

CARBayesST: An R Package for Spatio-temporal Areal Unit Modelling with Conditional Autoregressive Priors

Duncan Lee
University of Glasgow

Alastair Rushworth
University of Strathclyde

Gary Napier
University of Glasgow

Abstract

This is a vignette for the R package **CARBayesST** version 2.1, which describes the class of models that can be implemented by the package.

Keywords: Bayesian inference, conditional autoregressive priors, R package **CARBayesST**.

1. Introduction

Areal unit data are a type of spatial data where the observations relate to a set of K contiguous but non-overlapping areal units, such as electoral wards or census tracts. Each observation relates to an entire areal unit, and thus is typically a summary measure such as an average, proportion or total of the quantity being measured throughout the unit. This type of data can be total yields in sectors in an agricultural field trial (Besag and Higdon 1999), the proportion of people who are Catholic in lower super output areas in Northern Ireland (Lee *et al.* 2015), the average score on SAT college entrance exams across US states (Wall 2004), or the total number of cases of chronic obstructive pulmonary disease from populations living in counties in Georgia, USA (Choi and Lawson 2011). Such data have become increasingly available in recent times, due to the creation of databases such as Scottish Neighbourhood Statistics (<http://www.sns.gov.uk>), Health and Social Care Information Centre Indicator Portal (<http://www.hscic.gov.uk/indicatorportal>), and Surveillance Epidemiology and End Results programme (<http://seer.cancer.gov>). These databases provide data on a set of K areal units for T consecutive time periods, yielding a rectangular array of $K \times T$ spatio-temporal observations. The motivations for modelling these data are varied, and include estimating the effect of a risk factor on a response (see Wakefield 2007 and Lee *et al.* 2009), identifying clusters of contiguous areal units that exhibit an elevated risk of disease compared with neighbouring areas (see Charras-Garrido *et al.* 2012 and Anderson *et al.* 2014), and quantifying the level of segregation in a city between two or more different groups (see Lee *et al.* 2015).

The key statistical modelling challenge for these data is that of spatio-temporal autocorrelation, namely that observations from geographically close areal units and time periods tend to have more similar values than units and time periods that are further apart. Temporal autocorrelation occurs because the data relate to the same set of individuals over consecutive time periods, while the spatial autocorrelation can arise for a number of reasons. The

first is unmeasured confounding, which occurs when a spatially patterned risk factor for the response data being measured is not included in a regression model and hence its omission induces unmeasured spatial structure into the response. Other causes of spatial autocorrelation include neighbourhood effects, where the behaviours of individuals in an areal unit are influenced by individuals in adjacent units, and grouping effects where groups of people with similar behaviours choose to live together. A number of models have been developed to allow for such spatio-temporal autocorrelation, the majority of which utilise a set of autocorrelated random effects. Autoregressive (AR) priors are commonly used to represent temporal autocorrelation in the random effects, while conditional autoregressive (CAR) priors (Besag *et al.* 1991) are most often utilised for modelling spatial autocorrelation. Both these models are special cases of a Gaussian Markov Random Field (GMRF), and are typically implemented in a Bayesian setting using either Markov chain Monte Carlo (MCMC) simulation or Integrated Nested Laplace Approximations (INLA).

An array of freely available software can now implement purely spatial areal unit models, ranging from general purpose statistical modelling software such as BUGS (Lunn *et al.* 2009) and **R-INLA** (Rue *et al.* 2009), to specialist spatial modelling packages in the statistical software R (R Core Team 2013) such as **CARBayes**, **spatcounts** and **spdep**. However, due to the flexibility of general purpose software, implementing spatial models, in say BUGS, requires a degree of programming that is non-trivial for applied researchers. Specialist software for spatio-temporal modelling is much less well developed, with examples for geostatistical data including **spTimer** and **spBayes**. For areal unit data the **surveillance** package models epidemic data, the **plm** and **splm** packages model panel data, while the **nlme** and **lme4** packages have functionality to model spatial and temporal random effects structures. However, software to fit a range of spatio-temporal areal unit models with CAR type autocorrelation structures is not available, which has motivated us to develop the R package **CARBayesST**.

