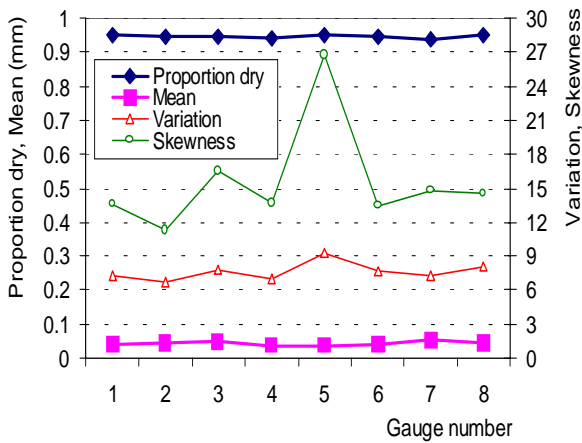




Figure 5.2: Marginal statistics of rainfall for the different raingauges for hourly (left) and daily (right) time scales and for months January (top) and July (bottom)

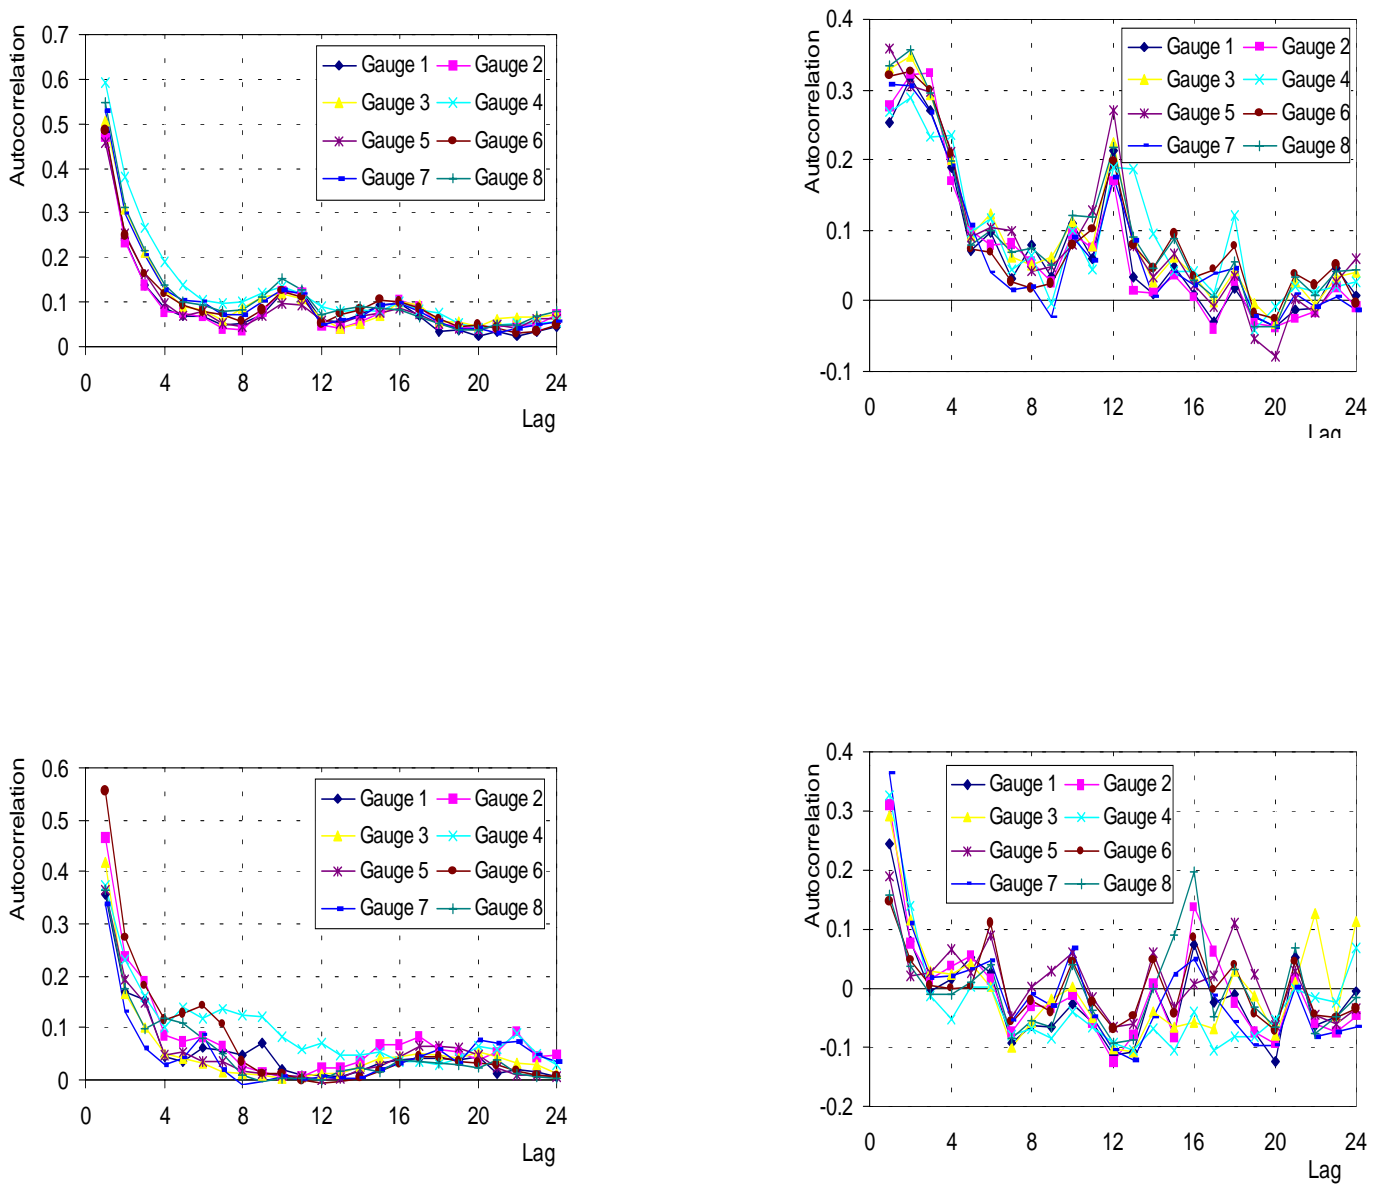


Figure 5.3: Autocorrelation functions of rainfall for the different raingauges for hourly (left) and daily (right) time scale and for months January (top) and July (bottom).

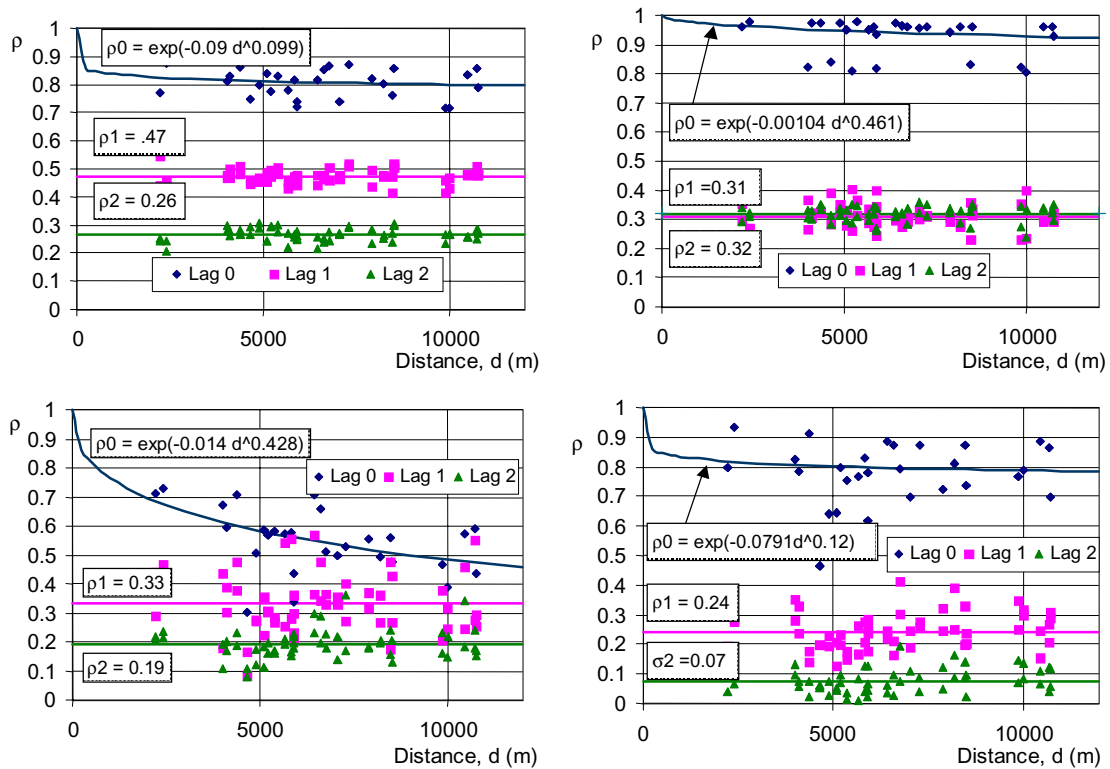


Figure 5.4: Cross-correlation coefficients of rainfall among the different raingauges for time lags 0, 1 and 2, versus the distance between gauges, for hourly (left) and daily (right) time scale and for months January (top) and July (bottom).

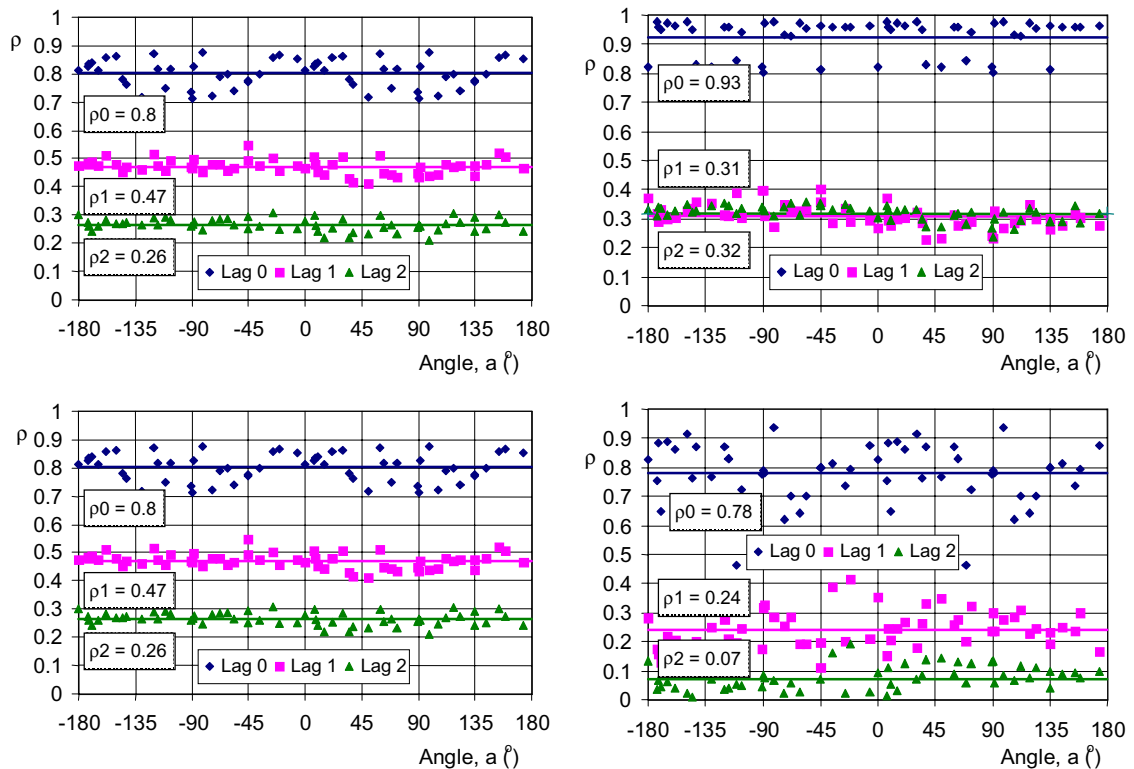


Figure 5.5: Cross-correlation coefficients of rainfall among the different raingauges for time lags 0, 1 and 2 versus the angular separation of the gauges, for hourly (left) and daily (right) time scales and for months January (top) and July (bottom).

section 4.2.2, where site altitude was found to have a statistically significant effect upon both rainfall occurrence and amounts at a daily timescale. However, in practical terms the effect is small (the GLM for rainfall amounts explains only 3% of the variance in daily rainfall, and altitude accounts for only a part of this), so for most purposes it may be considered adequate to assume spatial stationarity.

Isotropy of the rainfall field: In figure 5.5, where the cross-correlation coefficients among pairs of raingauges are plotted against the direction of the straight line joining the raingauges, no evidence of any effect of the orientation appears. Therefore, the spatial rainfall field can be regarded as isotropic. This also implies that there is no preferential velocity of storms, in a statistical sense. That is, velocities of different rainfall events must add up to a zero resultant. Additional evidence for this is provided by figure 5.4, which shows the cross-correlation coefficients of rainfall among the different raingauges for time lags 0, 1 and 2, versus the distance between gauges. Were a nonzero storm velocity \mathbf{V} affecting the rainfall field, this effect would be revealed in the lagged cross-covariance structure of the rainfall process. For example, considering two points 1 and 2 separated by a vector $\mathbf{L} = \mathbf{V}\tau$, whose point intensity processes are denoted by $X_1(t)$ and $X_2(t)$ respectively, the lagged correlation $\text{Corr}[X_1(t), X_2(t + \tau)]$ would be expected to be greater than the lag-zero cross-correlation $\text{Corr}[X_1(t), X_2(t)]$. However, as shown in figure 5.4, this is not the case for any of the pairs formed by the eight stations and for any lag $\tau > 0$. We must note, however, that the emerging insignificance of storm movement must be attributed to the consideration of all storms of a month as an ensemble (rather than as events with different velocities) and the adoption of rather coarse time scales (hourly and daily). An investigation was carried out using a finer time-scale (10 min) and the results obtained were similar. It is probable that at an even finer time scale or for a larger spatial extent, some effect of storm movement would become evident in the cross-correlation structure.

Cross-correlation structure: In figure 5.4 it is apparent that the lag-zero spatial correlation decreases slightly with increasing distance whereas lagged spatial correlations exhibit no significant variation with distance within the catchment. This has implications for model fitting, as will be discussed below.

Other properties Figure 5.4 shows that the spatial correlation of rainfall is impressively high, even at the hourly scale, a fact that, among others, indicates a good quality of the raingauge data. Both figures 5.3 and 5.4 show that the auto- and cross-correlations are higher in the wet season (January) than in the dry season (July).

