

Multivariate hierarchical analysis of car crashes data considering a spatial network lattice

Andrea Gilardi¹  | Jorge Mateu²  | Riccardo Borgoni³  |
Robin Lovelace⁴ 

¹Department of Statistics and Quantitative Methods (DISMEQ), University of Milan - Bicocca, Milan, Italy

²Department of Mathematics, University Jaume I, Castellon, Spain

³Department of Economics, Management and Statistics (DEMS), University of Milan - Bicocca, Milan, Italy

⁴Institute for Transport Studies (ITS), University of Leeds, Leeds, UK

Correspondence

Andrea Gilardi, Department of Statistics and Quantitative Methods, University of Milan - Bicocca, Milan, Italy.
Email: andrea.gilardi@unimib.it

Present address

Andrea Gilardi, Department of Economics, Management and Statistics (DEMS), University of Milan - Bicocca, Milan, Italy

Abstract

Road traffic casualties represent a hidden global epidemic, demanding evidence-based interventions. This paper demonstrates a network lattice approach for identifying road segments of particular concern, based on a case study of a major city (Leeds, UK), in which 5862 crashes of different severities were recorded over an 8-year period (2011–2018). We consider a family of Bayesian hierarchical models that include spatially structured and unstructured random effects to capture the dependencies between the severity levels. Results highlight roads that are more prone to collisions, relative to estimated traffic volumes, in the north-west and south of city centre. We analyse the modifiable areal unit problem (MAUP), proposing a novel procedure to investigate the presence of MAUP on a network lattice. We conclude that our methods enable a reliable estimation of road safety levels to help identify ‘hotspots’ on the road network and to inform effective local interventions.

KEYWORDS

Bayesian hierarchical models, car crashes data, MAUP, multivariate modelling, network lattice, spatial networks

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society.

1 | INTRODUCTION

Road casualties have been described as a global epidemic, representing the leading cause of death among young people worldwide (see, e.g. MacKay 1972; Nantulya & Reich, 2002). Car crashes and other types of collisions are responsible for more than 1 million deaths each year (1,250,000 in 2015, 17 deaths per 100,000 people), as reported by World Health Organization (2018). In high-income nations, such as the United Kingdom (UK), the roads are safer than the global average, but car crashes are still the cause of untold suffering. According to the statistics published by the UK's Department for Transport (DfT) in the *Annual report on Road Casualties in Great Britain* (Department for Transport, 2020), approximately 153,000 road traffic collisions resulting in casualties were recorded in 2019, 5% lower than 2018 and the lowest level since records began, in 1979. Nevertheless, the DfT estimates that approximately 33,648 people were killed or seriously injured (KSI) in 2019, and while this number is slightly lower than in 2018, the decaying rate has been getting lower and lower starting from 2010. These figures are worrisome considering that car occupant fatality rates are particularly high in the 17–24 age band (Department for Transport, 2020, p. 17)

To tackle the flattening trend in the KSI rate over the past decades, a range of interventions are needed, and analytical approaches can help prioritise them. This paper presents a statistical model to identify street sections with anomalously high car crashes rates, and to support police responses and cost-effective investments in traffic calming measures (PACTS, 2020).

Statistical models of road crashes have become more advanced over time. In the 1990s, models tended to consider only the discrete and heterogeneous nature of the data (Miaou, 1994; Miaou & Lum, 1993; Shankar et al., 1995), omitting spatial characteristics. More recent statistical models of crash data include consideration of crash location in two-dimensional space, with three main advantages for road safety research (El-Basyouny & Sayed, 2009). First, consideration of space allows estimating appropriate measures of risk (such as expected counts, rates or probabilities) at different levels of resolution and the subsequent ranking of geographical areas to support local interventions. Second, spatial dependence can be a surrogate for unknown, potentially unmeasured (or unmeasurable) covariates; adjusting for geographical location can reduce model misspecification (Cressie, 1993; Dubin, 1988). Third, the spatial dimension can be used to take advantage of autocorrelation in the relevant variables, borrowing strength from neighbouring sites and improving model parameter estimation.

Road crash data sets, which are typically available on a single accident basis, can be spatially aggregated in two main ways: administrative zones (such as cantons, census wards or regions) or street network features (either as contiguous segments or divided into corridors and intersections). In both cases, the spatial support is a lattice, that is a countable collection of geometrical units (polygons or lines, respectively), possibly supplemented by a neighbourhood structure. Several papers addressed the statistical modelling of crash frequencies at the areal level, based on available zoning systems in the study region (see, e.g. Aguero-Valverde & Jovanis, 2006; Boulieri et al., 2017; Miaou et al., 2003; Noland & Quddus, 2004). The second approach has gained in popularity in recent years, with a number of papers analysing road crash events aggregated to the street level (see, e.g. Aguero-Valverde & Jovanis, 2008; Miaou & Song, 2005; Wang et al., 2009), as detailed in several review papers (Lord & Mannering, 2010; Savolainen et al., 2011; Ziakopoulos & Yannis, 2020). In reference to street level data, we note that there has been a recent surge of research for spatial point patterns living on networks (Baddeley et al., 2021; Cronie et al., 2020; Rakshit et al., 2019).

Both zone and network level approaches have advantages, notably computational requirements for the former and spatially disaggregated results for the latter. Given that computational

resources are less of a constraint in the 2020s than they were in previous decades, and the fact that it is the nature of roads (not zones) that is responsible for crashes, we argue that road segments are the more appropriate aggregation units for the analysis of road crash data. Network analysis can be used to bring attention to specific segments, and, for these reasons, the models presented in the next sections were developed considering a network lattice.