The software can fit a number of different spatio-temporal models, which allow the user to answer different questions about their data. These models include a spatially varying linear time trends model similar to Bernardinelli *et al.* (1995), a spatial and temporal main effects and an interaction model similar to that proposed by Knorr-Held (2000), the spatially autocorrelated autoregressive time series model of Rushworth *et al.* (2014a), and a model with a common temporal trend but varying spatial surfaces. The software can also fit more complex spatio-temporal structures, including the adaptive smoothing model of Rushworth *et al.* (2014b) and a localised smoothing model which is a spatio-temporal extension of Lee and Sarran (2015). The software has the same syntax as the R package **CARBayes** (Lee 2013) for spatial areal unit modelling, and retains all of its easy-to-use features such as specifying the spatial adjacency information via a single matrix (unlike BUGS that requires 3 separate list objects), fitting models via a one-line function call, and compatibility with **CARBayes** which allows it to share the latter's model summary functionality for interpreting the results. All models available in this software can be fitted to binomial, Gaussian or Poisson data, and Section 2 in this vignette summarises the models that can be fitted. Finally, in Section 3 the software and functionality are discussed.

2. Spatio-temporal models for areal unit data

This section outlines the class of Bayesian hierarchical models that **CARBayesST** can fit, and in all cases inference is based on McMC simulation.

2.1. Data structure and likelihood

The study region comprises a set of $k = 1, \dots, K$ non-overlapping areal units $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_K\}$, and data are recorded for each unit for $t = 1, \dots, T$ consecutive time periods. Thus data are available for a $K \times T$ rectangular array with K rows (spatial units) and T columns (time periods). The response data are denoted by $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_T)$, where $\mathbf{Y}_t = (Y_{1t}, \dots, Y_{Kt})$ denotes the vector of observations for all K spatial units for time period t . Also available are a vector of known offsets $\mathbf{O} = (\mathbf{O}_1, \dots, \mathbf{O}_T)$, where $\mathbf{O}_t = (O_{1t}, \dots, O_{Kt})$ denotes the vector of offsets for time period t . Finally, $\mathbf{x}_{kt} = (x_{kt1}, \dots, x_{ktp})$ is a vector of p known covariates for areal unit k and time period t , and can include factors or continuous variables and a column of ones for the intercept term. **CARBayesST** models these data with a generalised linear mixed model, whose general form is:

$$\begin{aligned} Y_{kt} | \mu_{kt} &\sim f(y_{kt} | \mu_{kt}, \nu^2) \quad \text{for } k = 1, \dots, K, \quad t = 1, \dots, T, \\ g(\mu_{kt}) &= \mathbf{x}_{kt}^T \boldsymbol{\beta} + O_k + M_{kt}, \\ \boldsymbol{\beta} &\sim N(\boldsymbol{\mu}_\beta, \Sigma_\beta). \end{aligned} \tag{1}$$

The vector of regression parameters are denoted by $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$, and non-linear covariate effects can be incorporated into the above model by including natural cubic spline or polynomial basis functions into the design matrix. A multivariate Gaussian prior is assumed for $\boldsymbol{\beta}$, and the mean $\boldsymbol{\mu}_\beta$ and diagonal variance matrix Σ_β can be chosen by the user. The M_{kt} term is a latent component for areal unit k and time period t that captures spatio-temporal autocorrelation in these data, and **CARBayesST** can fit a number of different models for M_{kt} . The package can fit 3 special cases of the above model, for binomial, Gaussian and Poisson data, and their exact specifications are given below:

- **Binomial** - $Y_{kt} \sim \text{Binomial}(n_{kt}, \theta_{kt})$ and $\log(\theta_{kt}/(1 - \theta_{kt})) = \mathbf{x}_{kt}^T \boldsymbol{\beta} + O_k + M_{kt}$.
- **Gaussian** - $Y_{kt} \sim N(\mu_{kt}, \nu^2)$ and $\mu_{kt} = \mathbf{x}_{kt}^T \boldsymbol{\beta} + O_k + M_{kt}$.
- **Poisson** - $Y_{kt} \sim \text{Poisson}(\mu_{kt})$ and $\log(\mu_{kt}) = \mathbf{x}_{kt}^T \boldsymbol{\beta} + O_k + M_{kt}$.

2.2. Spatio-temporal models for M_{kt}

All the models in the package induce spatial autocorrelation into the response data \mathbf{Y} via the latent component M_{kt} , which is achieved by a $K \times K$ neighbourhood matrix $\mathbf{W} = (w_{kj})$. Typically, \mathbf{W} contains binary elements, where $w_{kj} = 1$ if areal units $(\mathcal{S}_k, \mathcal{S}_j)$ share a common border (i.e. are spatially close) and is zero otherwise. Additionally, $w_{kk} = 0$. This means that for spatially adjacent areal units $(\mathcal{S}_k, \mathcal{S}_j)$ (M_{kt}, M_{jt}) are spatially autocorrelated, where as values for non-neighbouring areal units are conditionally independent given the remaining

