

Structure Induced by a Multiple Membership Transformation on the Conditional Autoregressive Model

Marco Gramatica^{*}, Silvia Liverani^{*,†}, and Peter Congdon[‡]

Abstract. The objective of disease mapping is to model data aggregated at the areal level. In some contexts, however, (e.g. residential histories, general practitioner catchment areas) when data is arising from a variety of sources, not necessarily at the same spatial scale, it is possible to specify spatial random effects, or covariate effects, at the areal level, by using a multiple membership principle (MM) (Petrof et al., 2020; Gramatica et al., 2021). A weighted average of conditional autoregressive (CAR) spatial random effects embeds spatial information for a *spatially-misaligned* outcome and estimate relative risk for both frameworks (areas and memberships). In this paper we investigate the theoretical underpinnings of these application of the multiple membership principle to the CAR prior, in particular with regard to parameterisation, properness and identifiability. We carry out simulations involving different numbers of memberships as compared to number of areas and assess impact of this on estimating parameters of interest. Both analytical and simulation study results show under which conditions parameters of interest are identifiable, so that we can offer actionable recommendations to researchers. Finally, we present the results of an application of the multiple membership model to diabetes prevalence data in South London, together with strategic implications for public health considerations.

Keywords: Bayesian hierarchical modelling, conditional autoregressive prior, multiple membership, spatial misalignment.

1 Introduction

The scope for geo-referenced data analysis is vast, though particular methodological paradigms are more common. In the context of disease mapping, the information available is typically aggregated at areal level, thus constraining the range of viable modelling strategies. Bayesian hierarchical models have long been a widely popular choice for areal disease mapping (Lawson, 2018; Congdon, 2019; Blangiardo et al., 2015). Our interest focuses on the conditional autoregressive prior (CAR) (Besag, 1974) which has been employed extensively in the disease mapping literature to account for spatial autocorrelation which is a common occurrence in areal data (MacNab, 2010).

The most common framework in this context arises from the observations of an outcome of interest $\mathbf{y} = (y_1, \dots, y_n)$ and some explanatory covariates, both observed

^{*}School of Mathematical Sciences, Queen Mary University of London, London, UK, m.gramatica@qmul.ac.uk; s.liverani@qmul.ac.uk

[†]The Alan Turing Institute, The British Library, London, UK

[‡]School of Geography, Queen Mary University of London, London, UK, p.congdon@qmul.ac.uk

on the same set of n areas that partition a certain region of interest, modelled with a distribution that includes spatial random effects such as the CAR prior. In the present work we will investigate a different type of setup, where covariates are observed at the areal level but outcomes refer to a different spatial framework.

An instance of this non-standard setup can be found in Gramatica et al. (2021), where they investigate general practitioner (GP) level data of diabetes prevalence jointly modelled with areal level mortality data for the same condition. In this instance covariates of interest are available by area (i.e. census tracts) but not for practices. However, the patients of each practice do not necessarily reside in a unique set of areas. This gives rise to two misalignment problems: first, the population of each practice can be attributed to several non-contiguous areas and second, residents of a certain area can be registered at different practices. The former type of misalignment is more related to the (non) spatial structure of the data, given that spatial modelling strategies usually require continuity. On the other hand, the second misalignment problem concerns the *imperfect nesting* of areas in GP practices.

Another example of the same ilk can be found in Petrof et al. (2020), where the impact of residential history of people affected by a rare type of cancer caused by exposure to asbestos (mesothelioma), is investigated. For each patient the probability of contracting mesothelioma is modelled using age at death as a covariate together with a spatial random effect. Clearly, given that each individual might have resided in several different locations throughout the course of their life, their residential history has to be taken into account to correctly estimate the spatial risk. Hence, the random effect is specified in terms of a weighted average of the random effects of the areas where the individual has resided. The weights correspond to the share of time spent by the individual living in a certain area, hence the multiple membership of each area to different individuals' residential history.

Both of the examples above can be framed as applying the multiple membership principle (MM) (Browne et al., 2001) transformation to an areal framework. Hence, we will denote this new prior as CAR-MM. Given that this model is a weighted average of areal random effects, it can be seen as a linear transformation of a CAR prior. It is well known that any covariance matrix of a normal distribution can be represented as a CAR prior (Ver Hoef et al., 2018), as well as that the normal distribution is closed under affine transformations (Rao, 1973). Consequently, it is possible to use tools from Gaussian distribution theory to investigate the properties and assumptions underlying the CAR-MM prior, such as its identifiability.

Despite both Petrof et al. (2020) and Gramatica et al. (2021) including simulation studies to verify the performance of the model, they did not investigate the identifiability and the structure of the prior. In this paper, we develop the underpinning theory for this methodology, and provide a more general framework able to encompass a wider range of applications. In the case of Gramatica et al. (2021) as well as this paper, the weights are obtained by computing the share of population of a certain GP that resides in a specific area. This approach can take into account the fact that a single area can contribute to more than one membership (where each GP practice constitutes a membership) as well

that each membership is usually defined in terms of more than one area, thus addressing the misalignment issues mentioned above.

Instead of a *spatial misalignment* problem, defined as “different spatial data layers are collected at different scales” (Banerjee et al., 2014, p. 165), one could consider it as a *change of support* (COSP) problem. For instance, Bradley et al. (2016) start from the subregion $D \subset \mathbb{R}^d$, the domain of interest, and partition it according to $L \geq 1$ potentially non-nesting areal frameworks called *source supports* i.e. the spatial support for which data is available, as opposed to *target supports* where inference should take place. Due to the uniqueness requirement, this approach is not be feasible given the structure of the data that we work with. In fact, it would require the region of interest to be partitioned into smaller areas, each of which would have to be assigned uniquely to a specific GP. If one were to follow this approach, they would fail to take into account the actual practice population structure. Obviously, with individual level data this problem could be easily overcome, however these are rarely available in disease mapping applications.

The paper is organised as follows: in Section 2 we introduce the CAR model and its variant, the ICAR. In Section 3 we review the CAR-MM model, its definition and distributional form, we compare the two proposals by Petrof et al. (2020) and Gramatica et al. (2021) and introduce a more compact notation in matrix form. The main contribution of this paper is in Section 4 which (i) introduces a more general framework for the MM models and discusses results of both uniqueness of the covariance matrix and CAR specification for the CAR-MM prior, (ii) presents two different parameterisations for the CAR-MM model and (iii) uses these results to investigate the issue of identifiability in a Bayesian framework. Section 5 is mostly concerned with introducing the methodologies used for the simulation study that will follow and the model fit and model comparison methods that will be used in the data analysis. In Section 6 simulation results are introduced in reference to the considerations of Section 4. Lastly, in Section 7 we consider an example on a real South London diabetes prevalence dataset providing an application of the CAR-MM prior. Three approaches are compared: one where no spatial prior accounting for spatial autocorrelation is added, and two others employing the CAR-MM and ICAR-MM priors.

2 Generalised Linear Models (GLM) and CAR priors

When modelling area referenced data, the region of interest is generally divided into n areas for which the outcome of interest is observed. For each area $i = 1, \dots, n$ a random variable Y_i is used to model the observed data, whose mean is generally specified as follows:

$$g(E(Y_i|\boldsymbol{\theta}, \mathbf{X})) = \eta_i(\boldsymbol{\theta}_i, \mathbf{X}) \quad \forall i = 1, \dots, n, \quad (2.1)$$

where η_i is a linear function of the parameters $\boldsymbol{\theta}_i$ and the p covariates \mathbf{X} , and g is a continuous monotonic function called the *link-function* (McCullagh et al., 1983). The considerations in the remainder of this paper are valid for all models such as in (2.1) however, given that for areal data of prevalence or mortality counts Poisson distribution

is the most commonly assumed, for ease of presentation, we will work under the following framework:

$$\begin{aligned} Y_i &\sim \text{Poisson}(E_i \rho_i) \\ \log(E_i \rho_i) &= \log E_i + \gamma + \mathbf{x}_i^T \boldsymbol{\beta} + \phi_i, \end{aligned} \quad (2.2)$$

where the link function in (2.1) is the natural logarithm, E_i is the expected number of cases for area i and the relative risk ρ_i is modelled as a linear function of the parameters $\boldsymbol{\theta}_i = (\gamma, \boldsymbol{\beta}, \phi_i)$. In turn, the random effects ϕ_i are modelled using a CAR prior which is a random vector $\boldsymbol{\phi} = (\phi_1, \dots, \phi_n)^T \in \mathbb{R}^n$ defined on an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$.