Aggregation, for example number of crashes per road segment, enables comparison between different road segments. However, spatial aggregation also leads to a well-known problem in geographical analysis, the modifiable areal unit problem (MAUP), first described in Openshaw and Taylor (1981): the size of the spatial units impacts on the statistical analysis, influencing and possibly biasing modelling choices and results. Hence, conclusions drawn at one scale of spatial aggregation might not necessarily hold at another scale or be somehow different. The MAUP has been mainly ignored in the road safety literature and, as reported by Xu et al. (2018) and Ziakopoulos and Yannis (2020, p. 21), it is mentioned only in a handful of recent papers (Abdel-Aty et al., 2013; Briz-Redón et al., 2019b; Ukkusuri et al., 2012; Zhai et al., 2019), which explore the impact of changing the areal zoning system (e.g. TAZ, block groups and census tracts) on parameter estimates, significance and hotspot detection. Only one early paper (Thomas, 1996) could be found exploring the impact of MAUP on crash data at the street segment level, albeit only in terms of summary statistics of aggregated counts. To assess the MAUP effect on network data modelling, we employed an algorithm to modify the structure of a road network, merging contiguous segments in the same corridor while preserving the geometrical properties of the network (Padgham, 2019). Then, we compared the results obtained with the two different network configurations. To the best of our knowledge, this is the first attempt at exploring and estimating the presence and the magnitude of MAUP in models that consider a network lattice.

Finally, we note that systems of collision classification present a multivariate nature (Kirk et al., 2020). The occurrences of different severity degrees can be correlated to each other, and their spatial dynamics can be potentially interdependent. Hence, it is necessary to account for correlations between crashes counts at different levels of severity. We consider two types of accidents: *slight* and *severe*. The severe class is very sparse in the data set at hand, hence modelling both types of accidents simultaneously allows to borrow strength from the existing correlations and improves estimates. We underline that the methodology for classifying the severity level of a car crash in the UK has been modified starting from 2016, adopting the injury-based systems called CRASH and COPA (Braunholtz & Elliott, 2019). All police forces are gradually adopting these new reporting systems in England, and the Office for National Statistics (ONS) developed a logistic regression model to correct the severity levels between different years and classification systems. The data used in this paper have been adjusted using the procedures developed by ONS (Department for Transport, 2020, p. 38–41).

Following ideas introduced in Barua et al. (2014), we consider a range of competing models, developed in a full Bayesian hierarchical paradigm. This approach allows one to encompass complex structures of spatial dependence in a quite natural way. Spatially structured random effects are defined using both intrinsic multivariate conditional autoregressive (IMCAR) and proper multivariate conditional autoregressive (PMCAR) priors (Besag, 1974; Mardia, 1988; Martínez-Beneito & Botella-Rocamora, 2019; Palmí-Perales et al., 2021). Lord and Mannering (2010), Savolainen et al. (2011), and Ziakopoulos and Yannis (2020) review other examples of multivariate car crashes models on a network.

The case study is the metropolitan area of Leeds (population 800,000) in North England. We accessed Ordnance Survey data on major roads (3661 segments, total length 450 km), creating a spatial network substantially larger than previous studies, many of which report findings on only a few roads, with Borgoni et al. (2020) representing a notable exception albeit with a

simpler spatial structure. We present results for an entire metropolitan area, approximating more closely the level at which road policing activities and investments in road safety interventions are prioritised. The scale of the case study presented several computational challenges, and, in terms of Bayesian parameter estimation, we used the computationally efficient integrated nested laplace approximation (INLA) approach instead of Markov chain Monte Carlo (MCMC) sampling (Lindgren & Rue, 2015; Rue et al., 2017).

The rest of the paper is organised as follows. In Section 2, the data sources are described. In Section 3, the statistical methodology adopted in this paper is discussed in detail. In Section 4, the main results of the paper are presented whereas model criticism and further model discussion, such as MAUP analysis, are provided in Section 5. Conclusions, in Section 6, end the paper.

2 | DATA

The data sets analysed in this paper came from several different sources and required a number of preprocessing steps before they could be made into a structure suitable for a statistical analysis. The study region was defined as the middle super output area (MSOA) zones within the local authority of Leeds. The City of Leeds was selected because it is a car-dependent city with a large network of major roads that approach the city centre (the city was dubbed the 'motorway city of the 70s') and would therefore be expected to be a place where road safety could be improved. Leeds is part of West Yorkshire and accounts for approximately 40% of all car crashes in the region.

Origin-destination data from the 2011 UK Census were used to estimate traffic volumes, to provide an estimate of exposure, with traffic volumes used as part of the denominator of the statistical models presented in Section 3. The road network was obtained from Ordnance Survey, covering all major roads in Leeds. We matched the network and the MSOAs using an overlay operation. We associated all car crashes that occurred in the city of Leeds from 2011 to 2018 with the nearest point on the road network, counting the occurrences in each street segment. Finally, a set of socio-economic variables obtained from 2011 UK Census data were included in the statistical models as fixed effects.

2.1 | MSOA zones

There are 6791 MSOAs in England, 299 of which belong to the West Yorkshire region and 107 of which constitute Leeds. These were accessed from the github-page¹ of Propensity Cycle Tool (Lovelace et al., 2017). The MSOAs represent the starting point for all the following steps, and they are mapped in Figure 1 as grey polygons for the West-Yorkshire, and as dark-green polygons for the City of Leeds. The inset map is used to locate the study-area in the British territory.

2.2 | Traffic flow

The *traffic flow* data represent the *commuting journeys* from home to workplace using several modes of transport, such as train, bus, bike and motorcycle. The data were collected during the 2011 Census at the individual level, and then aggregated at the MSOA level. The UK Data Service shares the flow data through the WICID interface as cross-tables reporting the flows between all

¹<https://github.com/npct/pct-outputs-regional-R/tree/master/commute>, last access on 06/2020.

