Stochastic Environmental Research and Risk Assessment manuscript No.
(will be inserted by the editor)

Spatial-temporal rainfall models based on Poisson cluster processes

Nanda R. Aryal · Owen D. Jones

Received: date / Accepted: date

Abstract We fit stochastic spatial-temporal models to high-resolution rainfall radar data using Approximate Bayesian Computation (ABC). We consider models constructed from cluster point-processes, starting with the model of Cox, Isham and Northrop, which is the current state of the art. We then generalise this model to allow for more realistic rainfall intensity gradients and for a richer covariance structure that can capture negative correlation between the intensity and size of localised rain cells.

The use of ABC is of central importance, as it is not possible to fit models of this complexity using previous approaches. We also introduce the use of Simulated Method of Moments (SMM) to initialise the ABC fit.

Keywords Rainfall; spatial-temporal; spatiotemporal; Approximate Bayesian Computation.

1 Introduction

Our interest in spatial-temporal rainfall models comes from the use of rainfall simulators to understand the responses of hydrological systems to rainfall events. It has been argued by many authors (for example Wheater et al. [26], Segond et al. [22], Chander et al. [6]) that using simulated rainfall with realistic spatial and temporal variation is an effective way of understanding the range and frequency of responses from any given hydrological system. In particular, extreme responses such as flash floods and debris flows are regularly triggered by storm cells with an area of 4 km^2 and a duration of 30 minutes [10], so we want our models to have this level of detail.

Spatial-temporal rainfall models can typically be broken down into two main parts: a model for the frequency, duration and extent of rainfall events, and a separate model for the spatial and temporal structure of a single event.

The University of Sydney, Australia · Cardiff University, UK

There is no precise definition of what constitutes a rainfall event, and in practice it is somewhat model dependent. A rainfall event will not in general exhibit rainfall contiguous in space and time, but formally we imagine that during an event any point in space will be rained on more often than not. Within an event we do see contiguous convex patches of rain which we call rain cells. In this paper our rainfall events have a timescale of hours and a spatial scale of thousands of km^2; we required that during an event at least 20% of the area of interest should receiving rain at any given time.

We will concern ourselves solely with models for a single event, for which we consider stochastic cluster-type models. These models are constructed to mimic the cell-like structure of rainfall events, and are straight-forward to simulate. We start with the model of Cox & Isham [7] as extended by Northrop [15], which we call the CIN model, and then introduce a number of extensions. The CIN model and our extensions of it are all stationary models for the interior of a rainfall event; definitions are given in Section 2. To fit them we use high-resolution rainfall radar data, which gives rainfall intensity on a 1 km^2 grid every 6 minutes. In what follows we use a single rainfall event covering an area of 180×180 km^2 for a duration of 4 hours, described below. The thesis of Aryal [2] considers other rainfall events at the same and different locations and obtains similar results.

The CIN model assumes that rain cells have constant intensity and does not allow dependence between the intensity and size of a cell. Our first extension—the CIN-1 model—incorporates more realistic rainfall intensity gradients and gives a quantitatively better fit. This mirrors recent results for point rainfall models [17]. Our second extension—the CIN-2 model—is designed to capture the negative correlation between the intensity and size of localised rain cells that has previously been observed in point rainfall models [11]. The results here are mixed, with moderate evidence of the negative correlation but without a better fit overall. The implications of this are discussed in the Conclusions section.

Because it has an intractable likelihood function, in the past the CIN model has been fitted using the Generalized Method of Moments (GMM) [26]. GMM fitting matches theoretical and observed moments of the process, and thus is restricted to moments for which you have an analytic expression. However we do not have analytic expressions for the moments of our extended processes, so GMM fitting is no longer an option and we instead use Approximate Bayesian Computation (ABC). ABC fitting compares the observed process to simulations, and places no restrictions on the statistics used to compare them. It also has the advantage of providing credible intervals for the estimated parameters. We give a brief description of ABC in Section 3, before dealing with the specifics of the case in hand. In [3] the authors give a comparison of GMM and ABC fitting for the Bartlett-Lewis point rainfall model, demonstrating the advantages of ABC.

There is an extensive literature on modelling and simulating rainfall at a single point. In particular there are many models based on cluster point processes, going back to the rectangular pulse model of Rodriguez-Iturbe, Cox and

Isham [21]. The CIN model that we start with is a natural spatial-temporal extension of these models, and is the most sophisticated spatial-temporal model of this type to have appeared in the literature to date. The surveys of Onof et al. [16] and Wheater et al. [25] provide a good summary of the current state of the art for these models.

There are two main alternative approaches to spatial-temporal rainfall modelling. The first, dating back to Le Clerc and Schaake [12], is to extend point models into space using techniques such as the depth area reduction factor (DARF). These approaches are often used to simulate a single extreme rainfall event, but by their construction they ignore the internal spatial structure.

The second alternative is to use random fields rather than cluster processes. Such models do capture the spatial structure of a rainfall event though lack the physical intuition of processes built from cluster point processes. Recent examples of this type of model can be found in the papers of Paschalis et al. [18] and Benoit et al. [5]. These types of model have also been successfully incorporated into more general weather generators [1,19].

The data

The CIN model and our generalisations are stationary and are used to model the "interior" of a rainfall event. We will suppose that we have observations of the rainfall in some finite space-time window $A \times [0, T]$, where T is chosen so that the leading and trailing edges of the rainfall event are not observed. For this study we used radar data collected at Laverton, Melbourne, on 24th September 2016 from 12:54 to 16:48 hours, calibrated by the Australian Bureau of Meteorology using rain-gauge data. The data gives rainfall depth in $10^{-2}mm$ averaged over 1 km^2 pixels every 6 minutes for a period of 4 hours. The radar covers a circular region of 128 km radius, but we restrict ourselves to a square study area of size 180×180 km^2. To reduce the noise in the data, rainfall depths of less than 0.01 mm for a six minute interval are rounded down to 0.

A contour plot of the spatial rainfall intensity at a single time-point is given in Figure 1(a). Looking at the spatial maximum and mean over the study period, and the percentage of the study area covered by rain (Figure 1(b–d)), it is reasonable to consider this space-time region as being in the interior of a rainfall event.

2 The Cox-Isham-Northrop model and extensions

The Cox-Isham-Northrop (CIN) rainfall model is a spatial-temporal stochastic model for a rainfall event, constructed using a cluster point process. The cluster process is constructed by taking a primary point process, called the storm arrival process, and then attaching to each storm center a finite secondary point process, called a cell process. To each cell center we then attach a rain

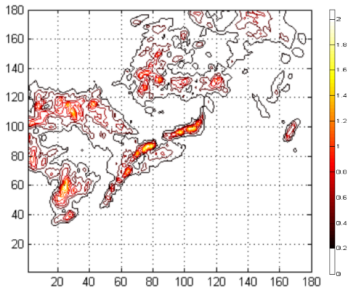

(a) Sample rainfall for a single time point.

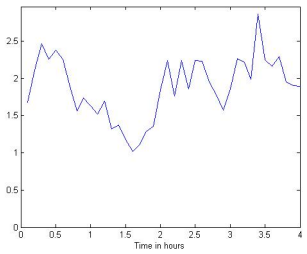

(b) Spatial maximum rainfall (mm) over time.