More specifically, using matrices \mathbf{M} and \mathbf{C} that are specified below in terms of the \mathcal{G} , we define the prior as:

$$\boldsymbol{\phi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma} = (\mathbf{I} - \mathbf{C})^{-1} \mathbf{M}). \quad (2.3)$$

Ver Hoef et al. (2018) state four conditions to ensure the positive definiteness of the covariance matrix $\boldsymbol{\Sigma}$

- C1:** $(\mathbf{I} - \mathbf{C})$ has positive eigenvalues
- C2:** \mathbf{M} is diagonal with positive diagonal elements
- C3:** $(\mathbf{C})_{ii} = 0$ for all $i = 1, \dots, n$
- C4:** $(\mathbf{C})_{ij}/(\mathbf{M})_{ii} = (\mathbf{C})_{ji}/(\mathbf{M})_{jj}$ for all $i, j = 1, \dots, n$

In practice, a very common specification that fulfils **C1–C4**, imposes $\mathbf{C} = \alpha \mathbf{D}^{-1} \mathbf{W}$, with $\alpha \in [0, 1)$ and \mathbf{W} the adjacency matrix such that $(\mathbf{W})_{ij} = 1$ iff i and j are neighbours and 0 otherwise, and $\mathbf{M} = \tau^{-1} \mathbf{D}^{-1}$ with $\tau \in \mathbb{R}^+$ and $(\mathbf{D})_{ii} = \sum_{j=1}^n (\mathbf{W})_{ij}$ for all i thus indicating the number of neighbours of area i . Finally, we can write the CAR prior as:

$$\boldsymbol{\phi} \sim \text{CAR}(\alpha, \tau, \mathbf{W}) \quad , \quad \boldsymbol{\Sigma}^{-1} = \mathbf{Q} = \tau(\mathbf{D} - \alpha \mathbf{W}). \quad (2.4)$$

Ver Hoef et al. (2018) note that the representation of the CAR covariance matrix in (2.3) is rather general, and in fact they prove it can represent any positive-definite covariance matrix:

Theorem 2.1. *Any positive-definite covariance matrix $\boldsymbol{\Sigma}$ can be expressed as the covariance matrix of a CAR model $(\mathbf{I} - \mathbf{C})^{-1} \mathbf{M}$, for a unique pair of matrices \mathbf{C} and \mathbf{M} .*

For a proof of this result see Ver Hoef et al. (2018)

A special case of the CAR prior in (2.4), arises by setting the parameter $\alpha = 1$. This prior is called *intrinsic conditional autoregressive* (ICAR) (Besag et al., 1991) and is widely used in practical applications. We will employ this prior for the data analysis in Section 7. Its precision matrix \mathbf{Q} is singular consequently leading to an improper density. In our implementation of the ICAR prior we impose the constraint $\sum_i \phi_i = 0$ at every MCMC iteration, as is commonly done in the literature (Eberly et al., 2000) to improve convergence and avoid intercept identifiability issues (Goicoa et al., 2018).

3 Multiple membership and the CAR-MM prior

Both Petrof et al. (2020) and Gramatica et al. (2021) proposed a GLM model with a CAR spatial prior and MM to deal with misaligned data. We will now lay down the model that they developed independently.

Let us consider a region partitioned in n areas and define a CAR prior on it as per (2.4) and let $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_m)$ be an outcome observed on m units called *memberships*. Typically no areal spatial framework would be available for the memberships, as the main goal of this modelling approach is to use the n areas as a latent spatial structure.

For each membership j we define a set of n weights $0 \leq h_{ji} \leq 1$, corresponding to each area, such that $\sum_{i=1}^n h_{ji} = 1$ for all $j = 1, \dots, m$. For example, the h_{ji} might be the proportions of the population of GP practice j living in areas i . We can thus re-define the relative risk term of the GLM in (2.2) for outcome \tilde{Y} so that it incorporates the latent areal structure we defined earlier:

$$\begin{aligned} \tilde{Y}_j &\sim \text{Poisson}(\tilde{E}_j \tilde{\rho}_j), \\ \log(\tilde{\rho}_j) &= \sum_{i=1}^n h_{ji} \log(\rho_i) = \sum_{i=1}^n h_{ji} (\gamma + \mathbf{X}_i^T \boldsymbol{\beta} + \phi_i), \end{aligned} \tag{3.1}$$

where \tilde{E} indicate membership level offsets. This model defines each membership level relative risk $\tilde{\rho}_j$ as a weighted average of the areal log relative risks $\log(\boldsymbol{\rho})$. In this case each area can contribute to multiple memberships with different weights.

We can more compactly represent the log relative risks using matrix notation, for a $m \times n$ matrix \mathbf{H} we define it element-wise using the weights in (3.1) so that $h_{ji} = (\mathbf{H})_{ji}$. Additionally, we assume \mathbf{H} to be of full column rank and that for each area i there is at least one membership j for which $h_{ji} > 0$, i.e. all areas contribute to at least one membership.

The second line in (3.1) can then be rewritten as:

$$\log(\tilde{\boldsymbol{\rho}}) = \mathbf{H} \log(\boldsymbol{\rho}) = \gamma \mathbf{H} \mathbf{1}_m + \mathbf{H} \mathbf{X} \boldsymbol{\beta} + \mathbf{H} \boldsymbol{\phi} = \gamma \mathbf{1}_m + \tilde{\mathbf{X}} \boldsymbol{\beta} + \tilde{\boldsymbol{\phi}}, \tag{3.2}$$

in line with Gramatica et al. (2021) memberships corresponded to General Practitioner (GP) practices, on which an outcome of interest was observed, and areas to English census areal subdivisions (Middle Super-Output Area; MSOA) where the covariates of interest were recorded. The MM matrix \mathbf{H} was obtained by computing the share of patients registered at a GP practice that resided in a specific MSOA. Petrof et al. (2020) instead defined individual patients as memberships and used membership level covariates, while the MM weights corresponded to the share of time that patient j resided in a particular area i . Nevertheless, both approaches can be represented with (3.2), by simply observing that the intercept term γ is unchanged by the linear transformation \mathbf{H} and so are coefficients $\boldsymbol{\beta}$.

4 Relationship between CAR and CAR-MM models

We will now focus on the CAR-MM prior and study its properties. Firstly, we can observe that since the CAR-MM is a linear transformation of normally distributed terms ϕ , $\tilde{\phi} = \mathbf{H}\phi$, it is also normally distributed (Rao, 1973). In fact:

$$\tilde{\phi} \sim \mathcal{N}_m(\mathbf{0}, \tilde{\Sigma} = \mathbf{H}\Sigma\mathbf{H}^T), \quad (4.1)$$

where Σ is as in (2.4). Additionally, in the case where $m \leq n$, since we assumed full column rank for \mathbf{H} , the resulting covariance matrix $\tilde{\Sigma}$, with the corresponding precision matrix $\tilde{\mathbf{Q}}$, will also be full rank (see Corollary 8.3.3 in Harville (2008)). Instead, when $m > n$, $\tilde{\Sigma}$ will be of rank n and thus singular.

In the remainder of this section we explore the corresponding uniqueness properties of the CAR-MM. As mentioned in Theorem 2.1, Section 2, Ver Hoef et al. (2018) show that any zero-mean Gaussian distribution on a finite set of points, $\mathbf{Z} \sim N(0, \Sigma)$, can be written with a covariance matrix parameterised as a CAR model.

Firstly, let us establish three conditions that will be common to the results presented below:

M1: Let Σ be an $n \times n$ covariance matrix as defined in (2.4)

M2: Let \mathbf{H} be a $m \times n$ matrix such that:

- (a) for each row $j = 1, \dots, m$ of \mathbf{H} : $\sum_{i=1}^n (H)_{ji} = 1$, $0 \leq (H)_{ji} \leq 1$ for all i and j
- (b) it is of full rank
- (c) there is no column i such that $(\mathbf{H})_{.i} = \mathbf{0}$

M3: Let $\tilde{\Sigma}$ be the covariance matrix of the corresponding CAR-MM model such that $\tilde{\Sigma} = \mathbf{H}\Sigma\mathbf{H}^T$

where $(\mathbf{H})_{.i} = \mathbf{0}$ indicates the i -th column of matrix \mathbf{H} .

We remark that, although our approach is mainly focused on the CAR prior, the following considerations directly apply to any Gaussian model. Therefore, the spatial random effects in (3.1) could also be modelled through Gaussian distributions with arbitrary covariance structure, for instance one could use a Matérn covariance function using the Euclidean distances between areas centroids (Martinez-Beneito et al., 2019, Example 6.1).

Above, following Theorem 2.1, we hinted at the fact that the CAR-MM can be seen as a CAR prior. We now introduce a corollary (proved in the Supplemental material (Gramatica et al., 2023)) to that theorem, that draws a parallel between the two matrices \mathbf{M} and \mathbf{C} that determine the CAR prior structure (C1–C4 in Section 2), and the CAR-MM.

Corollary 1. *Let Σ , \mathbf{H} , $\tilde{\Sigma}$ be as defined in M1–M3, then iff $m \leq n$, $\tilde{\Sigma}$ can be expressed as the covariance matrix of a CAR model $(\mathbf{I} - \tilde{\mathbf{C}})^{-1}\tilde{\mathbf{M}}$ for a unique pair of $m \times m$ matrices $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{M}}$.*

Since \mathbf{C} contains the off-diagonal non-zero elements of the CAR precision matrix, it determines the partial correlation structure of the prior, i.e. its zero entries correspond to partially uncorrelated components (Rue et al., 2005, Theorem 2.2). Therefore, the same could be said for $\tilde{\mathbf{C}}$. However, this matrix is not straightforward to obtain. In fact, the MM transformation \mathbf{H} is applied to the covariance matrix Σ so that to obtain the precision matrix $\tilde{\mathbf{Q}}$ we could assume $m = n$, **M1–M3** and, starting from (2.4) we could obtain:

$$\tilde{\mathbf{Q}} = (\mathbf{H}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{M}\mathbf{H}^T)^{-1} = (\mathbf{H}^T)^{-1}\mathbf{M}^{-1}(\mathbf{I} - (\mathbf{C}\mathbf{H}^{-1} - \mathbf{H}^{-1} + \mathbf{I})). \quad (4.2)$$

Given that the resulting matrices $\tilde{\mathbf{M}}$ and $\tilde{\mathbf{C}}$ would have to fulfil **C1–C4** (Section 2) it is not immediately clear how they would look in terms of \mathbf{M} , \mathbf{C} and \mathbf{H} . In fact, if we took for example $\tilde{\mathbf{M}}^{-1} = (\mathbf{H}^T)^{-1}\mathbf{M}^{-1}$ and $\tilde{\mathbf{C}} = \mathbf{C}\mathbf{H}^{-1} - \mathbf{H}^{-1} + \mathbf{I}$ we would not have any guarantee that $\tilde{\mathbf{M}}^{-1}$ would be diagonal as per **C2**, or $\tilde{\mathbf{C}}$ fulfils **C3** (null diagonal elements) or **C4** ($(\tilde{\mathbf{C}})_{ij}/(\tilde{\mathbf{M}})_{ii} = (\tilde{\mathbf{C}})_{ji}/(\tilde{\mathbf{M}})_{jj}$ for all $i, j = 1, \dots, n$). Therefore, differently from the spatial CAR model, we have no way of obtaining analytically the partial correlation structure of the CAR-MM prior.