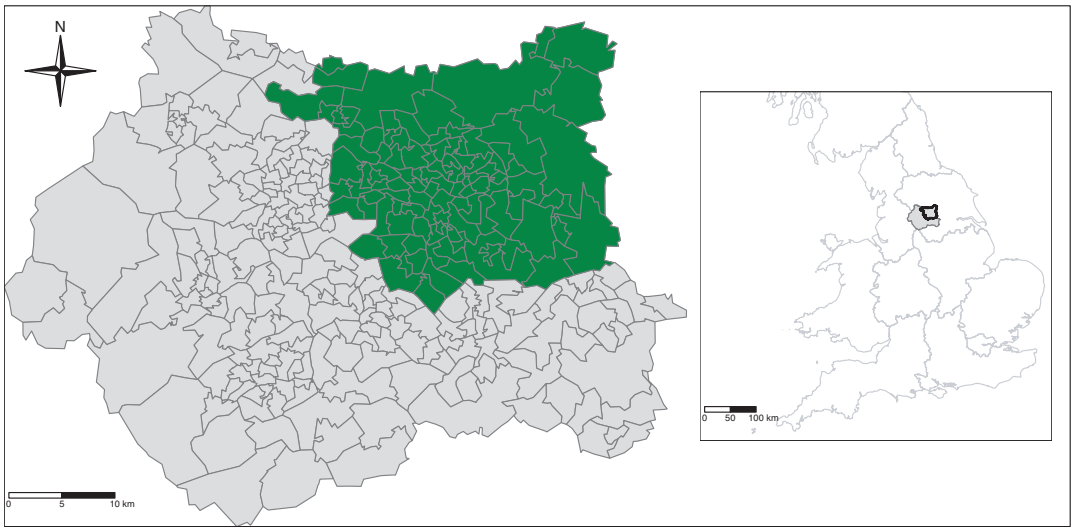
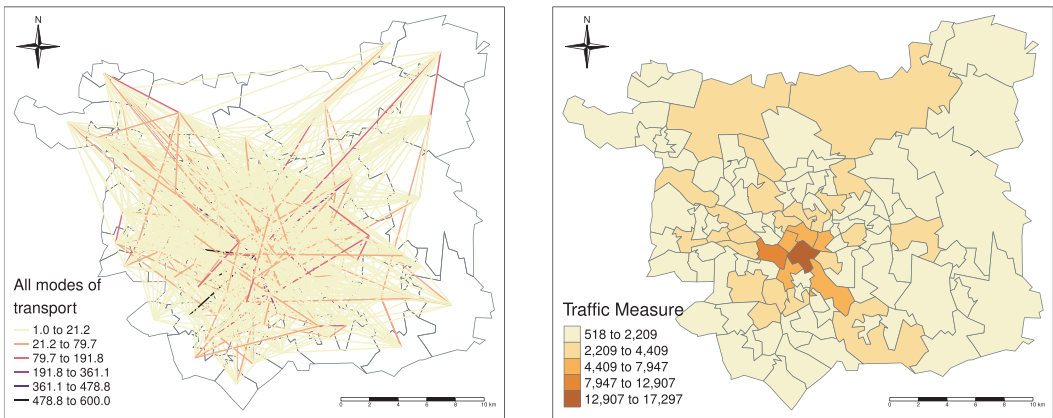


FIGURE 1 The grey polygons show the middle super output areas in West-Yorkshire region, while the dark-green area highlights the city of Leeds. The inset map locates the position of the study-area with respect to England



(a) Random sample of one thousand Origin-Destination data representing the daily flows between MSOAs in Leeds.

(b) Estimates of the new traffic measure. The white star represents Leeds City Centre.

FIGURE 2 Raw and modified traffic flows in the area of Leeds. The map on the right highlights several contiguous middle super output areas that correspond to the arterial thoroughfares that are used to reach the City Centre

pairs of a predefined set of MSOAs (UK Data Service Census Support, 2014). We considered the commuting flows in the region of Leeds for all possible modes of transport. Figure 2a shows a random sample of 1000 *traffic flows* (out of 10,536 in total) between the centroids of the MSOAs in Leeds, coloured according to the number of daily commuters.

Raw WICID data, however, ignore that people may travel to their workplace through several MSOAs. For this reason, we calculated a new traffic measure using the following procedure. Starting from the MSOAs, we defined a graph where the vertices are the centroids of each area, and the

edges connect neighbouring areas. Then, we estimated the shortest path for all commuting journeys downloaded from WICID and assigned to each MSOA a value that is equal to the number of all raw traffic measures going through the area. These values represent the new traffic measures and are displayed in Figure 2b. A similar approach was also adopted in Boulieri et al. (2017), and we refer to the references therein for more details. The raw data flows and the MSOAs polygons were downloaded using the R package `pct` (Lovell & Hama, 2020).

2.3 | Road network

The *road network* was built using data downloaded from Ordnance Survey (OS), an agency that provides digital maps and other services for location-based products (Ordnance Survey, 2020). We downloaded the `Vector OpenMap Local` data for a geographical region covering Leeds, selected the `Roads` and `Tunnels` layers, and filtered the streets that belong to the city. Ordnance Survey represents all the streets of a road network as the union of a finite set of segments, and it includes additional fields such as the road name or the street type. These segments represent the elementary units for the statistical analysis described in Section 3.

The road network downloaded from OS is composed of approximately 50,000 segments. The OS network was simplified using the following procedure. We first selected only the major roads, such as the *Motorways*, *Primary Roads* and *A roads*. They represent 3668 segments, that is less than 10% of the total road network, but more than 50% of all car crashes registered during 2011–2018 occurred in their proximity. The output of this procedure is a road network composed by a big cluster of connected streets, displayed in Figure 4a, and several isolated segments of small groups of road segments (which are also called *islands*), created by the exclusion of their links to the other roads. These small clusters can be problematic from a modelling perspective since they produce a not-fully connected network (see Freni-Sterrantino et al., 2018; Hodges et al., 2003 and the properties of ICAR and IMCAR distributions explained in Section 3), so we implemented an algorithm to further simplify the road network and remove them. This algorithm is based on the dual representation of a road network as a geographical entity, composed by points and lines, and a graph object, with nodes and edges (Gilardi et al., 2020; Marshall et al., 2018; Porta et al., 2006). More precisely, we created a graph whose vertices correspond to the street segments of the road network, and we defined an edge for each pair of spatial units sharing a point at their boundaries. This graph uniquely determines a (sparse) adjacency matrix amongst the spatial units (i.e. the road segments) that summarises the graph dimension of the road network. We sketched a toy example in Figure 3, representing the idea behind the dual representation of a road network and the definition of the adjacency matrix.