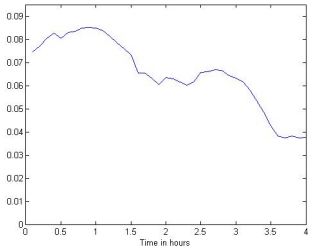

(c) Spatial mean rainfall (mm) over time.

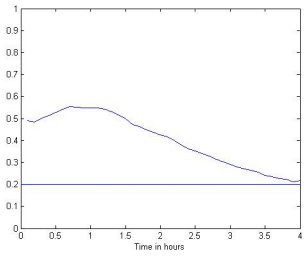

(d) Percentage of sample area covered by rainfall over time.

Fig. 1: Rainfall on 24th September 201 from 12:54 to 16:48 hours at Laverton, Melbourne.

cell, with an associated area, duration and intensity. The storm and cell centers all share a common velocity. The total rainfall intensity at point (x, y) and time t is then the sum of the intensity at (x, y) of all cells active at time t [7,15].

The storm arrival process is taken to be a Poisson process in $\mathbb{R}^2 \times [0, \infty)$ with homogeneous rate λ_s. Let $\mathbf{v} = (v_x, v_y)$ be the velocity of the rainfall event, so if a storm center arrives at $(\mathbf{u}, s)$ then at time $s + t$ it will be at $(\mathbf{u} + t\mathbf{v}, s + t)$. Storm durations are random with an $\exp(\gamma_s)$ distribution.

While a storm is active it produces cells at a rate λ_c in time, starting with a cell at the moment the storm begins. If the storm arrives at $(\mathbf{u}, s)$ and produces a cell at time $s + t$, the cell will be centered at $\mathbf{u} + t\mathbf{v} + \mathbf{w}$, where $\mathbf{w}$ comes from a Gaussian distribution with mean $\mathbf{0}$ and covariance Σ. The cell centre then also moves with velocity $\mathbf{v}$. We parameterise Σ using the storm diameter d_s, eccentricity e and orientation ω. d_s^{-1} has a gamma distribution with mean μ_{1/d_s} and coefficient of variation CV_{1/d_s}.

Individual cells have random durations, distributed as $\exp(\gamma_c)$, and random diameters d_c (the major axis). Rain cells are elliptical, with the same eccentricity e and orientation ω as the storms. d_c^{-1} has a gamma distribution with mean μ_{1/d_c} and coefficient of variation CV_{1/d_c}. It is convenient to use

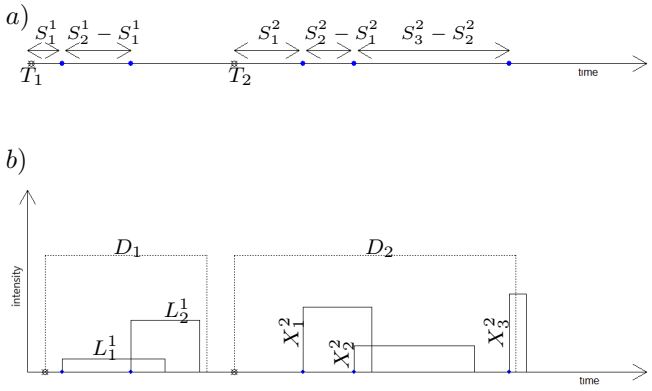

Fig. 2: Schematic of the temporal structure of the CIN model: a) Storm process (storm origins at T_i) and cell processes (the solid points show cell origins). The S_i^j note the times from the storm origins to cell origins; b) Storm durations D_i, cell durations L_i^j, and cell intensities X_i^j.

μ_A, the expected area of a rain cell, instead of CV_{1/d_c}, where

$$\mu_A = \pi\sqrt{1-e^2}\mu_{1/d_c}(1+CV_{1/d_c}^{-2}).$$

For the CIN model the intensity of a rain cell is constant over the area and duration of the cell, with an exponential distribution mean μ_X. The displacement, duration, diameter, and intensity of a cell are all independent, and independent of other cells. We give a schematic of the temporal structure of the CIN model in Figure 2 and of the spatial structure in Figure 3.

All together the CIN model has 13 parameters: velocity $\mathbf{v} = (v_x, v_y)$; eccentricity e; orientation ω; storm rate λ_s; mean storm duration $1/\gamma_s$; storm diameter given by μ_{1/d_s} and CV_{1/d_s}; cell rate λ_c; mean cell duration $1/\gamma_c$; cell diameter given by μ_{1/d_c} and μ_A; and mean cell intensity μ_X.

Generalisations of the CIN model

To date the CIN model has been the most detailed spatial-temporal rainfall model available using a stochastic cluster process construction. This is not because of a lack of scientific imagination, but because of the difficulty in obtaining theoretical moment expressions for these types of models, which have in the past been required for model fitting. We present here two novel extensions of the CIN model, incorporating increased flexibility and realism. We extend the model in two stages; in both cases the temporal structure of the process is unchanged and we refine the spatial structure of the rain cells.

The first stage we call the CIN-1 model; we incorporate two changes to give more realistic rainfall intensity gradients and provide a better match with

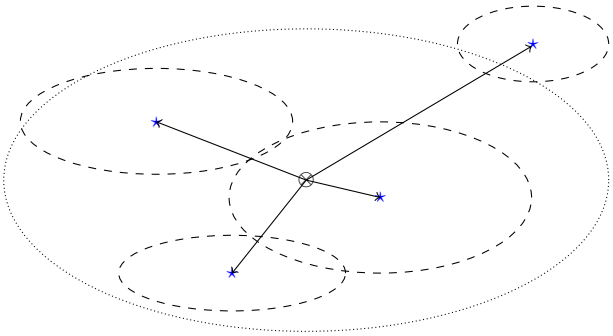

Fig. 3: Schematic diagram of the spatial structure of the CIN model. The centre point is a storm centre (which has a constant velocity). Stars are cell centres and the lines indicate the displacement of cell centres from the storm centre. The dashed curves indicate the cell areas, and the dotted curve gives a 95% prediction region for the storm area.

observations. Firstly we suppose that the rainfall intensity decreases continuously from the centre of a cell to the edge, rather than acting as a step function. If a and b are the lengths of the semi-major and semi-minor axis of a rain cell, and X is the intensity at the cell centre (c_x, c_y), then we model the intensity at (x, y) as

$$X\sqrt{1-\frac{(x-c_x)^2}{a^2}-\frac{(y-c_y)^2}{b^2}}.$$

Secondly we seek to capture an observed variation in the eccentricity of the cells: some cells appear circular in shape while some are long and thin. Accordingly we suppose that cell eccentricity has a normal distribution with mean μ_e and variance σ_e^2, truncated to $[0, 1]$.