Despite these difficulties regarding the precision matrix, we can consider the covariance matrix, together with its transformation \mathbf{H} , and prove a crucial result for model identifiability. The following lemma will help evaluate under which circumstances the MM transformation of the CAR prior is injective, and will also be used in the subsequent proofs.

Lemma 1. *Let \mathbf{H} be a matrix as defined in **M1–M3**. \mathbf{H} has a left inverse \mathbf{L} iff $m \geq n$, and this matrix is unique iff $m = n$.*

It is worth mentioning that any left inverse \mathbf{L} of \mathbf{H} will satisfy the same constraint, i.e the columns adding up to the $\mathbf{1}$ vector, in fact: $\mathbf{H}\mathbf{1}_n = \mathbf{1}_m \Leftrightarrow \mathbf{L}\mathbf{H}\mathbf{1}_n = \mathbf{L}\mathbf{1}_m \Leftrightarrow \mathbf{1}_n = \mathbf{L}\mathbf{1}_m$ (proof in Section 1 of the Supplemental material.).

Theorem 4.1. *Let Σ , \mathbf{H} , $\tilde{\Sigma}$ be as defined in **M1–M3**, then given $\tilde{\Sigma}$ and \mathbf{H} , Σ is uniquely defined iff $m \geq n$.*

See Section 1 in the Supplemental material for a proof of both results.

4.1 Parameterisation

In light of the results presented above, we can devise two different parameterisations for the random effects in (3.1): the *post* parameterisation and the *inverse* parameterisation. In fact vectors ϕ and $\tilde{\phi}$ are linked deterministically by the linear transformation \mathbf{H} (see (3.2)), so the model has to specify a prior only for one of the two.

The *post*-transform parameterisation consists of using the CAR prior directly to sample random effects ϕ^* from the spatial structure (2.4), and then multiply them by \mathbf{H} to obtain the MM equivalents using $\mathbf{H}\phi^* = \tilde{\phi}^*$.

The *inverse* parameterisation is defined as follows. We specify a prior on $\tilde{\phi}^\dagger$ directly from the model in (4.1) and obtain the corresponding spatial random effects using the inverse MM transformation $\mathbf{H}^{-1}\tilde{\phi}^\dagger = \phi^\dagger$.

In terms of notation, all quantities associated with the *post*-transform parameterisation will be denoted with an asterisk (*) superscript, while for the *inverse* parameterisation a dagger (†) will be used.

In order to be able to specify a proper CAR prior, we need the covariance matrix to be positive definite. This means that due to Theorem 4.1, the *inverse* parameterisation is available only in the case where $m \leq n$. Since when m is strictly smaller than n by Lemma 1 there is no left inverse that would allow us to obtain $\mathbf{H}^{-1}\tilde{\phi}^\dagger = \phi^\dagger$. Instead, the *post* parameterisation can be implemented in all cases, given that it is only necessary to specify a proper CAR prior on the areal structure as in (2.4).

It should be noted that whenever $m \neq n$ the specification of the CAR-MM prior can be seen as a case of Low-rank representation of a spatial process (Gelfand et al., 2010, Chapter 8), and specifically *fixed rank kriging* (Cressie et al., 2008). In fact, if we consider the *post* parameterisation for instance, given that $\mathbf{H}\phi^* = \tilde{\phi}^*$, for $m > n$ ϕ^* is a lower rank representation of $\tilde{\phi}^*$, and vice versa for $m < n$, as one random effect can be specified in terms of a deterministic transformation of a lower rank distribution.

Low-rank representations can be particularly useful for geostatistical models as they can significantly reduce the computational burden that is generally required for spatial prediction. In this context matrix \mathbf{H} takes the name of *basis* and its selection, and in some cases estimation, is generally arbitrary and depends on the problem at hand. This is different from the CAR-MM case where the MM matrix \mathbf{H} is entirely determined by the population structure of the areas and memberships. Additionally, contrary to most geostatistical applications in this context where the lower rank process is simply used as a latent factor, for the CAR-MM both the membership and the areal level, i.e. both the lower and higher rank process can be of inferential interest.

4.2 Identifiability

We can now investigate the identifiability of the model by introducing two definitions of non-identifiability in a Bayesian context and evaluate each one in the case of CAR-MM random effects. We will not provide explicit considerations for the hyperparameters α and τ , as they are generally not parameters of interest.

Firstly, according to Dawid (1979):

Definition 1 (Likelihood non-identifiability). *For a Bayesian model $M(\boldsymbol{\theta}, \mathbf{y})$ with data \mathbf{y} , likelihood L and prior f , if the parameter vector $\boldsymbol{\theta}$ is partitioned into $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ then $\boldsymbol{\theta}_2$ is non-identifiable if:*

$$f(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, \mathbf{y}) = f(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1). \quad (4.3)$$

Gelfand et al. (1999) remark that Definition 1 is equivalent to the classical notion of non-identifiability of the likelihood (see Definition 1 Section 3.1.1 in Cole (2020))

since:

$$f(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1, \mathbf{y}) \propto L(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \mid \mathbf{y})f(\boldsymbol{\theta}_2 \mid \boldsymbol{\theta}_1)f(\boldsymbol{\theta}_1) \tag{4.4}$$

is not identifiable according to Definition 1 iff $L(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \mid \mathbf{y})$ can be re-parameterised so that it becomes free of $\boldsymbol{\theta}_2$.

In addition, it is worth recalling that Catchpole et al. (1997) define *parameter redundancy* as

Definition 2 (Parameter redundancy). *A model $M(\boldsymbol{\theta}, \mathbf{y})$ is parameter redundant if we can write $\boldsymbol{\theta}$ as a function of just $\boldsymbol{\beta}$, where $\boldsymbol{\beta} = g(\boldsymbol{\theta})$ and $\dim(\boldsymbol{\beta}) < \dim(\boldsymbol{\theta})$, where \dim is length of the vector.*

They also prove that parameter redundancy implies non-identifiability in the likelihood which is then equivalent to Definition 1 (Theorem 4 in Catchpole et al. (1997)). We refer back to the model in (4.1) and conditions **M1–M3** imposed on \mathbf{H} in Section 4 and evaluate the two parameterisations under the definitions we just presented.

In the case of a *post*-parameterised model, it is evident that when $m > n$ the MM random effects $\tilde{\boldsymbol{\phi}}^*$ are parameter redundant as they can be written as a function of the sole areal random effects $\boldsymbol{\phi}^*$ which are n dimensional. More specifically, there is an injective function g such that $g(\boldsymbol{\phi}^*) = \mathbf{H}\boldsymbol{\phi}^* = \tilde{\boldsymbol{\phi}}^*$, where $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$.

This can also be seen more explicitly, by partitioning the model parameters in (3.1) into $(\boldsymbol{\phi}, \tilde{\boldsymbol{\phi}}, \boldsymbol{\theta})$, where $\boldsymbol{\theta} = (\gamma, \boldsymbol{\beta}, \alpha, \tau)$, we can write the posterior distribution, up to a constant, as

$$\begin{aligned} L(\tilde{\mathbf{y}} \mid \tilde{\boldsymbol{\phi}}^*, \boldsymbol{\phi}, \boldsymbol{\theta}, \mathbf{H}, \mathbf{X})f(\tilde{\boldsymbol{\phi}}^*, \boldsymbol{\phi}^*, \boldsymbol{\theta}) &= L(\tilde{\mathbf{y}} \mid g(\boldsymbol{\phi}^*), \boldsymbol{\phi}, \boldsymbol{\theta}, \mathbf{H}, \mathbf{X})f(g(\boldsymbol{\phi}^*)\boldsymbol{\phi}^*, \boldsymbol{\theta}) \\ &= L(\tilde{\mathbf{y}} \mid g(\boldsymbol{\phi}^*), \boldsymbol{\theta}, \mathbf{H}, \mathbf{X})f(\boldsymbol{\phi}^*, \boldsymbol{\theta}), \end{aligned} \tag{4.5}$$

where the second equality comes from the fact that the $\tilde{\boldsymbol{\phi}}^*$ are a deterministic function exclusively of the $\boldsymbol{\phi}^*$ and thus do not depend on either $\boldsymbol{\theta}$ or $\{\mathbf{H}, \mathbf{X}\}$. Consequently, in (4.5) we can see that the MM random effects $\tilde{\boldsymbol{\phi}}$ are non-identifiable in the likelihood (Definition 1). In the case of the *post* parameterisation, this holds true for all m and n .

With regard to the *inverse* parameterisation, where a prior is specified on the MM random effects $\tilde{\boldsymbol{\phi}}^\dagger$, in the case $m \geq n$ by Lemma 1 there is a left inverse such that $\boldsymbol{\phi}^\dagger = \mathbf{L}\tilde{\boldsymbol{\phi}}^\dagger$. Thus the same considerations hold as for the *post* parameterisation made above. However in this case the areal and MM random effects exchange statuses in terms of identifiability and redundancy. To see this, suffices to swap $\boldsymbol{\phi}^*$ for $\tilde{\boldsymbol{\phi}}^\dagger$ and k for g in (4.5), where $\mathbf{L}\tilde{\boldsymbol{\phi}}^\dagger = \boldsymbol{\phi}^\dagger = k(\tilde{\boldsymbol{\phi}}^\dagger)$ and $k : \mathbb{R}^m \rightarrow \mathbb{R}^n$. In the $m < n$ case for the *inverse* parameterisation, by Theorem 4.1 the areal random effects covariance matrix $\boldsymbol{\Sigma}^\dagger$ is not uniquely defined and thus the $\boldsymbol{\phi}^\dagger$ are not identifiable in the likelihood.