5.1.5 Model fitting

Given the analyses of the previous section, the GDSTM can be simplified for the present application by ignoring (setting to zero) four of its parameters, i.e., the storm eccentricity e , the event orientation θ and the two velocity components V_x and V_y . The remaining parameters are estimated using a method of moments: λ , μ_D , μ_L , μ_A , μ_S and μ_C are estimated using second order properties of the process, and the remaining parameter μ_X is estimated from the mean intensity (equation (5.2)).

In order to estimate the spatial structure of the model from second-order properties, it is necessary to incorporate lag zero cross-correlations into the fitting procedure. However, with data

available only from a discrete network of sites, it is unclear how this can best be achieved. One approach would be to calculate separate lag zero cross-correlation coefficients for each pair of sites ($S(S-1)/2$ coefficients in total where S sites are involved), and include all of these coefficients in the fitting so as to obtain as close a match as possible to the observed cross-correlation structure at all inter-site distances. However, given the variability in observed cross-correlations from pairs of gauges separated by similar distances (see figure 5.4), this would probably not be worth the computational expense involved. A more promising strategy is to fit a curve to the observed cross-correlations, treat the fitted curve as an observed cross-correlation function, and parameterise the GDSTM using the fitted cross-correlation at some specified spatial lag. In the work reported here, the following relationship has been used:

$$\rho(d) = \exp(-ad^b) \quad (5.6)$$

where $\rho(d)$ is the cross-correlation between sites separated by a distance d . The parameters a and b are determined by least squares. The relationship (5.6) corresponds to a semi-variogram of the form

$$\Gamma(d) = \text{Var}[Y_i(x, y) - Y_i(0, 0)] = C(0)[1 - \rho(d)] = C(0)[1 - \exp(-ad^b)] \quad (5.7)$$

where $d = \sqrt{x^2 + y^2}$, and $C(0) = \text{Var}[Y_i(0, 0)] = \text{Var}[Y_i(x, y)]$ is the process variance.

It should be noted that the autocorrelation structure defined at (5.6) is probably not exactly the same as that implied by the GDSTM which we are fitting. However, this does not cause any problems in practice, since the only role of (5.6) in this application is to provide a simple and automatic interpolation of the observed cross-correlations. It provides a very reasonable fit in all cases, as illustrated in figure 5.4.

We are now in a position to describe the estimation of the parameters λ , μ_D , μ_L , μ_A , μ_S and μ_C . The properties used in the fitting were as follows:

- The variance, lag one autocorrelation, and lag zero cross-correlation at a distance of 5km, at both hourly and daily timescales. The spatial separation of 5km was chosen as it is equal to about the mean distance between the stations in our case study. The six second-order statistics used here are regarded as ‘primary’ fitting statistics.
- The lag two autocorrelation and the lag zero cross-correlation at a distance of 10 km, again at both hourly and daily timescales. These four properties are regarded as ‘secondary’ fitting statistics.

Of course, the mean of the process should also be regarded as of primary importance in model fitting: however, recall that the mean is subsequently used to estimate μ_X using equation (5.2), so an exact fit to the observed mean is guaranteed.

The historical values of the variance and the autocorrelations were estimated as averages of point values at the eight raingauges, whereas the historical lag zero cross-correlations were estimated from equation (5.6). The theoretical values are given in terms of the unknown parameters by equations (5.2) and (5.5). The objective function to be minimized has the form

$$f(\lambda, \mu_D, \mu_L, \mu_A, \mu_S, \mu_C) = \sum_i w_i \left(\frac{m_i - h_i}{h_i} \right)^2 \quad (5.8)$$







Figure 5.6: Seasonal variation in fitted model parameters, for simplified GDSTM applied to hourly raingauge data in the Brue catchment.

Parameter description	Parameter symbol	Parameter set 1	Parameter set 2
Rate of event arrival ($km^{-2} h^{-1}$)	λ	1×10^{-5}	2.66×10^{-6}
Mean storm duration (h)	μ_D	0.6022	0.6164
Mean event duration (h)	μ_L	14.64	14.79
Mean storm area (km^2)	μ_A	499.9	615.4
Mean event area (km^2)	μ_S	313.3	1542.6
Mean number of storms/event	μ_C	47.4	137.4
Mean storm intensity (mm/h)	μ_X	0.56	0.57
Value of objective function		0.0236	0.0178

Table 5.1: Two different sets of parameter values for March, for simplified GDSTM fitted to raingauge data from the Brue catchment

where h_i and m_i denote historical and modelled (theoretical) values respectively for the i th fitting property, and the $\{w_i\}$ are weights which were set equal to 1 for each of the six ‘primary’ statistics, and 0.1 for each of the ‘secondary’ statistics. The optimisation was performed numerically using the tools described in Appendix C. The fitted parameter values are shown in figure 5.6.

Generally, the objective function in equation (5.8) has several local optima. Thus, different sets of parameter values may result in an acceptable preservation of the rainfall process statistics. As an example, two different sets of parameter values for March are given in table 5.1. It is seen that both sets yield very similar values of the objective function. In terms of the reproduction of these statistics, either parameter set may therefore be chosen. Discrimination between them would require that other statistics be considered. We note that for the first parameter set, μ_A is larger than μ_S , which means that most of the storm centres are clustered over an area which is on average smaller than the average extent of a storm. As a result of this lack of physical realism, coupled with the slightly lower objective function value achieved by the second parameter set, it is the second parameter set which has been used for the work reported below.

5.2 Continuous simulation performance with raingauge data

To assess the model performance we now compare several properties of synthetic point rainfall series, generated by simulating the fitted model, with their corresponding historical properties.

5.2.1 Simulation model

The simulation model developed for this study, and described above, first generates events and storms for a certain spatial region and time period and then calculates the rainfall time series at points of interest, together with their statistical characteristics.

The simulation area is wider than the region of interest and the simulation time period is longer than the period of interest. Let us assume that the region of interest is a rectangle with sides l_x and l_y , and the time period of interest is an interval of length d (see figure 5.7). Events located outside




Figure 5.7: Schematic of the simulation space of the model. The inner parallelepiped is the space (i.e. area and time period) of interest and the outer parallelepiped is the actual simulation space.

this rectangle (but nearby) can cover all or part of the region. Similarly, events generated before the time period of interest may continue into that period. Therefore, to simulate accurately the rainfall process in the inner parallelepiped of figure 5.7 with dimensions $l_x \times l_y \times d$, it is necessary to simulate event generation in a much larger spatial-temporal region. We shall consider all events generated in the outer parallelepiped of figure 5.7, which has dimensions $L_x \times L_y \times D$. We take

$$L_x = l_x + 2\Delta l_x \quad (5.9)$$

$$L_y = l_y + 2\Delta l_y \quad (5.10)$$

$$D = d + \Delta d, \quad (5.11)$$

and

$$\Delta l_x = 2.5\sqrt{\mu_s + \mu_A} \quad (5.12)$$

$$\Delta l_y = \Delta l_x \sigma_y / \sigma_x \quad (5.13)$$

$$\Delta d = 5\mu_L \quad (5.14)$$

where σ_x and σ_y are the standard deviations of the displacements of storm from event centres. They are determined in terms of the model parameters in section 5.1.2. These choices were determined using statistical reasoning and simulation experiments. Note that in the present case study l_x and l_y are of the order of 10 km whereas Δl_x and Δl_y are of the order of 100 km, so that the volume ratio of the outer to the inner parallelepiped is of the order of 100-1000. This large ratio indicates that the computer time for simulation can be significant even when the area of interest is small.

The simulation program, after generating the series of events in the outer parallelepiped of figure 5.7, generates the storms belonging to those events; note that storm centres may be located outside the outer parallelepiped. The program then determines the rainfall time series at 5 points of the inner parallelepiped over the duration d (namely, the four corners and the centre of the rectangle, at the hourly time scale) and evaluates certain statistics of these time series (namely, the dry proportions, means, standard deviations, coefficients of skewness, temporal correlation coefficients for lags one and two, and cross-correlation coefficients for zero time lag).

5.2.2 Statistical checking criteria

The model is checked at different points, using both the hourly and daily time scales. Three categories of statistical properties are used for checking:

1. Primary properties used in the fitting procedure:
 - Marginal statistics: Means and variances
 - Temporal structure: Lag one covariances
 - Spatial structure: Lag zero covariances at a characteristic distance
2. Secondary properties used in the fitting procedure:
 - Temporal structure: Lag two covariances
 - Spatial structure: Lag zero covariances at a distance greater than the characteristic one
3. Properties not used in the fitting procedure:
 - Wet/dry properties: Proportion of dry hours and days
 - Marginal statistics: Coefficients of skewness
 - Statistics of extremes: Rainfall depths for characteristic return periods

5.2.3 Results

Twenty years of synthetic point rainfall, at five locations, were generated using the monthly parameter sets presented in figure 5.6. The results are presented in a sequence of figures. Figure 5.8 shows a comparison of simulated, modelled (i.e. theoretically expected) and historical statistics for each month, for the primary fitting properties. We observe a very good agreement of simulated statistics with the modelled ones. These in turn are almost identical to the historical ones, the slight differences reflecting the model's inability to fit perfectly to historical values.

Figures 5.9 and 5.10 show similar comparisons, respectively for secondary fitting properties and for properties not used in the fitting. In figure 5.9 we observe a very good agreement of the simulated statistics with the modelled ones, which, however, depart somewhat (especially in the case of cross-correlations for a distance of 10 km) from the historical ones. In figure 5.10 we find a very good agreement of simulated and historical proportions of dry hours and days and an acceptable approximation of skewness at both hourly and daily scales, although these properties were not used in the model fitting procedure.







Figure 5.8: Comparison of simulated, modelled and historical statistics for each month: mean intensity, coefficient of variation, lag one autocorrelation and lag zero cross correlation. These are the primary statistics used in model fitting. Lag zero cross-correlation refers to an inter-site distance of 5 km.

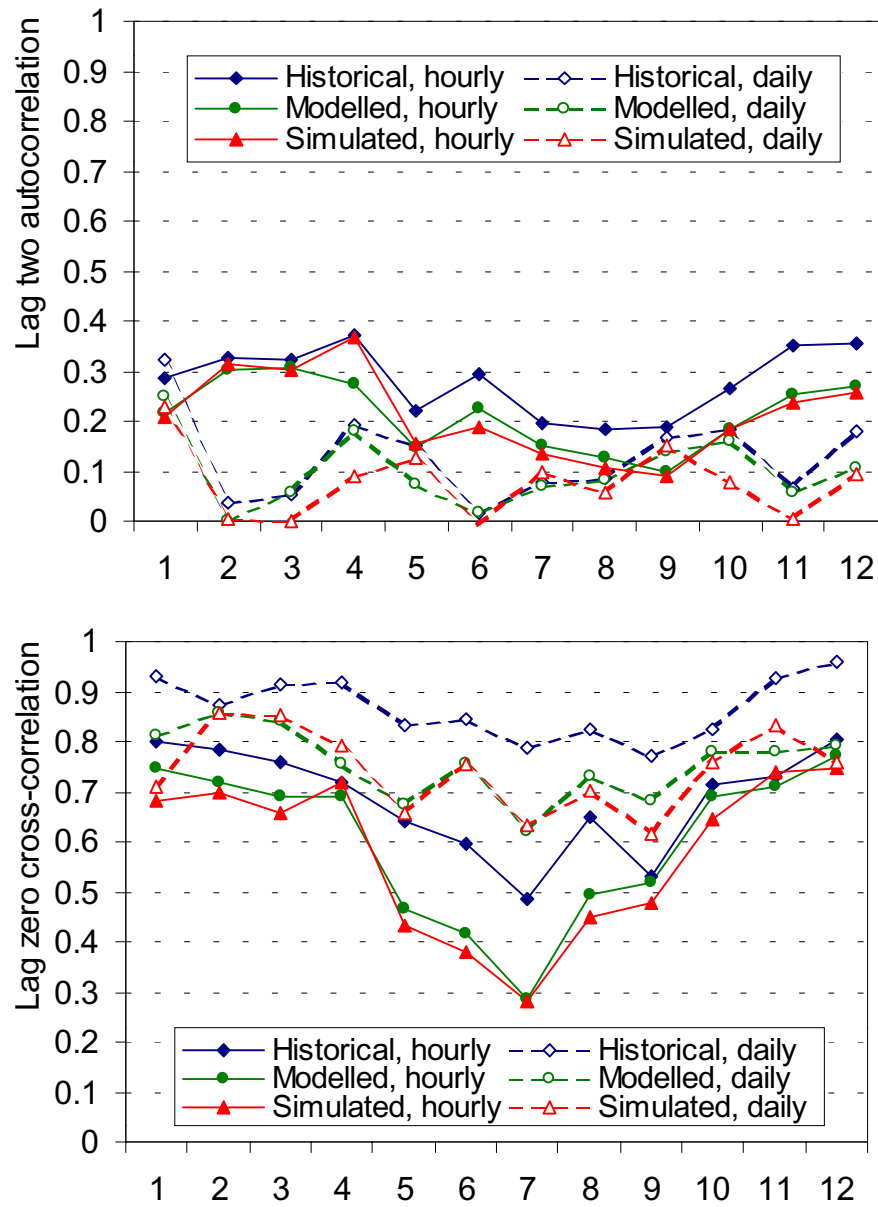


Figure 5.9: Comparison of simulated, modelled and historical statistics for each month: lag two autocorrelation and lag zero cross correlation. These are the secondary statistics were used in the model fitting. Lag zero cross-correlation refers to an inter-site distance of 10 km.