Using the graph and the adjacency matrix, we excluded all road segments that did not belong to the main cluster. In particular, we dropped seven segments spread across different parts of the city with a total length of approximately 380 m (out of 450 km). That also implied that we removed from the analysis three car crashes that occurred in those segments (see below). However, given the extremely small fraction of records discarded from the data set, we expect that excluding these observations does not influence the final results. It should be stressed that this procedure creates a fully connected network, which has relevant benefits on the rank-deficiency problem of the ICAR models described in Section 3.

In the end, the road network is composed of 3661 units, and it is shown in Figure 4a, where the segments are coloured according to their road types. The segments have different lengths, ranging from 0.1 m to 2597 m, with an average value of 118 m ($sd = 178$ m). As explained in





FIGURE 3 Graphical example showing the dual nature of a road network. Left: map showing the geographical dimension. Each segment is coloured and labelled using a different ID and colour respectively. Right: adjacency matrix of the graph associated with the road network. Each vertex corresponds to a segment whereas an edge connects two vertices if they share one boundary point. For example, segments 1 and 2 are not neighbours since they do not share any point at the boundaries, even if they intersect each other. This situation may occur at bridges or overpasses



(a) Road network and car crashes in Leeds.

(b) Choropleth map of severe counts.

FIGURE 4 The map on the left represents the road network in Leeds. Each segment is coloured according to its OS classification. The black dots represent the car crashes. On the right, we report a choropleth map displaying severe car crashes counts

the next Section, these lengths are used in the exposure parameter of the statistical models in order to guarantee a comparable rate amongst the network units. Finally, since the street network and the MSOAs are spatially misaligned, they were matched using an overlay operation: each road segment was assigned to the MSOA that intersects the largest fraction of the segment. This procedure allows us to assign a traffic estimate to each road segment, which will be used in the exposure parameter (along with the segments' lengths) in the statistical models considered below.

2.4 | Road traffic collision data

We analysed all road traffic collisions that occurred between 1 January 2011 and 31 December 2018 in the MSOAs pertaining to the City of Leeds, which involved at least one personal injury, occurred on public roads and became known to the Police forces within thirty days of the occurrence. The geographical coordinates of the crashes are provided at 10m or less resolution in the UK's official coordinate reference system (the Ordnance Survey National Grid, EPSG code 27700) depending upon the year to which the data refer to (Department for Transport, 2011). Hence, we adopted this value as a threshold to account for the potential misalignment between the event locations and the network and excluded all those events that occurred farther than ten meters from the closest segment in the simplified road network since they might be related to other streets. The data were downloaded from the UK's official road traffic casualty database, called *STATS19*, using the homonym R package (Lovelace et al., 2019).

STATS19 data also report the severity level of each casualty using one of three possible categories: *fatal*, *serious* or *slight*. In particular, a collision is classified as *fatal* if it involved a human casualty whose injuries caused his death less than thirty days after the accident; *severe* if at least one person was hospitalised after the accident or recorded a particular type of injury (like concussions or severe cuts), and *slight* in all the other cases. We harmonised the severity levels for different years and police forces using the CRASH methodology (Braunholtz & Elliott, 2019; Department for Transport, 2020), and we decided to aggregate *serious* and *fatal* levels since *fatal* crashes represent approximately 1% of the total number of car accidents. Henceforth, we will refer to *serious* or *fatal* crashes as *severe* accidents.

As mentioned above, we focussed on car crashes that occurred in a public highway and involved personal injuries. In general, however, there is no obligation for people to report all personal injury accidents to the police and, for this reason, a proportion of non-fatal and no-injury casualties remain unknown to the police (Department for Transport, 2020). Hence, we acknowledge that *slight* accidents counts considered hereinafter may suffer under-reporting to some extent. We refer to Savolainen et al. (2011, section 2.1) and references therein for more details.

The final sample is composed of 5862 events, and they are reported as black dots in Figure 4a. Then, we projected all crashes to the nearest point on the road network, and we counted the number of *slight* or *severe* occurrences on all street segments. We decided to ignore the temporal dimension since *severe* crashes counts present an extreme sparsity, with more than 80% of zero counts during 2011–2018. Moreover, 40% of all segments registered no car crashes during the study period, while another 40% reported two or more crashes. These numbers highlight a common temporal trend between the 8 years, and we refer the interested reader to the supplementary material for a space-time representation. The map in Figure 4b shows the spatial distribution of *severe* crashes counts.

2.5 | Socio-economic covariates

The statistical models described in Sections 3 and 4 include two socio-economic covariates that were obtained from the 2011 UK Census data and downloaded from Nomis website.² In particular, we consider the population density (given by the ratio of the number of inhabitants in a given region and its area in squared metres) and the employment rate (given by the ratio of employed

²<https://www.nomisweb.co.uk/>. Last access: 06/2021.

population between 16 and 64 and total population between 16 and 64). Both covariates were obtained at the LSOA level (a more detailed aggregation level than MSOA) and, for the same reasoning as before, they were matched with the street segments using an overlay operation.

3 | STATISTICAL METHODOLOGY

We first focus on the definition of a three-level hierarchical model structure which is shared among all the alternative specifications considered below. Then, we introduce two baseline models that serve as benchmarks and starting points for the other specifications. Thereafter, four different extensions to the baseline models are introduced. Finally, some techniques used for model comparison are discussed. The common theme behind all the six alternatives is the presence of spatially structured and unstructured multivariate random effects.

Let Y_{ij} , $i = 1, \dots, n$, represent the number of car crashes that occurred in the i -th road segment with severity level j , $j = 1, \dots, J$. In this paper, we consider two possible severity levels, a car crash being either *severe*, $j = 1$, or *slight*, $j = 2$.