$\{M_{it}\}$ values. This binary specification of \mathbf{W} based on sharing a common border is the most commonly used for areal data, but the only requirement by **CARBayesST** is for \mathbf{W} to be symmetric and contain non-negative elements. Similarly the model `ST.CARanova()` uses a binary $T \times T$ temporal neighbourhood matrix $\mathbf{D} = (d_{tj})$, where $d_{tj} = 1$ if $j = t - 1, t + 1$ and $d_{tj} = 0$ otherwise. **CARBayesST** can fit the following models:

- `ST.CARlinear()` - fits a model similar to the spatially varying linear time trends model proposed by [Bernardinelli *et al.* \(1995\)](#).
- `ST.CARanova()` - fits a model similar to the spatial and temporal main effects and space-time interaction model proposed by [Knorr-Held \(2000\)](#).
- `ST.CARsepspatial()` - extends `ST.CARanova()` by fitting an overall temporal trend and separate spatial surfaces for each time period. Note, this model can only be applied to binomial or Poisson data.
- `ST.CARar()` - fits the spatially autocorrelated autoregressive time series model of [Rushworth *et al.* \(2014a\)](#).
- `ST.CARadaptive()` - fits the localised smoothing model of [Rushworth *et al.* \(2014b\)](#) that is an extension of [Rushworth *et al.* \(2014a\)](#).
- `ST.CARlocalised()` - fits the localised smoothing and clustering model that is an extension of [Rushworth *et al.* \(2014a\)](#) and [Lee and Sarran \(2015\)](#). Note, this model can only be applied to binomial or Poisson data.

Full details of each model are given below.

ST.CARlinear()

The model is a modification of that proposed by [Bernardinelli *et al.* \(1995\)](#) and is given by

$$\begin{aligned}
 M_{kt} &= \beta_1 + \phi_k + (\alpha + \delta_k) \frac{(t - \bar{t})}{T}, & (2) \\
 \phi_k | \phi_{-k}, \mathbf{W} &\sim \text{N} \left(\frac{\rho_\phi \sum_{j=1}^K w_{kj} \phi_j}{\rho_\phi \sum_{j=1}^K w_{kj} + 1 - \rho_\phi}, \frac{\tau_\phi^2}{\rho_\phi \sum_{j=1}^K w_{kj} + 1 - \rho_\phi} \right), \\
 \delta_k | \delta_{-k}, \mathbf{W} &\sim \text{N} \left(\frac{\rho_\delta \sum_{j=1}^K w_{kj} \delta_j}{\rho_\delta \sum_{j=1}^K w_{kj} + 1 - \rho_\delta}, \frac{\tau_\delta^2}{\rho_\delta \sum_{j=1}^K w_{kj} + 1 - \rho_\delta} \right), \\
 \tau_\phi^2, \tau_\delta^2 &\sim \text{Inverse-Gamma}(a, b), \\
 \rho_\phi, \rho_\delta &\sim \text{Uniform}(0, 1), \\
 \alpha &\sim \text{N}(\mu_\alpha, \sigma_\alpha^2),
 \end{aligned}$$

where $\bar{t} = (1/T) \sum_{t=1}^T t$ and thus the modified linear temporal trend covariate is $t^* = (t - \bar{t})/T$ and runs over a centered unit interval. Each areal unit k has its own linear time trend, with a spatially varying intercept $\beta_1 + \phi_k$ and a spatially varying slope $\alpha + \delta_k$. Note, the β_1 term

comes from the covariate component $\mathbf{x}_{kt}^\top \boldsymbol{\beta}$ in (1). Each set of random effects $\boldsymbol{\phi} = (\phi_1, \dots, \phi_K)$ and $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K)$ are modelled as spatially autocorrelated by the CAR prior proposed by Leroux *et al.* (1999), and are mean centered. Here (ρ_ϕ, ρ_δ) are spatial dependence parameters, with values of one corresponding to strong spatial smoothness that is equivalent to the intrinsic CAR prior proposed by Besag *et al.* (1991), while values of zero correspond to independence. Flat uniform priors on the unit interval are specified for the spatial dependence parameters (ρ_ϕ, ρ_δ) , while conjugate inverse-gamma and Gaussian priors are specified for the random effects variances $(\tau_\phi^2, \tau_\delta^2)$ and the overall slope parameter α respectively. The corresponding hyperparameters $(a, b, \mu_\alpha, \sigma_\alpha^2)$ can be chosen by the user, and default values are $(a = 0.001, b = 0.001, \mu_\alpha = 0, \sigma_\alpha^2 = 1000)$.