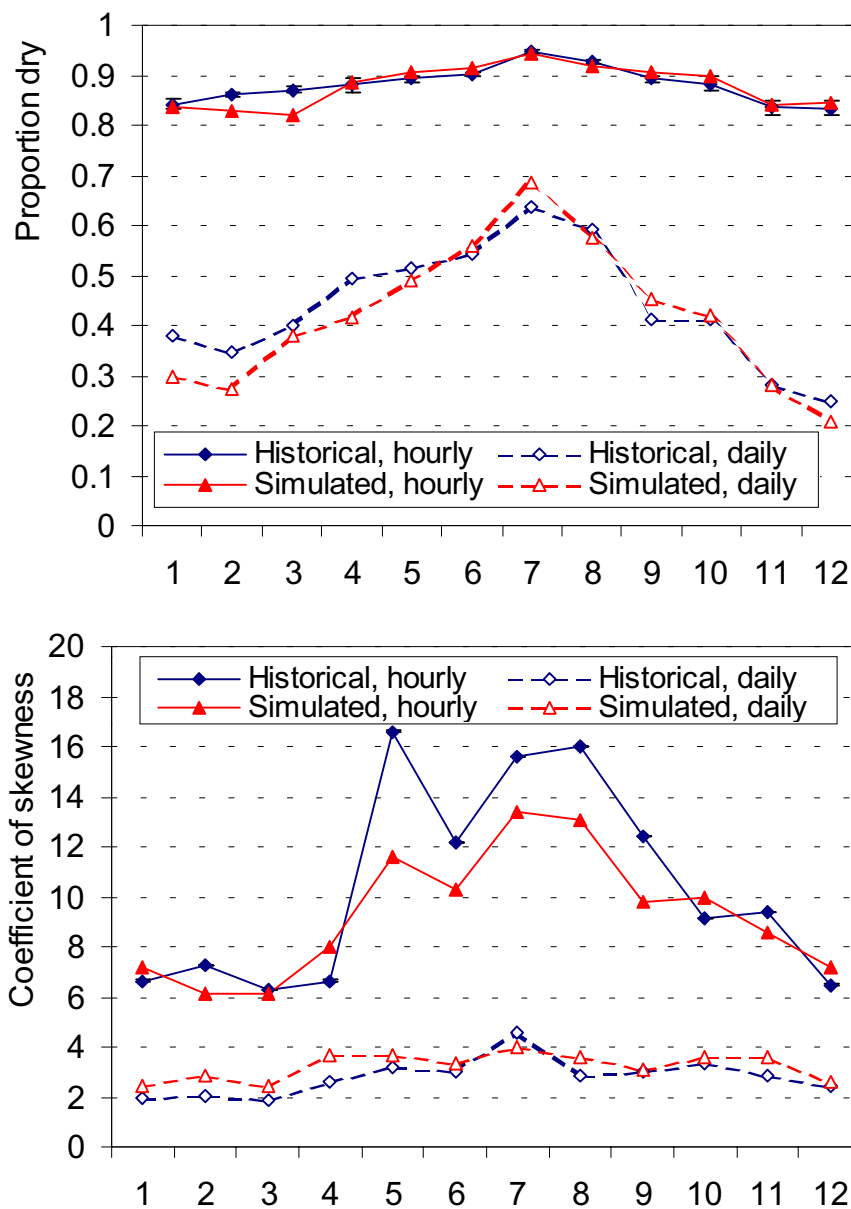


Figure 5.10: Comparison of simulated, modelled and historical statistics for each month: proportion of dry hours and days and coefficients of skewness. These statistics were not used in the model fitting.

5.2.4 Extreme-value performance

Validation data To evaluate of the ability of the proposed model to reproduce the extreme-value behaviour of rainfall in the given area, adequate validation data are required. The Brue raingauge data set is limited to under four full years of data. However, the fact that a large number of gauges are available can partially compensate for this very limited size. Although the methodology presented below requires independence of the annual maxima, this condition can be approximately realised by pooling data from different gauges since many such maxima are the result of very local rainfall peaks.

A more satisfactory validation data set would have a longer record. Although this is not available over the Brue catchment, the gauge of Boscombe Down is located some 50 km away to the East and has a record of more than 18 years (1980-1998) of hourly rainfalls.

If longer term data are required, the regionalised information of the Flood Studies Report (NERC 1975) can be utilised. In this report, maps give the 2 day rainfall with a 5 year return period, i.e. $R(2 \text{ day}, 5 \text{ year})$, as well as the ratio $R(1 \text{ hour}, 5 \text{ year}) / R(2 \text{ day}, 5 \text{ year})$. Using tables, one can then deduce $R(D, T)$ for any duration D and return period T . For the sake of comparison, we have therefore examined the extreme values obtained over the Brue and for Boscombe Down on the basis of the Flood Studies Report.

Methodology By simulating 5 sets of 18 years of data, we shall have 5 different scenarios of rainfall to compare with Boscombe Down's historical data set and with sets of 18 years of data put together on the basis of combining data from several Brue raingauges (as before, spatial dependence has not been accounted for). For each of these data sets, an extreme-value analysis is performed on the annual maxima at hourly and daily timescales, using the methods already described in section 3.3.3.

Results for hourly data: Figure 5.11 shows that the January extremes are very well reproduced by the model; the five lines representing the estimated rainfalls for any return period are very close to the extreme-value line extrapolated from the Boscombe Down observed maxima (denoted "Fitted EV-I"). This is a remarkable agreement, given that the modelling is based on only 5 years of data. Such good agreement is not however the case for all months, and over the whole year we find that the model underestimates the extremes (seen figure 5.12). This applies to the extremes estimated from both the Flood Studies Report and the Boscombe Down data set.

When using the pooled Brue raingauge maxima to fit the EV-I distribution however, we find a very good fit of the model to the data (figure 5.13). This indicates that the model is indeed able to reproduce the hourly extreme value distribution of the data used in the fitting. That the latter is significantly different from that of a longer term raingauge indicates that the maxima for the period 1993 – 1998 are not representative of those of the 1970-1998 period at Boscombe Down. This is however not a surprising result, given recent climate variability (Mayes 1995, DETR 1999).

Results for daily data: Here, the 5 simulations provide an excellent reproduction of the historical extremes extrapolated from the Boscombe Down data set (figure 5.14). The performance is also good with respect to the pooled data from the Brue (figure 5.15).

Hourly Extreme Values

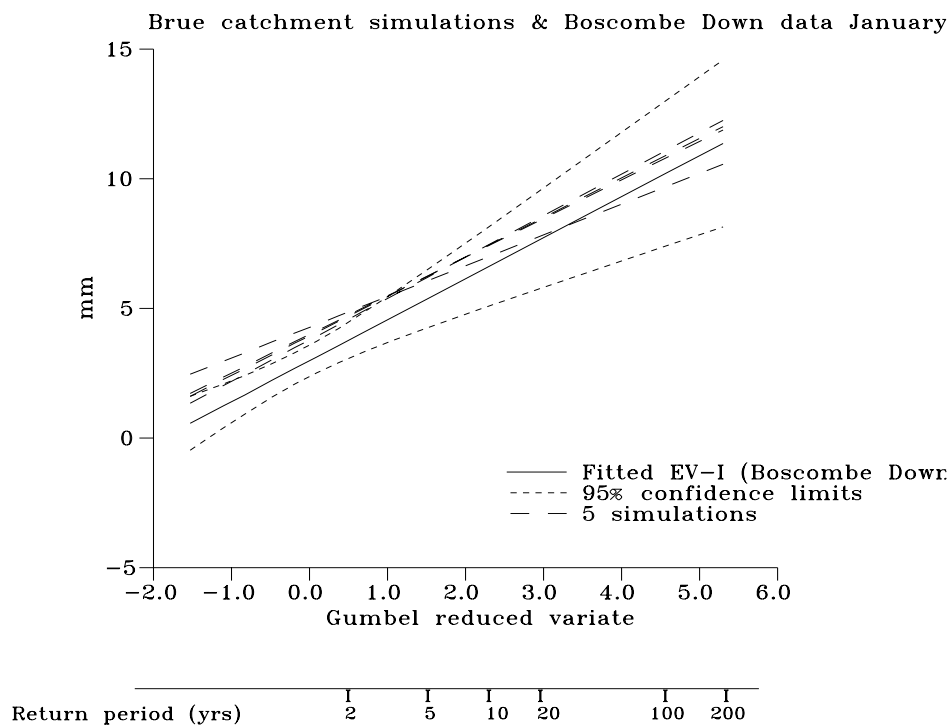


Figure 5.11: January hourly extreme values for simulated and Boscombe Down data.

These figures also indicate the discrepancy between the long term Flood Studies Report extremes and those based upon these shorter data sets. Thus, although the Boscombe Down series is sufficient to form a good estimate of the longer term hourly maxima, this is not the case for the daily time-scale. At that time-scale, there is insufficient data in the observed 18 years to allow extrapolation to maxima of larger return periods.

5.2.5 Conclusions

Typical densities of raingauges are evidently insufficient, without additional information, to identify the parameters of the full spatial-temporal model considered in chapter 2. Analysis of a network of raingauges has shown that the spatial structure of cells cannot be resolved, and that even with data at fine temporal resolution, no consistent velocity can be identified.

Within the available time, only limited testing of the reduced model has been possible. However, results were generally impressive. The spatial and temporal properties of observed rainfall were well reproduced and, given only a very limited period of record (5 years), extreme value performance was excellent at both hourly and daily timescale. The good extreme value performance held in comparison both with the data used in fitting and, at least for daily values, with an 18-year raingauge record from nearby.




Figure 5.12: Hourly annual extreme values for simulated and Boscombe Down data




Figure 5.13: Hourly annual extreme values for simulated and pooled Brue data

Daily Extreme Values




Figure 5.14: Daily extreme values for simulated and Boscombe Down data

Daily Extreme Values




Figure 5.15: Daily extreme values for simulated and pooled Brue data

5.3 Spatial-temporal rainfall disaggregation

Very frequently, there appear situations even poorer than those examined in the previous sections in terms of availability of rainfall data. For example, radar data may not exist and raingauge data may be available only at a daily time scale at most locations of interest. However, there often exist raingauge data at a finer time scale (e.g. hourly) at a neighbouring site. The question then arises whether we could utilise the available single-site fine scale rainfall information, in conjunction with the daily data, to generate spatially consistent rainfall series.

This can be considered as a particular case of a general multivariate spatial-temporal rainfall disaggregation problem. In another (commonly occurring) guise, this problem involves the use of observed fine-scale data from a single site to disaggregate historical multivariate daily series. Although there is substantial experience in multi-site disaggregation of rainfall from annual to monthly time scale, and in single-site disaggregation of rainfall to finer time scales, this multivariate fine-time-scale rainfall disaggregation problem has not been studied so far in the rainfall modelling literature. It presents significant differences from that of single-site disaggregation. Multiple sites certainly imply mathematical complexity. The spatial correlation (cross-correlation between different sites) must be maintained in the multivariate problem, whereas it does not appear at all in univariate problems. However, the spatial correlation can be turned to advantage since, in combination with the available single-site rainfall information, it enables realistic fine-scale rainfall series to be generated. Timings of rainfall events and maximum intensities at all sites can be guided by an observed or simulated fine-scale hyetograph.

5.3.1 Problem formulation

We standardise the problem that we examine throughout this section in the following manner, with reference to figure 5.16. We assume that we are given:

1. an hourly point rainfall series at point 1, as a result of either:
 - measurement by an autographic device or digital sensor;
 - simulation with a fine time scale point rainfall model such as a point process model;
 - simulation with a temporal point rainfall disaggregation model applied to a series of daily rainfalls.
2. several daily point rainfall series at neighbouring points (e.g. 2, 3, 4 and 5 in figure 5.16), as a result of either:
 - measurement by conventional rain gauges (daily observations), or
 - simulation with a multivariate daily rainfall model (such as the GLM).

We wish to produce series of hourly rainfall at points 2, 3, 4, and 5, so that:

1. their daily totals equal the given daily values;
2. their stochastic structure resembles that implied by the available historical data (see section 5.3.2).

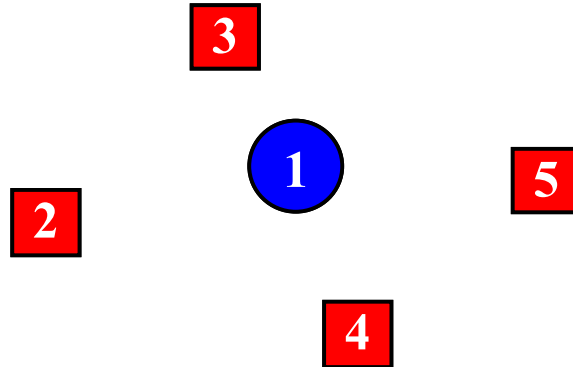


Figure 5.16: Schematic diagram illustrating the multivariate disaggregation problem. Squares indicate points with daily rainfall data and the circle represents a point with hourly rainfall data.

For the sake of simplicity we have assumed here that hourly information is available at one site only. In real world problems it could be the case that hourly rainfall is available at more than one location. This case can also be tackled with the modelling strategy described below, without any further difficulty apart from some generalisations of the computational algorithm. In fact, having more than one point with known hourly information would be advantageous for two reasons. First, it would allow a more accurate estimation of the spatial correlation of hourly rainfall depths (see section 5.3.2) or transformations of them (see sections 5.3.5 and 5.3.7). Second, it might reduce the residual variance of the rainfall process at each site, thus allowing for generated hyetographs closer to the real ones.

5.3.2 Estimation of the stochastic structure at the hourly level

The essential statistics that we wish to preserve in the generated hourly series are:

1. the means, variances and coefficients of skewness;
2. the temporal correlation structure (autocorrelations);
3. the spatial correlation structure (lag zero cross-correlations);
4. the proportions of dry intervals.

All these statistics, apart from the cross-correlation coefficients, can be estimated at the hourly time scale using the single hourly data set (gauge 1 in figure 5.16). To transfer these parameters to other locations, spatial stationarity of the process can be assumed. This may seem an oversimplification at first glance; however, it is not a problem in practice since possible spatial nonstationarities will manifest themselves in the available daily series. Therefore, despite the stationarity assumption made at this initial (parameter estimation) stage, the final hourly series, which is forced to respect the observed daily totals, will reflect these non-stationarities.

The most difficult statistics to estimate are the cross-correlations at the hourly level. Note that if more than one rainfall series were available at hourly level, then at least one cross-correlation coefficient could be estimated directly from these series. Then, by making reasonable assumptions about the spatial dependence, an expression of the cross-correlation as a function of distance could be established. This would then be used to estimate cross-correlations between all pairs of raingauges.

However, this method cannot be used when only one hourly rainfall series is available. In this case, a spatial temporal model can be used to infer indirectly the spatial correlation structure of the rainfall field. Here we have used the GDSTM (as in the previous sections) in the following manner:

1. The marginal statistics and temporal and spatial correlations at the daily level are estimated using the daily data sets.
2. The parameters of the spatial-temporal rainfall model are estimated (using the general method described in section 5.1) using the single-site statistics at the hourly level and the multi-site statistics at the daily level.
3. The spatial correlations at the hourly level are inferred from the spatial-temporal rainfall model.

It should be noted that the spatial-temporal rainfall model is used in this parameter estimation step only. To estimate cross-correlations in the manner described above, it suffices to calibrate the model and there is no need to run it. However, a model run may be necessary if statistics of transformations of rainfall depths are needed for subsequent modelling steps (see sections 5.3.5 and 5.3.7). Such statistics can be estimated from synthetic hourly series generated by the spatial-temporal rainfall model.

