Spatial and Spatio-temporal Epidemiology 26 (2018) 25–34

Contents lists available at ScienceDirect

Spatial and Spatio-temporal Epidemiology

journal homepage: www.elsevier.com/locate/sste

Original Research

A note on intrinsic conditional autoregressive models for disconnected graphs

Anna Freni-Sterrantino[a,*], Massimo Ventrucci[b], Håvard Rue[c]

[a] Small Area Health Statistics Unit, Department of Epidemiology and Biostatistics, Imperial College London, United Kingdom
[b] Department of Statistics, University of Bologna, Bologna, Italy
[c] CEMSE Division, King Abdullah University of Science and Technology, Saudi Arabia

ARTICLE INFO

Article history:
Received 6 November 2017
Revised 23 March 2018
Accepted 25 April 2018

Keywords:
CAR models
Disease mapping
Disconnected graph
Gaussian Markov random fields
Islands
INLA

ABSTRACT

In this note we discuss (Gaussian) intrinsic conditional autoregressive (CAR) models for disconnected graphs, with the aim of providing practical guidelines for how these models should be defined, scaled and implemented. We show how these suggestions can be implemented in two examples, on disease mapping.

1. Introduction

Conditional Autoregressive (CAR) models are widely used to represent local dependency between random variables, with numerous applications in spatial statistics, disease mapping (Lawson, 2013; Wakefield, 2007) and imaging (Besag et al., 1991). This paper discusses CAR models for disconnected graphs and provides specific recommendations in order to implement them in practice. Through this note, we will illustrates the proposed methods using a classic example on disease mapping, leaving the straightforward generalisation to the reader.

Disease mapping concerns the study of disease risk over a map of geographical regions. Let assume a study area partitioned in n non overlapping regions, indexed by $i = 1, \ldots, n$, with y_i the number of cases observed for a given disease in region i. If the disease is rare, y_i can be assumed as Poisson distributed with mean $\theta_i = E_i r_i$, where E_i is the expected number of cases (computed on the basis of the demographic characteristics of a reference population) and r_i is the relative risk associated with living in region i; $r_i = 1$ means no augmented risk with respect to the average in the whole study area, $r_i > 1$ ($r_i < 1$) indicates higher (lower) risk than average.

As an example, in ecological regression studies investigating the relationship between pollution, say z, and health, the risks might be modelled as

$$\log(r_i) = \alpha + \beta z_i + x_i$$

where α is the log baseline risk and β is the effect of pollution. The term x_i is a random effect capturing residual extra variability in i, possibly due to unobserved risk factors. Some of these unobserved may vary smoothly over space, inducing spatial structure in the residuals x_i's. This structure can be modelled using a CAR prior.

The definition of a CAR model starts by specifying a graph. The graph consists of a collection of nodes and edges representing, respectively, regions and neighbouring relationships between them. The number of nodes deter-

* Corresponding author.
E-mail address: a.freni-sterrantino@imperial.ac.uk (A. Freni-Sterrantino).

https://doi.org/10.1016/j.sste.2018.04.002
1877-5845/

mines the size of the graph. For the sake of a general definition, in this paper a graph is seen as a collection of nodes belonging to one or more *connected components*.

Within a connected component, each node is connected to at least another node through a path (i.e. a set of contiguous edges). Therefore, we say a graph is *connected* if it consists of one connected component of size larger than one (e.g. a 'mainland' component). A graph is *disconnected* if it is not connected, meaning that it consists of more than one connected component of any size (e.g. a 'mainland' component of size n plus an 'island' component of size 1).

Within a connected graph, specification of neighbouring relationships is clear, as each node has at least one neighbour. In this case, the definition of a CAR model follows straightforwardly from specification of the full conditionals (Besag et al., 1991)

$$x_i|\boldsymbol{x}_{-i}, \kappa \sim \mathcal{N}\left(\sum_{j:i\sim j} x_j/n_i, (n_i\kappa)^{-1}\right), \qquad i = 1, \ldots, n. \quad (1)$$

where $\boldsymbol{x}_{-i} = \{x_k, k \neq i\}$, $i \sim j$ means i and j are neighbouring nodes and n_i is the number of neighbours of i. The precision parameter κ regulates the degree to which x_i is shrunk to the local mean $\Sigma_{j:\, i\sim j} x_j/n_i$. Note that the variance $Var(x_i|\boldsymbol{x}_{-i})$ is inversely proportional to the number of neighbours n_i. Therefore, both the prior assigned to κ and the structure of the graph play a role in determining the shrinkage properties of the CAR prior.