For the second stage CIN-2 model we introduce correlation between cell intensity and diameter, again reflecting observed behaviour [11]. Specifically, for cell intensity X and cell diameter d_c, we suppose that

$$\log\begin{pmatrix}X\\ d_c\end{pmatrix} \sim N\left(\begin{pmatrix}\mu_X\\ \mu_{d_c}\end{pmatrix}, \begin{pmatrix}\sigma_X^2 & \rho_{X,d_c}\sigma_X\sigma_{d_c}\\ \rho_{X,d_c}\sigma_X\sigma_{d_c} & \sigma_{d_c}^2\end{pmatrix}\right).$$

Figure 4 represents the intensity of a single cell in (a) the CIN model and (b) the CIN-1 and CIN-2 models.

3 Model Fitting Using ABC

The CIN model has an intractable likelihood: calculating the rainfall intensity at any given point requires integration over all configurations of the underlying cluster point-process, which is impractical. Thus likelihood-based model fitting—such as maximum likelihood or Bayesian MCMC methods—can not be used. Moreover, for spatial-temporal data classical likelihood-free approaches

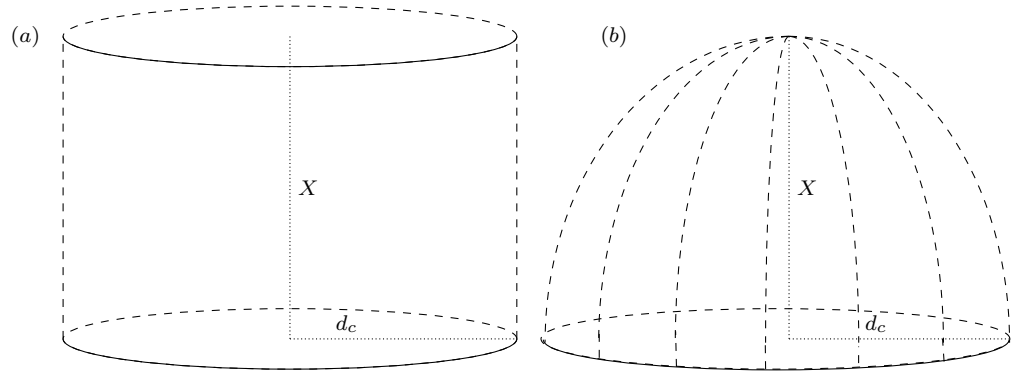

Fig. 4: Cell intensities for (a) CIN model (b) CIN-1 and CIN-2 models.

using ranks, permutations or bootstrapping also don't work. Wheater et al. [26] successfully used the Generalised Method of Moments (GMM) to fit the CIN model. However GMM requires analytical expressions for various moments and in this case their derivation relies on the strong independence assumptions present in the CIN model, so this approach does not carry over to the CIN-1 and CIN-2 models.

Simulation based model fitting has been increasing in popularity since the late 1980's, mirroring the increase in desktop computing power. The Simulated Method of Moments (SMM) of McFadden [14] uses simulation to estimate moments which are then compared to observed moments much like GMM, and can be used to fit the CIN-1 and CIN-2 models (see Gouriéroux and Monfort 1993 [8] for a review of SMM). However SMM is limited in that it only provides point estimates and does not allow model comparisons, so we only use it to provide initial estimates for our ABC procedure (details are given below). For estimation problems with a small number of parameters (not the case here) we can combine simulation with kernel density estimation to estimate the likelihood and hence obtain approximate maximum likelihood estimates (see for example Jones 2007 [9]). Alternatively the synthetic likelihood approach of Wood 2010 [27] is a more widely applicable method that uses simulation to form a Gaussian approximation to the likelihood, allowing likelihood-based inference. Rather than synthetic likelihood we have opted to use Approximate Bayesian Computation (ABC), which avoids any questions about whether it is appropriate to use a Gaussian approximation, and has all the benefits of Bayesian inference. ABC has already been successfully applied in many fields, but this is its first application to a spatial-temporal rainfall model.

ABC is a likelihood-free Bayesian inference technique, which uses simulations from the likelihood of interest in the absence of an analytic form. The technique developed from numerically intensive techniques for estimating population genetics models [20], and has since seen steadily increasing use in a variety of applications. The recent collection edited by Sisson, Fan and Beaumont [24] gives a comprehensive introduction to the subject. In what follows

we use ABC-MCMC, introduced by Marjoram et al. [13], which uses Markov Chain Monte Carlo to speed up the effective sampling rate of vanilla ABC.

We suppose that we have an observation D from some model $f(\cdot|\boldsymbol{\theta})$, depending on parameters $\boldsymbol{\theta}$, and that we are able to simulate from f (so $f(\cdot|\boldsymbol{\theta})$ is the likelihood). Let π be the prior distribution for $\boldsymbol{\theta}$ and $S = S(D)$ a vector of summary statistics for D, then ABC generates samples from $f(\boldsymbol{\theta}|\rho(S(D^*), S(D)) < \epsilon)$, where $D^* \sim f(\cdot|\boldsymbol{\theta})$, $\boldsymbol{\theta} \sim \pi$, and ρ is some distance function. If S is a sufficient statistic, then as $\epsilon \to 0$ this will converge to the posterior $f(\boldsymbol{\theta}|D)$. ABC-MCMC adds a proposal chain with density q and an additional rejection step, to generate a sample $\{\boldsymbol{\theta}_i\}$. Let $\boldsymbol{\theta}_0$ be some initial choice for $\boldsymbol{\theta}$ then the algorithm is as follows:

FOR $i = 1$ to N

1 Given current state $\boldsymbol{\theta}_i$ propose a new state $\boldsymbol{\theta}^*$ using $q(\cdot|\boldsymbol{\theta}_i)$
2 Put $\alpha = \min\{1, (\pi(\boldsymbol{\theta}^*)q(\boldsymbol{\theta}_i|\boldsymbol{\theta}^*))/(\pi(\boldsymbol{\theta}_i)q(\boldsymbol{\theta}^*|\boldsymbol{\theta}_i))\}$
3 Go to 4 with probability α, otherwise set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$ and return to 1
4 Simulate data $D^* \sim f(\cdot|\boldsymbol{\theta}^*)$
5 If $\rho(S(D^*), S(D)) \leq \epsilon$ then set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}^*$, otherwise set $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i$

END FOR

Unlike [13] we put the MCMC rejection step 3 before the ABC comparison step 5, to avoid unnecessarily running the simulation in step 4. Note that if we wish to use a non-uniform kernel in step 5 (see for example Sisson and Yan [23] §4.3: a non-uniform kernel allows us to assign weights to those $\boldsymbol{\theta}^*$ we accept, depending on $\rho(S(D^*), S(D))$, with smaller distances giving larger weights) then we can no longer separate the MCMC rejection step and the ABC comparison step, which increases the simulation burden. Sisson and Yan argue that using a kernel with unbounded support can improve the mixing of the Markov chain, however we did not find this to be a problem. In particular, using a regression adjustment [4] allows some relaxation of the threshold ϵ, to increase the acceptance rate without deliteriously impacting the posterior.