5.3.3 Modelling approach

The proposed approach to the multivariate fine-scale rainfall disaggregation problem involves the application of two separate models at the generation phase.

The first is a rather simplified multivariate model of hourly rainfall that can preserve the essential statistics of the multivariate rainfall process and, simultaneously, incorporate the available hourly information without any reference to the known daily totals at the other sites. The essential statistics considered here are the means, variances and coefficients of skewness, the lag one autocorrelation coefficients and the lag zero cross-correlation coefficients.

The second model is a transformation model that modifies the series generated by the first model, so that the daily totals are equal to the given ones. This uses a (multivariate) transformation, which does not affect the stochastic properties of the series. Both models are discussed further below.

Some questions may arise as to the adoption of a simplified multivariate model rather than a more sophisticated one which could preserve additional statistics of the rainfall process. While in the proposed modelling framework a simplified model is suggested, the adoption of a more

sophisticated approach is not excluded. For example, it may be possible to design the multivariate model so as to maintain a large number of autocovariance coefficients (for any lags). However, such complexity will generally be unnecessary since, for typical catchment sizes, spatial dependence between sites will be strong so that the primary control over the disaggregated output will come from the observed fine-scale series.

Specifically, assuming that sites 2 – 5 (figure 5.16) are all close to site 1 and highly spatially correlated, the given hourly series at site 1 can be used, with the simplified multivariate model, to:

- guide the generation of the hourly series at the sites with daily data, and act indirectly to preserve properties not modelled explicitly;
- properly locate the rainfall events in time;
- produce initial hourly rainfall series at the daily sites, whose departures from the actual hourly depths at those sites are not large (even though the known daily totals are not considered at all at this stage).

At a later stage, i.e. when the transformation model is applied, another source of information is additionally incorporated, that is the multi-site daily information. This results in preservation of additional properties, which are not captured by the statistics used. For example, as noted above, nonstationarities of the rainfall field (both in space and time) are reproduceable, even though the models used are both stationary.

The proportion of dry intervals, although considered as one of important properties to be preserved (page 142), is difficult to incorporate explicitly in either of the above described models, as it cannot be expressed in terms of statistical moments. However, it can be treated by an iterative procedure which will be discussed later (section 5.3.7).

5.3.4 Models involved

Several separate models are involved in the proposed disaggregation framework. These fall into three categories, as follows:

Category 1: The first category includes the models that are the core of this framework in the sense that they provide the required output (the hourly series). These are, as outlined in section 5.3.3 above,

- the simplified multivariate model for hourly rainfall, and
- the transformation model.

Category 2: This category contains models to provide the required input, if no observed series are available. These may include

- the GLM for providing daily rainfall depths;
- a single-site disaggregation model to disaggregate daily depths of one location into hourly depths (see section 5.4 below for a review of such procedures);

- a single-site Bartlett-Lewis model (Rodriguez-Iturbe et al. 1987) for providing hourly depths at one location.

The final example here may be appropriate if, under a stationary climate scenario, a single-site Bartlett-Lewis model has been simulated into the future, and a GLM has been simulated at neighbouring sites, conditional upon the daily totals from this model, so as to introduce spatial nonstationarity into the generated multivariate daily time series.

Category 3: The third category includes models that are able to provide some of the required parameters of the spatial-temporal rainfall process given the statistical properties that can be estimated from the available data. For example, in this case study, we used the GDSTM to provide the cross-correlation structure at the hourly level. Runs of this model may not be necessary unless it is required to estimate statistical properties of transformations of hourly rainfall depths.

We now describe the simplified multivariate model and the transformation model which have been used in this work.

5.3.5 The simplified multivariate rainfall model

Let the n -vector $\mathbf{X}_s = (X_s^1, X_s^2, \dots, X_s^n)^T$ represent the hourly rainfall at time (hour) s at n locations. We assume that the simplified multivariate rainfall model is an AR(1) (autoregressive-moving process of order 1) model, expressed by

$$\mathbf{X}_s = \mathbf{a}\mathbf{X}_{s-1} + \mathbf{b}\mathbf{V}_s \quad (5.15)$$

where \mathbf{a} and \mathbf{b} are $(n \times n)$ matrices of parameters and (\mathbf{V}_s) ($s = 0, 1, 2, \dots$) is an independent, identically distributed (i.i.d.) sequence of innovations (these are n -vectors of i.i.d. random variables, so that the innovations are both spatially and temporally independent). The time index s can take any integer value. The (\mathbf{X}_s) are not necessarily standardised to have zero mean and unit standard deviation, nor are they normally distributed. On the contrary, their distributions are quite skewed. Consequently, the distributions of (\mathbf{V}_s) are skewed too; a three-parameter gamma distribution is generally appropriate for the latter.

Alternatively, the model can be expressed in terms of some nonlinear transformation X_s^* of the hourly depths X_s , in which case (5.15) is replaced by

$$\mathbf{X}_s^* = \mathbf{a}\mathbf{X}_{s-1}^* + \mathbf{b}\mathbf{V}_s \quad (5.16)$$

It is natural to consider the power family of transformations here.

Equations to estimate the model parameters \mathbf{a} and \mathbf{b} and the moments of \mathbf{V}_s are given by Koutsoyiannis (1999) for the most general case. As there is no need to preserve lagged cross-covariances in the problem examined, the parameter matrix \mathbf{a} can be diagonal. The parameter matrix \mathbf{b} should be defined here as lower triangular. This is necessary in order to incorporate the known hourly rainfall at site 1. Thus, if \mathbf{b} is lower triangular then its first row will have only one nonzero item (b^1 say), so that from (5.15) we have

$$X_s^1 = a^1 X_{s-1}^1 + b^1 V_s^1 \quad (5.17)$$

This can be utilised to determine (rather than to generate) V_s^1 given the series of X_s^1 . This can be directly expanded to the case where several gauges with hourly information are available provided that \mathbf{b} is lower triangular.

5.3.6 The transformation model

Transformations that can modify a series generated by any stochastic process so as to satisfy some additive property (i.e. that the sum of the values of a number of consecutive variables be equal to a given amount), without affecting the first and second order properties of the process, have been studied by Koutsoyiannis (1994) and Koutsoyiannis & Manetas (1996). These transformations, more commonly known as adjusting procedures, are appropriate for univariate problems, although they can be applied to multivariate problems as well, but in an iterative framework. More recently, Koutsoyiannis (2000) has studied a true multivariate transformation of this type, which avoids any iteration, and also proposed a generalised framework for coupling stochastic models of different time scales.

This framework, specialised for the problem examined here, is depicted in figure 5.17. Here \mathbf{X}_s and \mathbf{Z}_p represent the *actual* hourly- and daily-level processes, related by

$$\sum_{s=(p-1)k+1}^{pk} \mathbf{X}_s = \mathbf{Z}_p, \quad (5.18)$$

whereas $\tilde{\mathbf{X}}_s$ and $\tilde{\mathbf{Z}}_p$ denote auxiliary processes, represented by the simplified rainfall model in our case, which also satisfy (5.18). k in this equation is the number of fine-scale timesteps within each coarse-scale step (24 for the current application).

The problem is, given a time series (\mathbf{z}_p) of the actual process (\mathbf{Z}_p), to generate a series (\mathbf{x}_s) of the actual process (\mathbf{X}_s). To this aim, we first generate another (auxiliary) time series ($\tilde{\mathbf{x}}_s$) using the simplified rainfall process ($\tilde{\mathbf{X}}_s$). The latter time series is generated independently of (\mathbf{z}_p) and, therefore, the $\tilde{\mathbf{x}}_s$ s do not add up to the corresponding \mathbf{z}_p s, as required by the additive property (5.18), but to some other quantities, denoted as ($\tilde{\mathbf{z}}_p$). Thus, in the next step, we modify the series ($\tilde{\mathbf{x}}_s$) to produce a series (\mathbf{x}_s) which is consistent with (\mathbf{z}_p) (in the sense that the \mathbf{x}_s s and \mathbf{z}_p s obey (5.18)) without affecting the stochastic structure of the $\tilde{\mathbf{x}}_s$ s. For this modification we use a linear transformation $f(\tilde{\mathbf{X}}_s, \tilde{\mathbf{Z}}_p, \mathbf{Z}_p)$ whose outcome is a process distributed identically to (\mathbf{X}_s) (so that we can write $\mathbf{X}_s = f(\tilde{\mathbf{X}}_s, \tilde{\mathbf{Z}}_p, \mathbf{Z}_p)$), and which is also consistent with (\mathbf{Z}_p) (it satisfies (5.18)).

Let $\mathbf{X}_{(p)}$ be the vector of hourly rainfall values for day p (for 5 locations, $\mathbf{X}_{(p)}$ contains $24 \times 5 = 120$ variables). Let also \mathbf{Y}_p be a vector containing:

- a. the daily values \mathbf{Z}_p ;
- b. the daily values \mathbf{Z}_{p+1} of the next day and
- c. the hourly values of the last hour of the previous day $p-1$ for all locations. This means that for 5 locations \mathbf{Y}_p contains $3 \times 5 = 15$ variables in total.

Items (b) and (c) of the vector \mathbf{Y}_p were included to ensure that the transformation preserves not only the covariance properties among the hourly values of each day, but the covariances with the

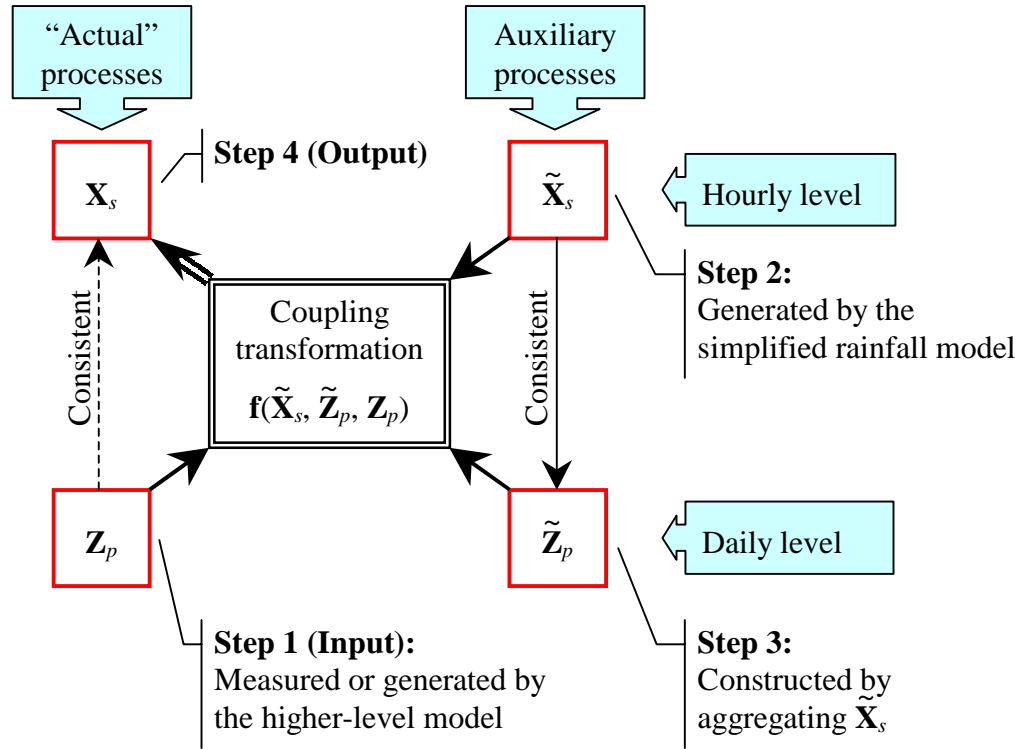


Figure 5.17: Schematic representation of actual and auxiliary processes, their links, and the steps followed to construct the actual hourly-level process from the actual daily-level process.

previous and next days as well. Note that at day p the hourly values of day $p - 1$ are known (therefore, in \mathbf{Y}_p we enter hourly values of the previous day) but the hourly values of day $p + 1$ are not known. This is why, in \mathbf{Y}_p , we enter daily values of the next day, which are known. In an identical manner, we construct the variables $\tilde{\mathbf{X}}_{(p)}$ and $\tilde{\mathbf{Y}}_p$ from vectors $\tilde{\mathbf{X}}_s$ and $\tilde{\mathbf{Z}}_p$.

Koutsoyiannis (2000) showed that there exists a matrix of coefficients \mathbf{h} such that if \mathbf{X} is generated using

$$\mathbf{X}_{(p)} = \tilde{\mathbf{X}}_{(p)} + \mathbf{h}(\mathbf{Y}_p - \tilde{\mathbf{Y}}_p) \quad (5.19)$$

then:

1. $\mathbf{X}_{(p)}$ has mean and variance-covariance matrix identical to those of $\tilde{\mathbf{X}}_{(p)}$, and joint covariance matrix with \mathbf{Y}_p identical to that of $\tilde{\mathbf{X}}_{(p)}$ and $\tilde{\mathbf{Y}}_p$;
2. any linear relationship which holds for $\tilde{\mathbf{X}}_{(p)}$ and $\tilde{\mathbf{Y}}_p$ and which can be written in the form

$$\mathbf{g}_X^T \tilde{\mathbf{X}}_{(p)} = \mathbf{g}_Y^T \tilde{\mathbf{Y}}_p \quad (5.20)$$

where \mathbf{g}_X and \mathbf{g}_Y are matrices of coefficients, also holds for $\mathbf{X}_{(p)}$ and \mathbf{Y}_p , that is

$$\mathbf{g}_X^T \mathbf{X}_{(p)} = \mathbf{g}_Y^T \mathbf{Y}_p \quad (5.21)$$

Note that the additive property (5.18) can be written in the matrix form (5.21) (for appropriately selected \mathbf{g}_X^T and \mathbf{g}_Y^T) and, therefore, its preservation by the transformation is ensured.

Details of how to determine \mathbf{h} in terms of covariance properties of $\mathbf{X}_{(p)}$ and \mathbf{Y}_p are given by Koutsoyiannis (2000).

5.3.7 Specific difficulties

A number of peculiarities of the rainfall process at a fine time scale cause specific difficulties that are examined in this section.

Negative values: Linear stochastic models such as those used here may generally generate negative values, which, of course, do not have any physical meaning. In practice, the probability of generating negative values depends upon the coefficient of variation of the variables. This probability is negligible when models operate at large time scales such as annual or monthly, because the coefficient of variation is small (usually smaller than 0.5–1.0). However, in the case of hourly rainfall this coefficient becomes as large as 3–4 and the probability of generating negative values becomes significant. This problem can be resolved by truncating the negative values generated, i.e., setting them to zero. This may have a beneficial effect in preserving the proportion of dry intervals (as shown in next paragraph) but it is also a potential source of bias to all statistical properties that are to be preserved. Specifically, it is anticipated to result in overprediction of cross-correlations as it is very probable that negative values are contemporary.