In the first stage of the hierarchy, we assume that

$$Y_{ij} | \lambda_{ij} \sim \text{Poisson}(E_i \lambda_{ij}),$$

where E_i is an exposure parameter and λ_{ij} represents the car crashes rate in the i th road segment for severity level j . As mentioned in the previous section, the exposure parameter, E_i , is given by the product of two quantities (Wang et al., 2009): the segment's length and the estimate of traffic flow (see Section 2). The exposure accounts for the fact that a longer street segment has a higher collision risk than a shorter one, guaranteeing that comparable rates amongst units (segments) are preserved. At the same time, the traffic flow estimates allow different segments of the network to be 'weighted' differently, being more exposed to accidents those segments with a higher traffic flow, all the rest being fixed.

At the second stage of the hierarchical model, a log-linear structure on λ_{ij} is specified. We assume that

$$\log(\lambda_{ij}) = \beta_{0j} + \sum_{m=1}^M \beta_{mj} X_{ijm} + \theta_{ij} + \phi_{ij},$$

where β_{0j} represents a severity-specific intercept, $\{\beta_{mj}\}_{m=1}^M$ is a set of coefficients, $(X_{ij1}, \dots, X_{ijM})$ is a collection of M covariates, ϕ_{ij} is a spatially structured random effect and θ_{ij} represents a normally distributed error component. The third stage that completes the hierarchical model is the specification of prior and hyperprior distributions. We assigned a vague $N(0, 1000)$ prior to β_{mj} , $m = 0, \dots, M$. The two random effects, namely θ_{ij} and ϕ_{ij} , represent the unstructured and structured spatial components and are defined differently in different models as discussed below. Hereafter, we follow the notation used in Martínez-Beneito and Botella-Rocamora (2019).

3.1 | Baseline models: independent spatial and unstructured effects

The two baseline models are defined considering multivariate spatial and unstructured random effects with independent components. More precisely, a bivariate Gaussian prior with independent components is assigned to $(\theta_{i1}, \theta_{i2})$ for both baseline models:

$$(\theta_{i1}, \theta_{i2}) \sim N_2 \left(\mathbf{0}, \begin{bmatrix} \sigma_{\theta_1}^2 & 0 \\ 0 & \sigma_{\theta_2}^2 \end{bmatrix} \right), \quad i = 1, \dots, n. \quad (1)$$

We assigned a Gamma hyperprior with parameters 1 (shape) and 0.00005 (inverse scale) to the inverse of $\sigma_{\theta_1}^2$ and $\sigma_{\theta_2}^2$, that is the precisions.

The spatially structured term in the first baseline model was defined using an *independent intrinsic multivariate conditional autoregressive* (IIMCAR) prior, whereas, for the second model, we adopted an *independent proper multivariate conditional autoregressive* (IPMCAR) prior. The IIMCAR and IPMCAR distributions are briefly introduced hereafter, starting from their classical univariate counterparts, namely the ICAR and PCAR distributions.

Univariate spatial random effects are traditionally modelled using a prior that belongs to the family of *conditional autoregressive* (CAR) distributions (Besag, 1974). Given a random vector $\boldsymbol{\phi} = (\phi_1, \dots, \phi_n)$, the *intrinsic conditional autoregressive* (ICAR) distribution, which is a particular case of the CAR family, is usually defined through a set of conditional distributions (Besag & Kooperberg, 1995):

$$\phi_i | \{\phi_{i'}, i' \in \partial_i\}; \sigma^2 \sim N \left(m_i^{-1} \sum_{i' \in \partial_i} \phi_{i'}, \frac{\sigma^2}{m_i} \right), \quad i = 1, \dots, n, \quad (2)$$

where ∂_i and m_i denote, respectively, the indices and the cardinality of the set of neighbours for spatial unit i . These quantities are defined through a sparse binary symmetric neighbourhood matrix \mathbf{W} with dimensions $n \times n$ that summarises the spatial relationships in the region of study. We built it taking advantage of the dual representation of a road network as a spatial and a graph object (see Porta et al. (2006) and Section 2). More precisely, \mathbf{W} is the adjacency matrix of a graph whose vertices correspond to the street segments of the road network and the edges identify a shared point at the boundaries of two spatial units. This procedure defines a *first-order* neighbourhood matrix. *Second-* and *third-order* neighbourhood matrices are defined iteratively in the same way.

It is possible to prove that the prior defined by Equation (2) suffers from rank-deficiency problems, that are usually fixed by imposing a set of sum-to-zero constraints on the vector $\boldsymbol{\phi}$, one for each group of connected segments in the graph of the road network (Hodges et al., 2003). In this paper, we deal with a fully connected road network (see the pre-processing procedures detailed in Section 2), so we always had to fix only one set of constraints.

The *proper conditional autoregressive* (PCAR) distribution is another member of the CAR family and it is usually defined as follows:

$$\phi_i | \{\phi_{i'}, i' \in \partial_i\}; \sigma^2, \rho \sim N \left(\rho \left(m_i^{-1} \sum_{i' \in \partial_i} \phi_{i'} \right), \frac{\sigma^2}{m_i} \right), \quad i = 1, \dots, n, \quad (3)$$

where ∂_i and m_i are defined as for the ICAR distribution and ρ is a parameter controlling the strength of spatial dependence, usually called *spatial autoregression coefficient* (Cressie, 1993). It is possible to prove that the joint distribution defined by Equation (3) is proper if $|\rho| < 1$, hence there is no need to set any sum-to-zero constraint in this case. The ICAR prior can be seen as a limit case of the PCAR distribution with $\rho \rightarrow 1$, analogously to the relationship between autoregressive and random-walk models in time series models (Botella-Rocamora et al., 2013).