Despite these issues with the likelihood, we can adopt a different perspective following a different definition, by Cole (2020) (Definition 12):

Definition 3 (Posterior non-identifiability). *A Bayesian model $M(\boldsymbol{\theta}, \mathbf{y})$ described by the posterior distribution, is globally identifiable if $f(\boldsymbol{\theta}_a \mid \mathbf{y}) = f(\boldsymbol{\theta}_b \mid \mathbf{y})$ implies $\boldsymbol{\theta}_a = \boldsymbol{\theta}_b$.*

A model is locally identifiable if there exists an open neighbourhood of any θ such that this is true. Otherwise the posterior is non-identifiable.

Under this perspective we can see that both membership and areal random effects are identifiable in the posterior for the post parameterisation when $m \geq n$. In fact, by setting $\Delta = \{\tilde{\mathbf{y}}, \mathbf{H}, \mathbf{X}\}$, we can investigate identifiability by writing down the posterior distribution as:

$$f(g(\phi_a^*) \mid \phi_a^*, \theta_a, \Delta) f(\phi_a^*, \theta_a \mid \Delta) = f(g(\phi_b^*) \mid \phi_b^*, \theta_b, \Delta) f(\phi_b^*, \theta_b \mid \Delta). \quad (4.6)$$

In the $m \geq n$ case, we can easily see that since the function g is injective, we have that if and only if $\phi_a^* = \phi_b^*$ then $g(\phi_a^*) = g(\phi_b^*)$. In turn, this entails that for the purpose of verifying Definition 3, (4.6) reduces to:

$$f(\phi_a^*, \theta_a \mid \Delta) = f(\phi_b^*, \theta_b \mid \Delta), \quad (4.7)$$

where the terms ϕ^* are simple CAR random effects.

Therefore, for $m \geq n$ (4.7) shows that a post parameterised CAR-MM model as in (3.1), is identifiable if and only if the underlying CAR model (2.2) is as well. Conversely, we remark that in the $m < n$ case, given that g is not injective, (4.6) does not hold in general, thus making the random effects non-identifiable in the posterior.

In the case of the *inverse* parameterisation for $m > n$, (4.6) remains the same with simply the ϕ^* swapped for $\tilde{\phi}^\dagger$ and k for g . However, in this instance the function k is not injective, thus leading to non-identifiability. Furthermore, for $m < n$ there is no function k , given that \mathbf{H} has no left inverse. Only the $m = n$ case is identifiable under both parameterisations.

With regard to the covariates coefficients β , it is apparent from (3.2) that the MM transformation of the log relative risks $\log(\rho)$ does not affect these parameters. Therefore, the question of identifiability of both the areal and membership relative risks hinges on the random effects ϕ and $\tilde{\phi}$ as discussed above.

To conclude, we recommend the use of the CAR-MM model only when a $m \geq n$ dataset is available. Only in this case, in fact, parameters of interest, such as the relative risks as well as risk factor coefficients will be identifiable.

5 Model validation and posterior checks

5.1 Simulation based calibration (SBC)

In order to both check the software implementation of our posterior sampler and verify that the model is able to correctly recover the true parameters, we devised a simulation study following a simulation based calibration (SBC) approach (Talts et al., 2018). In an SBC study, instead of comparing posterior samples to some fixed *ground truth* parameter, the procedure assesses the overall calibration of the posterior samples with

respect to the joint distribution $\pi(\mathbf{y}, \boldsymbol{\theta})$. In this way, it is possible to verify the performance of the posterior sampler for multiple *ground truths* $\tilde{\boldsymbol{\theta}}$ generated from the prior distribution.

In summary, the SBC procedure can be written as follows:

$$\tilde{\boldsymbol{\theta}} \sim \pi(\boldsymbol{\theta}), \quad \tilde{\mathbf{y}} \sim l(\mathbf{y}|\tilde{\boldsymbol{\theta}}), \quad \{\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(B)}\} \sim \pi(\boldsymbol{\theta}|\tilde{\mathbf{y}}). \tag{5.1}$$

Since we have $\pi(\mathbf{y}, \boldsymbol{\theta}) = \pi(\boldsymbol{\theta})l(\mathbf{y}|\boldsymbol{\theta})$, by sampling a parameter vector $\tilde{\boldsymbol{\theta}}$ from the prior and then generate new data from the likelihood $\tilde{\mathbf{y}} \sim l(\mathbf{y}|\tilde{\boldsymbol{\theta}})$, this would be equivalent to drawing a sample from the joint distribution $(\tilde{\mathbf{y}}, \tilde{\boldsymbol{\theta}}) \sim \pi(\mathbf{y}, \boldsymbol{\theta})$. Therefore, whenever we implement an algorithm that correctly samples from a posterior distribution, the $\{\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(B)}\}$ are going to have the same marginal distribution as $\tilde{\boldsymbol{\theta}}$ sampled from the prior (Cook et al., 2006).

On the basis of this, we can straightforwardly check that the implemented algorithm performs inference correctly by verifying that prior and posterior samples are distributed according to the same distribution (Talts et al., 2018): for any one-dimensional random variable $f : \Theta \rightarrow \mathbb{R}$, the *rank statistic* $r(\cdot)$ of the prior sample relative to the posterior samples:

$$r(\{f(\boldsymbol{\theta}^{(1)}), \dots, f(\boldsymbol{\theta}^{(B)})\}, f(\tilde{\boldsymbol{\theta}})) = \sum_{b=1}^B \mathbb{I} [f(\boldsymbol{\theta}^{(b)}) < f(\tilde{\boldsymbol{\theta}})] \in [0, B], \tag{5.2}$$

will be uniformly distributed across the integers $[0, B]$. Therefore, by sampling from the Bayesian joint distribution N times, and for each sample drawing B posterior samples, computing the rank statistic in (5.2) and finally checking it is actually uniformly distributed, we can build a validation procedure for any Bayesian model. We mention that in practice, given that we have both N and B very large, we divide the sampled ranks by $B + 1$ and verify calibration with the Uniform distribution on $[0, 1]$

5.2 Posterior checks and model comparison measures

Before illustrating the simulation results we introduce the methodologies used to assess both model fits and comparisons. Posterior predictive p-values (PPP) are a widely used tool in Bayesian analysis to compare the estimated posterior distribution with the observed data (Gelman et al., 2013). In the case of marginal predictive checks, model calibration can be verified in two steps: first generating new observations \mathbf{y}^{rep} using posterior samples and secondly by computing the marginal PPP's, for each observation $\mathbf{y} = (y_1, \dots, y_m)^T$:

$$p_j = P(T(y_j^{rep}) < T(y_j) | \mathbf{y}) + \frac{1}{2}P(T(y_j^{rep}) = T(y_j) | \mathbf{y}), \tag{5.3}$$

if \mathbf{y} is a discrete outcome, where $T(\cdot)$ is a generic test quantity; in our case we will take the identity function $T(y_j) = y_j$.

A computationally inexpensive alternative is the one from Marshall et al. (2003) where they propose a method to approximate this “leave-one-out” cross-validatory

model fit measure: $p_{j|-j} = P(y_j^{rep} \leq y_j | \mathbf{y}_{-j})$, which for continuous data will have a uniform distribution if the model is well calibrated. Referring to (2.2), we can simulate a new random effect for area i using $\phi_i \sim p(\phi_i^{rep} | \alpha, \tau, \phi_{-i})$ using posterior samples of α and τ and thus approximate the $p_{j|-j}$'s above. This approach can be directly extended to the CAR-MM model, by simply post-multiplying the quantities in the relative risk by the MM matrix \mathbf{H} as per (3.1).

With regard to model comparison, in the data analysis Section we compare different models fit using three criteria. The first one is Pareto smoothed importance-sampling leave-one-out cross-validation (PSIS-LOO) (Vehtari et al., 2017). Similarly to the approach described above for PPP, this method make use of a conditional independence assumption of data given model parameters $\boldsymbol{\theta}$, which in our case is implicit in (2.2) and (3.1). The PSIS-LOO approach allows to compute the *expected log pointwise predictive density* (elpd) to estimate out-of-sample predictive fit: $elpd_{loo} = \sum_{j=1}^m \log p(y_j | \mathbf{y}_{-j})$. Without the need to fit the model m times, estimates of $elpd_{loo}$ can be obtained using

$$\widehat{elpd}_{psis-loo} = \sum_{j=1}^m \log \left(\frac{\sum_{s=1}^S w_j^s p(y_j | \boldsymbol{\theta}^s)}{\sum_{s=1}^S w_j^s} \right), \quad (5.4)$$

where $s = 1, \dots, S$ indexes the posterior sample iterations, and w_j^s the Pareto smoothed importance weights. Model comparison can be carried out by computing the difference in \widehat{elpd}_{loo} between the two models, so for models A and B: $\text{diff}(\widehat{elpd}_{loo}^{A,B}) = \widehat{elpd}_{loo}^A - \widehat{elpd}_{loo}^B$, where a larger \widehat{elpd}_{loo} indicates better estimated predictive accuracy of the model. The standard error of the difference is thus:

$$se(\text{diff}(\widehat{elpd}_{loo-psis}^{A,B})) = \sqrt{m V_{j=1}^m \left(\text{diff}(\widehat{elpd}_{loo-psis}^{A,B}) \right)}, \quad (5.5)$$

where $V_{j=1}^m(a_j)$ is the sample variance of m values with mean \bar{a} : $V_{s=1}^S(a_s) = \frac{1}{m-1} \sum_{j=1}^m (a_j - \bar{a})^2$. The paired estimate of (5.5) relies on the two models being fitted on the same set of m data points.