Dry intervals: The proportion of dry intervals is an important characteristic of the rainfall process that must be preserved. This proportion cannot be preserved by the linear stochastic models of section 5.3.3 in an explicit and theoretically consistent manner. However, after rounding the generated values (e.g. to one decimal digit) a significant number of zero values emerge because of the high coefficient of skewness (of the order of 5–10 in the case study examined) of the rainfall process. Additional zero values result from the truncation of negative values. It cannot be expected however, that the proportion of such dry intervals produced in this way will match the proportion in the historical data. Usually, we expect the former to be lower than the latter. A simple practical technique may be to control the proportion of dry intervals. Specifically, a proportion π_0 of the very small positive values, chosen at random among the generated values that are smaller than a threshold l_0 (e.g. 0.1–0.3mm), are set to zero. The numbers π_0 and l_0 can be found by performing repetitions starting with different trial values until the proportion of dry intervals in the synthetic series matches that in the historical record.

Preservation of skewness: Although the coupling transformation preserves the first and second order statistics of the processes, it does not ensure the preservation of third order statistics. Thus, it is anticipated that it will result in underprediction of skewness. Iterative algorithms to remedy this problem have been studied by Koutsoyiannis (2000).

Homoscedasticity of innovations: By definition, the innovations \mathbf{V}_s in the simplified multivariate rainfall model are homoscedastic, in the sense that their variances are constant, independent of the values of rainfall depths \mathbf{X}_s . Therefore, if for instance we estimate (or generate) the value at location 2, given that at location 1, we assume that the conditional variance is constant and independent of the value at location 1. This, however, does not comply with reality: by examining simultaneous hyetographs at two locations we can observe that the variance is larger during the periods of high rainfall (peaks) and smaller for periods of low rainfall (heteroscedasticity). As a result of this inconsistency, synthesised hyetographs tend to have unrealistically similar peaks. To mitigate this problem we can apply a nonlinear transformation to rainfall depths.

The first candidate nonlinear transformation is the logarithmic one:

$$\mathbf{X}_s^* = \ln(\mathbf{X}_s + \mathbf{c}_s) \quad (5.22)$$

with constants $c_s^i > 0$ for all i (with a slight abuse of notation here — (5.22) should be read as an element-by-element transformation). This has the advantage that it may result in zero skewness in which case the transformed variables can be assumed to be normally distributed. Then, preservation of first and second order properties of the untransformed variables is equivalent to preservation of first and second order statistics of the transformed variables (Koutsoyiannis 2000). However, evidence from the examined data sets shows that the skewness of the transformed variables increases with increasing constants \mathbf{c}_s and it still remains positive even if very small values of these constants are chosen. This means that the lognormal assumption is not appropriate for hourly rainfall.

A second candidate is the power transformation

$$\mathbf{X}_s^* = \mathbf{X}_s^{(m)} \quad (5.23)$$

i.e. each component of \mathbf{X}_s is raised to the power m (item to item) where $m < 1$. The stationarity assumption implies that m is the same for all items. Note that the preservation of the statistics of the untransformed variables does not necessarily lead to the preservation of the corresponding statistics of the transformed variables. However, the discrepancies are expected to be low if m is close to unity.

5.3.8 Application to the Brue catchment

The methodology described above was applied to the Brue catchment using hourly data from one raingauge only and daily data from another four raingauges, shown in figure 5.18. Five years of data for the month of January were used for this case study.

The statistics estimated from the single-site hourly rainfall data (gauge 1) are shown in table 5.2. For comparison and verification, we have included the corresponding statistics for the other four raingauges (gauges 2–5) estimated from the hourly data of the same period, although the latter were not used at any phase of this case study. Using the statistics of gauge 1 at the hourly level, as well as the corresponding statistics at daily level together with the cross-correlations at daily level, we fitted the GDSTM in the manner already described in section 5.1. The cross-correlation coefficients resulting from this model at hourly level are given in table 5.3. For comparison, the historical cross-correlation coefficients among the five gauges are also shown, although these were

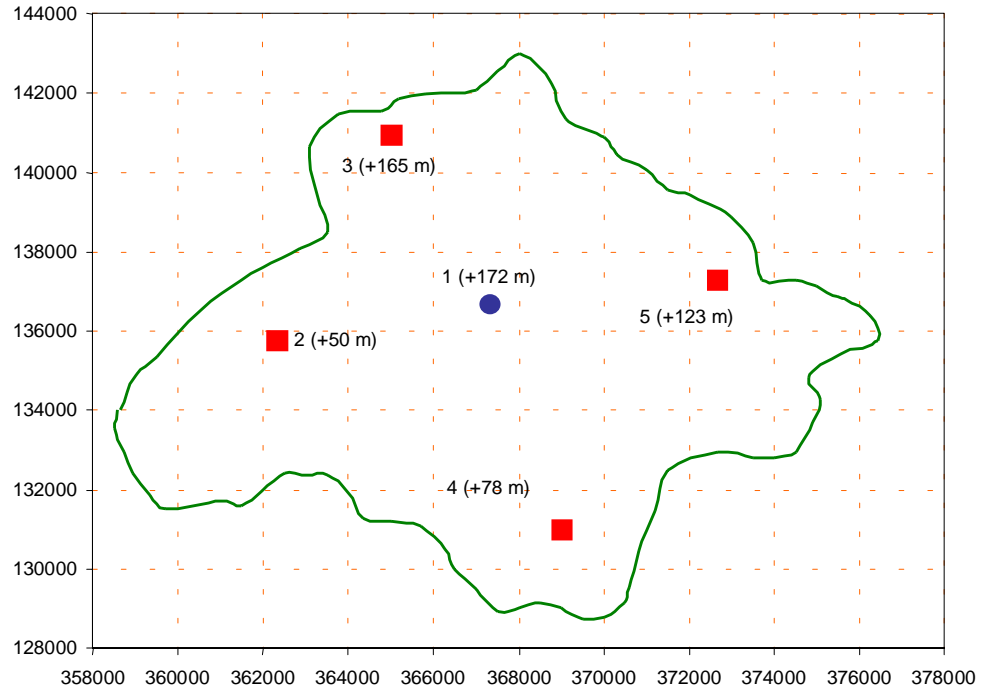


Figure 5.18: Schematic of the case study area and raingauges used for the multivariate disaggregation exercise. Each raingauge is labelled with its site number as used in this case study, and its altitude. Hourly data were used for gauge 1 (circle) whereas only daily values were used for other gauges (squares).

not used at all in any phase of the modelling. The differences between the two groups of values are less than 10%. Notably, the differences of the cross-correlation coefficients, predicted by the spatial-temporal rainfall model in this case study, from the values predicted again by the same model but using the full parameter set already found in section 5.1 (using both hourly and daily data at eight gauges) are less than 5.6%.

The simplified multivariate model was used in its form (5.16) along with the power transformation (5.23). The exponent m was chosen equal to 0.5, a value that was found (after trials) to prevent discrepancies between observed and simulated values of the statistics to be preserved (see section 5.3.7). The statistics of the transformed variables are shown in table 5.4 in a format similar to that of table 5.2. These statistics were used in the simplified rainfall model but the cross-correlations were unchanged as given in table 5.3. For the control of the proportion of dry intervals, the technique described in section 5.3.7 was used with $l_0 = 0.20\text{mm}$ and $\pi_0 = 0.40$. No special treatment was applied for the preservation of skewness.

Applying the disaggregation modelling framework, synthetic hourly rainfall series were produced for the five gauges, that of gauge 1 being identical to the historical series. The statistics of the synthetic series are compared to the historical and model statistics in tables 5.2 through 5.4. It can be observed that the statistics of the synthetic series are in good agreement with the historical and model statistics.

Property	Gauge				
	1	2	3	4	5
Proportion dry	0.84	0.85	0.84	0.85	0.84
	0.84	0.84	0.84	0.84	0.84
	0.84	0.82	0.80	0.82	0.80
Mean	0.10	0.10	0.12	0.11	0.12
	0.10	0.10	0.10	0.10	0.10
	0.10	0.10	0.12	0.11	0.12
Maximum value	8.00	7.60	5.80	4.80	4.80
	8.00	8.00	8.00	8.00	8.00
	8.00	6.70	7.10	6.90	7.20
Standard deviation	0.39	0.40	0.44	0.40	0.41
	0.39	0.39	0.39	0.39	0.39
	0.39	0.37	0.41	0.39	0.42
Skewness	7.60	7.23	5.81	6.03	5.25
	7.60	7.60	7.60	7.60	7.60
	7.60	6.90	6.89	7.91	7.22
Lag 1 autocorrelation	0.46	0.48	0.50	0.53	0.55
	0.46	0.46	0.46	0.46	0.46
	0.46	0.44	0.44	0.47	0.43

Table 5.2: Statistics of hourly rainfall depths at each gauge, obtained using multivariate disaggregation scheme. The top figure in each cell is the historical value, not used in the disaggregation model. The middle figure is the value used in the disaggregation model, which is the historical value from gauge 1. The bottom figure is the synthetic value.

A further comparison is given in figure 5.19 in terms of the autocorrelation function for higher lags, up to lag 10. Clearly, both models used, i.e. the GDSTM and the multivariate AR(1) model depart significantly from the historical autocorrelation functions (particularly the latter model). However, the synthetic autocorrelations which are dominated by the structure of the observed hourly series agree well with the historical ones.

As an additional means of comparison, some hyetographs are given in figure 5.20 (for relatively heavy rainfall) and figure 5.21 (for moderate and low rainfall). It can be seen that the disaggregation model predicted well the actual hyetographs, the only significant discrepancy being the long low intensity tails of the rainfall event of 17/01/95 (figure 5.20) generated during the period where the intensity was actually zero.

5.3.9 Conclusions

A promising new methodology for spatial-temporal disaggregation, with wide potential hydrological applicability, has been proposed. In the context of continuous simulation modelling, this provides a way to take GLM simulations of multivariate daily rainfall (incorporating spatial and temporal non-stationarity) and generate multivariate fields at fine temporal resolution. A minimum requirement for the methodology is a single temporally-disaggregated time series. As indicated in section 5.3.4

Gauge	1	2	3	4	5
1	1.00	0.84	0.80	0.72	0.83
	1.00	0.89	0.89	0.87	0.88
	1.00	0.93	0.92	0.84	0.86
2	0.84	1.00	0.82	0.80	0.83
	0.89	1.00	0.87	0.82	0.77
	0.93	1.00	0.90	0.84	0.83
3	0.80	0.82	1.00	0.79	0.86
	0.89	0.87	1.00	0.77	0.82
	0.92	0.90	1.00	0.81	0.83
4	0.72	0.80	0.79	1.00	0.87
	0.87	0.82	0.77	1.00	0.84
	0.84	0.84	0.81	1.00	0.79
5	0.83	0.83	0.86	0.87	1.00
	0.88	0.77	0.82	0.84	1.00
	0.86	0.83	0.83	0.79	1.00

Table 5.3: Cross-correlation coefficients for the five gauges at hourly level. The top figure in each cell is the historical value, not used in the disaggregation model. The middle figure is the value predicted by the spatial temporal model (used in the disaggregation model) The bottom figure is the synthetic value.

Property	Gauge				
	1	2	3	4	5
Proportion dry	0.84	0.85	0.84	0.85	0.84
	0.84	0.84	0.84	0.84	0.84
	0.84	0.82	0.80	0.82	0.80
Mean	0.11	0.11	0.13	0.11	0.13
	0.11	0.11	0.11	0.11	0.11
	0.11	0.12	0.14	0.13	0.14
Maximum value	2.83	2.76	2.41	2.19	2.19
	2.83	2.83	2.83	2.83	2.83
	2.83	2.59	2.66	2.63	2.68
Standard deviation	0.30	0.30	0.33	0.31	0.32
	0.30	0.30	0.30	0.30	0.30
	0.30	0.30	0.32	0.30	0.32
Skewness	3.13	3.18	3.00	3.13	2.90
	3.13	3.13	3.13	3.13	3.13
	3.13	2.90	2.73	2.88	2.85
Lag 1 autocorrelation	0.60	0.60	0.60	0.60	0.63
	0.60	0.60	0.60	0.60	0.60
	0.60	0.62	0.64	0.66	0.64

Table 5.4: Statistics for the power transformation of hourly rainfall depths at each gauge. The top figure in each cell is the historical value, not used in the disaggregation model. The middle figure is the value used in the disaggregation model, which is the historical value from gauge 1. The bottom figure is the synthetic value.




Figure 5.19: Comparison of autocorrelation functions as determined from historical (H1-H5 for gauges 1-5, respectively), or simulated (S2-S5 for gauges 2-5, respectively) series, or predicted from the AR(1) (Markov) and GDSTM models.

above, this may be obtained from a single-site point process model or from an existing single-site disaggregation scheme. Methods for obtaining such a series are reviewed in the following section.

Within the available time and resources, the spatial-temporal disaggregation methodology has necessarily been based on a number of assumptions, and only limited testing has been possible. The sensitivity to these assumptions should be addressed in further work. However, these preliminary results are extremely encouraging. There is considerable flexibility in the proposed scheme, and hence potential for further refinement.