ST.CARanova()

The model is a modification of that proposed by Knorr-Held (2000), and is given by

$$\begin{aligned} M_{kt} &= \phi_k + \delta_t + \gamma_{kt}, \\ \phi_k | \boldsymbol{\phi}_{-k}, \mathbf{W} &\sim \text{N} \left(\frac{\rho_\phi \sum_{j=1}^K w_{kj} \phi_j}{\rho_\phi \sum_{j=1}^K w_{kj} + 1 - \rho_\phi}, \frac{\tau_\phi^2}{\rho_\phi \sum_{j=1}^K w_{kj} + 1 - \rho_\phi} \right), \\ \delta_t | \boldsymbol{\delta}_{-t}, \mathbf{D} &\sim \text{N} \left(\frac{\rho_\delta \sum_{j=1}^T d_{tj} \delta_j}{\rho_\delta \sum_{j=1}^T d_{tj} + 1 - \rho_\delta}, \frac{\tau_\delta^2}{\rho_\delta \sum_{j=1}^T d_{tj} + 1 - \rho_\delta} \right), \\ \gamma_{kt} &\sim \text{N}(0, \tau_\gamma^2), \\ \tau_\phi^2, \tau_\delta^2, \tau_\gamma^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho_\phi, \rho_\delta &\sim \text{Uniform}(0, 1). \end{aligned}$$

Here the spatio-temporal autocorrelation is modelled by a common set of spatial random effects $\boldsymbol{\phi} = (\phi_1, \dots, \phi_K)$ and a common set of temporal random effects $\boldsymbol{\delta} = (\delta_1, \dots, \delta_T)$, and both are modelled by the CAR prior proposed by Leroux *et al.* (1999). Additionally, the model can incorporate an optional set of independent space-time interactions $\boldsymbol{\gamma} = (\gamma_{11}, \dots, \gamma_{KT})$, which can be specified by the argument `interaction=TRUE` (the default) in the function call. All sets of random effects are mean centered. Fixed uniform (ρ_ϕ, ρ_δ) or flexible conjugate $(\tau_\phi^2, \tau_\delta^2, \tau_\gamma^2)$ priors are specified for the remaining parameters, and the default specifications for the latter are $(a = 0.001, b = 0.001)$.

ST.CARsepspatial()

The model is given by

$$\begin{aligned} M_{kt} &= \phi_{kt} + \delta_t, \\ \phi_{kt} | \boldsymbol{\phi}_{-k,t}, \mathbf{W} &\sim \text{N} \left(\frac{\rho \sum_{j=1}^K w_{kj} \phi_{jt}}{\rho \sum_{j=1}^K w_{kj} + 1 - \rho}, \frac{\tau_{t\phi}^2}{\rho \sum_{j=1}^K w_{kj} + 1 - \rho} \right), \\ \delta_t | \delta_{t-1} &\sim \text{N}(\delta_{t-1}, \tau_\delta^2), \end{aligned}$$

$$\begin{aligned}\tau_{t_\phi}^2, \tau_\delta^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho &\sim \text{Uniform}(0, 1).\end{aligned}$$

Here $\phi_{-k,t}$, denotes the set of all ϕ random effects at time t except the k th. This model fits an overall temporal trend to the data $\boldsymbol{\delta} = (\delta_1, \dots, \delta_T)$, and a separate (uncorrelated) spatial surface $\boldsymbol{\phi} = (\phi_{1t}, \dots, \phi_{Kt})$ at each time period t . Each spatial surface is modelled by the CAR prior proposed by Leroux *et al.* (1999), with a common spatial dependence parameter ρ and a varying variance parameter $\tau_{t_\phi}^2$. Thus the collection $(\tau_{1_\phi}^2, \dots, \tau_{T_\phi}^2)$ allows us to examine the extent to which the magnitude of the spatial variation in the data has changed over time. As with all other models the random effects are zero mean centered, while flat and conjugate priors are specified for $(\rho, \tau_{t_\phi}^2)$ respectively with $(a = 0.001, b = 0.001)$ being the default values.

ST.CARar()

The model is that proposed by Rushworth *et al.* (2014a), and is given by

$$\begin{aligned}M_{kt} &= \phi_{kt}, \\ \boldsymbol{\phi}_t | \boldsymbol{\phi}_{t-1} &\sim \text{N}(\gamma \boldsymbol{\phi}_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}) \quad t = 2, \dots, T, \\ \boldsymbol{\phi}_1 &\sim \text{N}(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}), \\ \tau^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho, \gamma &\sim \text{Uniform}(0, 1).\end{aligned}\tag{3}$$

