

Spatio-temporal areal unit modelling in R with conditional autoregressive priors using the CARBayesST package

Duncan Lee
University of Glasgow

Alastair Rushworth
University of Strathclyde

Gary Napier
University of Glasgow

Abstract

Spatial data relating to non-overlapping areal units are prevalent in fields such as economics, environmental science, epidemiology and social science, and a large suite of modelling tools have been developed for analysing these data. Many utilise conditional autoregressive (CAR) priors to capture the spatial autocorrelation inherent in these data, and software such as **CARBayes** and **R-INLA** have been developed to make these models easily accessible to others. Such spatial data are typically available for multiple time periods, and the development of methodology for capturing temporally changing spatial dynamics is the focus of much current research. A sizeable proportion of this literature has focused on extending CAR priors to the spatio-temporal domain, and this article presents the **R** package **CARBayesST**, which is the first dedicated software for spatio-temporal areal unit modelling with conditional autoregressive priors. The software can fit a range of models focused on different aspects of space-time modelling, including estimation of overall space and time trends, and the identification of clusters of areal units that exhibit elevated values. This paper outlines the class of models that the software can implement, before applying them to one simulated and two real examples in the fields of epidemiology and housing market analysis.

Keywords: Bayesian inference, conditional autoregressive priors, **R** package, spatio-temporal areal unit modelling.

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, for the entire unit. Examples include the total yield 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, Minton, and Pryce 2015), the average score on SAT college entrance exams across US states (Wall 2004), and the total number of cases of chronic obstructive pulmonary disease from populations living in counties in Georgia, USA (Choi and Lawson 2011). Areal unit data such as these have become increasingly available in recent times, due to the creation of databases such as Scottish Statistics (<http://statistics.gov.scot/>), Health and Social Care Information Centre Indicator Portal (<http://www.hscic.gov.uk/indicatorportal>), and cancer registries such as

the Surveillance Epidemiology and End Results programme (<http://seer.cancer.gov>).

These databases provide data on a set of K areal units for N consecutive time periods, yielding a rectangular array of $K \times N$ 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, Ferguson, and Mitchell 2009), identifying clusters of contiguous areal units that exhibit an elevated risk of disease compared with neighbouring areas (see Charras-Garrido, Abrial, and de Goer 2012 and Anderson, Lee, and Dean 2014), and quantifying the level of segregation in a city between two or more different groups (see Lee *et al.* 2015). As a result different space-time structures have been proposed for modelling spatio-temporal data, which depend on the underlying question(s) of interest and the goals of the analysis.

However, a common challenge when modelling these data is that of spatio-temporal autocorrelation, namely that observations from geographically close areal units and temporally close 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 largely the same populations in consecutive time periods, while 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 variable is not included in a regression model, and hence its omission induces spatial structure into the residuals. 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 close together.

Predominantly, a Bayesian approach is taken to modelling these data in the statistical community, where the spatio-temporal structure is modelled via sets of autocorrelated random effects. Conditional autoregressive (CAR, Besag, York, and Mollié 1991) priors and spatio-temporal extensions thereof are typically assigned to these random effects to capture the autocorrelation, which are special cases of a Gaussian Markov Random Field (GMRF). A range of models have been proposed with different space-time structures, which can be used to answer different questions of interest about the data. For example, Knorr-Held (2000) proposed an ANOVA style decomposition of the data into overall spatial and temporal trends and a space-time interaction, which allows common patterns such as the region wide average temporal trend to be estimated. In contrast, Li, Best, Hansell, Ahmed, and Richardson (2012) developed a model for identifying areas that exhibited unusual trends that were different from the overall region wide trend, while Lee and Lawson (2016) presented a model for identifying spatio-temporal clustering in the data.

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, Spiegelhalter, Thomas, and Best 2009) and **R-INLA** (Rue, Martino, and Chopin 2009), to specialist spatial modelling packages in the statistical software R (R Core Team 2016) such as **CARBayes** (Lee 2013), **spatcounts** (Schabenberger 2009) and **spdep** (Bivand 2013). 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** (Bakar and Sahu 2015) and **spBayes** (Finley, Banerjee, and Gelfand 2015). For areal unit data the **surveillance** (Paul and Meyer 2016) package models epidemic data, the **plm** (Croissant and Millo 2008) and **splm** (Millo and Piras 2012) packages model panel data, while the **nlme** (Pinheiro, Bates, DebRoy, Sarkar, and R Core Team 2015) and **lme4** (Bates, Mächler, Bolker, and Walker 2015) 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**.

CARBayesST can fit models with different spatio-temporal structures, which allows the user to answer different questions about their data within a common software environment. A number of models can be implemented, including: (i) a spatially varying linear time trends model (similar to Bernardinelli, Clayton, Pascutto, Montomoli, Ghislandi, and Songini 1995); (ii) a spatial and temporal main effects and interaction model (similar to Knorr-Held 2000); (iii) a spatially autocorrelated autoregressive time series model (Rushworth, Lee, and Mitchell 2014); (iv) a model with a common temporal trend but varying spatial surfaces (similar to Napier, Lee, Robertson, Lawson, and Pollock 2016); (v) a spatially adaptive smoothing model for localised spatial smoothing (Rushworth, Lee, and Sarran 2016); and (vi) a spatio-temporal clustering model (Lee and Lawson 2016). The software has the same syntax and feel as the R package **CARBayes** for spatial areal unit modelling, and retains all of its easy-to-use features. These include 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** that allows it to share the latter's model summary functionality for interpreting the results. The models available in this software can be fitted to binomial, Gaussian (in some cases) or Poisson data, and Section 2 summarises the models that can be implemented. Section 3 provides an overview of the software and its functionality, while Section 4 gives an example of how to apply the software to simulated data to illustrate the correctness of the algorithms. Sections 5 and 6 present two worked examples illustrating how to use the software for epidemiological and housing market research, while Section 7 concludes with a summary of planned future development for this package.

2. Spatio-temporal models for areal unit data

This section outlines the class of Bayesian hierarchical models that **CARBayesST** can implement, and in all cases inference is based on Markov chain Monte Carlo (MCMC) simulation. The first part of this section outlines the general hierarchical model that can be fitted, while the second describes the range of space-time random effects structures that are available.

2.1. General Bayesian hierarchical model

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 $t = 1, \dots, N$ consecutive time periods. Thus data are available for a $K \times N$ rectangular array with K rows (spatial units) and N columns (time periods). The response data are denoted by $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_N)_{K \times N}$, where $\mathbf{Y}_t = (Y_{1t}, \dots, Y_{Kt})$ denotes the $K \times 1$ column vector of observations for all K spatial units for time period

t . Next, a vector of known offsets are denoted by $\mathbf{O} = (\mathbf{O}_1, \dots, \mathbf{O}_N)_{K \times N}$, where similarly $\mathbf{O}_t = (O_{1t}, \dots, O_{Kt})$ denotes the $K \times 1$ column 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. Note, non-linear covariate-response relationships can be included by adding transformations of covariates (e.g. squared) or spline basis functions (e.g. using `ns()`) to the covariate vector. **CARBayesST** can fit a generalised linear mixed model to these data, whose general form is:

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

The vector of covariate regression parameters are denoted by $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$, and a multivariate Gaussian prior is assumed with mean $\boldsymbol{\mu}_\beta$ and diagonal variance matrix Σ_β that can be chosen by the user. The ψ_{kt} term is a latent component for areal unit k and time period t encompassing one or more sets of spatio-temporally autocorrelated random effects, and the complete set are denoted by $\boldsymbol{\psi} = (\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_N)$, where $\boldsymbol{\psi}_t = (\psi_{1t}, \dots, \psi_{Kt})$. **CARBayesST** can fit a number of different spatio-temporal structures for ψ_{kt} , all of which are outlined in Section 2.2 below. The software can implement model (1) for binomial, Gaussian and Poisson data, and the exact specifications of each are given below:

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

In the binomial model (n_{kt}, θ_{kt}) respectively denote the number of trials and the probability of success in each trial in area k and time period t , while in the Gaussian model ν^2 is the observation variance. An inverse-gamma(a, b) prior is specified for ν^2 , and default values of ($a = 1, b = 0.01$) are specified by the software unless they are changed by the user.

2.2. Spatio-temporal random effects models

All models implementable in this package induce spatio-temporal autocorrelation into the response data \mathbf{Y} via the latent random effects component $\boldsymbol{\psi}$, using CAR type prior distributions and spatio-temporal extensions thereof. Spatial autocorrelation is controlled by a symmetric non-negative $K \times K$ neighbourhood, weight or adjacency matrix $\mathbf{W} = (w_{kj})$, where w_{kj} represents the spatial closeness between areal units $(\mathcal{S}_k, \mathcal{S}_j)$. Larger valued elements represent spatial closeness between the two areas in question, where as smaller or zero values correspond to areas that are not spatially close. Typically, \mathbf{W} is assumed to be binary, 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, $w_{kk} = 0$. Under this binary specification the values of (ψ_{kt}, ψ_{jt}) for spatially adjacent areal units (where $w_{kj} = 1$) are spatially autocorrelated, where as values for non-neighbouring areal units (where $w_{kj} = 0$) are conditionally independent given the remaining $\{\psi_{it}\}$ 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, non-negative, and for each row sum to be greater than zero. Similarly, the model **ST.CARanova()** defined below uses a binary $N \times N$ temporal neighbourhood matrix $\mathbf{D} = (d_{tj})$, where $d_{tj} = 1$ if $|j - t| = 1$ and $d_{tj} = 0$ otherwise. **CARBayesST** can fit the following models for ψ :

- **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()** - fits a model similar to the overall temporal trend and separate spatial surfaces model proposed by [Napier *et al.* \(2016\)](#). 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.* \(2014\)](#).
- **ST.CARadaptive()** - fits the localised smoothing model proposed by [Rushworth *et al.* \(2016\)](#).
- **ST.CARlocalised()** - fits the localised smoothing and clustering model proposed by [Lee and Lawson \(2016\)](#). Note, this model can only be applied to binomial or Poisson data.

This set of models allows users to fit different spatio-temporal structures to their data and thus examine different underlying hypotheses, and full details of each model are given below.

ST.CARlinear()

This model is a modification of that proposed by [Bernardinelli *et al.* \(1995\)](#), and estimates separate linear time trends for each areal unit. Thus it is appropriate if the goal of the analysis is to estimate which areas are exhibiting increasing or decreasing (linear) trends in the response over time. The full model specification is given below.

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

Here $\bar{t} = (1/N) \sum_{t=1}^N t$, and thus the modified linear temporal trend covariate is $t^* = (t - \bar{t})/N$ and runs over a centered unit interval. Thus 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, Lei, and Breslow (2000), and are mean centered, that is $\sum_{j=1}^K \phi_j = \sum_{j=1}^K \delta_j = 0$. Here (ρ_{int}, ρ_{slo}) 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 (for example if $\rho_{slo} = 0$ then $\delta_k \sim N(0, \tau_{slo}^2)$). Flat uniform priors on the unit interval are specified for the spatial dependence parameters (ρ_{int}, ρ_{slo}) , while conjugate inverse-gamma and Gaussian priors are specified for the random effects variances $(\tau_{int}^2, \tau_{slo}^2)$ and the overall slope parameter α respectively. The corresponding hyperparameters $(a, b, \mu_\alpha, \sigma_\alpha^2)$ can be chosen by the user, and the default values specified by the software are $(a = 1, b = 0.01, \mu_\alpha = 0, \sigma_\alpha^2 = 1000)$ which correspond to weakly informative prior distributions. Alternatively, the dependence parameters (ρ_{int}, ρ_{slo}) can be fixed at values in the unit interval $[0, 1]$ rather than being estimated in the model, by specifying arguments to the `ST.CARlinear()` function. For example, using the arguments `fix.rho.slo=TRUE`, `rho.slo=1` sets $\rho_{slo} = 1$ when fitting the model. Finally, missing (NA) values are allowed in the response data \mathbf{Y} for this model.