The \widehat{elpd}_{loo} (we drop ‘‘psis’’ in the subscript for conciseness) is easily computable using the R package `loo` (Vehtari et al., 2020) in conjunction with the RStan package (Carpenter et al., 2017) for sampling from the posterior distribution. The authors suggest checking the estimated shape parameters \hat{k} of the Pareto weights for all observation to assess reliability of estimates: when $\hat{k} > 0.7$ there might be stability issues.

Additionally, we also considered the mean Ranked Probability Score for count data (\overline{RPS}) (Czado et al., 2009)

$$\overline{RPS} = \frac{1}{m} \sum_{j=1}^m \left[\frac{1}{B} \sum_{i=1}^B |y_j^{rep(i)} - y_j| - \frac{1}{B} \sum_{i=1}^{B/2} |y_j^{rep(i)} - y_j^{rep(i+B/2)}| \right], \quad (5.6)$$

where $y_j^{rep(i)}$ is the i -th replicate of observation y_j . Finally, the Dawid-Sebastiani Score (\overline{DSS}) (Dawid et al., 1999) was also included in our comparisons

$$\overline{DSS} = \frac{1}{m} \sum_{j=1}^m \left[\left(\frac{y_j - \overline{y_j^{rep}}}{\sigma_{y_j^{rep}}} \right)^2 + 2 \log \sigma_{y_j^{rep}} \right], \quad (5.7)$$

where $\overline{y_j^{rep}}$ and $\sigma_{y_j^{rep}}$ are respectively the mean and standard deviations of the replicates for the j -th observation. For both \overline{RPS} and \overline{DSS} a lower value indicates better fit for the model.

6 Simulation study

We now turn to a simulation study designed to assess the considerations on identifiability of Section 4. We designed the simulations to evaluate the two parameterisations (*post* and *inverse*) and the number of memberships relative to areas, i.e. $m < n$, $m = n$ and $m > n$. We will first describe the implementation of the posterior sampler for each model, then the simulation setup in detail and finally the results.

6.1 Model implementation

All models are implemented in RStan with significant differences between the two parameterisations. For the *post* parameterised covariance matrices, we followed Joseph (2016)'s sparse representation combined with the approach suggested by Jin et al. (2005) for the determinant. The former relies on a sparse representation of the matrix \mathbf{W} whereby the terms corresponding to non-neighbouring areas are from the computation. With regard to the determinant, following Jin et al. (2005) we can avoid inverting the entire precision matrix \mathbf{Q} at each iteration. Taking $\tau = 1$, without loss of generality, it can be shown that: $\mathbf{Q} = \mathbf{D} - \alpha \mathbf{W} = \mathbf{D}^{1/2}(\mathbf{I} - \alpha \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}) \mathbf{D}^{1/2}$ then $\det(\mathbf{Q}) = \det(\mathbf{D}) \prod_{i=1}^n (1 - \alpha \lambda_i)$, where the λ_i 's are the eigenvalues of $\mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$. This representation significantly speeds up the computation of $\det(\mathbf{Q})$ since the λ_i 's do not depend on either α and τ and thus have to be calculated only once.

Conversely, for the *inverse* parameterisation we specified the entire $m \times m$ precision matrix $\tilde{\mathbf{Q}}$. We remark that in this instance, the sampler can only be applied to the $m \leq n$ case, as it is the only case where the covariance matrix is proper. Under this parameterisation, in order to obtain the areal random effects ϕ^\dagger we used the Moore-Penrose inverse of the \mathbf{H} matrix.

The code and some of the functions for both the simulation study and the RStan models can be found at <https://github.com/markgrama/CAR-MM> and in the R package CARME, available on CRAN (Gramatica et al., 2023).

With regard to the model, we used the following priors (see (3.1)): a uniform on $(0, 1)$ for α , a *Gamma*(2, 0.2), parameterised in terms of shape and rate, for the marginal precision τ , while for γ , β_1 and β_2 three independent $\mathcal{N}(0, 0.7^2)$ distributions were used.

In terms of computational time, with the same number of iterations (see Section 6.2 for details), the *post* parameterisation performed significantly better than the *inverse*, with an average runtime of 2 minutes across membership sizes, against the 10 minutes for

the $m = 70$ case and 29 minutes for $m = 100$ (Table 1 in the Supplemental Material). Our hypothesis is that the *post* parameterisation runs more quickly due to the fact that the sampling of the CAR random effect is done through the sparse representation mentioned above, and the linear transformation \mathbf{H} is only evaluated in the likelihood with RStan optimised matrix operations.

Conversely, the *inverse* parameterisation requires the evaluation of the entire precision matrix, since no adjacency matrix is directly available. This is due to the problem in identifying the partial correlation structure of the CAR-MM mentioned in Section 4. Additionally, even if we were able to determine partial independence relationships between components, there is no guarantee that the resulting precision matrix will be sparse.

6.2 Setup

In this simulation study we use the *post* and *inverse* parameterisations to both generate the data and sample from the posterior, thus obtaining 4 different scenarios. In order to refer to a particular scenario we might write *post-inverse*, where the first parameterisation refers to the data generation (*post*), while the second indicates the MCMC sampler parameterisation (*inverse*). For each scenario we consider three membership sizes (70, 100, 130) so that we can correspondingly evaluate the three cases $m < n$, $m = n$, and $m > n$. However, as we remark in Section 4.1, the *inverse* parameterisation is available only for $m \leq n$, therefore we simulate the $m = 130$ case only for the *post* parameterisation.

The spatial structure used in the study is a 10×10 grid, hence $n = 100$. Expected rates for each memberships (\tilde{E}_j in (3.1)) were simulated from m independent *Poisson*(20) distributions. For both the MM matrix and the $\tilde{\mathbf{E}}$ vector, we simulated the case $m = 130$ and removed the extra rows or elements to obtain the other membership sizes.

With regard to the MM matrix \mathbf{H} , this was simulated starting from each area's neighbourhood structure, encoded in the adjacency matrix \mathbf{W} , up to second order neighbours. For each area a new membership was created by first generating random weights from an uniform distribution to all first and second order neighbours. Then those quantities were normalised so that the area itself would amount to 50% of the total weights, the sum of all first order neighbours would be attributable 35% of the total weight, and the remaining 15% would be taken up by the second order neighbours. More formally, the simulated $m \times n$ matrix \mathbf{H} is obtained by first randomly drawing samples:

$$\begin{aligned} (G)_{ji} &\stackrel{iid}{\sim} \text{Uniform}(0, 1) \quad \text{for } i = 1, \dots, n \text{ and } j = 1, \dots, m \\ \text{s.t. } &\{w_{ji} = 1\} \vee \{w_{il} = 1 \wedge w_{lj} = 1\}, \end{aligned} \quad (6.1)$$

for $l = 1, \dots, n$. Therefore the elements of the matrix \mathbf{H} are defined as follows

$$(H)_{ji} = \begin{cases} 0.5 & \text{if } i = j, \\ \frac{(G)_{ji}}{\sum_{j:w_{ji}=1}(G)_{ji}} \cdot 0.35 & \text{if } w_{ji} = 1, \\ \frac{(G)_{ji}}{\sum_{j:\{w_{il}=1 \wedge w_{lj}=1\}}(G)_{ji}} \cdot 0.15 & \text{if } w_{il} = 1 \text{ and } w_{lj} = 1. \end{cases} \quad (6.2)$$

In order to obtain a 130 membership matrix, the extra rows were created by simply repeating the same process and generating different random weights. Before deleting the bottom rows to obtain the varying number of memberships, in order to ensure that all areas were represented as uniformly as possible in the final matrix its rows were randomly shuffled. Finally, the obtained 70 membership matrix was checked to ensure that each area was represented with at least one weight. In our case the area with the least weights had 2, so it contributed to 2 memberships.

A total of 10^4 datasets were simulated for the memberships structure using the *post* parameterisation, each membership size was obtained by simply truncating the simulated outcome vector accordingly, thus increasing the total of datasets used to 9000. Covariates \mathbf{X} were kept fixed across all datasets and they were generated using two independent standard normal samples and min-max normalised. Both the MM matrix and offsets were also the same in all simulated datasets.

Following the approach described in Section 5.1, hyperparameters for the CAR prior were sampled from the same priors described in Section 6.1 as well as *Hamiltonian Monte Carlo* (HMC) (Betancourt, 2018) sampling which was run using RStan (Carpenter et al., 2017) on Queen Mary's Apocrita High Performance Computing facility, supported by QMUL Research-IT (King et al., 2017), using five 10^4 iterations chains and 50% warm-up. To reduce autocorrelation, as suggested in Talts et al. (2018), samples were thinned by 5 leading to $L = 5000 + 1$ posterior samples (see Section 5.1).

Additionally, in order to ensure convergence of the posterior sampler, we only keep those simulations for which no parameter report a \hat{R} statistic over 1.01, following Vehtari et al. (2021). In Table 2 in the Supplemental material we report the number of simulations that were excluded. Convergence was satisfactory overall, since we always retained at least 98% of the simulations.