In this model $\boldsymbol{\phi}_t = (\phi_{1t}, \dots, \phi_{Kt})$ is the vector of random effects for time period t , which evolve over time via a multivariate first order autoregressive process with temporal autoregressive parameter γ . The temporal autocorrelation is thus induced via the mean $\gamma \boldsymbol{\phi}_{t-1}$, while spatial autocorrelation is induced by the variance $\tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}$. This precision matrix corresponds to the CAR prior proposed by Leroux *et al.* (1999) and is given by

$$\mathbf{Q}(\mathbf{W}, \rho) = \rho[\text{diag}(\mathbf{W}\mathbf{1}) - \mathbf{W}] + (1 - \rho)\mathbf{I},$$

where $\mathbf{1}$ is the $K \times 1$ vector of ones while \mathbf{I} is the $K \times K$ identity matrix. As with all other models the random effects are zero mean centered, while flat and conjugate priors are specified for (ρ, γ, τ^2) respectively with $(a = 0.001, b = 0.001)$ being the default values.

ST.CARadaptive()

The model is that proposed by Rushworth *et al.* (2014b), and is an extension of ST.CARar() proposed by Rushworth *et al.* (2014a). It has the same autoregressive random effects structure as ST.CARar(), namely:

$$\begin{aligned}M_{kt} &= \phi_{kt}, \\ \boldsymbol{\phi}_t | \boldsymbol{\phi}_{t-1} &\sim \text{N}(\gamma \boldsymbol{\phi}_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}) \quad t = 2, \dots, T,\end{aligned}\tag{4}$$

$$\begin{aligned}\phi_1 &\sim \text{N}(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}), \\ \tau^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho, \gamma &\sim \text{Uniform}(0, 1).\end{aligned}$$

However, this random effects structure assumes there is a single level of spatial dependence in the data, which is controlled by ρ . Thus all pairs of adjacent areal units will have strongly autocorrelated random effects if ρ is close to one, while no such spatial dependence will exist anywhere if ρ is close to zero. However, real data may exhibit spatially varying dependences, as two adjacent areal units may exhibit similar values suggesting a value of ρ close to one, while another pair may exhibit different values suggesting a value of ρ close to zero.

We allow for localised residual spatial autocorrelation by allowing spatially neighbouring random effects to be correlated (inducing smoothness) or conditionally independent (no smoothing), which is achieved by modelling the non-zero elements of the neighbourhood matrix \mathbf{W} as unknown parameters rather than fixed constants equal to one. These adjacency parameters are collectively denoted by $\mathbf{w}^+ = \{w_{kj} | k \sim j\}$, where $k \sim j$ means areas (k, j) share a common border. Estimating $w_{kj} \in \mathbf{w}^+$ as equal to zero means (ϕ_{kt}, ϕ_{jt}) are conditionally independent for all t given the remaining random effects, while estimating it close to one means they are correlated. The adjacency parameters in \mathbf{w}^+ are each modelled on the unit interval, by assuming a multivariate Gaussian prior distribution on the transformation $\mathbf{v}^+ = \log(\mathbf{w}^+ / (\mathbf{1} - \mathbf{w}^+))$. This prior is a shrinkage model with a constant mean μ and a diagonal variance matrix with variance parameter ζ^2 , and is given by

$$\begin{aligned}f(\mathbf{v}^+ | \zeta^2, \mu) &\propto \exp \left[-\frac{1}{2\zeta^2} \left(\sum_{v_{ik} \in \mathbf{v}^+} (v_{ik} - \mu)^2 \right) \right], \\ \zeta^2 &\sim \text{Inverse-Gamma}(a, b).\end{aligned}\tag{5}$$

The prior distribution for \mathbf{v}^+ assumes that the degree of smoothing between pairs of adjacent random effects is not spatially dependent, which results from the work of [Rushworth *et al.* \(2014b\)](#) that shows poor estimation performance when \mathbf{v}^+ (and hence \mathbf{w}^+) is assumed to be spatially autocorrelated. Under small values of ζ^2 the elements of \mathbf{v}^+ are shrunk to μ , and here we follow the work of [Rushworth *et al.* \(2014b\)](#) and fix $\mu = 15$ because it avoids numerical issues when transforming between \mathbf{v}^+ and \mathbf{w}^+ and implies a prior preference for values of w_{kj} close to 1. That is as $\zeta^2 \rightarrow 0$ the prior becomes a global smoothing model `ST.CARar()`. As with the other models the default values for the inverse-gamma prior for ζ^2 are $(a = 0.001, b = 0.001)$. For further details see [Rushworth *et al.* \(2014b\)](#).