ST.CARanova()

The model is a modification of that proposed by Knorr-Held (2000), and decomposes the spatio-temporal variation in the data into 3 components, an overall spatial effect common to all time periods, an overall temporal trend common to all spatial units, and a set of independent space-time interactions. Thus this model is appropriate if the goal is to estimate overall region wide time trends and spatial patterns. The model specification is given below.

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

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 $\boldsymbol{\delta} = (\delta_1, \dots, \delta_N)$, and both are modelled by the CAR prior proposed by Leroux *et al.* (2000). Additionally, the model can incorporate an optional set of independent space-time interactions $\boldsymbol{\gamma} = (\gamma_{11}, \dots, \gamma_{KN})$, 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_S, ρ_T) or conjugate $(\tau_S^2, \tau_T^2, \tau_I^2)$ priors are specified for the remaining parameters, and the default specifications for the latter are $(a = 1, b = 0.01)$. Alternatively, in common with the `ST.CARlinear()` function the

dependence parameters (ρ_S, ρ_T) can be fixed at values in the unit interval $[0, 1]$ rather than being estimated in the model, for full details see the help file for this function. Finally, missing (NA) values are allowed in the response data \mathbf{Y} for this model.

ST.CARsepspatial()

The model is a generalisation of that proposed by Napier *et al.* (2016), and represents the data by two components, an overall region-wide temporal trend, and separate spatial surfaces for each time period which share a common spatial dependence parameter but have different spatial variances. This model is appropriate if the goal is to estimate both a common overall temporal trend and the extent to which the spatial variation in the response has changed in magnitude over time. The model specification is given below.

$$\begin{aligned} \psi_{kt} &= \phi_{kt} + \delta_t, \\ \phi_{kt} | \phi_{-kt}, \mathbf{W} &\sim \text{N} \left(\frac{\rho_S \sum_{j=1}^K w_{kj} \phi_{jt}}{\rho_S \sum_{j=1}^K w_{kj} + 1 - \rho_S}, \frac{\tau_t^2}{\rho_S \sum_{j=1}^K w_{kj} + 1 - \rho_S} \right), \\ \delta_t | \delta_{-t}, \mathbf{D} &\sim \text{N} \left(\frac{\rho_T \sum_{j=1}^N d_{tj} \delta_j}{\rho_T \sum_{j=1}^N d_{tj} + 1 - \rho_T}, \frac{\tau_T^2}{\rho_T \sum_{j=1}^N d_{tj} + 1 - \rho_T} \right), \\ \tau_1^2, \dots, \tau_N^2, \tau_T^2 &\sim \text{Inverse-Gamma}(a, b), \\ \rho_S, \rho_T &\sim \text{Uniform}(0, 1), \end{aligned}$$

where $\phi_{-k,t} = (\phi_{1,t}, \dots, \phi_{k-1,t}, \phi_{k+1,t}, \dots, \phi_{K,t})$. This model fits an overall temporal trend to the data $\boldsymbol{\delta} = (\delta_1, \dots, \delta_N)$ that is common to all areal units, which is augmented with a separate (uncorrelated) spatial surface $\boldsymbol{\phi}_t = (\phi_{1,t}, \dots, \phi_{K,t})$ at each time period t . The overall temporal trend and each spatial surface are modelled by the CAR prior proposed by Leroux *et al.* (2000), and the latter have a common spatial dependence parameter ρ_S but a temporally-varying variance parameter τ_t^2 . Thus the collection $(\tau_1^2, \dots, \tau_N^2)$ allows one to examine the extent to which the magnitude of the spatial variation in the data has changed over time. Note that here we fix ρ_S to be constant in time as it is not orthogonal to τ_t^2 , thus if it varied then any changes in τ_t^2 would not directly correspond to changes in spatial variance over time. As with all other models the random effects are zero mean centered, while flat and conjugate priors are specified for (ρ_S, ρ_T) and $(\tau_T^2, \tau_1^2, \dots, \tau_N^2)$ respectively with $(a = 1, b = 0.01)$ being the default values. Alternatively, in common with the ST.CARlinear() function the dependence parameters (ρ_S, ρ_T) can be fixed at values in the unit interval $[0, 1]$ rather than being estimated in the model.

ST.CARar()

The model is that proposed by [Rushworth *et al.* \(2014\)](#), and models the spatio-temporal structure as a multivariate first order autoregressive process with a spatially correlated precision matrix. This model is appropriate if one wishes to estimate the evolution of the spatial response surface over time without forcing it to be the same for each time period. The model specification is given below.

$$\begin{aligned}
\psi_{kt} &= \phi_{kt}, \\
\phi_t | \phi_{t-1} &\sim N(\rho_T \phi_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho_S)^{-1}) \quad t = 2, \dots, N, \\
\phi_1 &\sim N(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho_S)^{-1}), \\
\tau^2 &\sim \text{Inverse-Gamma}(a, b), \\
\rho_S, \rho_T &\sim \text{Uniform}(0, 1).
\end{aligned} \tag{3}$$

In this model $\phi_t = (\phi_{1t}, \dots, \phi_{Kt})$ 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 ρ_T . The temporal autocorrelation is thus induced via the mean $\rho_T \phi_{t-1}$, while spatial autocorrelation is induced by the variance $\tau^2 \mathbf{Q}(\mathbf{W}, \rho_S)^{-1}$. The corresponding precision matrix $\mathbf{Q}(\mathbf{W}, \rho_S)$ was proposed by [Leroux *et al.* \(2000\)](#) and corresponds to the CAR models used in the other models used above. The algebraic form of this matrix is given by

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

where $\mathbf{1}$ is the $K \times 1$ vector of ones while \mathbf{I} is the $K \times K$ identity matrix. In common with all other models the random effects are zero mean centered, while flat and conjugate priors are specified for (ρ_S, ρ_T) and τ^2 respectively with $(a = 1, b = 0.01)$ being the default values. In common with the `ST.CARanova()` function the dependence parameters (ρ_S, ρ_T) can also be fixed at values in the unit interval $[0, 1]$ rather than being estimated in the model. Finally, missing (NA) values are allowed in the response data \mathbf{Y} for this model.

ST.CARadaptive()

The model is that proposed by [Rushworth *et al.* \(2016\)](#), and is an extension of `ST.CARar()` proposed by [Rushworth *et al.* \(2014\)](#) to allow for spatially adaptive smoothing. It is appropriate if one believes that the residual spatial autocorrelation in the response after accounting for the covariates is consistent over time but has a localised structure. That is, it is strong in some parts of the study region but weak in others. The model has the same autoregressive random effects structure as the previous model `ST.CARar()`, namely:

$$\begin{aligned}
\psi_{kt} &= \phi_{kt}, \\
\phi_t | \phi_{t-1} &\sim N(\rho_T \phi_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho_S)^{-1}) \quad t = 2, \dots, N, \\
\phi_1 &\sim N(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho_S)^{-1}), \\
\tau^2 &\sim \text{Inverse-Gamma}(a, b), \\
\rho_S, \rho_T &\sim \text{Uniform}(0, 1).
\end{aligned} \tag{4}$$

However, the random effects from `ST.CARar()` have a single level of spatial dependence that is controlled by ρ_S . All pairs of adjacent areal units will have strongly autocorrelated random effects if ρ_S is close to one, while no such spatial dependence will exist anywhere if ρ_S 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_S close to one, while another adjacent pair may exhibit very different values suggesting a value of ρ_S close to zero.

This model allows for localised 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 assuming they are fixed constants. 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 time periods 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 transformed scale $\mathbf{v}^+ = \log(\mathbf{w}^+ / (\mathbf{1} - \mathbf{w}^+))$. 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

$$f(\mathbf{v}^+ | \zeta^2, \mu) \propto \exp \left[-\frac{1}{2\tau_w^2} \left(\sum_{v_{ik} \in \mathbf{v}^+} (v_{ik} - \mu)^2 \right) \right], \quad (5)$$

$$\tau_w^2 \sim \text{Inverse-Gamma}(a, b).$$

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.* \(2016\)](#) that shows poor estimation performance when \mathbf{v}^+ (and hence \mathbf{w}^+) is assumed to be spatially autocorrelated. Under small values of τ_w^2 the elements of \mathbf{v}^+ are shrunk to μ , and here we follow the work of [Rushworth *et al.* \(2016\)](#) 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 $\tau_w^2 \rightarrow 0$ the prior becomes the global smoothing model `ST.CARar()`, where as when τ_w^2 increases both small and large values in \mathbf{w}^+ are supported by the prior. As with the other models the default values for the inverse-gamma prior for τ_w^2 are $(a = 1, b = 0.01)$. Alternatively, it is possible to fix ρ_S using the `rhofix` argument, e.g. `rhofix=1` fixes $\rho_S = 1$, so that globally the spatial correlation is strong and is altered locally by the estimates of \mathbf{w}^+ . For further details see [Rushworth *et al.* \(2016\)](#).

`ST.CARlocalised()`

The model was proposed by [Lee and Lawson \(2016\)](#), and augments the smooth spatio-temporal variation in `ST.CARar()` with a piecewise constant intercept component. This model is appropriate when the aim of the analysis is to identify clusters of areas that exhibit elevated (and reduced) values of the response compared with their geographical and temporal neighbours. Thus this model is similar to `ST.CARadaptive()`, in that both relax the restrictive assumption that if two areas are close together then their estimated random

effects must be similar. This model captures any step-changes in the response via the mean function, whereas `ST.CARadaptive()` captures them via the correlation structure (via \mathbf{W}). Model `ST.CARlocalised()` is given by

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

The random effects $\phi = (\phi_1, \dots, \phi_T)$ are modelled by a simplification of the `ST.CARar()` model with $\rho_S = 1$, which corresponds to the intrinsic CAR model proposed by Besag *et al.* (1991). Note, for this model the inverse $\mathbf{Q}(\mathbf{W})^{-1}$ does not exist as the precision matrix is singular. This simplification is made so that the random effects capture the globally smooth spatio-temporal autocorrelation in the data, allowing the other component to capture localised clustering and step-changes. This second component 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 the extreme low and high risk classes if supported by the data. Thus G is the maximum number of distinct intercept terms allowed in the model, and is not the actual number of intercept terms estimated in the model. The allocation prior is independent across areal units but correlated in time, and is given by:

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

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. The default value is $m = 10$, and full details of this model can be found in Lee and Lawson (2016).