5.4 Temporal disaggregation

As outlined in chapter 1 of this report, an original intention of this study was to simulate the evolution of a spatial rainfall field in continuous time by using the continuous spatial-temporal model (GDSTM) to disaggregate the predictions of daily rainfall totals given by the Generalised Linear Model (GLM). More specifically, we would simulate rainfall from the spatial-temporal model in such a manner that the daily rainfall totals accumulated at a network of sites were consistent with the predictive distribution given by the GLM. In this way, the simulation would benefit from the fine spatial-temporal resolution of the stochastic GDSTM while having the capability of the GLM to incorporate spatial nonstationarities and long-term climate changes. We emphasise that,

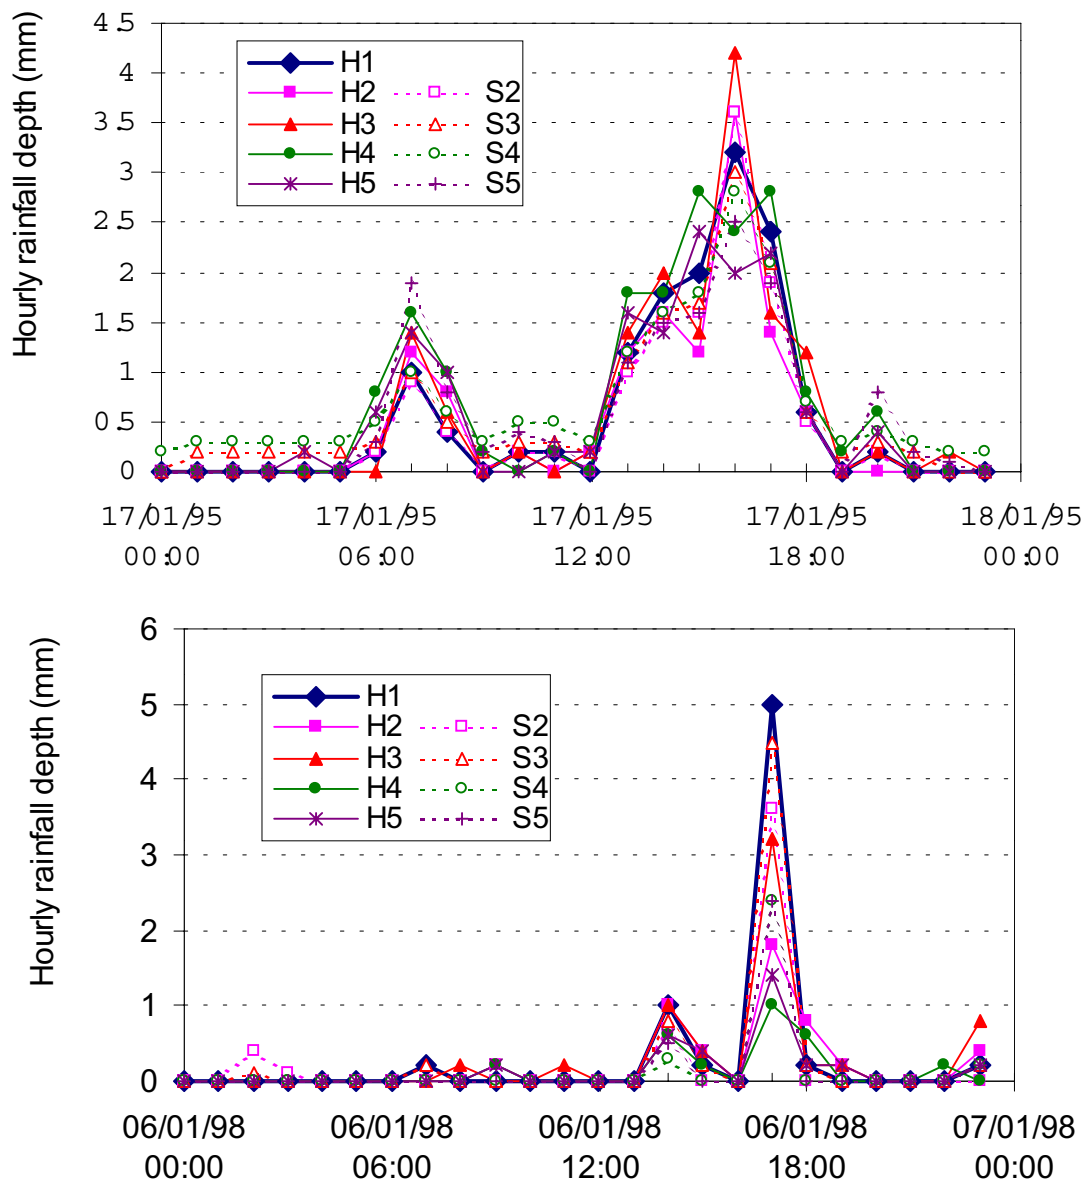


Figure 5.20: Comparison of historical (H1-H5 for gauges 1-5, respectively) and simulated (S2-S5 for gauges 2-5, respectively) hyetographs for two days with high rainfall (average daily rainfall depths 14.3 mm at 17/01/95 and 4.8 mm at 06/01/98)

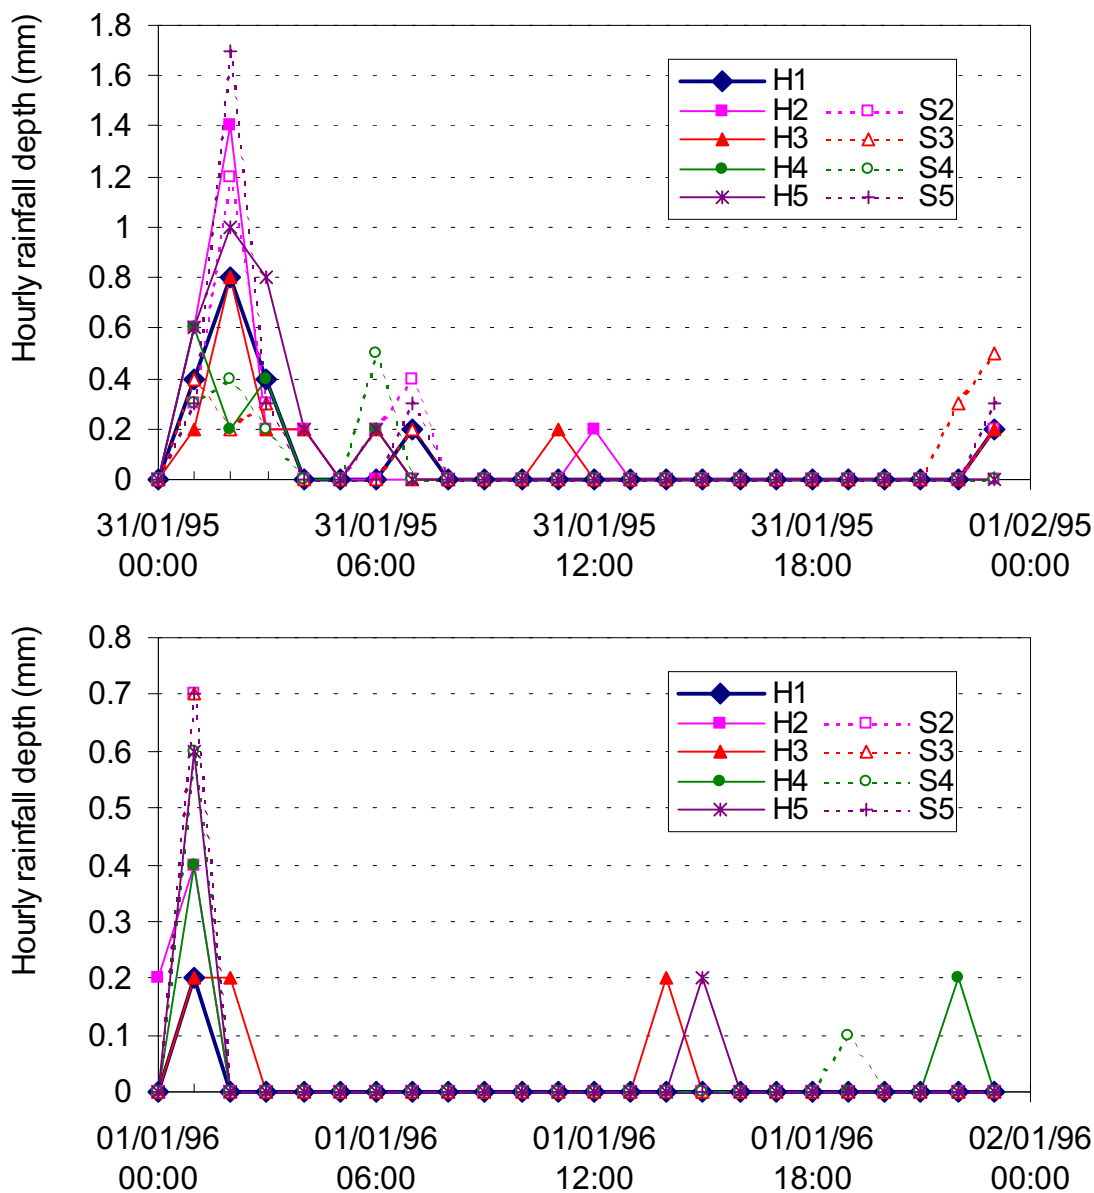


Figure 5.21: Comparison of historical (H1-H5 for gauges 1-5, respectively) and simulated (S2-S5 for gauges 2-5, respectively) hyetographs at two days with moderate and low rainfall (average daily rainfall depths 2.2 mm at 31/1/95 and 0.6 mm at 01/01/96).

in contrast with the alternative schemes for spatial-temporal disaggregation described in section 5.3 above, the idea is not to disaggregate fixed (or simulated) sequences of daily values at a set of sites (raingauges). Instead, we aim directly to provide fine time-scale sequences whose daily totals have the *joint distribution* predicted by the GLM. The proposed technique is based on the rejection sampling method of stochastic simulation (von Neumann 1951, Devroye 1986).

As a prelude to tackling the full spatial-temporal problem, we have investigated the simpler task of purely temporal disaggregation. Here the aim is to provide a fine resolution sequence at a single site, whose daily totals are consistent with the probability distributions forecast by a GLM. While this exercise is of interest in itself, its primary purpose here is to investigate the general feasibility of the approach and to indicate whether extension to a network of raingauges is practicable.

5.4.1 Disaggregation of fixed daily totals

The single-site disaggregation problem is a common one in hydrology since, for many applications, hourly rainfall totals are required when only daily totals are available. Various methods of disaggregating daily rainfall totals into realistic hourly sequences have been proposed, many of them based on models such as the Bartlett-Lewis Rectangular Pulses Model (BLRPM) of Rodriguez-Iturbe et al. (1987). An ideal solution to the problem would provide hourly totals from the (multivariate) distribution of hourly totals conditional upon the daily total. However, it is not possible to derive this distribution for models such as the BLRPM and approximate methods are necessary. For example, Glasbey, Cooper & McGechan (1995) use a long sequence of simulated rainfall data from a variant of the BLRPM to disaggregate an historical record of daily rainfall totals. For each sequence of wet days in the historical record they find the sequence of wet days in the simulated data archive that provides the best match to the observations. An alternative approach is to develop a procedure for simulating from the chosen rainfall model such that certain important sub-daily statistical properties are at least approximately preserved and the hourly totals sum to the daily total in the historical record — see, for example, Koutsoyiannis & Manetas (1996). Either of these methods could be used to simulate hourly data from a GLM for daily rainfall. We would simply simulate a sequence of daily rainfall totals from the GLM and treat them as if they were historical daily totals. Such methods may be used to obtain a single disaggregated rainfall sequence for application with the multivariate disaggregation methodology described above in section 5.3.

5.4.2 Direct generation of fine-resolution data

In contrast to the methods described above, the approach investigated here has been to condition upon the *distribution* of the daily rainfall total predicted by the GLM, rather than upon a particular realised value from this distribution. Subject to the simplifying approximation given in section D.5.2 of Appendix D, the hourly totals generated by this procedure are from the (multivariate) conditional distribution of hourly totals, given the distribution of the daily total specified by the GLM.

The methodology developed for temporal disaggregation in this study is described in Appendix D. In addition, the methodology has been applied to data from the Elmdon raingauge near Birmingham, in a student project at UCL (Keheyian 1999). For the continuous time model of rainfall, we use the BLRPM, (Rodriguez-Iturbe et al. 1987, Onof & Wheater 1993). We can, in principle,

extend the approach detailed in this report to a network of sites rather than a single site. However, two particular difficulties were encountered in the single site case: the derivation of an approximate density for the conditional distribution of positive daily rainfall totals under the BLRPM, and the fact that this density is not compatible (in terms of its intended use as a proposal density in the standard rejection sampling detailed in section D.4) with the corresponding density under the GLM in all cases. Both difficulties have been overcome for the case of temporal disaggregation, but present serious analogous problems in the spatial temporal case. There, we must derive an approximation to the conditional (multivariate) distribution of daily rainfall totals at a network of sites given the state of the rainfall process at the start of the day, or possibly a univariate summary of this distribution. This is anticipated to be extremely difficult. An additional consideration is that simulation from the spatial-temporal model is more computationally expensive than from the BLRPM. Thus, given the success of the alternative disaggregation scheme described in section 5.3, this method will not be pursued further for the moment.

5.5 Summary of work on limited data applications

5.5.1 Spatial-temporal modelling using raingauge data

Given the fact that radar data are at present limited with respect to availability, record lengths and quality, it is important to investigate the potential use of raingauge data to identify the spatial-temporal model. In the present study, data from 8 raingauges from the Brue network were used to evaluate the potential of the model for UK rainfall. The data were analysed at hourly and daily resolution.

The data were insufficient to allow the identification of variation of event properties between events within any month. The analysis showed that the rainfall fields, after allowing for monthly seasonality, can be regarded as isotropic and having no preferred velocity. The raingauge data are also highly correlated spatially, with some seasonal variability (higher in Winter, lower in Summer, as expected). Given that non-zero storm velocity could not be detected from the raingauges at the space and time scales considered, a reduced form of the model was required. Two levels of event structure (of storms within events) were generated, with no lower level of cell structure.

Having fitted the model to the available short length of record (4 years), simulation results were examined. In general, the simulation performance was excellent. Hourly and daily means, standard deviations, lag 1 and 2 autocorrelations and lag zero cross-correlation were very well reproduced, as were the proportion of dry hours and days. Only a limited investigation of extreme value properties was possible, given the available record lengths. For hourly and daily extreme values, excellent agreement with the Brue raingauge data was obtained, demonstrating that the model was able to reproduce the extreme value characteristics of the calibration data. The simulations lay below extreme values determined from longer records for the higher return period events, but this is primarily due to the limited record lengths used in model fitting.

5.5.2 Spatial-temporal rainfall disaggregation

The Generalized Linear Model has the capability to represent spatial non-stationarity and long term temporal climate non-stationarity, but at a daily time step. Preliminary work has been undertaken

to develop a new spatial-temporal disaggregation procedure. Here, we consider a GLM simulated field of daily rainfall and require a consistent set of sub-daily time series at selected locations which preserve the spatial structure of the simulated field. A single-site sub-daily time series is required, and can be produced by one of a number of single site disaggregation methods. We then apply a simplified multivariate model of hourly rainfall that preserves the basic statistical properties of the daily multivariate process and incorporates the available hourly time-series at the single site. Finally a transformation is applied which modifies the multivariate time-series to match the daily values generated by the GLM.

Results have been presented for a trial situation, based on Brue data, of 4 daily raingauges and a single sub-daily rainfall recorder. Hourly rainfall sequences are produced for each of the locations, and these preliminary results are most encouraging. Individual gauge statistics are well reproduced, including mean, standard deviation, skewness, proportion dry and the autocorrelation function, while the daily values are maintained as observed. The method needs more evaluation, but these preliminary studies suggest that the method has considerable promise.

Chapter 6

Summary and Recommendations

6.1 Summary of methods and results

6.1.1 Spatial-temporal modelling using radar data

The Poisson-based spatial-temporal model has the flexibility to allow application to event-based or continuous simulation modelling at any required space and time-step, including fine resolution in both space and time. Thus it can represent local structure in rainfall fields which may be important for particular hydrological applications, for example the response to convective rainfall systems Samuel (1999).