`ST.CARlocalised()`

The model extends `ST.CARar()` proposed by [Rushworth *et al.* \(2014a\)](#) and the localised smoothing and clustering model of [Lee and Sarran \(2015\)](#). In common with `ST.CARadaptive()` this model allows for localised spatio-temporal autocorrelation, in that some pairs of observations from spatially or temporal adjacent areal units will have similar values (correlation) while others will have large differences between their values (step-changes, no correlation).

This model captures these step-changes via the mean function, where as `ST.CARadaptive()` captured then via the correlation structure (via \mathbf{W}). Model `ST.CARlocalised()` is given by

$$\begin{aligned} M_{kt} &= \lambda_{Z_{kt}} + \phi_{kt}, \\ \phi_t | \phi_{t-1} &\sim \text{N}(\gamma \phi_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}) \quad t = 2, \dots, T, \\ \phi_1 &\sim \text{N}(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}), \\ \tau^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho, \gamma &\sim \text{Uniform}(0, 1), \end{aligned} \tag{6}$$

where the random effects $\phi = (\phi_1, \dots, \phi_T)$ are as specified in `ST.CARar()` and capture spatio-temporal autocorrelation in the data. The other component in the model is a piecewise constant clustering or intercept component $\lambda_{Z_{kt}}$. Thus spatially and temporally adjacent data points (Y_{kt}, Y_{js}) will be similar (autocorrelated) if they are in the same cluster or intercept, that is if $\lambda_{Z_{kt}} = \lambda_{Z_{js}}$, but exhibit a step-change if they are estimated to be in different clusters, that is if $\lambda_{Z_{kt}} \neq \lambda_{Z_{js}}$. The piecewise constant intercept or clustering component comprises at most G distinct levels, making this component a piecewise constant intercept term. The G levels are ordered via the prior specification:

$$\lambda_j \sim \text{Uniform}(\lambda_{j-1}, \lambda_{j+1}) \quad \text{for } j = 1, \dots, G, \tag{7}$$

where $\lambda_0 = -\infty$ and $\lambda_{G+1} = \infty$. Here $Z_{kt} \in \{1, \dots, G\}$ and controls the assignment of the (k, t) th data point to one of the G intercept levels. A penalty based approach is used to model Z_{kt} , where G is chosen larger than necessary and a penalty prior is used to shrink it to the middle intercept level. This middle level is $G^* = (G + 1)/2$ if G is odd and $G^* = G/2$ if G is even, and this penalty ensures that Z_{kt} is only in risk class 1 or G if supported by the data. The allocation prior is independent across areal units but correlated in time, and is given by:

$$\begin{aligned} f(Z_{kt} | Z_{k,t-1}) &= \frac{\exp(-\delta[(Z_{kt} - Z_{k,t-1})^2 + (Z_{kt} - G^*)^2])}{\sum_{r=1}^G \exp(-\delta[(r - Z_{k,t-1})^2 + (r - G^*)^2])} \quad \text{for } t = 2, \dots, T, \\ f(Z_{k1}) &= \frac{\exp(-\delta(Z_{k1} - G^*)^2)}{\sum_{r=1}^G \exp(-\delta(r - G^*)^2)}, \\ \delta &\sim \text{Uniform}(1, M = 100). \end{aligned} \tag{8}$$

Temporal autocorrelation is induced by the $(Z_{kt} - Z_{k,t-1})^2$ component of the penalty, while the $(Z_{kt} - G^*)^2$ component penalises class indicators Z_{kt} towards the middle risk class G^* . The size of this penalty and hence the amount of smoothing that is imparted on \mathbf{Z} is controlled by δ , which is assigned a uniform prior on a large range.

2.3. Inference

All models in this package are fitted in a Bayesian setting using Markov chain Monte Carlo (MCMC) simulation. A combination of Gibbs sampling (when the appropriate full conditional distributions are standard statistical distributions) and Metropolis / Metropolis-Hastings

steps are used, and the majority of the latter use simple random walk proposals. The overall functions that implement the MCMC algorithms are written in R, while the computationally intensive updating steps are written in computationally efficient C++ routines using the R package **Rcpp** (?). Additionally, the sparsity of the neighbourhood matrices **W** and **D** is utilised via their triplet form within the MCMC algorithms, to make the software more computationally efficient.

3. Obtaining and using the software

3.1. Obtaining the software

CARBayesST is a package for the R statistical software, and can be downloaded from CRAN (<http://cran.r-project.org/>) for Windows, Linux and Apple platforms. The package requires R ($\geq 3.0.0$) and depends on packages **MASS**, and **Rcpp** ($\geq 0.11.5$). It also imports functionality from the **coda**, **spam** and **truncdist** packages. Once installed it can be loaded using the command

```
> library(CARBayesST)
```

The packages listed above are automatically loaded by the above call, but a complete spatial analysis beginning with reading in and formatting shapefiles and data, creating the neighbourhood matrix **W**, and plotting the results requires a number of other packages. Thus you may also find the following packages useful: **CARBayes**, **maptools**, **shapefiles**, **sp**, **spdep**.