Finally we evaluate bias and root mean square error (RMSE) across simulations by computing the following quantities. For a generic parameter θ with true value $\hat{\theta}_k$ for simulated dataset k , $k = 1, \dots, 10^4$ we compute

$$\text{Bias}_k = \frac{1}{B} \sum_{b=1}^B \theta_k^{(b)} - \hat{\theta}_k; \text{Abs bias}_k = \frac{1}{B} \sum_{b=1}^B \left| \theta_k^{(b)} - \hat{\theta}_k \right|; \text{RMSE}_k = \sqrt{\frac{1}{B} \sum_{b=1}^B (\theta_k^{(b)} - \hat{\theta}_k)^2}, \quad (6.3)$$

where $\theta_k^{(b)}$ is the b^{th} HMC sample for the k^{th} dataset and $B = 2.5 \cdot 10^4$.

6.3 Results

In Table 1 we report the coverage of the 95% intervals of the expected variation of the SBC rank statistics against the uniform distribution.¹ For the $m = 100, 130$ cases, we observe high coverage for all the parameters in all scenarios, with the exception of the α parameter, which is generally not of inferential interest. We believe that the relatively

¹The plots for all parameters can be found in Section 2.2 in the Supplemental material.

	Data generation: Post					Data generation: Inverse			
	MCMC: Post			MCMC: Inverse		MCMC: Post		MCMC: Inverse	
	70	100	130	70	100	70	100	70	100
α	87.35	66.62	80.61	37.75	73.45	43.16	76.44	95.80	71.72
τ	68.52	93.80	94.43	8.92	93.36	30.82	90.85	91.47	91.54
γ	98.38	99.42	98.60	99.70	99.23	99.96	99.52	99.95	99.57
β_1	96.15	96.14	83.34	99.79	97.76	99.99	96.41	99.96	97.25
β_2	99.97	99.96	99.75	99.96	99.77	99.95	99.52	99.81	99.46
ϕ	94.63	94.33	94.22	50.72	94.53	57.72	94.53	93.83	94.54
ρ	94.45	94.37	93.51	26.35	94.68	59.30	94.46	27.95	94.17
$\tilde{\rho}$	95.37	95.78	94.72	94.32	95.78	94.38	95.72	95.16	95.88

Table 1: Coverage of the 95% expected variation intervals of the SBC obtained posterior samples of the CAR-MM parameters for each scenario and membership size.

lower coverage of β_1 in the $m = 130$ scenario is not a cause of concern for two reasons: first, as will be seen below, posterior means are unbiased across the simulation study, and second, this fact does not seem to hinder the good calibration of the relative risks ρ and $\tilde{\rho}$.

Overall, in the $m = 70$ case, we have satisfactory coverage for the linear term parameters γ , β_1 and β_2 in all scenarios. Despite this, we observe poor calibration for the spatial random effects ϕ and the areal relative risks ρ , the only exception being the *post-post* case. These results are consistent with the considerations on identifiability that we discussed in Section 4 where we remarked the identifiability only of the $m = n$ case under both parameterisation and the $m > n$ for the *post* one. Although theoretically the $m = 70$ should not be identifiable, the calibration in the *post-post* case could be due to the coincidence between data generation and the MCMC parameterisation. Conversely, in the *inverse-inverse* case, despite the two parameterisations being the same, we remark that the generalised inverse used to obtain areal relative risks (ρ) from the sampled membership risks ($\tilde{\rho}$) is not a right inverse (Lemma 1). This leads to the poor calibration of the ρ .

All in all, these simulations suggest that in the *post* parameterisation case the posterior samples are well calibrated and support the considerations on identifiability of the CAR-MM model discussed above. Overall these results lead us to recommend using the *post* parameterisations in all applications, not only due to its computational efficiency as mentioned in Section 6.1, but also due to its overall better calibration.

However, despite the positive results for the $m < n$ case, we remind the reader that in Section 4.2 we showed its non identifiability in the posterior, under both parameterisations. The good calibration results in the *post-post* (both data and MCMC sampler are *post* parameterised) case could simply be an artefact of the simulation setup. In fact, in that scenario the function that maps the areal relative risks to the membership risks, i.e. the MM matrix \mathbf{H} is the same both when simulating the data and sampling from the posterior distribution. Conversely, whenever an inverse parameterisation is involved, either in the data simulation stage or during MCMC sampling, given that a left inverse

does not exist (Lemma 1), the data does not contain enough information to correctly recover the spatial random effects at the areal level. Finally, high calibration for all parameters across different parameterisations for the two $m \geq n$ cases is evidence of the reliability of the CAR-MM model.

7 Data analysis

7.1 Data structure and modelling

We present an application of this methodology to primary care catchment areas. As argued above, the CAR-MM approach is of utility in situations where data are observed for non-spatial units (memberships) but where an underlying spatial framework can be conceived, and where an area to membership population cross-referencing is available (Health and Social Care Information Centre, 2019). In particular, UK primary care data on disease prevalence (which is recorded at GP practice level) offers a suitable application. Specifically, we consider diabetes prevalence for a set of GP practices in South East London (SEL), and to the MSOAs (small areas) defining the catchment areas of these GP practices.

The data pertain to 153 GP practices which together constitute what are known as Clinical Commissioning Groups or CCGs: the 153 practices are those for Bexley, Lewisham, Greenwich and Bromley CCGs. Based on the choice of GPs, a total of 152 MSOAs were included so that at least 90% coverage of each GPs population was reached.

Prevalence data for GPs refers to October 2015, as does the cross reference information providing the number of each MSOA residents registered at each GP practice (Health and Social Care Information Centre, 2016). Age structure information was also available at GP level, so that expected rates could be easily computed using England and Wales national rates of diabetes prevalence.

For MSOA level covariates, two were included: from the 2011 Census the share of South Asian residents (\mathbf{X}).₁ i.e. Indian, Pakistani and Bangladeshi, see Table KS201EW in the NOMIS database (Office for National Statistics, 2016), and the 2015 Index of Multiple deprivation (IMD) (\mathbf{X}).₂. The two variables were included as a control of socio-economic status and due to the higher prevalence of diabetes among the South Asian population (Hall et al., 2010). Both variables were min-max normalised.

7.2 Results

Initial runs using the Poisson likelihood model (3.1) returned concerning PSIS-LOO weights \hat{k} which might be a symptom of misspecification of the model. In fact, Vehtari et al. (2017) note that high \hat{k} mean that the full posterior distribution is a poor approximation of the leave-one-out posterior.

Since the possible misspecification might be attributable to overdispersion, i.e. when the variance of the outcome is larger than its average (Coly et al., 2019) we therefore

turned to a Negative Binomial likelihood for the model in (3.1), that allows for direct modelling of overdispersion.

The density of the parameterisation we used is the following:

$$f_Y(y_j; \mu_j, \psi) = \binom{y_j + \psi - 1}{y_j} \left(\frac{\mu_j}{\mu_j + \psi}\right)^{y_j} \left(\frac{\psi}{\mu_j + \psi}\right)^\psi, \tag{7.1}$$

where $y_j \in \mathbb{R}$ for all $j = 1, \dots, m$; $E(Y_j) = \mu_j \in \mathbb{R}^+$, $\psi \in \mathbb{R}^+$ and $Var(Y_j) = \mu_j + \mu_j^2/\psi$. Therefore the mean of each Y_j can be modelled directly through the μ_j 's. (3.1) then becomes:

$$\tilde{Y}_j \sim NegBin(\tilde{E}_j \tilde{\rho}_j, \psi) \quad , \quad \log(\tilde{\rho}) = \log(\mathbf{H}\boldsymbol{\rho}) = \mathbf{H}(\gamma \mathbf{1}_n + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\phi}), \tag{7.2}$$

so the overdispersion parameter ψ is assumed common across all GP practices (i.e. memberships). The priors used were the same as for the simulation study in Section 6.

In the analysis we compared three models: the CAR-MM as described above, the ICAR-MM and a GLM-MM that simply consists in removing the spatial random effects $\boldsymbol{\phi}$ from (7.2). We include the ICAR since it is one of the most widely used prior in the disease mapping literature (Martinez-Beneito et al., 2019). All models were coded in RStan (see Section 6.1 and Morris et al. (2019) for the ICAR prior) and run for five 10^5 iterations chains with 50% warm-up. With regard to the prior for the added overdispersion parameter ψ for the Negative Binomial likelihood, $Gamma(2, 0.2)$, parameterised in terms of shape and rate, was used.

Table 2 shows posterior estimates for the three models estimated on the SEL dataset. In general estimates are very similar across the three models, with the exceptions of marginal CAR precision τ and β_1 coefficient, the former larger and the latter smaller in the ICAR-MM. As expected, given the fact that no term is directly accounting for residual spatial autocorrelation, the GLM-MM model estimates slightly smaller credible intervals for β_1 and β_2 than the two CAR based models.

	α		τ		γ			β_1			β_2			ψ			
mean	-	0.52	-	18.42	26.02	-0.75	-0.74	-0.69	0.75	0.69	0.54	1.02	1.02	0.99	31.61	33.85	34.71
sd	-	0.28	-	7.89	9.01	0.05	0.06	0.07	0.14	0.17	0.20	0.07	0.09	0.11	4.05	4.59	4.63
2.5%	-	0.03	-	7.13	12.26	-0.84	-0.86	-0.83	0.49	0.34	0.13	0.88	0.85	0.78	24.35	25.73	26.35
50%	-	0.54	-	17.04	24.71	-0.75	-0.74	-0.69	0.75	0.69	0.54	1.02	1.02	0.99	31.37	33.55	34.46
97.5%	-	0.96	-	37.64	47.06	-0.66	-0.62	-0.56	1.02	1.02	0.94	1.16	1.19	1.19	40.20	43.68	44.58
width CI	-	0.93	-	30.51	34.80	0.18	0.24	0.28	0.53	0.68	0.80	0.28	0.34	0.41	15.85	17.94	18.23
\hat{R}	-	1.00	-	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
ESS	-	3414	-	14048	13034	10821	7586	5325	12741	9850	6496	12570	14058	5937	14984	27963	28552

Table 2: Posterior results for SEL data analysis comparing Negative Binomial GLM-MM models (white columns), CAR-MM (light grey columns) and ICAR-MM (dark grey columns). β_1 is the coefficient, associated to the share of South Asian population, and β_2 for the IMD (both min-max normalised).