Practically, if $\boldsymbol{\theta}_0$ has very low posterior probability then ABC-MCMC can fail to accept any new sample points. Previous authors have suggested using a separate ABC step (without MCMC) to find a $\boldsymbol{\theta}_0$ with large posterior probability; we found that using Simulated Method of Moments (SMM) instead requires much less computation time. SMM is a variant of the Generalised Method of Moments (GMM) that uses Monte-Carlo estimates of moments, rather than analytic expressions (McFadden 1989 [14]). Thus, like ABC, using SMM we have much more freedom in the choice of moments used to fit the model to the data, and we found that it worked well using the same summary statistics S that we use for the ABC fitting.

Following Wheater et al. [26], for the CIN model the velocity $\mathbf{v}$, eccentricity e and orientation ω were all estimated using temporal and spatial autocovariance estimates, and then fixed. For the CIN-1 and CIN-2 models we used our

estimate of e for μ_e but also need to estimate σ_e^2. To do this we divided the study region spatially into four equal parts, then estimated the eccentricity for each part at each time point, giving four time-series of estimates for μ_e. Treating each series as an AR(1) process with mean μ_e, we can correct for the autocorrelation to get the usual moment-based estimate for σ_e^2.

The remaining parameters are all estimated using ABC. It is possible to include $\mathbf{v}$, e and ω in the ABC fitting rather than estimate them separately, however we found that doing so had no appreciable impact on the fit of the other parameters while it did increase the time required to run the ABC procedure.

We used the same set of 23 summary statistics for the CIN, CIN-1, and CIN-2 models:

- The overall mean and standard deviation of rainfall, taken over all pixels and all times.
- The spatial-temporal auto-correlation, with lags of (x, y, t), where x and y are measured in pixels and t is in units of 6-minutes. We take $t = 0$, $x \in \{-1, 0, 1\}$, $y \in \{-1, 0, 1\}$, and $t = 1$, $x \in \{-1, 0, 1\} + v_x$, $y \in \{-1, 0, 1\} + v_y$. Here v_x and v_y are the velocity components, in units of pixels per 6-minutes. Note that the lag $(0, 0, 0)$ auto-correlation is just the variance and so has already been included.
- The probability of an arbitrary pixel and time being dry.
- The ratio of dry/wet area and mean and standard deviation of wet area, averaged over time.

For the distance function ρ we used a weighted sum of squares $\rho(S(D^*), S(D)) = \sum_i w_i (S^*(i) - S(i))^2$, where $S^*(i)$ and $S(i)$ are the i-th components of $S(D^*)$ and $S(D)$ respectively. We found empirically that a good choice for w_i is to take it inversely proportional to the variance of $S^*(i)$ conditioned on using a $\boldsymbol{\theta}$ with high posterior probability. Given that $\boldsymbol{\theta}_0$ was chosen using SMM to have maximal posterior probability, we used a separate sample of $S(D^*)$ given $\boldsymbol{\theta}_0$ to calculate the w_i.

The choice of summary statistics S and distance metric ρ plays a large part in the performance of ABC. Ideally S should be sufficient, but certainly it should reflect those aspects of the real process considered most important. However choosing S too large reduces the efficiency of ABC, though this can be mitigated to some extent by using the approach of Beaumont et al. [4], which uses local linear regression to correct for bias in the posterior that can result from conditioning on $\rho(S(D^*), S(D)) \leq \epsilon$ instead of on $D^* = D$.

The remaining parameters were transformed to reduce dependence and skewness, and mapped to $(-\infty, \infty)$ (given as $\theta(1), \ldots, \theta(9)$ and $\psi(1), \ldots, \psi(11)$ below). This makes it easier for the proposal chain to spend its time in regions of high posterior probability. Vague normal priors are used for all the transformed parameters, and for the proposal chain we used a random walk with $N(0, 0.2^2 I)$ steps.

For the CIN and CIN-1 model our new ABC-parameters are

$$\begin{aligned}
\theta(1) &= \log(\lambda_s\gamma_s^{-1}), & \theta(2) &= \log(\lambda_s\gamma_s),\\
\theta(3) &= \log(\lambda_c\gamma_c^{-1}), & \theta(4) &= \log(\lambda_c\gamma_c),\\
\theta(5) &= \log(\mu_X\mu_A), & \theta(6) &= \log(\mu_X\mu_A^{-1}),\\
\theta(7) &= \log(\mu_{1/d_c}), & \theta(8) &= \log(\mu_{1/d_s}),\\
\theta(9) &= \log(CV_{1/d_s}).
\end{aligned}$$

For the CIN-2 model we replace parameters μ_{1/d_c} and μ_A by μ_{d_c} and $\sigma^2_{d_c}$, and gain parameters σ^2_X and ρ_{X,d_c}.

$$\begin{aligned}
\psi(1) &= \log(\lambda_s\gamma_s^{-1}), & \psi(2) &= \log(\lambda_s\gamma_s),\\
\psi(3) &= \log(\lambda_c\gamma_c^{-1}), & \psi(4) &= \log(\lambda_c\gamma_c),\\
\psi(5) &= \log(\mu_X), & \psi(6) &= \log(1/\sigma^2_X),\\
\psi(7) &= \log(\mu_{d_c}), & \psi(8) &= \log(1/\sigma^2_{d_c}),\\
\psi(9) &= \log\left(\frac{\rho_{X,d_c}+1}{1-\rho_{X,d_c}}\right), & \psi(10) &= \log(\mu_{1/d_s}),\\
\psi(11) &= \log(CV_{1/d_s}).
\end{aligned}$$

Comparison of Fitted Models

Using the spatial autocorrelation function we estimated $\mathbf{v} = (v_x, v_y) = (0.10, 29.9)$, $e = \mu_e = 0.86$, $\sigma^2_e = 0.04$ and $\omega = 39°$.

In the Appendix Figures 9–11 plot the posterior sample traces for models CIN, CIN-1 and CIN-2. In each case the chains appear stationary and exhibit good mixing.

For the CIN model using a threshold of $\epsilon = 10$ the overall acceptance rate was approximately 4.1%. Increasing ϵ improves the acceptance rate and the degree of mixing, at the expense of reduced accuracy for the posterior approximation. For the CIN-1 and CIN-2 models the acceptance rates were approximately 5% and 11% respectively, using thresholds of $\epsilon = 14$ and 8. For ABC-MCMC rather than the acceptance rate—which is commonly reported for vanilla ABC sampling—a better measure of sampling efficiency is the "effective sample rate" for each parameter, which combines the efficiency of the MCMC sampling and the ABC rejection step and is given by the effective sample size for the sampling chain divided by the number of simulations required to produce them. For the CIN, the effective sample rate varied from 0.002 to 0.004 depending on the parameter. For the CIN1 and CIN2 models the effective sample rate varied from 0.001 to 0.003 and from 0.001 to 0.005 respectively.