3.2. Using the software

The software can fit six main models: **ST.CARlinear()**, **ST.CARanova()**, **ST.CARsepspatial()**, **ST.CARar()**, **ST.CARadaptive()** **ST.CARlocalised()**, and full details of the arguments required for each function are given in the helpfiles. However, the main arguments needed are as follows.

- **formula** - A formula for the covariate part of the model using the syntax of the **lm()** function. Offsets can be included here using the **offset()** function. The response and each covariate should be vectors of length $KT \times 1$, where each vector is ordered so that the first K data points are the set of all K spatial locations at time 1, the next K are the set of spatial points for time 2 and so on.
- **family** - The likelihood model which must be one of 'binomial', 'Gaussian' or 'Poisson'.
- **trials** - A vector the same length as the response containing the total number of trials for each area and time period. Only used if **family**='binomial'.
- **W** - A $K \times K$ neighbourhood matrix, which must be symmetric and non-negative. Typically a binary specification is used, where the kj th element equals one if areas ($\mathcal{S}_j, \mathcal{S}_k$) are spatially close (e.g. share a common border) and is zero otherwise. This matrix can be created from a shapefile and data frame using the functionality from the **CARBayes** and **spdep** packages, see Lee (2013) for details.

- `burnin` - The number of MCMC samples to discard as the burnin period.
- `n.sample` - The number of MCMC samples to generate.

When a model has run (see the helpfiles for how to run a model), **CARBayesST** has a `print()` function that will print a model summary to the screen. The **CARBayes** function `summarise.samples` can also be applied to the models in **CARBayesST** to summarise the results. Each model in this package returns a list object with the following components.

- `summary.results` - A summary table of selected parameters that is presented in the print function. This table includes the posterior median (`Median`) and 95% uncertainty interval (2.5%, 97.5%), the number of samples generated (`n.sample`), the acceptance rate for the Markov chain (`% accept`), the effective number of independent samples using the function `effectiveSize()` from the `coda` package (`n.effective`), and the convergence Z-score diagnostic (convergence is suggested by the statistic being within the range (-1.96, 1.96)) proposed by Geweke (1992) and implemented in the `coda` package (`Geweke.diag`).
- `samples` - A list containing the MCMC samples from the model. Each element in the list is a matrix, where each column is the set of samples for a single parameter. The names of the elements in the list correspond to the parameter names in this vignette. For `ST.CARlinear` the (`tau2`, `rho`) elements of the list have columns ordered as $(\tau_\phi^2, \tau_\delta^2)$ and $(\rho_\phi^2, \rho_\delta^2)$ respectively. For `ST.CARanova` the (`tau2`, `rho`) elements of the list have columns ordered as $(\tau_\phi^2, \tau_\delta^2, \tau_\gamma^2)$ (the latter only if `interaction=TRUE`) and $(\rho_\phi^2, \rho_\delta^2)$ respectively. Finally, each model returns samples from the posterior distribution of the fitted values for each data point (`fitted`).
- `fitted.values` - A vector of fitted values for each area and time period in the same order as the data `Y`.
- `residuals` - A vector of residuals for each area and time period in the same order as the data `Y`.
- `modelfit` - Model fit criteria including the Deviance Information Criterion (DIC, Spiegelhalter *et al.* (2002)), the effective number of parameters in the model (p.d), and the Log Marginal Predictive Likelihood (LMPL, Congdon (2005)).
- `accept` - The acceptance probabilities for the parameters.
- `localised.structure` - This element is NULL except for the models `ST.CARadaptive()` and `ST.CARlocalised()`. For `ST.CARadaptive()` this element is a list with 2 $K \times K$ matrices, `Wmn` and `W99` summarising the estimated adjacency relationships. `Wmn` contains the posterior median for each w_{kj} element estimated in the model for adjacent areal units, while `W99` contains $\mathbb{P}(w_{jk} < 0.5 | \mathbf{Y})$. For both matrices, elements corresponding to non-adjacent pairs of areas have NA values. For `ST.CARlocalised()` this element is a vector of length KT , and gives the posterior median class (Z_{kt} value) that each data point is assigned to. This vector is in the same order as the data `Y`.
- `formula` - The formula (as a text string) for the covariate and offset part of the model.

- `model`- A text string describing the model that has been fitted.
- `X` - The design matrix of covariates inherited from the `formula` argument.