2.3. Inference

All models in this package are fitted in a Bayesian setting using Markov chain Monte Carlo simulation. A combination of Gibbs sampling (when the appropriate full conditional distributions are proportional to 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 as computationally efficient C++ routines using the R package **Rcpp** (Eddelbuettel and Francois 2011). Additionally, the sparsity of the neighbourhood matrices \mathbf{W} and \mathbf{D} are utilised via their triplet forms when updating the random effects within the algorithms, which increases the computational efficiency of the software.

3. Loading and using the software

3.1. Loading the software

CARBayesST is a package for the R (R Core Team 2016) statistical software and can be downloaded from the Comprehensive R Archive Network (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** (Venables and Ripley 2002), and **Rcpp** ($\geq 0.11.5$). Additionally, it imports functionality from the **CARBayesdata** (Lee 2016), **coda** (Plummer, Best, Cowles, and Vines 2006), **dplyr** (Wickham and Francois 2015) **matrixcalc** (Novomestky 2012), **sp** ((Bivand, Pebesma, and Gomez-Rubio 2013)), **spam** (Furrer and Sain 2010), **spdep**, **stats**, **truncdist** (Novomestky and Nadarajah 2012) and **utils** packages. Once installed it can be loaded using the command `library(CARBayesST)`.

The packages listed above for use in **CARBayesST** 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 \mathbf{W} , and plotting the results requires a number of other packages. Thus the worked examples in this paper utilise functionality from the following packages: **CARBayes**, **maptools** (Bivand and Lewin-Koh 2015), **sp** and **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 common to all the functions that are required for a baseline analysis (for example using default priors) are as follows.

- **formula** - A formula for the covariate part of the model using the same syntax used in 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** - This is only needed if **family**=‘binomial’, and is a vector the same length and in the same order as the response containing the total number of trials for each area and time period.
- **W** - A $K \times K$ symmetric and non-negative neighbourhood matrix whose row sums must all be positive. Typically a binary specification is used, where the kj th element w_{kj} 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 specified by hand or created from a shapefile and data frame using functionality from the **CARBayes** and **spdep** packages.
- **burnin** - The number of MCMC samples to discard as the burn-in period.
- **n.sample** - The number of MCMC samples to generate.

When a model has run **CARBayesST** has its own **print()** function that prints a summary of the fitted model to the screen, including both parameter summaries and convergence diagnostics for the MCMC run. Additionally, the **CARBayes** functions **summarise.samples()** and **summarise.lincomb()** can be applied to **CARBayesST** models to summarise the results. The software updates the user on its progress to the R console, which allows the user to monitor the function’s progress. However, using the **verbose=FALSE** option will disable this feature. Once run, each model in this package returns a list object with the following components.

- **summary.results** - A summary table of selected parameters that is presented when using the **print()** function. The table includes the posterior median (**Median**) and 95% credible 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 **effectiveSize()** function from the **coda** package (**n.effective**), and the convergence diagnostic proposed by Geweke (1992) and implemented in the **coda** package (**Geweke.diag**). This diagnostic takes the form of a Z-score, so that convergence is suggested by the statistic being within the range (-1.96, 1.96).
- **samples** - A list containing the MCMC samples from the model, where each element in the list is a matrix. The names of these matrix objects correspond to the parameters defined in Section 2 of this paper, and each column contains the set of samples for a single parameter. For example, for **ST.CARlinear()** the (**tau2**, **rho**) elements of the list have columns ordered as $(\tau_{int}^2, \tau_{slo}^2)$ and $(\rho_{int}^2, \rho_{slo}^2)$ respectively. Similarly, for **ST.CARanova()** the (**tau2**, **rho**) elements of the list have columns ordered as $(\tau_S^2, \tau_T^2, \tau_I^2)$ (the latter only if **interaction=TRUE**) and (ρ_S^2, ρ_T^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 based on posterior medians for each area and time period in the same order as the data **Y**.

- **residuals** - A vector of raw residuals based on posterior medians for each area and time period in the same order as the data \mathbf{Y} .
- **modelfit** - Model fit criteria including the Deviance Information Criterion (DIC, Spiegelhalter, Best, Carlin, and Van der Linde (2002)) and its corresponding estimated effective number of parameters (p.d), the Watanabe-Akaike Information Criterion (WAIC, Watanabe (2010)) and its corresponding estimated number of effective parameters (p.w), and the Log Marginal Predictive Likelihood (LMPL, Congdon (2005)). The best fitting model is one that minimises the DIC and WAIC but maximises the LMPL.
- **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 indicator variables for whether $\mathbb{P}(w_{jk} < 0.5 | \mathbf{Y}) > 0.99$. 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 \mathbf{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.

The remainder of this paper illustrates the **CARBayesST** software via a small simulation study to illustrate the correctness of the MCMC algorithms, and two worked examples, the latter of which utilise spatio-temporal data to answer important questions in public health and the housing market.

4. Simulation example

This section presents a short simulation example to illustrate both how to use the software to fit a model and the correctness of the model fitting algorithm. The example here is for the `ST.CARanova()` model, and we note that the correctness of the algorithms from the `ST.CARsepspatial()` (Napier *et al.* 2016), `ST.CARar()` (Rushworth *et al.* 2014), `ST.CARadaptive()` (Rushworth *et al.* 2016), and `ST.CARlocalised()` (Lee and Lawson 2016) models have been assessed in the accompanying papers where the models were developed. We generate data from a binomial logistic model because the two example analyses that follow in Sections 5 and 6 utilise Poisson log-linear models. However, similar simulated examples can easily be created for the other models by adapting the code below.

4.1. Data generation

Consider a spatial region comprising $K = 400$ areal units on a regular 20×20 grid and $N = 20$ consecutive time periods. Such a grid can be constructed from the code

```
R> x.easting <- 1:20
R> x.northing <- 1:20
R> Grid <- expand.grid(x.easting, x.northing)
R> K <- nrow(Grid)
R> N <- 20
R> N.all <- N * K
```

A binary 400×400 spatial neighbourhood matrix \mathbf{W} can be constructed for this spatial region based on (rook, in chess) adjacency using the code

```
R> W <-array(0, c(K,K))
R>   for(i in 1:K)
+   {
+     for(j in 1:K)
+     {
+       temp <- (Grid[i,1] - Grid[j,1])^2 + (Grid[i,2] - Grid[j,2])^2
+       if(temp==1) W[i,j] <- 1
+     }
+   }
```

whilst a binary 20×20 temporal neighbourhood matrix \mathbf{D} can be constructed in a similar manner using

```
R> D <-array(0, c(N,N))
R>   for(i in 1:N)
+   {
+     for(j in 1:N)
+     {
+       if(abs((i-j))==1) D[i,j] <- 1
+     }
+   }
```

From this the precision matrix can be computed for the multivariate Gaussian representation of the spatial random effects ϕ from (Leroux *et al.* 2000) as follows:

```
R> Q.W <- 0.99 * (diag(apply(W, 2, sum)) - W) + 0.01 * diag(rep(1,K))
```

where $\rho_S = 0.99$. This matrix can then be inverted and a sample of random effects generated (assuming $\tau_S^2 = 0.01$) using the code

```
R> Q.W.inv <- solve(Q.W)
R> library(MASS)
R> phi <- mvrnorm(n=1, mu=rep(0,K), Sigma=(0.01 * Q.W.inv))
R> phi.long <- rep(phi, N)
```

Here the last line repeats the spatial random effects N times, as the `ST.CARanova()` model assumes that there is the same set of spatial random effects for each time period. The temporal random effects under the `ST.CARanova()` model have the same functional form but depend on \mathbf{D} rather than \mathbf{W} , and thus a realisation can be generated analogously using the code

```
R> Q.D <- 0.99 * (diag(apply(D, 2, sum)) - D) + 0.01 * diag(rep(1,N))
R> Q.D.inv <- solve(Q.D)
R> delta <- mvrnorm(n=1, mu=rep(0,N), Sigma=(0.01 * Q.D.inv))
R> delta.long <- kronecker(delta, rep(1,K))
```

Again, the final line repeats the temporal random effects for each spatial unit. Next we generate space-time interactions and a covariate x , both of which are generated independently from Gaussian distributions.

```
R> x <- rnorm(n=N.all, mean=0, sd=1)
R> gamma <- rnorm(n=N.all, mean=0, sd=0.1)
```

Finally, assuming an intercept term of $\beta_1 = -0.1$, a regression coefficient of $\beta_2 = 0.1$ and the number of trials for the binomial likelihood in each area and time period being $n_{kt} = 50$, we generate the response variable via the code below. Here LP denotes the linear predictor which contains an intercept term, a covariate and three sets of random effects (spatial, temporal, and interactions).

```
R> beta1 <- -0.1
R> beta2 <- 0.1
R> n <- rep(200, N.all)
R> LP <- beta1 + beta2 * x + phi.long + delta.long + gamma
R> theta.true <- exp(LP) / (1 + exp(LP))
R> Y <- rbinom(n=N.all, size=n, prob=theta.true)
```

4.2. Running the model

The `ST.CARanova()` model can then be applied to these data using the following code.

```
R> library(CARBayesST)
R> model <- ST.CARanova(formula=Y~x, family="binomial", trials=n, W=W, burnin=20000,
+ n.sample=220000, thin=50)
```

In the code above inference is based on 4,000 MCMC samples, which were generated from a single Markov chain that was run for 220,000 iterations with a 20,000 burn-in period and subsequently thinned by 50 to remove the temporal autocorrelation. The estimated proportions $\hat{\theta}_{kt}$ (posterior medians) are plotted in panel (a) of Figure 1 against the true values `theta.true` that generated the data, while panel (b) displays the sample proportions Y_{kt}/n_{kt} against the true values for comparison purposes. In both cases the red line is the line of equality. The code to generate the plot is below. As you can see the fitted values are close to the true values from panel (a), and show markedly less variation than the sample proportions as expected.

```
R> par(mfrow=c(1,2))
R> plot(theta.true, model$fitted.values / n, pch=19, xlab="True values",
+ ylab="Estimated proportions from the model", main="(a)")
R> abline(0,1, col="red", lwd=2)
```

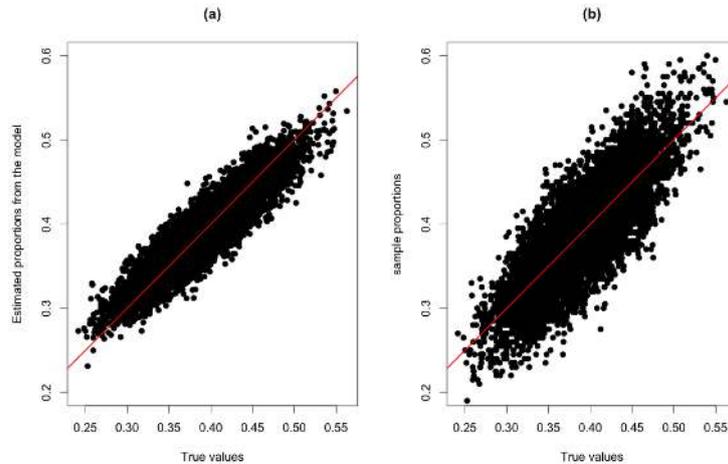


Figure 1: Scatterplots of the true proportions against the estimated proportions from the model (panel (a)) and the raw sample proportions (panel (b)). The red line is that of equality.

```
R> plot(theta.true, Y/n, pch=19, xlab="True values", ylab="sample proportions",
+       main="(b)")
R> abline(0,1, col="red", lwd=2)
```

The accuracy of individual components of the model can also be assessed graphically, and the code below plots the true and estimated values for the spatial random effects surface ϕ . Again, the true and estimated values are close, and the resulting scatterplot is shown in Figure 2. Note, as the estimated random effects have been mean centred the true values are also mean centred to aid comparison.

```
R> par(mfrow=c(1,1))
R> plot(phi - mean(phi), apply(model$samples$phi,2,median), pch=19, xlab="True values",
+       ylab="Posterior median")
R> abline(0,1,col="red")
```

Finally, the posterior distributions of the regression parameters (β_1, β_2) can be visualised by the code below, the result of which is shown in Figure 3.

```
R> colnames(model$samples$beta) <- c("beta1", "beta2")
R> plot(model$samples$beta)
```

The figures show the Markov chains have converged, and that the posterior distribution of β_2 is centred close to the true value of β_2 . Note, due to the centering of the random effects in the model fitting but not in the data generation, the posterior distribution of β_1 is not centred on the true value, as the estimates also incorporates non-zero mean levels from the 3 sets of random effects. The next two sections give two worked examples of applying the models to real data to address important questions in epidemiology and housing market research.