Posterior plots for models CIN, CIN-1 and CIN-2 are given in Figures 5–7. In all three cases we get nice peaks on our posteriors. Posterior summaries

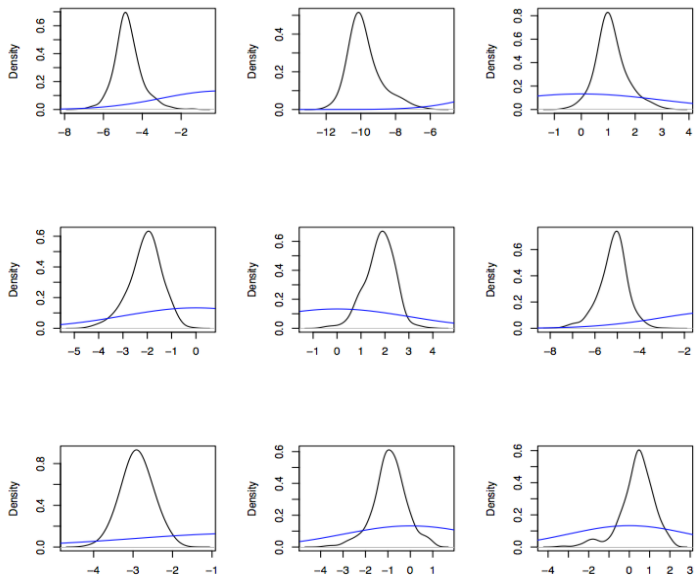

Fig. 5: CIN model: posteriors for $\theta(i), i = 1, 2, ..., 9$. Priors are given by the dashed lines.

for the original (untransformed) parameters are given in Tables 1–3 in the Appendix.

To compare the performance of our three models we use posterior predictive probabilities to judge how close each model is to the original data, as measured by our distance ρ. That is, we compare for each model the distribution of $\rho(S(D^*), S(D))$ where D is the original data and D^* is generated by the model when $\boldsymbol{\theta}$ is distributed according to the posterior. We can sample from this distribution by sampling $\boldsymbol{\theta}$ from the posterior, generating D^* from $\boldsymbol{\theta}$ and the model, and then calculating $\rho(S(D^*), S(D))$. We can then estimate the c.d.f. of $\rho(S(D^*), S(D))$ using the e.d.f. of a suitably large sample.

Figure 8 plots the estimated c.d.f. of the posterior for $\rho(S(D^*), S(D))$ under models CIN and CIN-1. We see that the CIN-1 and model gives a better fit than the CIN model, that is, it more likely to produce data close to our observation. The CIN-2 model does not give an improved fit, however it does show moderate evidence of negative correlation between cell intensity and diameter: the posterior mean for ρ_{X,d_c} is -0.69 with a 95% credible interval of $(-0.93, 0.10)$. We discuss the implications of this in the Conclusions section.

We can also construct posterior predictive distributions for $(S^*(i) - S(i))^2$. That is, we can compare the fit of each model as measured by individual components of the summary statistic S. They are given for the CIN and CIN-1 models in Figures 12 and 13 in the Appendix. Generally the CIN-1 model

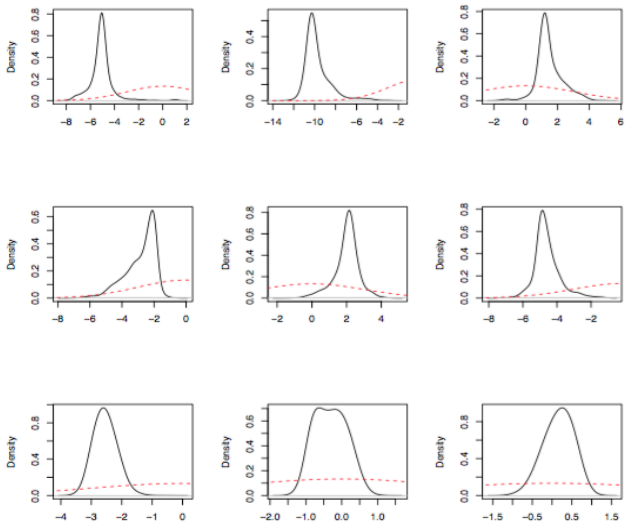

Fig. 6: CIN-1 model: posteriors for $\theta(i), i = 1, 2, ..., 9$. Priors are given by the dashed lines.

performs better than the CIN model, though not uniformly, but without any obvious pattern to the exceptions.

4 Conclusions

By using Approximate Bayesian Computation (ABC) we have been able to fit extensions of the spatial-temporal model of Cox, Isham and Northrop [7, 15] (CIN model), allowing for more realistic rainfall intensity gradients and correlation between the intensity and size of localised rain cells. Using rainfall radar data for a rainfall event in Melbourne, Australian, and looking at the posterior predictive distribution of the distance between the observed and simulated data, we showed that the CIN-1 model gave a better fit than the CIN model. We also demonstrated moderate evidence for negative correlation between the intensity and size (diameter) of rain cells.

ABC is a very flexible methodology, and the approach demonstrated here could be easily applied to further generalisations of the CIN model. When introducing correlation between the cell intensity X and cell diameter d_c we used a multivariate lognormal distribution as this was a parsimonious way to do so, however it appears that while there is moderate evidence that the correlation is non-zero, qualitatively the lognormal marginals did a poor job of matching observed cell diameters, producing too many small cells. This suggests that we could improve the model further by using a copula to capture the dependence

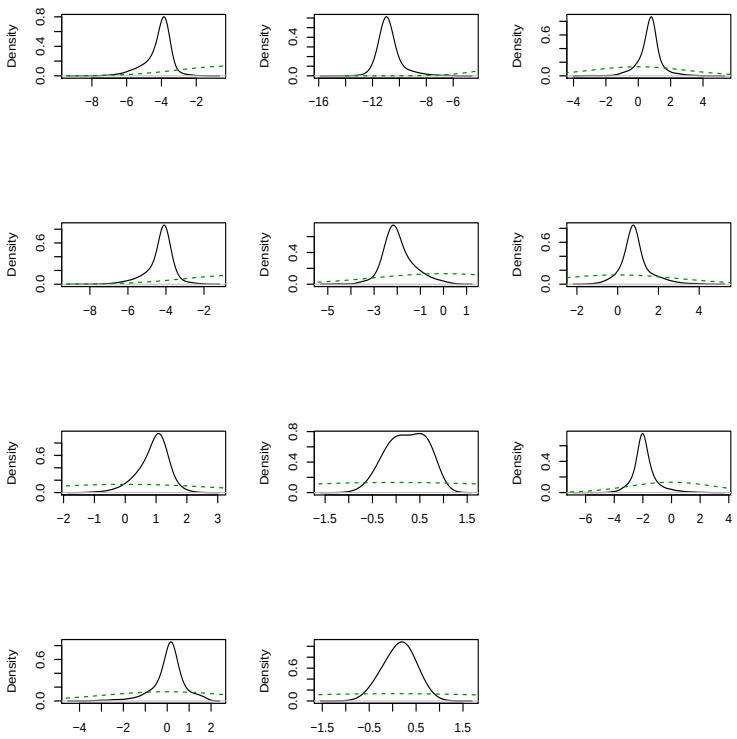

Fig. 7: CIN-2 model: posteriors for $\psi(i), i = 1, 2, ..., 11$. Priors are given by the dashed lines.

between X and d_c and experimenting with marginals other than the lognormal (the CIN and CIN-1 models use the exponential and gamma distribution for X and d_c respectively). A further extension would be to parametrise the cell intensity profile illustrated in Figure 4 to interpolate between the two models

More generally, by considering different components of the distance between observed and simulated data we can see where one model outperforms another, which has the potential to inform future efforts at model refinement. Moreover additional components can be added to the summary S to focus on specific aspects of the model. For example if the maximum rainfall intensity produced by the model was considered particularly important, then this could be included in S and hence targeted when fitting the model.