The family of *multivariate conditional autoregressive* (MCAR) distributions was firstly introduced by Mardia (1988), extending the ideas of Besag (1974) to the multivariate case. Given a random matrix $\Phi = (\phi_{ij})$, which is defined for $i = 1, \dots, n$ units and $j = 1, \dots, J$ levels, the *intrinsic multivariate conditional autoregressive* (IMCAR) distribution is a particular case of the MCAR family, defined through a set of multivariate conditional distributions (Martínez-Beneito & Botella-Rocamora, 2019):

$$\Phi_i | \text{vec}(\Phi_{-i}); \Omega \sim N_J \left(m_i^{-1} \sum_{i' \in \partial_i} \Phi_{i'}^T; m_i^{-1} \Omega^{-1} \right). \quad (4)$$

The terms Φ_i and Φ_{-i} denote, respectively, the i th row of Φ and the matrix obtained by excluding the i th row from Φ . The vec operator is used for row-binding the columns of a matrix, meaning that $\text{vec}(\Phi_{-i}) = (\Phi_{-i1}^T, \dots, \Phi_{-iJ}^T)^T$. The elements m_i and ∂_i are defined as before, through the adjacency matrix \mathbf{W} of the graph associated to the road network, and they represent the spatial dimension of the IMCAR distribution. The $J \times J$ precision matrix Ω is used to model the associations between pairs of levels in the same road segment i , and it acts as a multivariate extension of the parameter σ^2 in Equation (2).

This distribution suffers from the same rank-deficiency problems as its univariate counterpart, which are usually solved by imposing appropriate sum-to-zero constraints. The number of restrictions is equal to the number of clusters in the graph of the road network times the number of levels in the multivariate setting. The pre-processing operations that we performed on the network data (see Section 2) imply that we always have to set only J sum-to-zero constraints.

The IIMCAR distribution is a particular case of (4), which is obtained by setting $\Omega^{-1} = \text{diag}(\sigma_{\phi_1}^2, \dots, \sigma_{\phi_J}^2)$. More precisely, if we assume $J = 2$ as we do in this paper, then IIMCAR is defined by the following set of multivariate conditional distributions:

$$\Phi_i | \text{vec}(\Phi_{-i}); \sigma_{\phi_1}^2, \sigma_{\phi_2}^2 \sim N_2 \left(m_i^{-1} \sum_{i' \in \partial_i} \Phi_{i'}^T; m_i^{-1} \begin{bmatrix} \sigma_{\phi_1}^2 & 0 \\ 0 & \sigma_{\phi_2}^2 \end{bmatrix} \right). \quad (5)$$

In Equation (5), we are assuming independence between the 2 levels, and this implies that the IIMCAR distribution is equivalent to two independent ICAR distributions, one for each level.

Analogously to the univariate case, the *proper multivariate conditional autoregressive* (PMCAR) distribution is a particular case of the MCAR family characterised by the following set of multivariate conditional distributions:

$$\Phi_i | \text{vec}(\Phi_{-i}); \rho, \Omega \sim N_J \left(m_i^{-1} \rho \sum_{i' \in \partial_i} \Phi_{i'}^T; m_i^{-1} \Omega^{-1} \right). \quad (6)$$

The strength of the spatial dependence is controlled by ρ (as for the univariate PCAR distribution) and all the other parameters are defined as before. It can be proved that the joint distribution defined by Equation (6) is proper if $|\rho| < 1$, although we restricted ourself to $\rho \in (0, 1)$ to avoid some counter-intuitive behaviour of the PMCAR distribution (Miaou & Song, 2005; Wall, 2004).

The IPMVAR distribution is defined as a particular case of Equation (6) with $\Omega^{-1} = \text{diag}(\sigma_{\phi_1}^2, \dots, \sigma_{\phi_J}^2)$. More precisely, if we assume $J = 2$, then IPMVAR is defined through the following set of multivariate conditional distribution:

$$\Phi_{i \cdot} | \text{vec}(\Phi_{-i \cdot}); \rho, \sigma_{\phi_1}^2, \sigma_{\phi_2}^2 \sim N_2 \left(m_i^{-1} \rho \sum_{i' \in \partial_i} \Phi_{i'}^T; m_i^{-1} \begin{bmatrix} \sigma_{\phi_1}^2 & 0 \\ 0 & \sigma_{\phi_2}^2 \end{bmatrix} \right). \quad (7)$$

For the same reasoning as in Equation (5), the IPMCAR distribution is equivalent to J independent PCAR distributions.

Now we can characterise the random effects for the two baseline models. The first model was defined by considering unstructured random effects with a bivariate independent Gaussian prior (1), and spatial random effects with an IIMCAR prior (5). The second one was defined analogously to the first baseline model, but assuming an IPMCAR distribution for the spatial random effects (7). These models assume independence between the two levels both in the spatial and unstructured components, so they were used as benchmarks. In the next sections, we will also refer to the two baseline models using, respectively, the codes (A) and (B). We assigned an improper prior to σ_1^2 and σ_2^2 , the variances in Ω , defined on \mathbb{R}^+ , and a Uniform(0,1) prior to ρ .

Hereafter we introduce two increasingly complex sets of extensions that generalise the baseline models. The first one is characterised by the removal of the independence assumption from the spatially structured random effects, whereas in the second set of extensions, we also relax the independence assumption from the unstructured random effects.

3.2 | Model extensions

3.2.1 | First set of extensions

Starting from the baselines, we defined two new models replacing the IIMCAR and IPMCAR priors with their non-independent multivariate counterparts, the generic IMCAR and PMCAR defined above. If we assume $J = 2$, then the variance-covariance matrix Ω^{-1} in (4) and (6) can be written as

$$\Omega^{-1} = \begin{bmatrix} \sigma_{\phi_1}^2 & \rho_{\phi} \sigma_{\phi_1} \sigma_{\phi_2} \\ \rho_{\phi} \sigma_{\phi_1} \sigma_{\phi_2} & \sigma_{\phi_2}^2 \end{bmatrix}$$

where σ_1^2 and σ_2^2 represent the conditional variances and ρ_{ϕ} represents the correlation coefficient between the two levels in the same spatial unit. These models represent a generalisation of the baselines since we are now taking into account the correlations between different levels in the same road segment. We will also refer to them using, respectively, the codes (C) and (D). Following Palmí-Perales et al. (2021), we assigned a Wishart hyperprior to Ω^{-1} with parameters 2 and \mathbf{I}_2 , that is the identity matrix of size two. The prior distributions on the unstructured random effects were left unchanged with respect to the baselines.