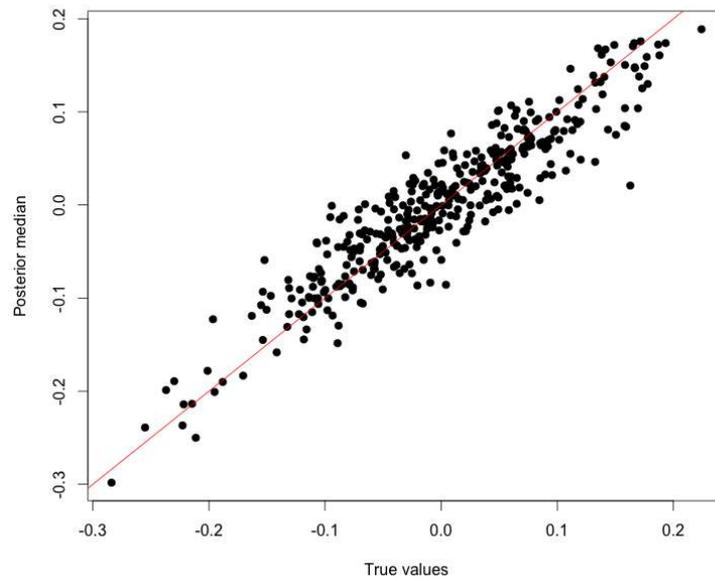


Figure 2: Scatterplot of the true spatial random effects against the estimates from the model. The red line is that of equality.

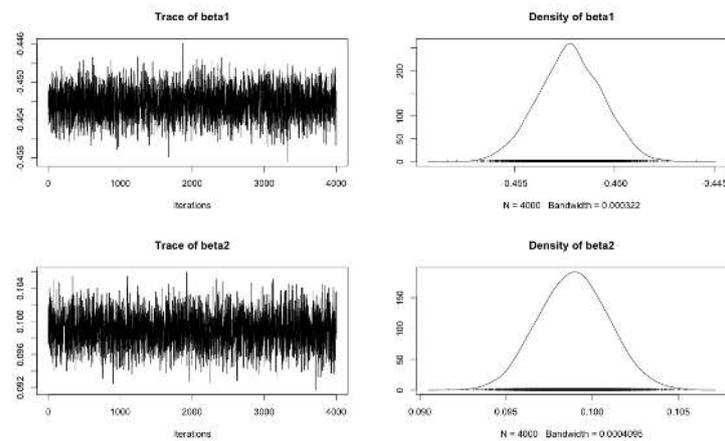


Figure 3: Posterior distributions of the regression parameters (β_0, β_1) . The left panel contains traceplots while the right panel are density estimates.

5. Example 1 - Quantifying the effect of air pollution on human health

This first example is an ecological regression problem, whose aim is to estimate the effect that air pollution concentrations have on respiratory disease risk.

5.1. Data and exploratory analysis

For the purposes of delivering health care Scotland is split into 14 health boards, and this study focuses on the Greater Glasgow and Clyde health board, which contains the city of Glasgow and has a population of around 1.2 million people during the 2007 to 2011 study period. This health board is split into $K = 271$ intermediate geographies (IG), which are a key geography for the distribution of small-area statistics in Scotland and contain populations of between 2,468 and 9,517 people. The aim of this study is to quantify the impact of particulate matter air pollution concentrations on respiratory ill health, and we have yearly data for $N = 5$ years (2007 to 2011) for the $K = 271$ IGs. The disease and covariate data are freely available from Scottish Statistics (<http://statistics.gov.scot/>), while the particulate matter pollution concentrations are available from the Department for the Environment, Food and Rural Affairs (DEFRA, <https://uk-air.defra.gov.uk/data/pcm-data>).

The respiratory disease data are population level counts of the numbers of admissions to hospital in each IG and year with a primary diagnosis of respiratory disease, which corresponds to the International Classification of Disease tenth revision (ICD-10) codes J00-J99 and R09.1. However, the observed numbers of admissions in an IG and year depends on the size and demographic structure (e.g. age and sex profile) of the population living there, which is adjusted for using indirect standardisation. This involves computing the number of admissions that would be expected in each IG and year if national age and sex specific admissions rates applied. The observed and expected numbers of respiratory hospital admissions in the k th IG and t th year are denoted by (Y_{kt}, E_{kt}) respectively, and the Poisson model, $Y_{kt} \sim \text{Poisson}(E_{kt}R_{kt})$ is typically used to model these data. Here R_{kt} is the risk, relative to E_{kt} , of disease in IG k and year t , and a value of 1.2 corresponds to a 20% increased risk of disease. Operationally, the E_{kt} is included as an offset term in the model on the natural log scale, that is $O_{kt} = \ln(E_{kt})$ in (1).

The pollution data we utilise are yearly average modelled concentrations of particulate matter less than 10 microns (PM_{10}), which come from both anthropogenic (e.g. particles in car exhaust fumes) and natural (e.g. sea salt) sources. These data are estimates on a 1 kilometre square grid produced by a numerical simulation model, and full details can be found in [RICARDO-AEA \(2015\)](#). These 1 kilometre square estimates are spatially misaligned with the irregularly shaped polygonal IGs at which the disease and covariate data are available, and thus simple averaging is used to produce IG level PM_{10} estimates. Specifically, the median value of PM_{10} over the set of 1 kilometre grid squares having centroids lying within each IG was computed, and if an IG was too small to contain a grid square centroid, the nearest grid square was used as the concentration.

Finally, the data set contains 2 potential confounders that will be included in the model, both of which are proxy measures of socio-economic deprivation (poverty). The main confounder in spatio-temporal air pollution and health studies is smoking rates, but such smoking data are unavailable. However, smoking rates are strongly linked to socio-economic deprivation, and

thus existing studies such as [Haining, Li, Maheswaran, Blangiardo, Law, Best, and Richardson \(2010\)](#) control for smoking effects using deprivation based proxy measures. Here we have two measures of socio-economic deprivation, the average property price in each IG and year (in hundreds of thousands), and the proportion of the working age population who are in receipt of job seekers allowance (JSA), the latter being a benefit paid to individuals who are unemployed and seeking employment.

These data are available in the **CARBayesdata** package in the object `pollutionhealthdata`, and the package also contains the spatial polygon information for the Greater Glasgow and Clyde health board study region in the object `GGHB.IG`. These data can be loaded using the following commands.

```
R> library(CARBayesdata)
R> library(sp)
R> data(GGHB.IG)
R> data(pollutionhealthdata)
```

The structure of `pollutionhealthdata` is shown below

```
R> head(pollutionhealthdata)
```

| | IG | year | observed | expected | pm10 | jsa | price |
|---|-----------|------|----------|-----------|----------|------|-------|
| 1 | S02000260 | 2007 | 97 | 98.24602 | 14.02699 | 2.25 | 1.150 |
| 2 | S02000261 | 2007 | 15 | 45.26085 | 13.30402 | 0.60 | 1.640 |
| 3 | S02000262 | 2007 | 49 | 92.36517 | 13.30402 | 0.95 | 1.750 |
| 4 | S02000263 | 2007 | 44 | 72.55324 | 14.00985 | 0.35 | 2.385 |
| 5 | S02000264 | 2007 | 68 | 125.41904 | 14.08074 | 0.80 | 1.645 |
| 6 | S02000265 | 2007 | 24 | 55.04868 | 14.08884 | 1.25 | 1.760 |

where the first column labelled `IG` is the set of unique identifiers for each IG. An exploratory measure of disease risk is the standardised morbidity ratio (SMR), which for the k th IG and t th year is computed as $SMR_{kt} = Y_{kt}/E_{kt}$. However, due to the natural log link function in the Poisson model, the covariates are related in the model to the natural log of the SMR. Therefore the code below adds the SMR and the natural log of the SMR to the data set and produces a `pairs()` plot showing the relationship between the variables.

```
R> pollutionhealthdata$SMR <- pollutionhealthdata$observed / pollutionhealthdata$expected
R> pollutionhealthdata$logSMR <- log(pollutionhealthdata$SMR)
R> pairs(pollutionhealthdata[,c(9, 5:7)], pch=19, cex=0.5, lower.panel=NULL,
+       labels=c("ln(SMR)", "PM10", "JSA", "Price (*100,000)"))
```

The pairs plot shown in Figure 4 shows respectively positive and negative relationships between the natural log of SMR and the two deprivation covariates `jsa` and `price`, in both cases suggesting that increasing levels of poverty are related to an increased risk of respiratory hospitalisation. There also appears to be a weak positive relationship between $\log(\text{SMR})$ and PM_{10} , while the only relationship that exists between the covariates is a negative non-linear one between `jsa` and `price`. Next it is of interest to visualise the average spatial

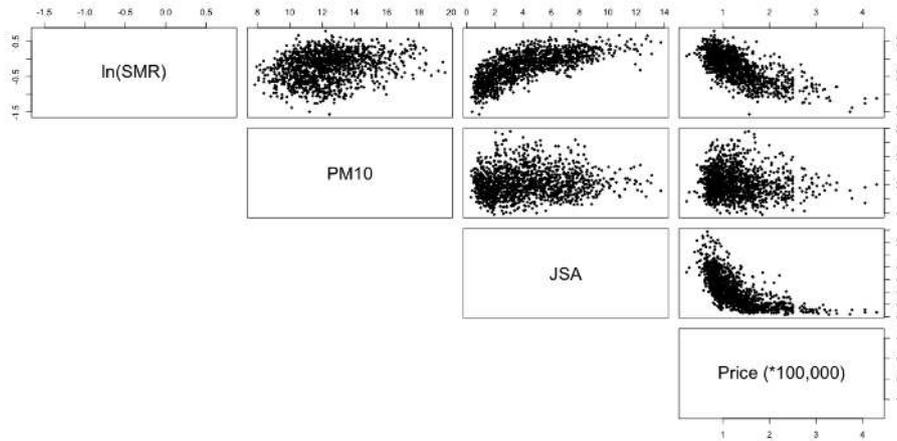


Figure 4: Scatterplot of the disease, pollution and covariate data.