Finally we note that from a technical point of view we showed that ABC-MCMC could be improved by separating the ABC and MCMC rejection steps and by using the Simulated Method of Moments (SMM) to initialise the proposal chain.

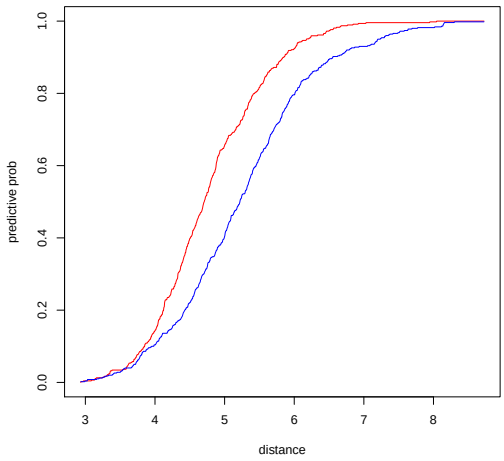

Fig. 8: Posterior distribution of $\rho(S(D^*), S(D))$ for the CIN model (blue) and CIN-1 model (red). The CIN-1 model is more likely to produce simulations D^* close to D, as measured by ρ, that is, it gives a better fit than the CIN model.

Acknowledgements Datasets for this research are available from the Australian Bureau of Meteorology, GPO Box 1289, Melbourne VIC 3001. http://www.bom.gov.au/climate/data-services/data-requests.shtml. The Bureau charges according to the Australian Government Cost Recovery Guidelines.

Appendix

Here we provide additional details of the model comparisons discussed in Section 3.

Figures 9–11 plot the posterior sample traces for models CIN, CIN-1 and CIN-2. In each case the chains appear stationary and exhibit good mixing.

Tables 1–3 give posterior summaries for the parameters of the CIN, CIN-1 and CIN-2 models. We notice that for the CIN-2 model some credible intervals are wider than the CIN model and the CIN-1 model estimations, particularly mean storm duration γ_s^{-1}, mean cell duration γ_c^{-1}.

Figures 12 and 13 give posterior predictive distributions for $(S^*(i) - S(i))^2$, for the CIN and CIN-1 models. That is, we can compare the fit of each model as measured by individual components of the summary statistic S. Generally the CIN-1 model performs better, though not uniformly, but without any obvious pattern to the exceptions.

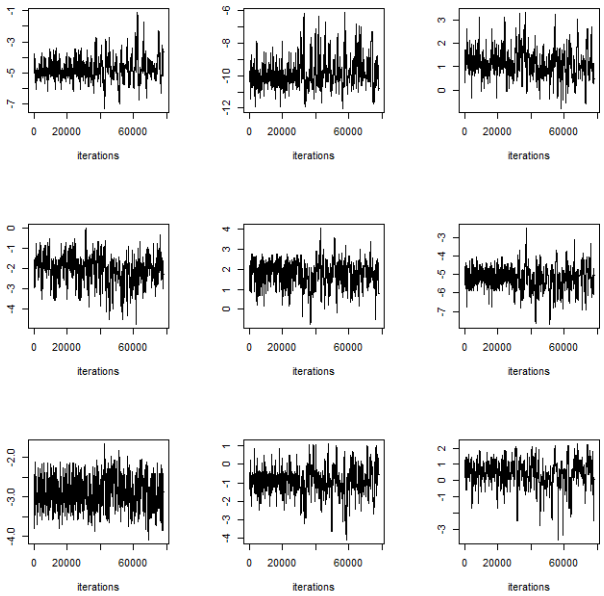

Fig. 9: ABC-MCMC fitting for CIN model: posterior chains for $\theta(i)$, $i = 1, 2, \ldots, 9$.

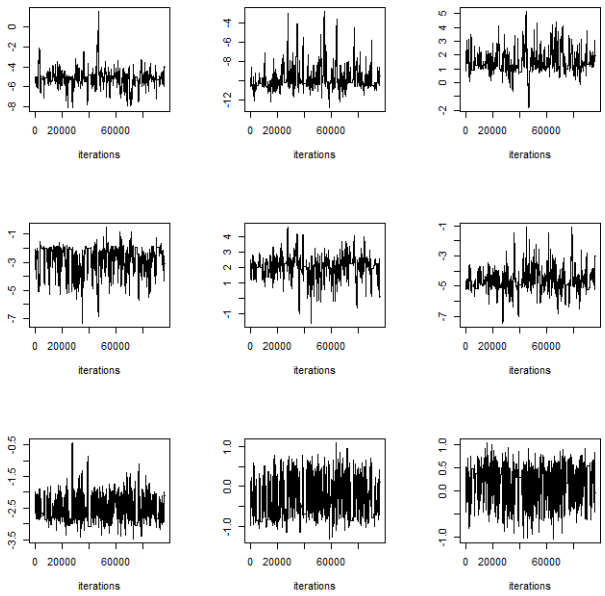

Fig. 10: ABC-MCMC fitting for CIN-1 model: posterior chains for $\theta(i)$, $i = 1, 2, \ldots, 9$.

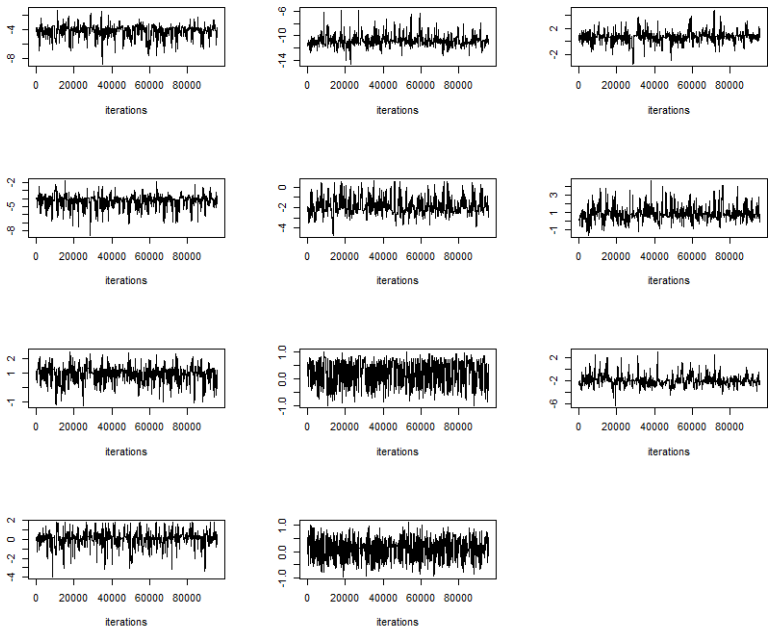

Fig. 11: ABC-MCMC fitting for CIN-2 model: posterior chains for $\psi(i)$,$i = 1, 2, \ldots, 11$.