3.2.2 | Second set of extensions

In these models, the independence assumption of the spatially unstructured random effects is removed. More precisely, assuming $J = 2$, we assign a generic bivariate Gaussian prior to the unstructured random effects:

$$(\theta_{i1}, \theta_{i2}) \sim N_2 \left(\mathbf{0}, \begin{bmatrix} \sigma_{\theta_1}^2 & \rho_{\theta} \sigma_{\theta_1} \sigma_{\theta_2} \\ \rho_{\theta} \sigma_{\theta_1} \sigma_{\theta_2} & \sigma_{\theta_2}^2 \end{bmatrix} \right).$$

TABLE 1 Summary of the prior distributions assigned to the random effects in the models introduced in Section 3

ID	Model	Unstructured effect	Spatial effect
(A)	Baseline 1	Independent Gaussian	Independent IMCAR
(B)	Baseline 2	Independent Gaussian	Independent PMCAR
(C)	Extension 1 - Model 1	Independent Gaussian	IMCAR
(D)	Extension 1 - Model 2	Independent Gaussian	PCAR
(E)	Extension 2 - Model 1	Correlated Gaussian	IMCAR
(F)	Extension 2 - Model 2	Correlated Gaussian	PCAR

Parameters $\sigma_{\theta_1}^2$ and $\sigma_{\theta_2}^2$ represent the marginal variances of the unstructured random error, whereas ρ_θ represents their correlation. These models will also be identified using the codes (E) and (F). We assigned a Wishart hyperprior to the variance-covariance matrix with parameters 2 and \mathbf{I}_2 .

Finally, following the ideas introduced in Gelfand and Vounatsou (2003), we also tested an extension of model (F), named model (G), which is characterised by a generalisation of the PMCAR distribution that introduces a separate spatial autoregression coefficient, ρ_j , for each level in Φ . We found that this extension did not improve over model (F). Hence, we will not add more details here and refer the interested reader to the supplementary materials.

The prior distributions adopted for the random effects in the two baselines and their extensions are summarised in Table 1. We also included the IDs that will be used to identify each model in subsequent Tables and Sections.

3.3 | Model comparison

The models proposed in the previous paragraphs were compared using Deviance information criterion (DIC) (Spiegelhalter et al., 2002) and Watanabe-Akaike information criterion (WAIC) (Gelman et al., 2014; Watanabe, 2010). These criteria represent a measure for the adequacy of a model, penalised by the number of effective parameters. In both cases, the lower is the value of the index, the better is the fitting of the model.

4 | RESULTS

We estimated the models previously described using the software INLA (Gómez-Rubio, 2020; Lindgren & Rue, 2015; Rue et al., 2017), interfaced through the homonymous R package (R Core Team, 2020). We used the *simplified Laplace* strategy for approximating the posterior marginals and the *central composite design* strategy for determining the integration points. The code behind the definition of multivariate ICAR and PCAR random effects is defined in the package INLAMSM (Palmí-Perales et al., 2021). It took approximately 30–45 min to estimate each model using a virtual machine with an Intel Xeon E5-2690 v3 processor, six cores, and 32GB of RAM.

4.1 | Fixed effects

The models listed in Table 1 share a common structure for the fixed effect component, with a severity-specific intercept and a set of covariates representing some social and physical characteristics of the road segments. In particular, we considered five severity-specific covariates, namely the two socio-economic variables mentioned in Section 2, a dummy variable recording whether a road segment lays in a dual carriageway street or not, the road type (either *Motorway*, *Primary Road* or *A Road*, according to OS definition), and the edge betweenness centrality measure, which reflects the number of shortest paths traversing each segment (Kolaczyk & Csárdi, 2014). The last covariate can be considered as a proxy for the average vehicle miles travelled (VMT) (Briz-Redón et al., 2019a), whereas the employment rates may represent an indirect measure of local economic activities level. Finally, to improve the stability of INLA algorithms, we scaled all numerical variables to zero mean and unit variance.

Table 2 shows the posterior means and standard deviations for the fixed effects. We first notice that the estimates are stable among the models. The intercept for severe car crashes, β_{01} , is found slightly smaller than β_{02} . This is not surprising since severe accidents are rarer than the slight ones. The coefficients of edge betweenness centrality measures are found close to zero for all models, and their 95% credible interval (not reported in the table) always include the value zero. The Road type parameters represent relative differences with respect to the reference category (i.e. *A Roads*), hence *Motorways* are found less prone to severe and slight car crashes than *A roads*. A similar finding was also reported by Boulieri et al. (2017) for UK data. An analogous interpretation applies to *Primary Roads*. In agreement with other studies (see, e.g. Huang et al. (2017) and references therein), we found that population density significantly correlates to both slight and severe accidents, although its effect is stronger for the former type. On the other hand, the second socio-economic variable was not found significant in any of the outcomes considered in this paper. Unfortunately, more direct measures for local economic activity are not available and are difficult to construct at such detailed spatial resolution. More work is needed to assess the relationship between road safety and local economic levels. Finally, dual carriageway roads have been found significantly less prone to slight car accidents, whereas no impact has been found for severe car crashes. An analogous result was also reported by Castle and Lynam (2008).

4.2 | Random effects

The posterior means and standard deviations for all hyperparameters are reported in Table 3, which reflects the models' nested structure, also summarised in Table 1. We started from two baselines, (A) and (B), with independent random effects, and generalised them until model (F), that presents multiple autocorrelation parameters between the two severity levels.