pattern in the SMR over all five years, and the data can be appropriately aggregated using the `summarise()` function from the **dplyr** package using the code below. The aggregation is done by the second line, while the final line adds the aggregated averages to the `GGHB.IG` `SpatialPolygonsDataFrame` object.

```
R> library(dplyr)
R> SMR.av <- summarise(group_by(pollutionhealthdata, IG), SMR.mean=mean(SMR))
R> GGHB.IG@data$SMR <- SMR.av$SMR.mean
```

Finally, the `spplot()` function from the **sp** package can be used to draw a map of the average SMR over time using the code below:

```
R> l1 = list("SpatialPolygonsRescale", layout.north.arrow(), offset = c(220000,
+ 647000), scale = 4000)
R> l2 = list("SpatialPolygonsRescale", layout.scale.bar(), offset = c(225000,
+ 647000), scale = 10000, fill=c("transparent","black"))
R> l3 = list("sp.text", c(225000,649000), "0")
R> l4 = list("sp.text", c(230000,649000), "5000 m")
R> breakpoints <- seq(min(SMR.av$SMR.mean)-0.1, max(SMR.av$SMR.mean)+0.1, length.out=11)
R> spplot(GGHB.IG, "SMR", sp.layout=list(l1, l2, l3, l4),
+ col="transparent", xlab="Easting", ylab="Northing",
+ scales=list(draw = TRUE), at=breakpoints,
+ col.regions=terrain.colors(n=length(breakpoints)-1))
```

The first four lines in the code creates a northarrow and scalebar to add to the map, the fifth line creates the breakpoints for the colour scheme into 10 equally sized intervals, while the last line draws the map. The resulting map is shown in Figure 5, where the green shaded areas are low risk ($SMR < 1$) while the orange to silver areas exhibit elevated risks ($SMR > 1$). The map shows that the main high-risk areas are in the east-end of Glasgow in the east of the study region and the Port Glasgow area in the far west of the region on the lower bank of

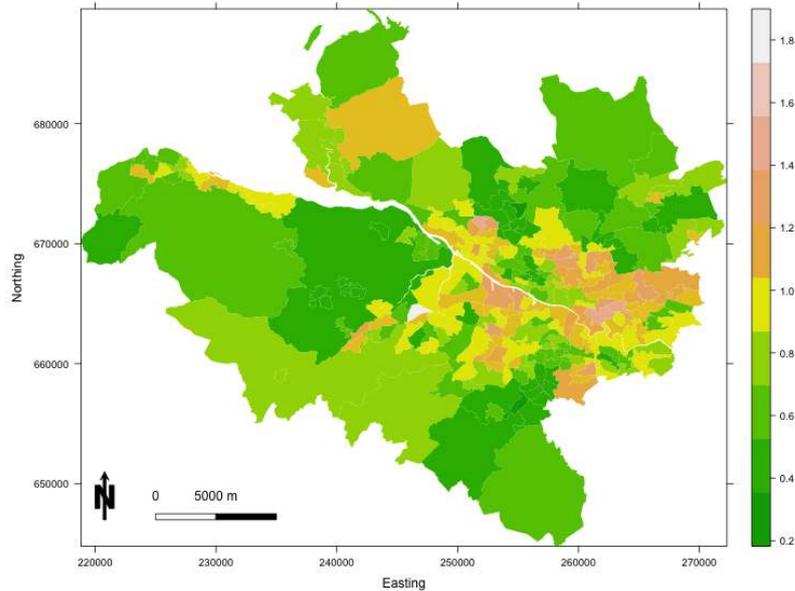


Figure 5: Map showing the average SMR over all five years from 2007 to 2011.

the river Clyde (the white line running north west to south east). The analysis that follows requires us to compute the neighbourhood matrix \mathbf{W} and a `listw` object variant of the same spatial information, the latter being used in a hypothesis test for spatial autocorrelation. Both of these quantities can be computed from the `SpatialPolygonsDataFrame` object using functionality from the `spdep` package, and code to achieve this is shown below.

```
R> library(spdep)
R> W.nb <- poly2nb(GGHB.IG, row.names = SMR.av$IG)
R> W.list <- nb2listw(W.nb, style="B")
R> W <- nb2mat(W.nb, style="B")
```

Here \mathbf{W} is a binary $K \times K$ neighbourhood matrix computed based on sharing a common border, and `W.list` is the `listw` object variant of this spatial information.

5.2. Assessing the presence of spatial autocorrelation

The spatio-temporal models in `CARBayesST` allow for spatio-temporal autocorrelation via random effects, which capture the remaining autocorrelation in the disease data after the effects of the known covariates have been accounted for. Therefore, we assess the presence of spatial autocorrelation in our data set by first computing the residuals from a simple overdispersed Poisson log-linear model that incorporates the covariate effects. This model is fitted using the code:

```
R> formula <- observed~offset(log(expected)) + jsa + price + pm10
R> model1 <- glm(formula=formula, family="quasipoisson", data=pollutionhealthdata)
R> resid.glm <- residuals(model1)
R> summary(model1)$coefficients
```

| | Estimate | Std. Error | t value | Pr(> t) |
|-------------|-------------|-------------|-----------|--------------|
| (Intercept) | -0.59752496 | 0.054333524 | -10.99735 | 5.287385e-27 |
| jsa | 0.06041994 | 0.003231475 | 18.69732 | 1.467196e-69 |
| price | -0.28293191 | 0.018292049 | -15.46748 | 8.225472e-50 |
| pm10 | 0.04174701 | 0.003282156 | 12.71938 | 4.344434e-35 |

```
R> summary(model1)$dispersion
```

```
[1] 4.399561
```

The results show significant effects of all three covariates on disease risk, as well as substantial overdispersion with respect to the Poisson equal mean and variance assumption (over dispersion parameter equal to around 4.40). To quantify the presence of spatial autocorrelation in the residuals from this model we can compute Moran's I statistic (Moran 1950) and conduct a permutation test for each year of data separately. The permutation test has the null hypothesis of no spatial autocorrelation and an alternative hypothesis of spatial autocorrelation (either positive or negative), and is conducted using the `moran.mc()` function from the `spdep` package. The test can be implemented for the first year of residuals (2007) using the code below.

```
R> moran.mc(x=resid.glm[1:271], listw = W.list, nsim = 10000)
```

```
Monte-Carlo simulation of Moran I
```

```
data: resid.glm[1:271]
weights: W.list
number of simulations + 1: 10001
```

```
statistic = 0.10358, observed rank = 9975, p-value = 0.0026
alternative hypothesis: greater
```

The estimated Moran's I statistic is 0.10358 and the p-value is less than 0.05, suggesting strong evidence of unexplained spatial autocorrelation in the residuals from 2007 after accounting for the covariate effects. Similar results were obtained for the other years and are not shown for brevity. We note that residual temporal autocorrelation could be assessed similarly for each IG, for example by computing the lag-1 autocorrelation coefficient, but with only 5 time points the resulting estimates would not be reliable. These results show that the assumption of independence is not valid for these data, and that spatio-temporal autocorrelation should be allowed for when estimating the covariate effects.

5.3. Spatio-temporal modelling with CARBayesST

We illustrate model fitting in **CARBayesST** by applying the `ST.CARar()` model to the data, details of which are defined in Section 2. This model has previously been used to account for spatio-temporal autocorrelation in an air pollution and health study, for details see [Rushworth *et al.* \(2014\)](#). The model can be fitted with the following one-line function call, and we note that all data vectors (response, offset and covariates) have to be ordered so that the first K

data points relate to all spatial units at time 1, the next K data points to all spatial units at time 2 and so on.

```
R> library(CARBayesST)
R> model2 <- ST.CARar(formula=formula, family="poisson", data=pollutionhealthdata,
+                     W=W, burnin=20000, n.sample=220000, thin=50)
```

In the above code the covariate and offset component defined by `formula` is the same as for the simple Poisson log-linear model fitted earlier, and the neighbourhood matrix is binary and defined by whether or not two areas share a common border. The `ST.CARar()` model is run for 220,000 MCMC samples, the first 20,000 of which are removed by the burn-in period. The samples are then thinned by 50 to largely remove the autocorrelation, resulting in 4,000 samples for inference. A summary of the model results can be visualised using the `print()` function developed for **CARBayesST**, which gives a very similar summary to that produced in the **CARBayes** package.

```
R> print(model2)
```

```
#####
#### Model fitted
#####
Likelihood model - Poisson (log link function)
Latent structure model - Autoregressive CAR model
Regression equation - Y ~ offset(log(E)) + jsa + price + pm10

#####
#### Results
#####
Posterior quantities for selected parameters and DIC

           Median    2.5%   97.5% n.sample % accept n.effective Geweke.diag
(Intercept) -0.6556 -0.8211 -0.4843   4000   30.1       717.9      -1.8
jsa          0.0673  0.0574  0.0770   4000   30.1       887.9       1.0
price       -0.1939 -0.2326 -0.1530   4000   30.1      1443.7       1.3
pm10        0.0334  0.0209  0.0444   4000   30.1       613.9       1.5
tau2        0.0591  0.0501  0.0694   4000  100.0      2848.5       0.6
rho.S       0.5701  0.4131  0.7380   4000   45.3      2142.6       1.0
rho.T       0.7555  0.6953  0.8130   4000  100.0      4000.0       0.9

DIC = 10385.78      p.d = 767.0432      LMPL = -4532.71
```

The Summary is presented in two parts, the first of which describes the model that has been fit. The second summarises the results, and includes the posterior median (**Median**) and 95% credible intervals (2.5%, 97.5%) for selected parameters (not the random effects), the convergence diagnostic proposed by [Geweke \(1992\)](#) (**Geweke.diag**) as a Z-score and the effective number of independent samples (**n.effective**). Also displayed are the actual number of

samples kept from the McMC run (`n.sample`) as well as the acceptance rate for each parameter (`% accept`). Note, parameters that have an acceptance rate of 100% have been Gibbs sampled due to their full conditional distributions being a standard distribution. Finally, the DIC and LMPL overall model fit criteria are displayed, which allows models with different space-time structures to be compared.