Model comparisons in Table 3 do not seem to show substantial differences between model fits: the ICAR shows the highest \widehat{elpd}_{loo} although with no significant difference from the other models. Despite the GLM-MM having the lowest \overline{RPS} and \overline{DSS} , the small differences and the result from other model comparison criteria, there is not

	\widehat{elpd}	$se(\widehat{elpd}_{loo})$	$\text{diff}(\widehat{elpd}_{loo}^{A,B})$	$se[\text{diff}(\widehat{elpd}_{loo}^{A,B})]$	\overline{RPS}	\overline{DSS}
ICAR-MM	-847.8	14.2	0	0	35.93	9.39
CAR-MM	-849.0	13.8	-1.2	1.1	35.11	9.32
GLM-MM	-849.1	13.0	-1.3	2.5	34.86	9.28

Table 3: Comparison of model check measures. ICAR-MM model is used as reference (model A) when computing differences $\text{diff}(\widehat{elpd}_{loo}^{A,B})$.

sufficient evidence to claim it provides a better fit. In fact, in Figure 10 (Supplemental material), we can see no discernible difference between posterior calibration of marginal PPP's. Therefore, we focus on the CAR-MM model as it is the one on which we focused our attention in this paper.

7.3 Strategic implications

Comparison of posterior relative risk estimates between the 152 MSOAs in the catchment area shows a wide contrast in estimated prevalence and hence in the need for diabetes care. Such contrasts show substantial ecological variation – that is, variation due to area context. Such variations support locality based or place-based interventions and allocation of resources for tackling diabetes (National Institute for Health and Care Excellence, 2014).

Considering the CAR-MM model, the ratio of the 90th percentile relative risk to the 10th percentile relative risk has posterior mean (95% Credible Interval) of 2.31 (2.06, 2.55), an over 2-fold variation in risk.

Figure 1b maps out the probabilities $P(\rho_i > 1 | \tilde{\mathbf{y}})$ that the relative risk exceeds 1 under the CAR-MM in the posterior samples (Richardson et al., 2004). There are 24 areas with over 95% probability of elevated risk, indicating highly elevated diabetes prevalence, and 14 areas with this probability exceeding 99%. By contrast, a considerably larger minority of 61 areas have under 5% probability of elevated risk, with 49 areas having exceedance probabilities below 0.01, indicating significantly low prevalence.

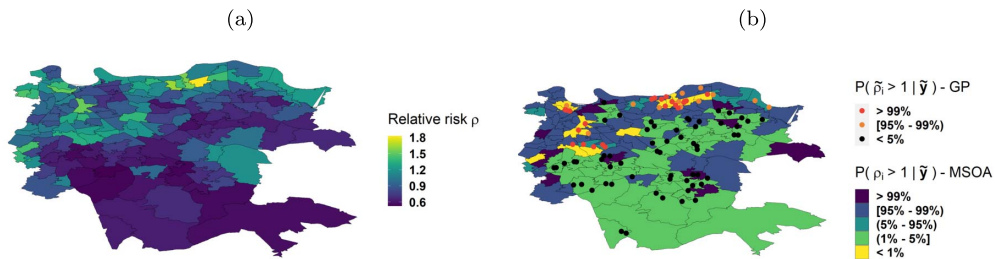


Figure 1: (a) Posterior means for each areal relative risk (ρ in (7.2)) for the CAR-MM model (b) Probabilities $P(\rho_i > 1 | \tilde{\mathbf{y}})$ that the relative risk in each exceeds 1 among posterior samples under the CAR-MM prior.

Figures 1a and 1b also show the spatial features of elevated diabetes levels in a minority of areas, particularly a spatial concentration in the north of the region, reinforcing findings such as (Green et al., 2003). The posterior mean (95% CI) of the Moran's I for the relative risks is 0.44 (0.35, 0.49). The clustering of high risk reflects residential concentration in observed area risk factors, namely deprived communities and in the South Asian community, as well as spatial clustering in unobserved influences on diabetes risk. Hence both area risk factors should figure in any area rankings for locality based intervention.

Table 4 disaggregates the 152 MSOAs into quintiles according to their deprivation scores and their percentages South Asian. If we consider the mean of posterior relative risks in the 30 MSOAs which are most deprived we see an average (95% CI) of 1.32 (1.15, 1.55), as compared to a much lower relative risk in the least deprived MSOAs. The posterior relative risk in the 30 MSOAs with the relatively largest Asian communities has mean (95% CI) 1.07 (0.64, 1.68) as compared to significantly lower relative risks in MSOAs with relatively small Asian communities. Such findings add to existing UK-based findings that area deprivation and ethnic mix significantly affects area diabetes prevalence and incidence variations (Moody et al., 2016; Nishino et al., 2015).

	IMD			South Asian population		
	mean	2.5%	97.5%	mean	2.5%	97.5%
1 st	0.62	0.54	0.79	0.92	0.54	1.39
2 nd	0.74	0.62	0.89	0.91	0.56	1.38
3 rd	0.90	0.75	1.10	0.87	0.55	1.45
4 th	1.13	0.96	1.64	0.93	0.58	1.45
5 th	1.32	1.15	1.55	1.07	0.64	1.68

Table 4: Relative risk profiles of posterior means by covariate quintiles.

Finally, so far we analysed the areal relative risks ρ however, as mentioned in Section 4, the membership relative risks $\tilde{\rho}$ are also identifiable in the posterior (since $m > n$ in our case) and can be of inferential interest in many applications. In fact, in our case each $\tilde{\rho}_j$ corresponds to the relative risk of a single GP practice and can be of significant epidemiological and public health interest.

Therefore, we can replicate the analysis of the probabilities of a relative risk exceeding 1 as we did above. In the case of GP practices, we can see that: 41 practices have over 95% probability, 29 with over 99%, and 68 are under 5%, with 62 under 1%. In Figure 1b we also show the locations of those practices on the map.

As expected, in the SEL region, these results are similar to the ones observed for areal relative risks, given that we have a similar number of practices and areas. However, these membership level relative risk can give insight in which GP's could be facing significant public health challenges. In fact, although most of the practices with either high or low probability are in accordance located in areas with similar risk profiles, there are some GP's for which the opposite is true, i.e. high or low risk practices in respectively low or high risk areas. A purely areal analysis of prevalence would have not been able to identify these practices.

8 Discussion

In this work we have studied and evaluated the theoretical underpinnings of the CAR-MM model devising a general framework to evaluate its structure and implications to spatial and relative risk modelling. We compared our approach to that of Petrof et al. (2020), provided a more general framework able to encompass a potentially wide range of applications, and showed how this could apply to the parameterisation and evaluation of a CAR-MM model. Building on that, we considered identifiability in both the likelihood and the posterior for different membership sizes compared to areas and both parameterisations.

The simulation study highlights how this approach leads to well calibrated posterior samples and small bias for estimated parameters. We showed how the two parameterisations, although theoretically identical in the $m = n$ case, can lead to significantly different results in terms of both computational time and identifiability when $m \neq n$. More specifically, we recommend the use of the *post* parameterisation in all cases although we warn of the possible identifiability problems in the $m < n$ case. In this regard, the present paper supports the previous analyses and findings of both Gramatica et al. (2021) and Petrof et al. (2020), where $m > n$ datasets were considered, and where the models were *post* parameterised.

Finally, the data analysis allowed us to show how this approach can be used in a real world application where the MM principle and CAR frameworks merge naturally, thus allowing for a direct implementation of the CAR-MM prior. Datasets from different institutions, in this case the National Health Service (NHS) and the Census, can be directly integrated leading to a modelling strategy that not only is able to pinpoint areas with increased relative risks, but also specific practices, by making use of the membership level modelling. These evaluations can be of interest for both epidemiological research and public health considerations.

Supplementary Material

Supplemental material (DOI: [10.1214/23-BA1370SUPP](https://doi.org/10.1214/23-BA1370SUPP); .pdf).