This vignette is for version 2.1 of **CARBayesST**, and future version of the software (and hence the vignette) will come out in due course. These future versions will contain a larger suit of spatio-temporal areal unit models, a more comprehensive vignette with fully worked examples, and more functionality for summarising the results.

References

- Anderson C, Lee D, Dean N (2014). “Identifying clusters in Bayesian disease mapping.” *Biostatistics*, **15**, 457–469.
- Bernardinelli L, Clayton D, Pascutto C, Montomoli C, Ghislandi M, Songini M (1995). “Bayesian Analysis of Space-Time Variation in Disease Risk.” *Statistics in Medicine*, **14**, 2433–2443.
- Besag J, Higdon D (1999). “Bayesian Analysis of Agricultural Field Experiments.” *Journal of the Royal Statistical Society Series B*, **61**, 691–746.
- Besag J, York J, Mollié A (1991). “Bayesian Image Restoration with Two Applications in Spatial Statistics.” *Annals of the Institute of Statistics and Mathematics*, **43**, 1–59.
- Charras-Garrido M, Abrial D, de Goer J (2012). “Classification method for disease risk mapping based on discrete hidden Markov random fields.” *Biostatistics*, **13**, 241–255.
- Choi J, Lawson AB (2011). “Evaluation of Bayesian spatial-temporal latent models in small area health data.” *Environmetrics*, **22**, 1008–1022.
- Congdon P (2005). *Bayesian models for categorical data*. 1st edition. John Wiley and Sons.
- Geweke J (1992). “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments.” In *IN BAYESIAN STATISTICS*, pp. 169–193. University Press.
- Knorr-Held L (2000). “Bayesian modelling of inseparable space-time variation in disease risk.” *Statistics in Medicine*, **19**, 2555–2567.
- Lee D (2013). “CARBayes: An R Package for Bayesian Spatial Modelling with Conditional Autoregressive Priors.” *Journal of Statistical Software*, **55**, 13.
- Lee D, Ferguson C, Mitchell R (2009). “Air Pollution and Health in Scotland: A Multicity Study.” *Biostatistics*, **10**, 409–423.
- Lee D, Minton J, Pryce G (2015). “Bayesian inference for the dissimilarity index in the presence of spatial autocorrelation.” *Spatial Statistics*, **11**, 81–95.
- Lee D, Sarran C (2015). “Controlling for unmeasured confounding and spatial misalignment in long-term air pollution and health studies.” *Environmetrics*, **to appear**.

- Leroux B, Lei X, Breslow N (1999). *Estimation of Disease Rates in Small Areas: A New Mixed Model for Spatial Dependence*, chapter Statistical Models in Epidemiology, the Environment and Clinical Trials, Halloran, M and Berry, D (eds), pp. 135–178. Springer-Verlag, New York.
- Lunn D, Spiegelhalter D, Thomas A, Best N (2009). “The BUGS Project: Evolution, Critique and Future Directions .” *Statistics in Medicine*, **28**, 3049–3082.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.
- Rue H, Martino S, Chopin N (2009). “Approximate Bayesian Inference for Latent Gaussian Models Using Integrated Nested Laplace Approximations (with discussion).” *Journal of the Royal Statistical Society Series B*, **71**, 319–392.
- Rushworth A, Lee D, Mitchell R (2014a). “A spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in Greater London.” *Spatial and Spatio-temporal Epidemiology*, **10**, 29–38.
- Rushworth A, Lee D, Sarran C (2014b). “An adaptive spatio-temporal smoothing model for estimating trends and step changes in disease risk.” *arXiv*, **1411.0924**, 1–27.
- Spiegelhalter D, Best N, Carlin B, Van der Linde A (2002). “Bayesian measures of model complexity and fit.” *Journal of the Royal Statistical Society Series B*, **64**, 583–639.
- Wakefield J (2007). “Disease Mapping and Spatial Regression with Count Data.” *Biostatistics*, **8**, 158–183.
- Wall M (2004). “A Close Look at the Spatial Structure Implied by the CAR and SAR Models.” *Journal of Statistical Planning and Inference*, **121**, 311–324.

Affiliation:

Duncan Lee
School of Mathematics and Statistics
15 University Gardens
University of Glasgow
Glasgow
G12 8QQ, Scotland
E-mail: Duncan.Lee@glasgow.ac.uk
URL: <http://www.gla.ac.uk/schools/mathematicsstatistics/staff/duncanlee/>

Alastair Rushworth
Mathematics and Statistics
26 Richmond Street
University of Strathclyde
Glasgow
G1 1XH, Scotland
E-mail: alastair.rushworth@strath.ac.uk