The output from the `print()` function shows that all three covariates exhibit relationships with disease risk, as none of the 95% credible intervals contain zero. Furthermore, the spatial (`rho.S`) and temporal (`rho.T`) dependence parameters exhibit relatively high values in the unit interval, suggesting that both spatial and temporal autocorrelation are present in these data even after adjusting for the covariate effects. The model object `model2` is a list, and details of its elements are described in Section 3 of this paper. The main components are vectors of fitted values (`model2$fitted.values`) and residuals (`model2$residuals`), and a list object containing the McMC samples for each individual parameter and the fitted values. These McMC samples are stored in `model2$samples`, and each element of the list corresponds to a different group of parameters. Each element is a `mcmc` object from the `coda` package. Applying the `summary()` function to this object yields:

```
R> summary(model2$samples)
```

| | Length | Class | Mode |
|--------|---------|-------|---------|
| beta | 16000 | mcmc | numeric |
| phi | 5420000 | mcmc | numeric |
| rho | 8000 | mcmc | numeric |
| tau2 | 4000 | mcmc | numeric |
| fitted | 5420000 | mcmc | numeric |
| Y | 1 | mcmc | logical |

Here the `Y` object is `NA` as there are no missing Y_{kt} observations in this data set. If there had been say m missing values, then the `Y` component of the list would have contained m columns, with each one containing posterior predictive samples for one of the missing observations. The key interest in this analysis is the effects of the covariates on disease risk, which for Poisson models are typically presented as relative risks. The relative risk for an ϵ unit increase in a covariate with regression parameter β_1 is given by the transformation $\exp(\epsilon\beta_1)$, and a relative risk of 1.02 corresponds to a 2% increased risk if the covariate increased by ϵ . The code below draws the posterior relative risk distributions for a one unit increase in each covariate, which are all realistic increases given the variation observed in the data in Figure 4.

```
R> colnames(model2$samples$beta) <- c("Intercept", "JSA", "Price", "PM10")
R> plot(exp(model2$samples$beta[, -1]))
```

These distributions are displayed in Figure 6, where the left side shows traceplots and the right side shows density estimates. All three parameters appear to have converged and have largely symmetric and unimodal posterior distributions. Posterior medians and 95% credible intervals for the relative risks can be computed using the `summarise.samples()` function from the `CARBayes` package using the code below:

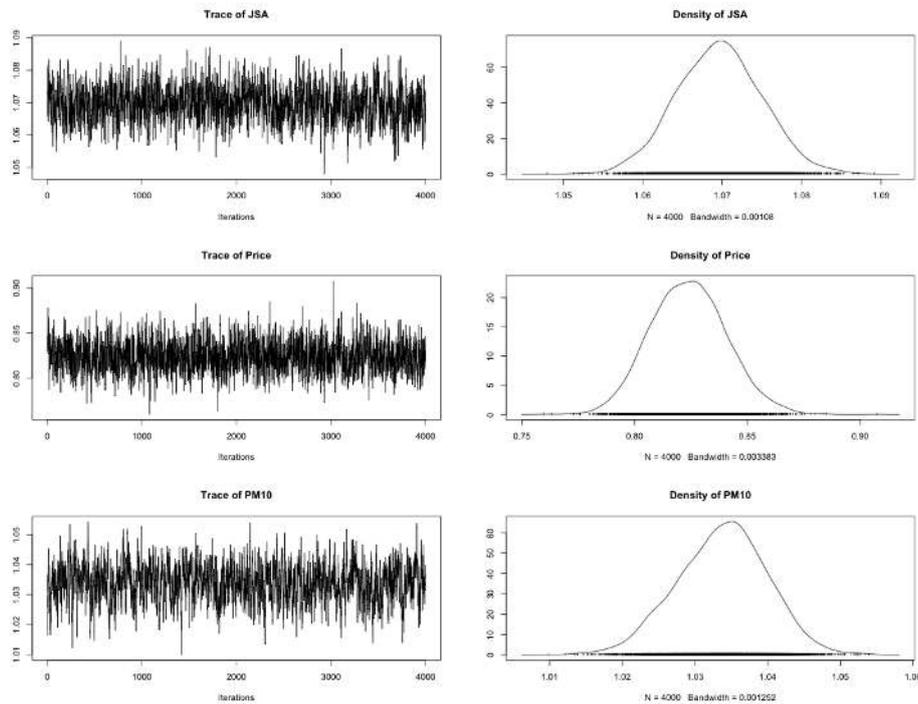


Figure 6: Posterior distributions for the covariate effects.

```
library(CARBayes)
R> parameter.summary <- summarise.samples(exp(model2$samples$beta[ , -1]),
+   quantiles=c(0.5, 0.025, 0.975))
R> round(parameter.summary$quantiles,3)
```

```
      0.5 0.025 0.975
[1,] 1.069 1.059 1.080
[2,] 0.823 0.789 0.857
[3,] 1.034 1.021 1.046
```

The output above shows that the posterior median and 95% credible interval for the relative risk of a $1\mu\text{gm}^{-3}$ increase in PM_{10} is 1.034 (1.021, 1.046), suggesting that such an increase corresponds to 3.4% additional hospital admissions. The corresponding relative risk for a one percent increase in JSA is 1.069 (1.059, 1.080), while for a one thousand pound increase in property price (the units for the property price data were in hundreds of thousands) the risk is 0.823 (0.789, 0.857). Thus, we find that air pollution is related, at this ecological level, to increased respiratory hospitalisation, while decreased socio-economic deprivation, as measured by both property price and JSA, is also related to decreased risks of hospital admission.

6. Example 2 - Monitoring the changing state of the housing market

This second example focuses on the state of the housing market, specifically property sales, and aims to quantify its changing trend over time in an era that encompasses the global

financial crisis that began in late 2007.

6.1. Data and exploratory analysis

The study region is the same as for the first example, namely the set of $K = 271$ intermediate geographies that make up the Greater Glasgow and Clyde health board. The data also come from the same source (Scottish Statistics, <http://statistics.gov.scot/>), and include yearly observations of house sales from 2003 to 2013 inclusive. The response variable is the number of property sales Y_{kt} in each IG (indexed by k) and year (indexed by t), and we additionally have the total number of properties n_{kt} in each IG and year that will be used in the model as the offset term. We use the following Poisson log-linear model for these data, $Y_{kt} \sim \text{Poisson}(n_{kt}\theta_{kt})$, where θ_{kt} is the rate of property sales as a proportion of the total number of properties. We note that we have not used a binomial model here as a single property could sell more than once in a year, meaning that each property does not constitute a Bernoulli trial. Thus θ_{kt} is not strictly the proportion of properties that sell in a year, but is on approximately the same scale for interpretation purposes.

These data are available in the **CARBayesdata** package in the object `salesdata`, as is the spatial polygon information for the Greater Glasgow and Clyde health board study region (in the object `GGHB.IG`). These data can be loaded using the following commands.

```
R> library(CARBayesdata)
R> library(sp)
R> data(GGHB.IG)
R> data(salesdata)
```

The `data.frame` `salesdata` contain 4 columns, the intermediate geography code (IG), the year the data relate to (`year`), the number of property sales (`sales`, Y_{kt}) and the total number of properties (`stock` n_{kt}). We visualise the temporal trend in these data using the code below, where the first line creates the raw rate of property sales as a proportion of the total number of properties.

```
R> salesdata$salesprop <- salesdata$sales / salesdata$stock
R> boxplot(salesdata$salesprop~salesdata$year, range=0, xlab="Year", ylab="Property sales
+       as a proportion of total properties", col="darkseagreen", border="navy")
```

This produces the boxplot shown in Figure 7, where the global financial crisis began in 2007. The plot shows a clear step-change in property sales between 2007 and 2008, as sales were increasing up to and including 2007, before markedly decreasing in subsequent years. Sales in the last year of 2013 show slight evidence of increasing relative to the previous 4 years, possibly suggesting the beginning of an upturn in the market. Also there appears to be a change in the level of spatial variation from year to year, with larger amounts of spatial variation observed before the global financial crisis. The spatial pattern in the average (over time) rate of property sales as a proportion of the total number of properties is shown in Figure 8, which was created using the code below.

```
R> library(dplyr)
R> salesprop.av <- summarise(group_by(salesdata,IG), salesprop.mean=mean(salesprop))
```

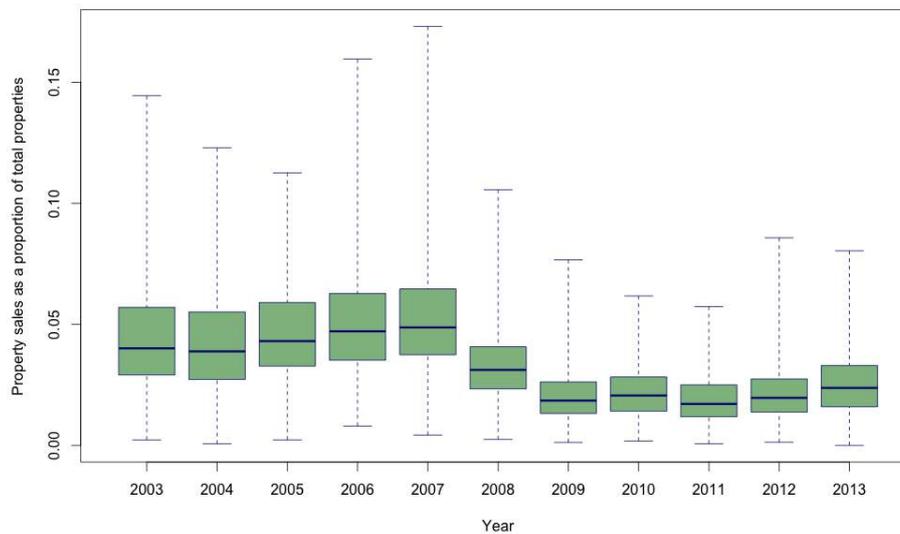


Figure 7: Boxplots showing the temporal trend in the raw rate of property sales as a proportion of the total number of properties between 2003 and 2013.

```
R> GGHB.IG@data$sales <- salesprop.av$salesprop.mean
R> l1 = list("SpatialPolygonsRescale", layout.north.arrow(), offset = c(220000,647000),
+          scale = 4000)
R> l2 = list("SpatialPolygonsRescale", layout.scale.bar(), offset = c(225000,647000),
+          scale = 10000, fill=c("transparent","black"))
R> l3 = list("sp.text", c(225000,649000), "0")
R> l4 = list("sp.text", c(230000,649000), "5000 m")
R> breakpoints <- c(0, quantile(salesprop.av$salesprop.mean, seq(0.1,0.9,0.1)), 0.1)
R> spplot(GGHB.IG, "sales", sp.layout=list(l1, l2, l3, l4),
+        col="transparent", xlab="Easting", ylab="Northing",
+        scales=list(draw = TRUE), at=breakpoints,
+        col.regions=terrain.colors(n=length(breakpoints)-1))
```

The first 3 lines of code create the IG specific temporal averages using the **dplyr** package and add it to the GGHB.IG `SpatialPolygonsDataFrame` object, both of which are very similar to code used in the first example. The remaining lines of code produce the map shown in Figure 8, again using code very similar to that in Example 1. The map and its colour key shows that the property sales data are skewed to the right, as the colour key chosen has unequal groups. Initially an equally-spaced colour scheme was created, but that showed little colour variation, hence the use of the unequal quantile based colour key. Secondly, the map shows a largely similar pattern to that seen for respiratory disease risk in Figure 5, with areas that exhibit relative high sales rates being the same as those that exhibit relatively low disease risk. Figures 7 and 8 highlight the change in temporal dynamics and the spatial structure in property sales in Glasgow, and we now apply the `ST.CARsepspatial()` model from **CARBayesST** to more formally quantify these features.