Parameter	Mean	Median	95 % Credible Interval
λ_s	0.0008	0.0006	(0.0003, 0.0026)
γ_s^{-1}	14.8159	13.2333	(2.9086, 38.5852)
λ_c	0.6968	0.6201	(0.2321, 1.7451)
γ_c^{-1}	5.3733	4.6462	(2.6768, 13.1802)
μ_{1/d_s}	0.5199	0.4006	(0.0738, 1.8749)
CV_{1/d_s}	1.9438	1.5953	(0.1467, 6.0694)
μ_X	0.1946	0.1917	(0.0792, 0.3332)
μ_A	37.2997	32.6773	(11.8394, 93.7327)
μ_{1/d_c}	0.0595	0.0548	(0.0255, 0.1211)

Table 1: Posterior estimates of the ABC-parameters for the CIN model.

Parameter	Mean	Median	95 % Credible Interval
λ_s	0.0008	0.0005	(0.0002 0.0036)
γ_s^{-1}	12.7340	12.3910	(1.0938 38.4967)
λ_c	0.6166	0.5722	(0.1171, 1.8787)
γ_c^{-1}	10.7880	7.6703	(4.0369, 35.9604)
μ_{1/d_s}	0.8022	0.7224	(0.3761, 1.6064)
CV_{1/d_s}	1.2535	1.2105	(0.5601, 2.1132)
μ_X	0.2850	0.2731	(0.1508, 0.4848)
μ_A	33.0017	31.5407	(4.6533, 92.1609)
μ_{1/d_c}	0.0840	0.0771	(0.0441, 0.1617)

Table 2: Posterior estimates of the ABC-parameters for the CIN-1 model.

Parameter	Mean	Median	95% Credible Interval
λ_s	0.0006	0.0006	(0.0002, 0.00150
γ_s^{-1}	31.797	32.462	(4.4944, 64.524)
λ_c	0.1887	0.1801	(0.0457, 0.4429)
γ_c^{-1}	13.883	11.407	(5.6338, 42.133)
μ_{1/d_s}	1.3331	1.1647	(0.1671, 4.3280)
CV_{1/d_s}	1.2102	1.1901	(0.6291, 1.9917)
μ_X	0.1882	0.1258	(0.0426, 0.7603)
σ_X^2	0.5235	0.4519	(0.0671, 1.5503)
μ_{d_c}	2.7478	2.7523	(0.7829 , 5.6306)
$\sigma_{d_c}^2$	0.8602	0.7859	(0.4335, 1.7103)
ρ_{X,d_c}	-0.6906	-0.7564	(-0.9327, 0.0996)

Table 3: Posterior estimates of the ABC-parameters for the CIN-2 model.

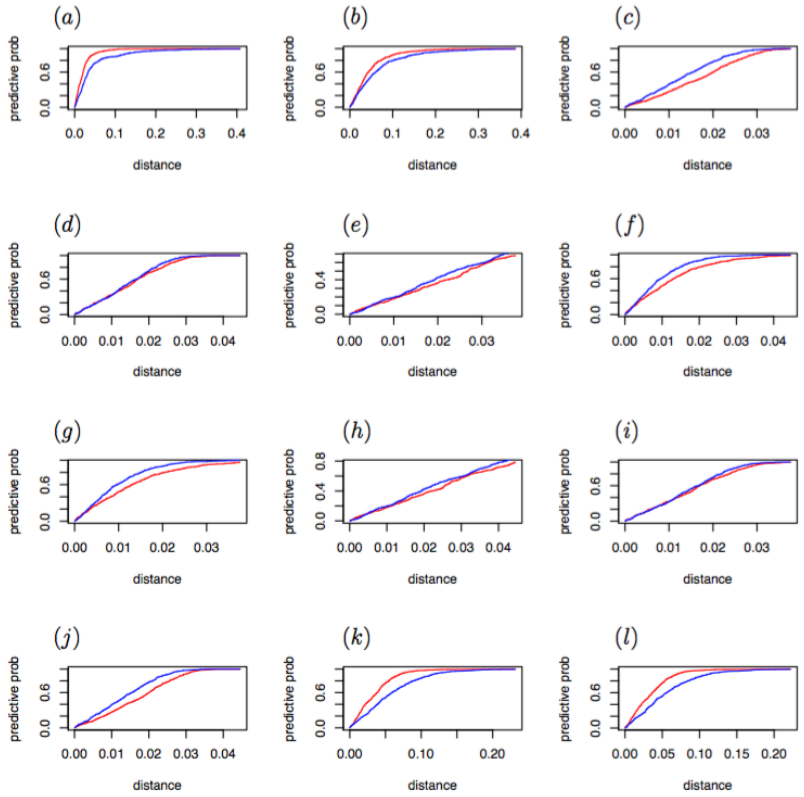

Fig. 12: Predictive probability for the CIN model (blue line) and CIN-1 model (red line). Plots (a) and (b) are for mean and standard deviation summaries. Plots (c) to (j) are spatial correlations $\rho(-1,-1,0), \rho(-1,0,0), \rho(1,1,0), \rho(0,-1,0), \rho(0,1,0), \rho(1,-1,0), \rho(1,0,0)$, and $\rho(1,1,0)$. Plots (k) and (l) are of $\rho(-1+v_x,-1+v_y,1)$, and $\rho(-1+v_x,0+v_y,1)$.