References

- Banerjee, S., B. P. Carlin, and A. E. Gelfand (2014). *Hierarchical modeling and analysis for spatial data*. Chapman and Hall/CRC, p. 562. [MR3362184](#). 727
- Besag, J. (1974). *Spatial Interaction and the Statistical Analysis of Lattice Systems*. 725
- Besag, J., J. York, and A. Mollié (1991). “Bayesian image restoration, with two applications in spatial statistics”. In: *Annals of the Institute of Statistical Mathematics* 43.1, pp. 1–20. [MR1105822](#). doi: <https://doi.org/10.1007/BF00116466>. 728
- Betancourt, M. (2018). “A Conceptual Introduction to Hamiltonian Monte Carlo”. In: arXiv: [1701.02434v2](https://arxiv.org/abs/1701.02434v2). 739

- Blangiardo, M. and M. Cameletti (2015). *Spatial and Spatio-temporal Bayesian Models with R – INLA*. en. John Wiley & Sons, pp. 1–308. [MR3364017](#). 725
- Bradley, J. R., C. K. Wikle, and S. H. Holan (2016). “Bayesian Spatial Change of Support for Count-Valued Survey Data With Application to the American Community Survey”. In: *Journal of the American Statistical Association* 111.514, pp. 472–487. [MR3538680](#). doi: <https://doi.org/10.1080/01621459.2015.1117471>. 727
- Browne, W. J., H. Goldstein, and J. Rasbash (2001). “Multiple membership multiple classification (MMMC) models”. In: *Statistical Modelling: An International Journal* 1.2, pp. 103–124. 726
- Carpenter, B. et al. (2017). “Stan: A Probabilistic Programming Language”. In: *Journal of Statistical Software* 76.1, pp. 1–32. 736, 739
- Catchpole, E. A. and B. J. Morgan (1997). “Detecting parameter redundancy”. In: *Biometrika* 84.1, pp. 187–196. [MR1450200](#). doi: <https://doi.org/10.1093/biomet/84.1.187>. 733
- Cole, D. (2020). *Parameter Redundancy and Identifiability*. Chapman & Hall/CRC Interdisciplinary Statistics. CRC Press. 732, 733
- Coly, S. et al. (2019). “Bayesian hierarchical models for disease mapping applied to contagious pathologies”. In: *bioRxiv*, p. 766071. 741
- Congdon, P. D. (2019). *Bayesian Hierarchical Models*. Chapman and Hall/CRC. 725
- Cook, S. R., A. Gelman, and D. B. Rubin (2006). “Validation of Software for Bayesian Models Using Posterior Quantiles”. In: *Journal of Computational and Graphical Statistics* 15.3, pp. 675–692. [MR2291268](#). doi: <https://doi.org/10.1198/106186006X136976>. 735
- Cressie, N. and G. Johannesson (2008). “Fixed rank kriging for very large spatial data sets”. In: *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 70, pp. 209–226. [MR2412639](#). doi: <https://doi.org/10.1111/j.1467-9868.2007.00633.x>. 732
- Czado, C., T. Gneiting, and L. Held (2009). “Predictive model assessment for count data”. In: *Biometrics* 65.4, pp. 1254–1261. [MR2756513](#). doi: <https://doi.org/10.1111/j.1541-0420.2009.01191.x>. 736
- Dawid, A. P. (1979). *Conditional Independence in Statistical Theory*. Tech. rep. 1, pp. 1–31. [MR0535541](#). 732
- Dawid, A. P. and P. Sebastiani (1999). “Coherent dispersion criteria for optimal experimental design”. In: *The Annals of Statistics* 27.1, pp. 65–81. [MR1701101](#). doi: <https://doi.org/10.1214/aos/1018031101>. 736
- Eberly, L. E. and B. P. Carlin (2000). “Identifiability and convergence issues for Markov chain Monte Carlo fitting of spatial models”. eng. In: *Statistics in Medicine* 19.17–18, pp. 2279–2294. 728

- Gelfand, A. E. et al. (2010). *Handbook of Spatial Statistics*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press. 732
- Gelfand, A. E. and S. K. Sahu (1999). “Identifiability, Improper Priors, and Gibbs Sampling for Generalized Linear Models”. In: *Journal of the American Statistical Association* 94.445, pp. 247–253. MR1689229. doi: <https://doi.org/10.2307/2669699>. 732
- Gelman, A. et al. (2013). *Bayesian data analysis*. CRC Press, p. 667. 735
- Goicoa, T. et al. (2018). “In spatio-temporal disease mapping models, identifiability constraints affect PQL and INLA results”. In: *Stochastic Environmental Research and Risk Assessment* 32.3, pp. 749–770. 728
- Gramatica, M., P. Congdon, and S. Liverani (2021). “Bayesian modelling for spatially misaligned health areal data: A multiple membership approach”. In: *Journal of the Royal Statistical Society. Series C: Applied Statistics* 70.3, pp. 645–666. MR4275840. doi: <https://doi.org/10.1111/rssc.12480>. 725, 726, 727, 729, 745
- Gramatica, M., S. Liverani, and P. Congdon (2023). *CARME: CAR-MM Modelling in Stan*. R package version 0.1. 737
- Gramatica, M., S. Liverani, and P. Congdon (2023). “Supplementary Material for “Structure Induced by a Multiple Membership Transformation on the Conditional Autoregressive Model”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/23-BA1370SUPP>. 730
- Green, C. et al. (2003). “Geographic analysis of diabetes prevalence in an urban area”. In: *Social Science & Medicine* 57.3, pp. 551–560. 744
- Hall, L. M. L. et al. (2010). “Fat Oxidation, Fitness and Skeletal Muscle Expression of Oxidative/Lipid Metabolism Genes in South Asians: Implications for Insulin Resistance?” In: *PLoS ONE* 5.12. Ed. by J. A. L. Calbet, e14197. 741
- Harville, D. A. (2008). *Matrix Algebra From a Statistician’s Perspective*. Springer. MR1467237. doi: <https://doi.org/10.1007/b98818>. 730
- Health and Social Care Information Centre (2016). *Quality and Outcomes Framework (QOF) – 2015–16*. 741
- Health and Social Care Information Centre (2019). *Patients Registered at a GP Practice – January 2019*. 741
- Jin, X., B. P. Carlin, and S. Banerjee (2005). “Generalized hierarchical multivariate CAR models for areal data”. en. In: *Biometrics* 61.4, pp. 950–961. MR2216188. doi: <https://doi.org/10.1111/j.1541-0420.2005.00359.x>. 737
- Joseph, M. (2016). *Exact sparse CAR models in Stan*. URL: <https://mc-stan.org/users/documentation/case-studies/mbjoseph-CARStan.html> (visited on 01/28/2020). 737
- King, T., S. Butcher, and L. Zalewski (2017). *Apocrita – High Performance Computing Cluster for Queen Mary University of London*. 739

- Lawson, A. (2018). *Bayesian Disease Mapping: Hierarchical Modeling in Spatial Epidemiology, Third Edition*. en. CRC Press, p. 652. [MR2484272](#). 725
- MacNab, Y. C. (2010). “On Gaussian Markov random fields and Bayesian disease mapping:” in: <https://doi.org/10.1177/0962280210371561> 20.1, pp. 49–68. [MR2767372](#). doi: <https://doi.org/10.1177/0962280210371561>. 725
- Marshall, E. C. and D. J. Spiegelhalter (2003). “Approximate cross-validators predictive checks in disease mapping models”. In: *Statistics in Medicine* 22.10, pp. 1649–1660. 735
- Martinez-Beneito, M. A. and P. Botella-Rocamora (2019). *Disease Mapping: From Foundations to Multidimensional Modeling*. Chapman and Hall/CRC. 730, 742
- McCullagh, P. and J. Nedler (1983). *Generalized Linear Models*. CRC Press. [MR0727836](#). doi: <https://doi.org/10.1007/978-1-4899-3244-0>. 727
- Moody, A. et al. (2016). “Social inequalities in prevalence of diagnosed and undiagnosed diabetes and impaired glucose regulation in participants in the Health Surveys for England series”. In: *BMJ Open* 6.2, e010155. 744
- Morris, M. et al. (2019). “Bayesian hierarchical spatial models: Implementing the Besag York Mollie model in stan”. In: *Spatial and Spatio-temporal Epidemiology* 31, p. 100301. 742
- National Institute for Health and Care Excellence (2014). *Preventing Type 2 Diabetes: Population and Community-level Interventions*. 743
- Nishino, Y., S. Gilmour, and K. Shibuya (2015). “Inequality in Diabetes-Related Hospital Admissions in England by Socioeconomic Deprivation and Ethnicity: Facility-Based Cross-Sectional Analysis”. In: *PLOS ONE* 10.2, e0116689. 744
- Office for National Statistics National Records of Scotland, N. (2016). *2011 Census aggregate data*. 741
- Petrof, O. et al. (2020). “On the impact of residential history in the spatial analysis of diseases with a long latency period: A study of mesothelioma in Belgium”. In: *Statistics in Medicine* 39.26, pp. 3840–3866. [MR4164140](#). doi: <https://doi.org/10.1002/sim.8697>. 725, 726, 727, 729, 745
- Rao, C. R., ed. (1973). *Linear Statistical Inference and its Applications*. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc. 726, 730
- Richardson, S. et al. (2004). “Interpreting Posterior Relative Risk Estimates in Disease-Mapping Studies”. In: *Environmental Health Perspectives* 112.9, p. 1016. 743
- Rue, H. and L. Held (2005). *Gaussian Markov random fields: theory and applications*. Chapman & Hall/CRC, p. 263. [MR2130347](#). doi: <https://doi.org/10.1201/9780203492024>. 731
- Talts, S. et al. (2018). “Validating Bayesian Inference Algorithms with Simulation-Based Calibration”. In: arXiv: [1804.06788](#). 734, 735, 739

- Vehtari, A., A. Gelman, and J. Gabry (2017). “Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC”. In: *Statistics and Computing* 27.5, pp. 1413–1432. MR3647105. doi: <https://doi.org/10.1007/s11222-016-9696-4>. 736, 741
- Vehtari, A. et al. (2020). *loo: Efficient leave-one-out cross-validation and WAIC for Bayesian models*. R package version 2.4.1. 736
- Vehtari, A. et al. (2021). “Rank-Normalization, Folding, and Localization: An Improved \hat{R} for Assessing Convergence of MCMC”. In: *Bayesian Analysis* 16.2, pp. 667–719. MR4298989. doi: <https://doi.org/10.1214/20-ba1221>. 739
- Ver Hoef, J. M., E. M. Hanks, and M. B. Hooten (2018). “On the relationship between conditional (CAR) and simultaneous (SAR) autoregressive models”. In: *Spatial Statistics* 25, pp. 68–85. MR3809256. doi: <https://doi.org/10.1016/j.spasta.2018.04.006>. 726, 728, 730