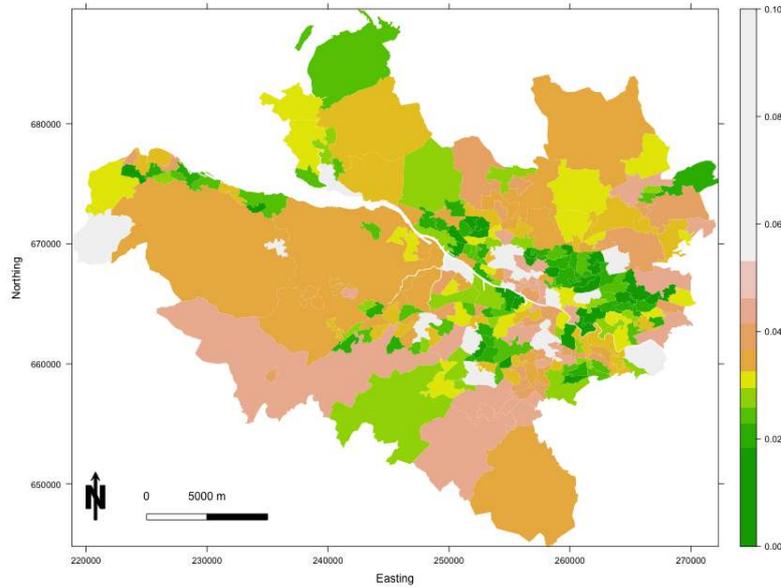


Figure 8: Map showing the average (between 2003 to 2013) raw rate of property sales as a proportion of the total number of properties.

6.2. Quantifying the changing temporal trends and spatial patterns in sales rates

The extent to which the region-wide average level of sales and its spatial variation and spatial structure changes over time can be assessed by applying the model proposed by [Napier *et al.* \(2016\)](#) to the data, which can be implemented by the `ST.CARsepspatial()` function. Before fitting this model we need to create the neighbourhood matrix using the following code:

```
R> library(spdep)
R> W.nb <- poly2nb(GGHB.IG, row.names = salesprop.av$salesprop.mean)
R> W <- nb2mat(W.nb, style="B")
```

Then the model can be fitted using the code below, where inference is again based on 4,000 post burn-in and thinned MCMC samples.

```
R> library(CARBayesST)
R> formula <- sales ~ offset(log(stock))
R> model1 <- ST.CARsepspatial(formula=formula, family="poisson", data=salesdata,
+                             W=W, burnin=20000, n.sample=220000, thin=50)
```

A summary of the model fit can be obtained using the `print()` function, and the output is similar to that shown in example 1 and is not shown for brevity. The model fitted represents the estimated rate of property sales by

$$\theta_{kt} = \exp(\beta_1 + \phi_{kt} + \delta_t),$$

which is the sum of an overall intercept term β_1 , a space-time effect ϕ_{kt} with a time period specific variance, and a region-wide temporal trend δ_t . The mean and standard deviation of $\{\theta_{kt}\}$ over space for each year is computed by the following code, which produces the posterior median and a 95% credible interval for each quantity for each year.

```
R> trend.mean <- array(NA, c(11, 3))
R> trend.sd <- array(NA, c(11, 3))
R> for(i in 1:11)
+   {
+     posterior <- exp(model1$samples$phi[ ,((i-1)*271 + 1):(i*271)] +
+                   matrix(rep(model1$samples$beta + model1$samples$delta[ ,i], 271),
+                           ncol=271, byrow=FALSE))
+     trend.mean[i, ] <- quantile(apply(posterior, 1, mean), c(0.5, 0.025, 0.975))
+     trend.sd[i, ] <- quantile(apply(posterior, 1, sd), c(0.5, 0.025, 0.975))
+   }
```

These temporal trends in the average rate of property sales and its level of spatial variation can be plotted by the following code, and the result is displayed in Figure 9.

```
R> par(mfrow=c(2,1))
R> plot(2003:2013, trend.mean[ ,1], col="red", type="l", xlab="Year", main="(a)",
+       ylab="Average property sales rate", ylim=c(0, 0.055))
R> lines(2003:2013, trend.mean[ ,2])
R> lines(2003:2013, trend.mean[ ,3])

R> plot(2003:2013, trend.sd[ ,1], col="red", type="l", xlab="Year", main="(b)",
+       ylab="Spatial standard deviation", ylim=c(0, 0.03))
R> lines(2003:2013, trend.sd[ ,2])
R> lines(2003:2013, trend.sd[ ,3])
```

The figure shows that both the region-wide average (panel (a)) and the level of spatial variation (as measured by the spatial standard deviation, panel (b)) in property sales rates show the same underlying trend, with maximum values just before the global financial crisis in 2007, and then sharp decreases afterwards. This provides some empirical evidence that the global financial crisis negatively affected the housing market in Greater Glasgow, with average sales rates dropping from just under 5.3% in 2007 to 3.3% in 2008. The spatial standard deviation also dropped from 0.024 to 0.014 over the same two-year period, suggesting that the global financial crisis had the effect of reducing the disparity in sales rates in different regions of Glasgow. We note that when measuring the spatial standard deviation we have not simply presented the posterior distribution of τ_t^2 , because this relates to the linear predictor scale and thus the results change after the inverse logit transformation to the $\{\theta_{kt}\}$ scale due to a non-constant mean level (due to the δ_t term).

The posterior median sales rate $\{\theta_{kt}\}$ is computed and then plotted for the 6 odd numbered years using the code below, where the colour scale is the same as for Figure 8. The first 4 lines create a `data.frame` object of estimated sales rates, while the fifth line a `SpatialPolygonsDataFrame` object. Finally, the last line plots the estimated sales rates for odd numbered years.

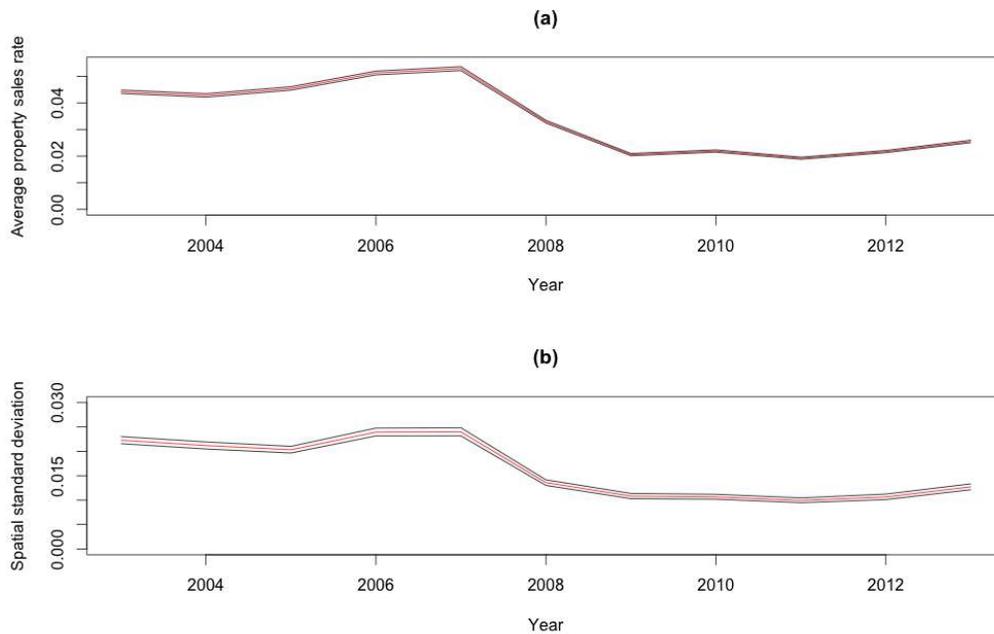


Figure 9: Posterior median (red) and 95% credible interval (black) for the temporal trend in: (a) region-wide average property sales rates; and (b) spatial standard deviation in property sales rates.

```
R> rate.est <- matrix(model1$fitted.values / salesdata$stock, nrow=nrow(W), byrow=FALSE)
R> rate.est <- as.data.frame(rate.est)
R> colnames(rate.est) <- c("r2003", "r2004", "r2005", "r2006", "r2007", "r2008",
+   "r2009", "r2010", "r2011", "r2012", "r2013")
R> GGHB.IG@data <- data.frame(GGHB.IG@data, rate.est)
R> spplot(GGHB.IG, c("r2011", "r2013", "r2007", "r2009", "r2003", "r2005"),
+   names.attr=c("Rate 2011", "Rate 2013", "Rate 2007", "Rate 2009", "Rate 2003",
+   "Rate 2005"), sp.layout=list(11, 12, 13, 14), col="transparent",
+   xlab="Easting", ylab="Northing", scales=list(draw = TRUE), at=breakpoints,
+   col.regions=terrain.colors(n=length(breakpoints-1)))
```

The maps are displayed in Figure 10, and show the clear changing spatial pattern in sales rates over time. The spatial rates for 2003 to 2007 are largely consistent, with a correlation coefficient of 0.76 and a mean absolute difference of 0.012 on approximately (as a property could sell more than once in a year) the proportion scale. However, a clear step-change is evident between 2007 and 2009, which incorporates the start of the global financial crisis, with a correlation coefficient of 0.65 and mean absolute difference of 0.032. The figure shows that the downturn in sales rates continues into 2011 (the mean absolute difference compared to 2009 is 0.0050) but that the property market is beginning an upturn by 2013. So in conclusion, Figures 9 and 10 show that the global financial crisis in 2007 resulted in a downturn in both the regionwide rate of sales and the level of spatial variation in sales across Glasgow, but that areas of high sales, such as the west-end of Glasgow (the thin strip of orange shaded areas north of the river in 2009), always remained higher than other parts of the study region.