One limitation is the assumption of spatial stationarity. At this stage of model development this precludes the modelling of orographic effects (although such features could be incorporated in subsequent work). Temporal stationarity is also assumed; the model can incorporate monthly seasonal effects, but within-month temporal variability of the stochastic parameters (for example to differentiate between rainfall types) and longer term climatic variability have not been represented. This does not, of course, rule out between-event variability, which is intrinsic to the model.

From the test data sets investigated, it has been concluded that in general, availability of radar data is required to identify the full model structure for UK conditions. However, radar data quality has been a source of difficulty in model calibration and performance assessment. In particular, due to the general occurrence of artefacts within the data set, it was not possible clearly to identify dry periods, which has limited the comparisons between observed and simulated sequences. Limited available record lengths have also restricted our ability to fit the models; with a data set of less than four years, there is a shortage of information on event sequences from which to identify storm arrivals and, in particular extreme values. In addition, it should be noted that no additional radar calibration has been carried out; comparisons are therefore primarily between model simulated and radar-derived properties, which may not be wholly consistent with raingauge data.

Despite the data restrictions, overall the model has performed extremely well in reproducing basic statistical properties over different space and time-scales, including temporal and spatial autocorrelations. The proportions of wet pixels (averaged over space and time) are also very well reproduced across the range of space and time-scales examined. Only limited evaluation of extreme value properties has been possible, but the results, in terms of intercomparison of the model and

the radar data set, are most encouraging. The model reproduces well the observed (radar) hourly extreme values at $2 \times 2 \text{ km}^2$ and $16 \times 16 \text{ km}^2$. Daily values show an increasing underestimation at high return periods, but lie well within 95% confidence limits. They agree rather better with observed raingauge data. It must be emphasised that within the timescale of the project there has been no opportunity to refine the model parameters or distributions used to define rainfall event properties, or the sampling strategy by which event arrivals are associated with particular parameter combinations. We note that the assumption of spatial stationarity does not appear to have been a significant limitation for the location and areal scale investigated.

For this project, we have fitted the full spatial-temporal model to a single data set, dominated by frontal rainfall systems. However, we consider that the model has widespread applicability to different rainfall regimes. As noted above, it has also been found to reproduce extremely localised convective rainfall (Arizona), and the underlying temporal structure, as implemented in single site models, has now been widely tested for different locations around the world (including S.Africa, Colorado, Arizona, and several UK and other European locations).

The future for the model is highly promising, but further testing on other data sets is required. The comparisons here have been affected by radar data quality, and there is more work to be done on combining radar and raingauge data to clean up radar images and provide longer time-series for the identification of temporal properties and the spatial structure of extreme events. A possible extension to the basic spatial-temporal model could include spatial non-stationarity using a scaling approach. The model could also be readily adapted to represent climate change.

6.1.2 Spatial-temporal modelling using raingauge data

Given the fact that radar data are at present limited with respect to availability, record lengths and quality, it is important to investigate the potential use of raingauge data to identify the spatial-temporal model. Samuel (1999) considered the highly localised thunderstorm rainfall at Walnut Gulch, Arizona, observed using an exceptional density of raingauges (> 1 per 2 km^2), and found that a reduced form of the model, with circular cells, no movement and a reduced level of aggregation (cells and storms, with no higher level event structure) could be identified. In the present study, data from the Brue raingauge network were used to evaluate the potential of the model for UK rainfall. Although more than 50 raingauges are available, a subset of 8 was used as more representative of UK raingauge densities. The data were analysed at hourly and daily resolution.

A major issue is the reduced level of information available to identify spatial structure. For the raingauge data used, information was insufficient to allow the identification of variation of event properties between events within any month. The resulting analysis therefore showed that the rainfall fields, after allowing for monthly seasonality, can be regarded as isotropic and having no preferred velocity. This contrasts with the analysis of the larger-scale Wardon Hill radar field, which detected preferential velocities (and directions). The raingauge data are also highly correlated spatially, with some seasonal variability (higher in Winter, lower in Summer, as expected). Given that non-zero storm velocity could not be detected from the raingauge data at the space and time scales considered, a reduced form of the model was required. As in Samuel's study, only two levels of event structure (of storms within events) were generated, but here with no lower level of cell structure.

Having fitted the model to the available short length of record (4 years), simulation results

were examined. Due to time constraints, 5 points from the simulated field were taken as the basis of comparison. In general, the simulation performance was excellent. Hourly and daily means, standard deviations, lag 1 and 2 autocorrelations and lag zero cross-correlation were very well reproduced, as were the proportion of dry hours and days. Only a limited investigation of extreme value properties was possible, given the available record lengths. For hourly and daily extreme values, excellent agreement with the Brue raingauge data was obtained, demonstrating that the model was able to reproduce the extreme value characteristics of the calibration data. The simulations lay below extreme values determined from longer records for the higher return period events, but this is primarily due to the limited record length used in model fitting.

6.1.3 Generalized Linear Models (GLMs)

Generalized Linear Models (GLMs) have been shown to provide a powerful and flexible tool to identify spatial and temporal structure in daily rainfall fields for the analysis of raingauge data. Spatial non-stationarity can be detected, as illustrated by the results presented in chapter 4, including complex topographic and location effects. Given current concerns for the impact of climate change on rainfall, it is also important to note that long-term change can be identified, and readily included in simulations of future rainfall sequences. The methodology is quick and simple to use, and the software is freely available (see Appendix B.2).

The relative simplicity of the fitted models means that they can readily be used to generate large numbers of realisations of long simulated sequences, thus allowing definition of extreme rainfall frequencies, including the dependence of frequency on climate variability. Although the results presented above are based on analysis of raingauge data (i.e. from a limited number of raingauge sites), the model can simulate rainfall at any location of interest within the observation field.

The results from three contrasting areas have been presented. An extensive study of Irish rainfall data demonstrates that although the daily rainfall data are extremely noisy, subtle spatial and temporal effects can be detected and quantified. The capacity to simulate long sequences of daily rainfall under a changing climate regime is demonstrated. The model performance in simulating daily and 3-monthly rainfall is good. Forecasting the probability of rainfall occurrence at a site is excellent, and the basic statistical properties of point rainfall time-series, including autocorrelation and the proportion of wet days, are well reproduced. Considering simulation performance at pairs of sites, there is a slight under-representation of the proportion of days when both sites are wet, but the inter-site correlation for wet days is excellent.

A second example (the Brue) is considered where a subset of the available data is used as a test case for the situation in which data from the catchment of interest are limited. It has been demonstrated that additional gauges from sites outside the catchment can be used within the GLM scheme to improve simulation performance, and validation tests show good predictive performance; discrepancies fall within reasonable confidence limits. Independent tests of the occurrence model show excellent results; the rainfall amounts are slightly over-predicted, but this is probably due to the influence of the distant raingauges. Within the scope of the present study it has not been possible to explore this aspect further, but options to improve performance are suggested, in particular the use of indicator variables to increase the weighting given to data from the catchment itself.

The third example, the Blackwater, has an extensive data-set (44 gauges, over an area of 400km²), and a limited modelling exercise was carried out at the Institute of Hydrology. It is

demonstrated that good results can rapidly be obtained, and, as for the Irish example, that complex spatial effects are apparent in the data.

6.1.4 Hybrid approaches

The main limitation of the GLM methodology is that sub-daily rainfall cannot be easily modelled, hence the original concept for the project was to combine the strengths of the GLM and spatial-temporal model in a hybrid approach. Spatial non-stationarity and long term temporal climate non-stationarity could be represented using the GLM, with the sub-daily spatial-temporal structure produced by the spatial-temporal model, the two being combined using a rejection sampling approach. However, as explained in chapter 5, although the possibility was demonstrated for a single site model, practical implementation for the full spatial-temporal model presents substantial theoretical and computational problems which remain to be addressed.

As an alternative approach, preliminary work has been undertaken to develop a new spatial-temporal disaggregation procedure. The methodology is flexible, powerful, and has widespread potential utility. It could, for example, be applied to the common situation in which a number of daily raingauges are available, supplemented by one or two gauges with sub-daily data. Here, we consider a GLM simulated field of daily rainfall; we require a consistent set of sub-daily time series at selected locations which preserve an appropriate spatial structure of the simulated field. The method is, in principle, the same as that originally suggested. However, in the new approach, a simplified spatial-temporal model is used, with a subsequent stage of transformation.

The first requirement is a single-site sub-daily time series. This can be produced by the single site disaggregation method described in Appendix D, or using a recent development by Koutsoyiannis & Onof (2000), building on previous work by Koutsoyiannis & Manetas (1996) and Koutsoyiannis (1994). Using the structure of a single site Poisson cluster model, a resampling procedure is used to generate a subdaily realisation consistent with the daily value at the single site.

We then apply a simplified multivariate model of hourly rainfall that preserves the basic statistical properties of the daily multivariate process and incorporates the available hourly time-series at the single site. Finally a transformation is applied which modifies the multivariate time-series to match the daily values generated by the GLM.

Results are presented for a trial situation, based on Brue data, of 4 daily raingauges and a single sub-daily rainfall recorder. Hourly rainfall sequences are produced for each of the locations, and these preliminary results are most encouraging. Individual gauge statistics are well reproduced, including mean, standard deviation, skewness, proportion dry and the autocorrelation function, while the daily values are maintained as observed. The method needs more evaluation, but these preliminary studies suggest that the method has considerable promise.

6.2 Recommendations

The problem addressed in this project, namely the simulation of long temporal sequences of spatial rainfall, is extremely challenging, and the work described above is a first substantive attempt to develop an appropriate set of methodologies and subject them to a set of rigorous tests using a variety of UK rainfall data.

The full spatial-temporal model is the most powerful of these methods, providing simulated sequences in continuous space and time, and is thus suitable for representing both fine-scale local structure and large scale structure, albeit with the current limitation of spatial stationarity. Tests using Wardon Hill radar data have shown that the model, even without any opportunity for refinement of parameters or distributions, performs extremely well in preserving rainfall spatial and temporal properties, and in matching the extreme value distributions of the radar and local raingauge data. Further refinement of the model, based on this data set, can be expected to improve performance further. It would be beneficial to combine this with a radar recalibration exercise, to improve confidence in the identification of dry periods, as well as rainfall intensities. The next step for this model is to extend the analysis of performance to a different (and contrasting) radar data set, such as that currently available for the North-West of England.

The main restriction in general applicability of the full spatial-temporal model at present lies in the radar data. These are limited in record length and data quality. For widespread application, ideally, long records of error free radar data are required. These will become progressively available in the UK as archiving is extended and post-processing quality control is improved. For example, in the USA, extensive spatial data sets (e.g. 850000 km² of 4 × 4km² hourly post-processed NEXRAD data) are routinely available. In the absence of such data, further work is required to improve the quality (through improved calibration) of available radar data for model testing, and to combine the information available from long raingauge records.

If raingauge data only are available, certain features of the full spatial-temporal model cannot be readily identified. In both UK and US applications, a reduced level of model has been required. Given this simplification, results were generally excellent for the case examined (the Brue catchment). However, further test applications are required across a range of catchment scales before general UK use could be confidently recommended.

The GLM modelling approach has now been tested on 3 different data sets, and its ability to discriminate subtle spatial and temporal variability from extremely noisy daily data has been clearly demonstrated. It is demonstrably a powerful tool for data analysis. More work could be done to improve some underestimation of simultaneous rainfall at multiple sites, but the main limitation of this method for the current application (continuous simulation modelling for flood design) is the daily time step. An additional stage of modelling is required to produce sub-daily resolution spatial fields.

The development of a new spatial-temporal disaggregation procedure (chapter 5) appears extremely promising as a solution to the problem of extended use of the GLM procedure (incorporating spatial and temporal non-stationarity). However, further evaluation is required to gain confidence in the proposed methodology.

6.3 Strategic Priorities

The research reported above represents, in our view, a significant step forward in the development of new methods for the representation of spatial rainfall for hydrological modelling. A family of alternative, but inter-related, simulation tools has been developed and tested. Contrasting strengths and weaknesses have been identified, and for each approach considered, extremely encouraging performance has been reported. However, this work should be regarded as a feasibility study. The

GLM methodology can be applied immediately; we do not feel confident that the other methods developed can be taken directly into practice without further work.

6.3.1 Medium term applications

As remotely-sensed rainfall fields become more accurate and archives more extensive (say over the next 5 years in the UK), there will be extensive high quality data available to calibrate the full spatial-temporal model. This model has the major advantage that rainfall fields are simulated in continuous space and time, thus giving complete flexibility in spatial applications. The main issues are the definition of the event occurrence model, and characterisation of the regional variability of rainfall types. However, testing on a contrasting radar data set will provide short-term guidance concerning the transferability of parameters identified for South West England. A UK wide methodology could be in place in 5 years time.

Recommendation: Continuation of strategic research through model and data refinement using Wardon Hill and North West radar/raingauge data (2 years), followed by a full regionalisation study for national implementation.

6.3.2 Short term applications

Two approaches have been developed. It has been shown that a reduced order of spatial-temporal model can be fitted to raingauge data, with good results. Alternatively, the Generalized Linear Modelling approach can be combined with a new spatial-temporal disaggregation methodology. These methods can be implemented with currently available data. The latter offers greater flexibility for use with limited data (although accuracy of results is restricted by the data available), and can accommodate temporal and spatial heterogeneity, however was a development that emerged in the final stages of the current research programme. More rigorous testing and validation of assumptions is required for both of these simplified approaches, as well as testing over a range of scales. Although the GLM has been evaluated on contrasting scales and climates, the coupled GLM/disaggregation method has been tested on the Brue only. Intercomparison with the full spatial-temporal modelling approach would be beneficial to identify the loss of information in generated rainfall fields associated with the simplified methodologies.

Thus far, assessment of all methods has been based on a range of statistical properties of the modelled rainfall fields. There is an important underlying question of the relative importance of particular features of the modelled fields for flood estimation. Each of the methods has flexibility in parameter identification to enhance particular features of the simulated fields, usually at the expense of others. It is therefore recommended that a pilot study be undertaken, using distributed rainfall-runoff simulation models, to study runoff sensitivity to spatial rainfall properties.

Recommendation: a 12 month extension would be sufficient to undertake the required further testing of the simplified approaches. At that point, firm recommendations could be made for a national strategy for implementation. If combined with analysis of rainfall-runoff simulation performance, an 18 month study would be required to address all of these elements.