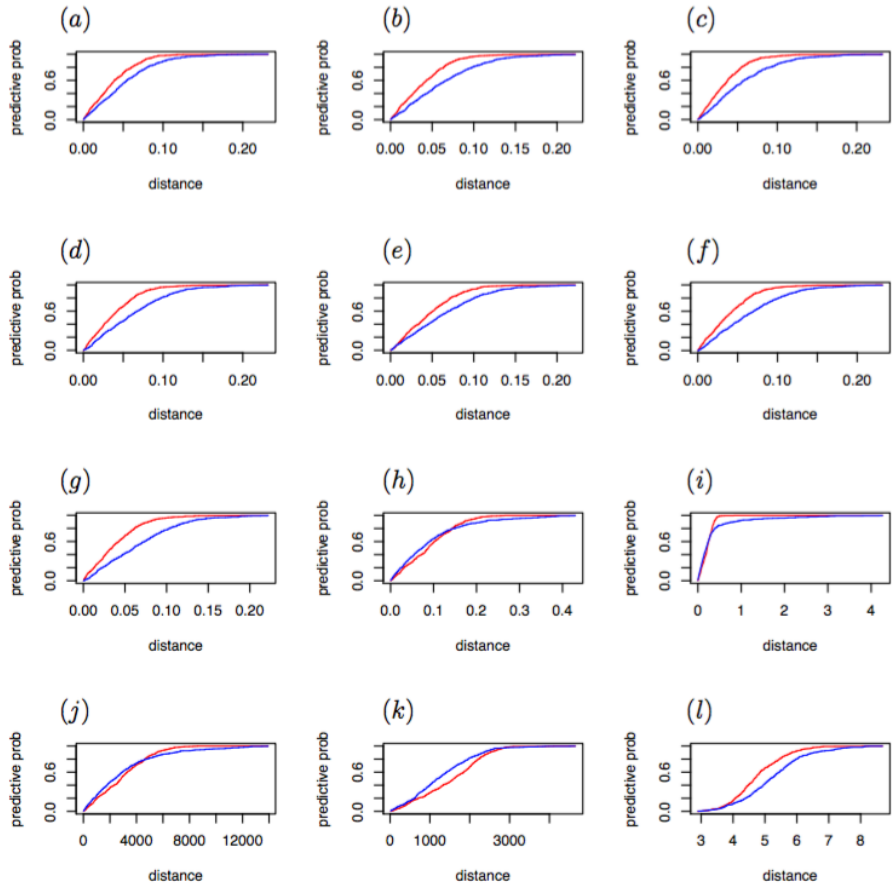

Fig. 13: Predictive probability for the CIN model (blue line) and CIN-1 model (red line). Plots (a) to (g) are of spatial autocorrelations $\rho(-1+v_x, 1+v_y, 1), \rho(0+v_x, -1+v_y, 1), \rho(0+v_x, 0+v_y, 1), \rho(0+v_x, 1+v_y, 1), \rho(1, +v_x, -1+v_y, 1), \rho(1, +v_x, 0+v_y, 1)$, and $\rho(1, +v_x, 1+v_y, 1)$. Plots (h) and (i) are of dry probability of an arbitrary pixel and dry and wet area ratio. Plots (j) and (k) are of mean wet area over time and standard deviation of wet area over time. Plot (l) is of total distance form all summaries.

References

1. Allard, D., Ailliot, P., Monbet, V., Naveau, P.: Stochastic weather generators: an overview of weather type models. Journal de la Société Française de Statistique **156**(1), 101–113 (2015)
2. Aryal, N.R.: Stochastic spatial-temporal models for rainfall processes. Ph.D. thesis, School of Mathematics and Statistics, The University of Melbourne (2018)

3. Aryal, N.R., Jones, O.D.: Fitting the Bartlett-Lewis rainfall model using Approximate Bayesian Computation. Mathematics and Computers in Simulation **175**, 153–163 (2020)
4. Beaumont, M.A., Zhang, W., Balding, D.J.: Approximate Bayesian Computation in population genetics. Genetics **162**(4), 2025–2035 (2002)
5. Benoit, L., Allard, D., Mariethoz, G.: Stochastic rainfall modeling at sub-kilometer scale. Water Resources Research **54**(6), 4108–4130 (2018)
6. Chandler, R., Isham, V., Northrop, P., Wheater, H., Onof, C., Leith, N.: Uncertainty in rainfall inputs. Applied Uncertainty Analysis for Flood Risk Management, edited by: Beven, KJ and Hall, JW, Imperial College Press: London pp. 101–152 (2014)
7. Cox, D.R., Isham, V.: A simple spatial-temporal model of rainfall. Proceedings of the Royal Society of London A, Mathematical and Physical Sciences **415**(1849), 317–328 (1988)
8. Gourieroux, C., Monfort, A.: Simulation-based inference. J. Econometrics **59**, 5–33 (1993)
9. Jones, O.D.: Modelling electricity power cuts in the UK. In: A. Stacey, B. Blyth, J. Shepherd, A.J. Roberts (eds.) Proceedings of the 7th Biennial Engineering Mathematics and Applications Conference, EMAC-2005, *ANZIAM J.*, vol. 47, pp. C603–C620 (2007)
10. Jones, O.D., Nyman, P., Sheridan, G.J.: Modelling the effects of fire and rainfall regimes on extreme erosion events in forested landscapes. Stochastic Environmental Research and Risk Assessment **28**(8), 2015–2025 (2014)
11. Kaczmarska, J., Isham, V., Onof, C.: Point process models for fine-resolution rainfall. Hydrological Sciences Journal **59**(11), 1972–1991 (2014)
12. Leclerc, G., Schaake, J.: Derivation of hydrologic frequency curves. Tech. Rep. 142, Mass. Inst. of Technol., Cambridge, MA, USA (1972)
13. Marjoram, P., Molitor, J., Plagnol, V., Tavaré, S.: Markov chain monte carlo without likelihoods. Proceedings of the National Academy of Sciences **100**(26), 15324–15328 (2003)
14. McFadden, D.: A method of simulated moments for estimation of discrete response models without numerical integration. Econometrica **57**(5), 995–1026 (1989)
15. Northrop, P.: A clustered spatial-temporal model of rainfall. Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **454**(1975), 1875–1888 (1998)
16. Onof, C., Chandler, R., Kakou, A., Northrop, P., Wheater, H., Isham, V.: Rainfall modelling using poisson-cluster processes: a review of developments. Stochastic Environmental Research and Risk Assessment **14**(6), 384–411 (2000)
17. Park, J., Cross, D., Onof, C., Chen, Y., Kim, D.: A simple scheme to adjust poisson cluster rectangular pulse rainfall models for improved performance at sub-hourly timescales. Journal of Hydrology **598**, 126296 (2021)
18. Paschalis, A., Molnar, P., Fatichi, S., Burlando, P.: A stochastic model for high-resolution space-time precipitation simulation. Water Resources Research **49**(12), 8400–8417 (2013)
19. Peleg, N., Fatichi, S., Paschalis, A., Molnar, P., Burlando, P.: An advanced stochastic weather generator for simulating 2-d high-resolution climate variables. Journal of Advances in Modeling Earth Systems **9**(3), 1595–1627 (2017)
20. Pritchard, J.K., Seielstad, M.T., Perez-Lezaun, A., Feldman, M.W.: Population growth of human y chromosomes: a study of y chromosome microsatellites. Molecular biology and evolution **16**(12), 1791–1798 (1999)
21. Rodriguez-Iturbe, I., Cox, D.R., Isham, V.: Some models for rainfall based on stochastic point processes. Proceedings of the Royal Society of London A, Mathematical and Physical Sciences **410**(1839), 269–288 (1987)
22. Segond, M.L., Wheater, H.S., Onof, C.: The significance of spatial rainfall representation for flood runoff estimation: a numerical evaluation based on the Lee catchment, UK. Journal of Hydrology **347**(1-2), 116–131 (2007)
23. Sisson, S., Fan, Y.: ABC samplers. In: Handbook of Approximate Bayesian Computation, pp. 87–123. Chapman and Hall/CRC (2018)
24. Sisson, S.A., Fan, Y., Beaumont, M. (eds.): Handbook of Approximate Bayesian Computation. Handbooks of Modern Statistical Methods. CRC Press (2018)

445 25. Wheater, H., Chandler, R., Onof, C., Isham, V., Bellone, E., Yang, C., Lekkas, D.,
446 Lourmas, G., Segond, M.L.: Spatial-temporal rainfall modelling for flood risk estimation.
447 Stochastic Environmental Research and Risk Assessment **19**(6), 403–416 (2005)
448 26. Wheater, H., Isham, V., Chandler, R., Onof, C., Stewart, E.: Improved methods for
449 national spatial-temporal rainfall and evaporation modelling for BSM. Tech. rep., Defra
450 (2006)
451 27. Wood, S.N.: Statistical inference for noisy nonlinear ecological dynamic systems. Nature
452 **466**(7310), 1102–1104 (2010)