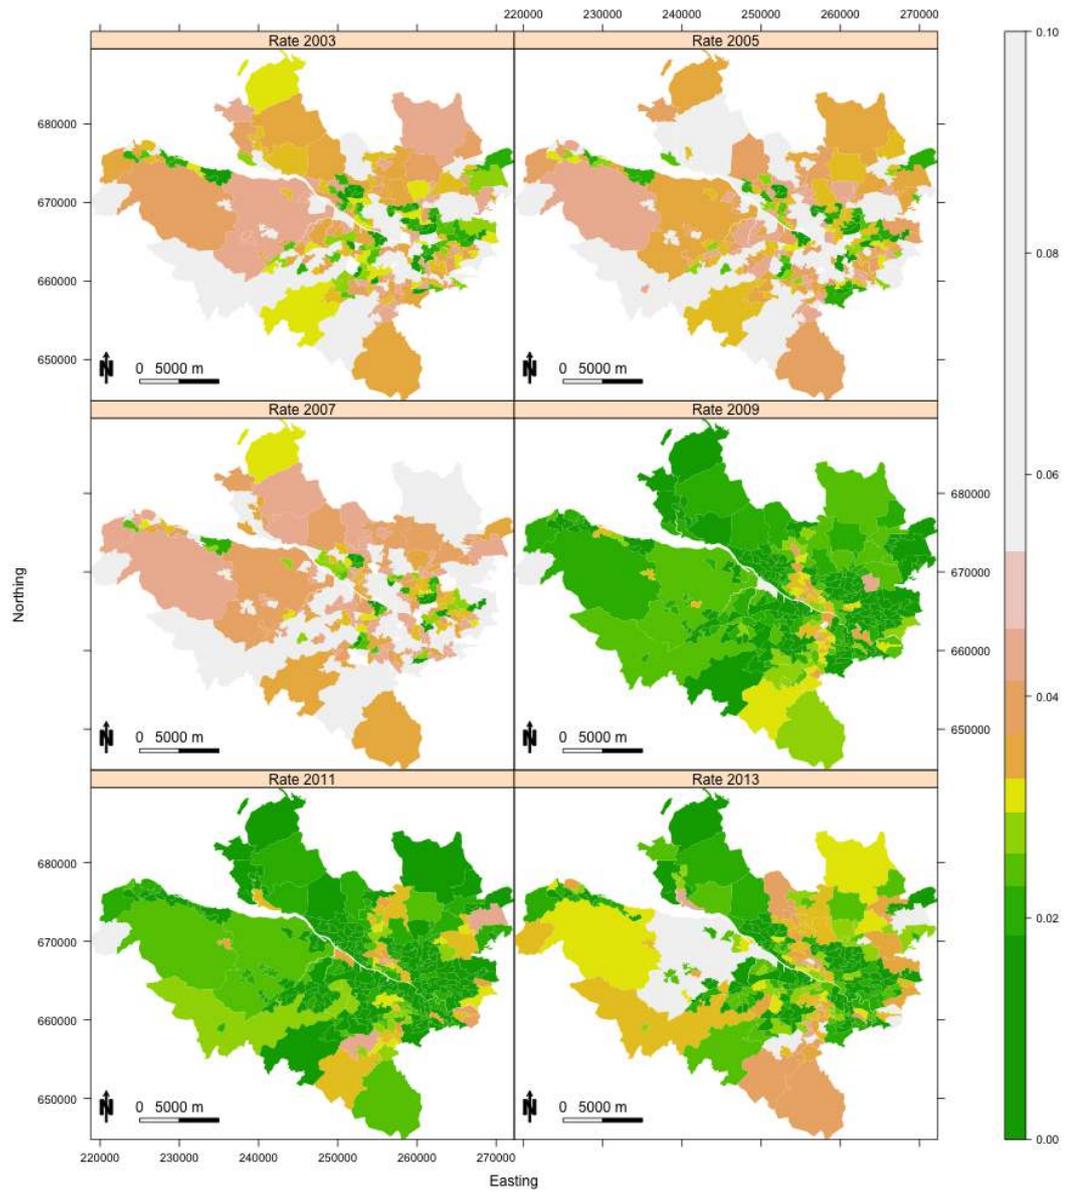


Figure 10: Maps showing the changing spatial evolution of the posterior median estimated sales rates $\{\theta_{kt}\}$ for odd numbered years.

7. Discussion

This paper has presented the **CARBayesST** software, which is the first software package dedicated to fitting spatio-temporal CAR type models to areal unit data. This paper has outlined the class of models that can be implemented by the software together with the Bayesian inferential framework, as well as the range of specific models that can be fitted. The key advantage of this package, compared to say implementing the models in **BUGS**, is its ease of use, which includes fitting models with a one-line function call, specifying the spatial adjacency information via a single matrix, and the ability to fit multiple models addressing different questions about the data in a common software environment. The paper has presented a small simulation example to illustrate the correctness of the algorithms, as well as two fully worked examples based on quantifying the effect of air pollution on human health and the changing nature of the housing market.

Future development of the software will be in two main directions. First, as the literature on spatio-temporal modelling advances we aim to increase the number of spatio-temporal models that can be implemented, giving the user an even wider set of modelling tools. Second, with the rapid increase in the availability of small-area data, we aim to develop a suite of multivariate space-time models (MVST). The development of MVST methodology for areal unit data is in its infancy, and the ability to jointly examine the spatio-temporal patterns in multiple response variables simultaneously allows one to address questions that cannot be addressed by single variable models. For example, in a public health context it allows one to estimate overall and disease specific spatio-temporal patterns in disease risk, allowing one to see which areas repeatedly signal at high risk for all diseases and which exhibit elevated risks for only one disease. In the housing context of example 2 a MVST approach would allow one to extend the analysis carried out by different property types, e.g. flats, terraced houses, etc, which would allow more insight to be gained about the exact nature of the housing market.

8. Acknowledgements

The Development of the **CARBayesST** software and this paper were supported by the UK Engineering and Physical Sciences Research Council (EPSRC, grant EP/J017442/1) and the UK Medical Research Council (MRC, grant MR/L022184/1). The data and shapefiles used in this paper were provided by the Scottish Government via their Scottish Statistics website (<http://statistics.gov.scot/>).

References

- Anderson C, Lee D, Dean N (2014). “Identifying clusters in Bayesian disease mapping.” *Biostatistics*, **15**, 457–469.
- Bakar K, Sahu S (2015). “spTimer: Spatio-Temporal Bayesian Modeling Using R.” *Journal of Statistical Software*, **63**(1). URL <https://www.jstatsoft.org/index.php/jss/article/view/v063i15>.
- Bates D, Mächler M, Bolker B, Walker S (2015). “Fitting Linear Mixed-Effects Models Using

- lme4.” *Journal of Statistical Software*, **67**(1). URL <https://www.jstatsoft.org/index.php/jss/article/view/v067i01>.
- 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.
- Bivand R (2013). “spdep: Spatial Dependence: Weighting Schemes, Statistics and Models. R package version 0.5-56, <http://CRAN.R-project.org/package=spdep>.”
- Bivand R, Lewin-Koh N (2015). *maptools: Tools for Reading and Handling Spatial Objects*. R package version 0.8-36, URL <https://CRAN.R-project.org/package=maptools>.
- Bivand R, Pebesma E, Gomez-Rubio V (2013). *Applied spatial data analysis with R*. 2nd edition. Springer, NY.
- 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.
- Croissant Y, Millo G (2008). “Panel Data Econometrics in R: The plm Package.” *Journal of Statistical Software*, **27**(1). URL <https://www.jstatsoft.org/index.php/jss/article/view/v027i02>.
- Eddelbuettel D, Francois R (2011). “Rcpp: Seamless R and C++ Integration.” *Journal of Statistical Software*, **40**, 8.
- Finley A, Banerjee S, Gelfand A (2015). “spBayes for Large Univariate and Multivariate Point-Referenced Spatio-Temporal Data Models.” *Journal of Statistical Software*, **63**(1). URL <https://www.jstatsoft.org/index.php/jss/article/view/v063i13>.
- Furrer R, Sain SR (2010). “spam: A Sparse Matrix R Package with Emphasis on MCMC Methods for Gaussian Markov Random Fields.” *Journal of Statistical Software*, **36**(10), 1–25. URL <http://www.jstatsoft.org/v36/i10/>.
- 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.
- Haining R, Li G, Maheswaran R, Blangiardo M, Law J, Best N, Richardson S (2010). “Inference from ecological models: estimating the relative risk of stroke from air pollution exposure using small area data.” *Spatial Spatio-temporal Epidemiol*, **1**, 123–131, DOI:10.1016/j.sste.2010.03.006.

- 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 (2016). *CARBayesdata: Data Used in the Vignettes Accompanying the CARBayes and CARBayesST Packages*. R package version 2.0.
- Lee D, Ferguson C, Mitchell R (2009). “Air Pollution and Health in Scotland: A Multicity Study.” *Biostatistics*, **10**, 409–423.
- Lee D, Lawson A (2016). “Quantifying the spatial inequality and temporal trends in maternal smoking rates in Glasgow.” *Annals of Applied Statistics*, -, to appear.
- Lee D, Minton J, Pryce G (2015). “Bayesian inference for the dissimilarity index in the presence of spatial autocorrelation.” *Spatial Statistics*, **11**, 81–95.
- Leroux BG, Lei X, Breslow N (2000). *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, chapter Estimation of Disease Rates in Small Areas: A new Mixed Model for Spatial Dependence, pp. 179–191. Springer New York, New York, NY. ISBN 978-1-4612-1284-3. doi:10.1007/978-1-4612-1284-3_4. URL http://dx.doi.org/10.1007/978-1-4612-1284-3_4.
- Li G, Best N, Hansell A, Ahmed I, Richardson S (2012). “BaySTDetect: detecting unusual temporal patterns in small area data via Bayesian model choice.” *Biostatistics*, **13**, 695–710.
- Lunn D, Spiegelhalter D, Thomas A, Best N (2009). “The BUGS Project: Evolution, Critique and Future Directions .” *Statistics in Medicine*, **28**, 3049–3082.
- Millo G, Piras G (2012). “splm: Spatial Panel Data Models in R.” *Journal of Statistical Software*, **47**(1), 1–38. URL <http://www.jstatsoft.org/v47/i01/>.
- Moran P (1950). “Notes on continuous stochastic phenomena.” *Biometrika*, **37**, 17–23, DOI:10.1093/biomet/37.1-2.17.
- Napier G, Lee D, Robertson C, Lawson A, Pollock K (2016). “A model to estimate the impact of changes in MMR vaccination uptake on inequalities in measles susceptibility in Scotland.” *Statistical Methods in Medical Research*, p. to appear.
- Novomestky F (2012). *matrixcalc: Collection of functions for matrix calculations*. R package version 1.0-3, URL <https://CRAN.R-project.org/package=matrixcalc>.
- Novomestky F, Nadarajah S (2012). *truncdist: Truncated Random Variables*. R package version 1.0-1, URL <https://CRAN.R-project.org/package=truncdist>.
- Paul M, Meyer S (2016). “hhh4: An endemic-epidemic modelling framework for infectious disease counts, <http://CRAN.R-project.org/package=surveillance>.”
- Pinheiro J, Bates D, DebRoy S, Sarkar D, R Core Team (2015). *nlme: Linear and Nonlinear Mixed Effects Models*. R package version 3.1-121, URL <http://CRAN.R-project.org/package=nlme>.

- Plummer M, Best N, Cowles K, Vines K (2006). “CODA: Convergence Diagnosis and Output Analysis for MCMC.” *R News*, **6**(1), 7–11. URL <http://CRAN.R-project.org/doc/Rnews/>.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.
- RICARDO-AEA (2015). “Technical report on UK supplementary assessment under the Air Quality Directive (2008/50/EC), the Air Quality Framework Directive (96/62/EC) and Fourth Daughter Directive (2004/107/EC) for 2013.” <https://uk-air.defra.gov.uk/data/pcm-data>.
- 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 (2014). “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 (2016). “An adaptive spatio-temporal smoothing model for estimating trends and step changes in disease risk.” *Journal of the Royal Statistical Society Series C*, -, DOI: 10.1111/rssc.12155.
- Schabenberger H (2009). *spatcounts: Spatial count regression*. R package version 1.1, URL <http://CRAN.R-project.org/package=spatcounts>.
- 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.
- Venables WN, Ripley BD (2002). *Modern Applied Statistics with S*. Fourth edition. Springer, New York. ISBN 0-387-95457-0, URL <http://www.stats.ox.ac.uk/pub/MASS4>.
- 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.
- Watanabe S (2010). “Asymptotic equivalence of the Bayes cross validation and widely applicable information criterion in singular learning theory.” *Journal of Machine Learning Research*, **11**, 3571–3594.
- Wickham H, Francois R (2015). *dplyr: A Grammar of Data Manipulation*. R package version 0.4.3, URL <https://CRAN.R-project.org/package=dplyr>.

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
URL: <https://www.strath.ac.uk/staff/rushworthalastairmr/>

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