Bibliography

- Abramowitz, M. & Stegun, I. (1965), *Handbook of mathematical functions: with formulas, graphs and mathematical tables*, Dover, New York.
- Ash, A. & Shwartz, M. (1999), ' R^2 : a useful measure of model performance when predicting a dichotomous outcome', *Statist. Med.* **18**, 375–384.
- Austin, P. & Houze, R. (1972), 'Analysis of the structure of precipitation patterns in New England', *J. Appl. Met* **11**, 926–935.
- Besag, J. E. (1974), 'Spatial interaction and the statistical analysis of lattice systems', *J. R. Statist. Soc., Series B* **36**, 192–236.
- Besag, J. & Green, P. J. (1993), 'Spatial statistics and bayesian computation', *J. R. Statist. Soc., series B* **55**, 25–37.
- Brooks, S. P. (1998), 'Markov chain Monte Carlo and its application', *The Statistician* **47**, 69–100.
- Calenda, G. & Napolitano, F. (1999), 'On parameter estimation of Neyman-Scott processes for temporal point rainfall simulation', *J. Hydrol.* **in press**.
- Calver, A., Lamb, R. & Morris, S. (1999), 'River flood frequency estimation using continuous runoff modelling', *Proc. Instn. Civil Engrs. Wat., Marit. & Energy* **136**, 225–234.
- Chandler, R. (1998*a*), Model checking, in P. Armitage & T. Colton, eds, 'Encyclopedia of Biostatistics', Wiley, Chichester.
- Chandler, R. (1998*b*), Orthogonality, in P. Armitage & T. Colton, eds, 'Encyclopedia of Biostatistics', Wiley, Chichester.
- Chandler, R. & Wheeler, H. (1998*a*), Climate change detection using Generalized Linear Models for rainfall — a case study from the West of Ireland. I. Preliminary analysis and modelling of rainfall occurrence, Technical report, no. 194, Department of Statistical Science, University College London. <http://www.ucl.ac.uk/Stats/research/abstracts.html>.
- Chandler, R. & Wheeler, H. (1998*b*), Climate change detection using Generalized Linear Models for rainfall — a case study from the West of Ireland. II. Modelling of rainfall amounts on wet days, Technical report, no. 195, Department of Statistical Science, University College London. <http://www.ucl.ac.uk/Stats/research/abstracts.html>.
- Coe, R. & Stern, R. (1982), 'Fitting models to daily rainfall.', *J. Appl. Meteorol.* **21**, 1024–1031.

- Collier, C. (1989), *Applications of weather radar systems: A guide to uses of radar data in meteorology and hydrology*, Ellis Horwood, Chichester.
- Cox, D. & Hinkley, D. (1974), *Theoretical Statistics*, Chapman & Hall, London.
- Cox, D. & Isham, V. (1980), *Point Processes*, Chapman and Hall, London.
- Cox, D. & Isham, V. (1988), 'A simple spatial-temporal model of rainfall', *Proc. R. Soc. Lond.* **A415**, 317–328.
- Cox, D. & Isham, V. (1994), Stochastic models of precipitation., in V. Barnett & K. Turkman, eds, 'Statistics for the Environment 2: Water related issues', Wiley, Chichester, pp. 3–18.
- Cox, D. & Snell, E. (1989), *Analysis of binary data (2nd edition)*, Chapman and Hall, London.
- Cox, D. & Wermuth, N. (1996), *Multivariate Dependencies: models, analysis and interpretation*, Chapman and Hall, London.
- Daly, D. (1992), 'A report on the flooding in the Gort-Ardrahan area', Geological Survey of Ireland Report, Dublin.
- Damien, P., Wakefield, J. & Walker, S. (1999), 'Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables', *J. R. Statist. Soc., series B* **61**, 331–344.
- Dawid, A. (1986), Probability forecasting, in S. Kotz & N. Johnson, eds, 'Encyclopedia of Statistical Sciences', Wiley, New York.
- DETR (1999), *Indicators of climate change in the UK*, Department of the Environment, Transport and the Regions, Centre for Ecology and Hydrology. Editors: M.G.R. Cannell, J.P. Palutikof and T.H. Sparks.
- Deutsch, C. & Journel, A. (1992), *GSLIB: Geostatistical software and user's guide*, Oxford University Press, New York.
- Devroye, L. (1986), *Non-uniform random variate generation*, Vol. New York, Springer-Verlag.
- Dobson, A. (1990), *An Introduction to Generalized Linear Models*, Chapman and Hall, London.
- DOE (1996), *UK Climate Change Impacts Review Group: Review of the potential effects of climate change in the UK*, Vol. Report No.2, Department of the Environment, HMSO.
- Edwards, R. G. & Sokal, A. D. (1988), 'Generalisation of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithms', *Phys. Rev. D* **38**, 2009–2012.
- Efron, B. & Tibshirani, R. (1993), *An Introduction to the Bootstrap*, Chapman and Hall, London.
- Emrich, L. & Piedmonte, M. (1991), 'A method for generating high-dimensional multivariate binary variates', *Amer. Statistician* **45**, no.4, 302–304.
- Fahrmeir, L. & Tutz, G. (1994), *Multivariate Statistical Modelling Based on Generalized Linear Models*, Springer-Verlag, New York.
- Fisher, R. A. (1915), 'frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population', *Biometrika* **10**, 507–521.

- Fisher, R. A. (1936), 'The use of multiple measurements in taxonomic problems', *Annals of Eugenics (London)* **7**, 179–188.
- Foufoula-Georgiou, E. (1998), On scaling theories of space-time rainfall: some recent results and open problems, in O. Barndorff-Nielsen, V. Gupta, V. Perez-Abreu & E. Waymire, eds, 'Stochastic Methods in Hydrology', World Scientific, pp. 25–72.
- Foufoula-Georgiou, E. & Lettenmaier, D. (1987), 'A markov renewal model for rainfall occurrences', *Water Resources Research* **23**, no. 5, 875–884.
- Geman, S. & Geman, D. (1984), 'Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images', *IEEE Trans. Pattern Anal. Machine Intel.* **PAMI-6**, 721–741.
- Glasbey, C., Cooper, G. & McGechan, M. (1995), 'Disaggregation of daily rainfall by conditional simulation from a point process model', *J. Hydrol.* **165**, 1–9.
- Green, P. (1984), 'Iteratively reweighted least squares for Maximum Likelihood estimation, and some robust and resistant alternatives', *J. R. Statist. Soc., Series B* **46**, No. 2, 149–192.
- Gupta, V. & Waymire, E. (1993), 'A statistical analysis of mesoscale rainfall as a random cascade', *J. Appl. Meteor.* **32**, 251–267.
- Guyen, A. K. (1951), 'The frequency distribution of the product-moment correlation coefficient in random samples of any size drawn from non-normal universes', *Biometrika* **38**, 219–247.
- Gyasi-Agyei, Y. & Willgoose, G. (1997), 'A hybrid model for point rainfall modelling', *Water Resources Research* **33**, no. 7, 1699–1706.
- Hansen, L. R. (1982), 'Large sample properties of generalized method of moments estimators.', *Econometrica* **50**, 1029–1054.
- Hougaard, P. (1982), 'Parametrizations of nonlinear models', *J. R. Statist. Soc., Series B* **44**, 244–252.
- Institute of Hydrology (1999), 'Flood estimation handbook (5 volumes)', Institute of Hydrology, Wallingford.
- Joe, H. (1997), *Multivariate models and dependence concepts*, Vol. 73 of *Monographs on Statistics and Applied Probability*, Chapman & Hall, London.
- Johnson, N. & Kotz, S. (1970), *Continuous univariate distributions*, Vol. I and II, Houghton Mifflin, Boston.
- Kakou, A. (1997), Point process based models of rainfall, PhD thesis, Department of Statistical Science, University College London.
- Keheyian, R. (1999), 'Rainfall modelling for hydrology', Final year undergraduate project, Department of Statistical Science, University College London.
- Kendall, M. & Ord, J. (1990), *Time Series (third edition)*, Edward Arnold.
- Khaliq, M. & Cunnane, C. (1996), 'Modelling point rainfall occurrences with the modified bartlett-lewis rectangular pulses model', *J. Hydrol.* **180**, 109–138.

- Kotegoda, N. (1980), *Stochastic Water Resource Technology*, Macmillan, London.
- Koutsoyiannis, D. (1994), 'A stochastic disaggregation method for design storm and flood synthesis', *Journal of Hydrology* **156**, 193–225.
- Koutsoyiannis, D. (1999), 'Optimal decomposition of covariance matrices for multivariate stochastic models in hydrology', *Water Resources Research* **35**, no.4, 1219–1229.
- Koutsoyiannis, D. (2000), 'Coupling stochastic models of different time scales', *Water Resources Research* **under review**.
- Koutsoyiannis, D. & Manetas, A. (1996), 'Simple disaggregation by accurate adjusting procedures', *Water Resources Research* **32**, no. 7, 2105–2117.
- Koutsoyiannis, D. & Onof, C. (2000), 'Rainfall disaggregation using adjusting procedures on a Poisson cluster model', In preparation.
- Krzanowski, W. (1988), *Principles of Multivariate Analysis*, Oxford University Press.
- Lasdon, L. & Smith, S. (1992), 'Solving large nonlinear programs using GRG', *ORSA Journal on Computing* **4**, no. 1, 2–15.
- Lunn, A. & Davies, S. (1998), 'A note on generating correlated binary variables', *Biometrika* **85**, No.2, 487–490.
- Marshall, R. (1996), 'Causal space-time multifractal processes: predictability and forecasting of rainfall fields', *J. Geophys. Res.* **26**, 333–346.
- Mason, J. (1986), 'Numerical weather prediction.', *Proc. R. Soc. Lond.* **A407**, 51–60.
- Mayes, J. (1995), 'Changes in the distribution of annual rainfall in the British Isles', *J. CIWEM* **9**, 531–539.
- McCullagh, P. & Nelder, J. (1989), *Generalized Linear Models (second edition)*, Chapman and Hall, London.
- Michaud, J. & Sorooshian, S. (1994), 'Effect of rainfall sampling errors on simulations of desert flash floods', *Water Resources Research* **30**, no. 10, 2765–2775.
- Naden, P. (1992), 'Spatial variability in flood estimation for large catchments: the exploitation of channel network structure', *Hydrological Sciences Journal* **37**, 53–71.
- NERC (1975), 'Flood studies report', Report of the Natural Environment Research Council (5 volumes).
- Ngirane-Katashaya, G. & Wheeler, H. (1985), 'Hydrograph sensitivity to storm kinematics', *Water Resources Research* **21**, no. 3, 337–345.
- Northrop, P. (1996), Modelling and statistical analysis of spatial-temporal rainfall fields, PhD thesis, Department of Statistical Science, University College London.
- Northrop, P. J. (1998), 'A clustered spatial-temporal model of rainfall.', *Proc. R. Soc. Lond.* **A454**, 1875–1888.

- Onof, C. & Wheater, H. (1993), 'Modelling of British rainfall using a random parameter Bartlett-Lewis rectangular pulse model', *J. Hydrol.* **149**, 67–95.
- Onof, C. & Wheater, H. (1994), 'Improvements to the modelling of british rainfall using a modified random parameter bartlett-lewis rectangular pulses model', *J. Hydrol.* **157**, 177–195.
- OPW (1998), *An investigation of the flooding problems in the Gort-Ardrahan area of South Galway, by Southern Water Global and Jennings O'Donovan and partners*, Vol. Final Report, Office of Public Works, Dublin.
- Over, T. & Gupta, V. (1996), 'A space-time theory of mesoscale rainfall using random cascades', *J. Geophys. Res.* **101 (D21)**, 319–331.
- Park, C., Park, T. & Shin, D. (1996), 'A simple method for generating correlated binary variables', *Am. Statistician* **50**, 306–310.
- Perica, S. & Foufoula-Georgiou, E. (1996), 'Linkage of scaling and thermodynamic parameters of rainfall: results from midlatitude mesoscale convective systems', *J. Geophys. Res.* **101(D3)**, 7431–7448.
- Petterssen, S. (1956), *Weather Analysis and Forecasting*, Vol. 2, McGraw Hill, New York.
- Pfister, C. (1992), Monthly temperature and precipitation in central Europe 1525–1979: quantifying documentary evidence on weather and its effects, in R. Bradley & P. Jones, eds, 'Climate since A.D. 1500', Routledge, London, pp. 118–142.
- Press, W., Teukolsky, S., Vetterling, W. & Flannery, B. (1992), *Numerical Recipes in FORTRAN (second edition)*, Cambridge University Press.
- Priestley, M. (1981), *Spectral Analysis and Time Series*, Academic Press.
- Rodriguez-Iturbe, I., Cox, D. & Isham, V. (1987), 'Some models for rainfall based on stochastic point processes.', *Proc. R. Soc. Lond.* **A410**, 269–288.
- Rodriguez-Iturbe, I., Cox, D. & Isham, V. (1988), 'A point process model for rainfall: further developments.', *Proc. R. Soc. Lond.* **A417**, 283–298.
- Samuel, C. (1999), Stochastic rainfall modelling of convective storms in Walnut Gulch, Arizona, PhD thesis, Department of Civil Engineering, Imperial College.
- Singh, V. (1997), 'Effect of spatial and temporal variability in rainfall and watershed characteristics on stream flow hydrograph', *Hydrological Processes* **11**, 1649–1669.
- Stern, R. & Coe, R. (1984), 'A model fitting analysis of rainfall data (with discussion).', *J. Roy. Stat. Soc.* **A147**, 1–34.
- Tessier, Y., Lovejoy, S. & Schertzer, D. (1993), 'Universal multifractals: Theory and observations for rain and clouds', *J. of Appl. Meteor.* **32**, 223–249.
- Velghe, T., Troch, P., de Troch, F. & Van de Velde, J. (1994), 'Evaluation of cluster-based rectangular pulses point process models for rainfall', *Water Resources Research* **30**, no. 10, 2847–2857.

- Verhoest, N., Troch, P. & de Troch, F. (1997), 'On the applicability of Bartlett-Lewis rectangular pulses models for calculating design storms at a point', *J. Hydrol.* **202**, 108–120.
- von Neumann, J. (1951), 'Various techniques used in connection with random digits, 'Monte Carlo Method'', *US Nat. Bur. Stand. Appl. Math. Ser.* **12**, 36–38.
- Wei, B.-C. (1997), *Exponential Family Nonlinear Models*, Vol. 130 of *Lecture Notes in Statistics*, Springer, Singapore.
- Wheater, H., Isham, V., Cox, D., Chandler, R., Kakou, A., Northrop, P., Oh, L., Onof, C. & Rodriguez-Iturbe, I. (1999), 'Spatial-temporal rainfall fields: modelling and statistical aspects', *Hydrological and Earth Systems Science* . To appear.