The graph structure is essentially defined through conditional independence assumptions between the nodes. In particular, prior (1) specifies that the x_i's are conditionally independent given the information in neighbourhood x_j, $i \sim j$. These assumptions are reflected in the precision structure of the joint distribution for $\boldsymbol{x} = (x_1, \ldots, x_n)^{\mathsf{T}}$, derived from the full conditionals in (1),

$$\pi(\boldsymbol{x}|\kappa) = \left(\frac{\kappa}{2\pi}\right)^{(n-1)/2} |\boldsymbol{R}|_*^{1/2} \exp\left(-\frac{\kappa}{2}\sum_{i\sim j}(x_i - x_j)^2\right), \quad (2)$$

where the summation is over the set of all pairs of neighbours, $i \sim j$, and $|\cdot|_*$ represents the generalised determinant, calculated as the product of the non zero eigenvalues. Model (2) is a multivariate Gaussian with zero mean and $n \times n$ precision $\boldsymbol{Q} = \kappa\boldsymbol{R}$, where $\boldsymbol{R}$ is a matrix representing the neighbourhood structure of the model:

$$R_{ij} = \begin{cases} n_i & i = j \\ -1 & i \sim j \\ 0 & \text{otherwise}, \end{cases} \quad (3)$$

The conditional independence assumptions between the nodes can be checked immediately looking inside $\boldsymbol{R}$: if $R_{ij} = 0$ then x_i and x_j are conditionally independent, given all the other variables $\{x_k: k \neq i, k \neq j\}$ (Rue and Held, 2005). The structure matrix in (3) has rank equal to $n - 1$, thus density (2) is improper. As a consequence of this, the overall mean is unspecified in Eq. (2) and can be identified only when adding a linear constraint, such as $\sum_i x_i = 0$. For this reason, this type of prior is referred to as intrinsic CAR. The purpose of using an intrinsic CAR prior is to borrow strength of information between neighbours, in the disease mapping example this yields a smoothed map for the risk.

In some applications the graph might be disconnected. A typical example is a graph made of two connected components, one of size larger than 1 (e.g. mainland) and the other of size 1 (e.g. island; components of size 1 will be denoted as singletons hereafter). In this case, a direct application of the intrinsic CAR (2) implies that $n_i = 0$, if i is a singletons. This yields a constant prior for x_i, i.e. the singleton random effect is, at prior, allowed infinite variance. This poses a general issue about interpretation of the prior for κ, that may be different in the various connected components of the graph. These issues will be discussed in details in Section 3, were we will introduce the general case of a disconnected graph, with several connected components of any size.

Motivated by the issues involved in a direct application of (2) to a disconnected graph, we propose a solution based on rescaling the precision matrix $\boldsymbol{Q}$ to have similar shrinkage properties, at prior, in each connected component. The rescaling procedure is drawn by ideas in Sørbye and Rue (2014). With our new definition the effect of the graph on the shrinkage properties (defined by the prior for κ) is marginalized out. Therefore, κ has a clear interpretation as a smoothing parameter, regulating the degree to which x_i will shrink to a local mean, if i has neighbours, and to a global mean, if i has no neighbour. In literature there is a lack of attention (Knorr-Held, 2002) on the definition and/or proprieties of a CAR for disconnected graphs, the only reference on this topic is Hodges et al. (2003) who discuss the form of the normalizing constant in (2). On the practical side, the GeoBUGS manual (Spiegelhalter et al., 2002) offers some guidelines in the case of a graph containing singletons, with a default option which is to set x_i to zero, if i is a singleton. Note that this is equivalent to enforcing a sum-to-zero constraint ($x_i = 0$) on each singleton random effect.

An alternative approach is to *correct* the graph, i.e. to remove the singletons by connecting islands to mainland. In our opinion neither of these strategies address the issue in a satisfactory way: the first one adds new constraints, the second one essentially corrupts the graph. We would like to stress that the definition of the graph is part of the modelling process, therefore editing new edges between islands and mainland is only appropriate in cases where borrowing strength of information between them is needed. Changing the graph is inappropriate every time the application at hand requires the original graph, i.e. when borrowing strength of information between islands and mainland is not a sensible choice.

The rest of the paper is organized as follows. Section 2 introduces formally the concept of an intrinsic CAR model defined with respect to a graph. In Section 3, we revise scaling of the precision matrix of an intrinsic CAR model defined with respect to a connected graph, following Sørbye and Rue (2014). In Section 4 we discuss the issues caused by direct application of (2) in the case of a disconnected graph. We then outline recommendations on how the model should be scaled in this case. Section 5 deals with linear constraints and computation