Models from (A) to (D), which assume independent unstructured random effects, exhibit a degenerate posterior distribution of $\sigma_{\theta_1}^2$, that is the variance of severe random component, and this is possibly due to the severe car crashes sparseness. This problem gets mitigated once the correlation parameter between the two severity levels is included in the model, suggesting that the estimation procedure benefits from the inclusion of a multivariate structure that allows borrowing strength from less rare events.

The estimates of $\sigma_{\theta_2}^2$ and ρ_{θ} are stable among the models, and the correlation parameter is estimated as high as 0.40, suggesting a positive and mildly strong relationship between the two random components. The posterior means for hyperparameters ρ in models (B), (D), and (F), always very close to one, which is not uncommon for this type of models (Carlin & Banerjee, 2003).

TABLE 2 Estimates for the posterior means and standard deviations, in round brackets, of the fixed effects included in the models described in Table 1

ID	Severe crashes						Slight crashes							
	β_{01}	Betw.	Motorways	Roads	Ratio Empl.	Pop. Dens.	Dual Carr.	β_{02}	Betw.	Motorways	Roads	Ratio Empl.	Pop. Dens.	Dual Carr.
(A)	-14.229 (0.092)	0.011 (0.058)	-0.850 (0.180)	0.327 (0.148)	0.009 (0.061)	0.181 (0.053)	0.040 (0.110)	-12.679 (0.061)	-0.004 (0.036)	-0.009 (0.117)	0.656 (0.101)	-0.005 (0.043)	0.265 (0.040)	-0.318 (0.072)
(B)	-14.348 (0.153)	-0.015 (0.059)	-0.834 (0.187)	0.336 (0.150)	0.026 (0.063)	0.157 (0.055)	0.031 (0.114)	-12.700 (0.132)	-0.023 (0.036)	-0.023 (0.123)	0.630 (0.104)	0.011 (0.045)	0.258 (0.042)	-0.324 (0.073)
(C)	-14.245 (0.092)	-0.051 (0.058)	-0.737 (0.185)	0.397 (0.149)	0.011 (0.063)	0.168 (0.056)	-0.023 (0.108)	-12.665 (0.062)	-0.031 (0.036)	0.011 (0.121)	0.636 (0.104)	-0.005 (0.045)	0.241 (0.041)	-0.317 (0.071)
(D)	-14.354 (0.12)	-0.069 (0.053)	-0.697 (0.181)	0.516 (0.131)	0.008 (0.059)	0.115 (0.055)	-0.111 (0.107)	-12.662 (0.098)	-0.038 (0.036)	-0.077 (0.124)	0.58 (0.101)	0.022 (0.045)	0.243 (0.042)	-0.3 (0.072)
(E)	-14.285 (0.092)	-0.005 (0.057)	-0.805 (0.183)	0.426 (0.143)	-0.007 (0.062)	0.209 (0.056)	-0.006 (0.11)	-12.665 (0.061)	-0.012 (0.036)	-0.021 (0.117)	0.639 (0.101)	-0.004 (0.043)	0.261 (0.04)	-0.299 (0.071)
(F)	-14.357 (0.138)	-0.043 (0.055)	-0.758 (0.183)	0.473 (0.136)	0.009 (0.06)	0.155 (0.056)	-0.053 (0.11)	-12.677 (0.119)	-0.026 (0.036)	-0.055 (0.121)	0.599 (0.102)	0.015 (0.044)	0.256 (0.041)	-0.299 (0.072)

TABLE 3 Estimates for the posterior means and standard deviations, in round brackets, of hyperparameters included in the models described in Table 1

ID	$\sigma_{\theta_1}^2$	$\sigma_{\theta_2}^2$	ρ_{θ}	ρ	$\sigma_{\phi_1}^2$	$\sigma_{\phi_2}^2$	ρ_{ϕ}
(A)	0.0001 (0.0003)	0.759 (0.049)			0.152 (0.028)	0.127 (0.019)	
(B)	0.0001 (0.0003)	0.578 (0.047)		0.996 (0.0008)	0.373 (0.052)	0.347 (0.045)	
(C)	0.0001 (0.0003)	0.594 (0.048)			0.308 (0.047)	0.249 (0.035)	0.898 (0.021)
(D)	0.0001 (0.0003)	0.376 (0.054)		0.986 (0.005)	0.747 (0.121)	0.690 (0.108)	0.904 (0.019)
(E)	0.589 (0.081)	0.752 (0.046)	0.413 (0.013)		0.132 (0.020)	0.134 (0.017)	0.789 (0.042)
(F)	0.484 (0.081)	0.643 (0.046)	0.405 (0.016)	0.997 (0.001)	0.276 (0.038)	0.269 (0.033)	0.829 (0.030)

The estimates of the posterior distributions for the two conditional variances, $\sigma_{\phi_1}^2$ and $\sigma_{\phi_2}^2$, are found less stable compared to the unstructured errors. The credible intervals of the two hyperparameters overlap in all models, indicating a similar spatial structure between the two kinds of severities. The posterior mean of ρ_{ϕ} , the correlation coefficient between the two severity levels, is found approximately equal to 0.9 (models C and D) and 0.8 (models E and F), indicating a strong multivariate nature for the spatial random component.

These results suggest that car crash data have a complex latent structure being the severity levels strongly correlated, and the spatially structured and unstructured effects statistically relevant.

4.3 | Model comparisons

We compared the models listed in Table 1 using DIC and WAIC criteria. The results are reported in Table 4. We first notice that PMCAR models (i.e. (B), (D) and (F)) are found to perform always better than their Intrinsic counterparts in terms of goodness of fit. They are somewhat unexplored in the road safety literature on spatial networks, Miao and Song (2005) being the only paper we found that analyse the importance of a spatial autocorrelation parameter. However, our results suggest that the PCAR distribution and its generalisations should deserve more attention.

Moving from (A) to (F) the model performance improves, indicating one more time the benefits of considering a correlated multivariate structure for the spatial and the unstructured components. In particular, model (F) is the best one according to both criteria; hence, hereafter, we focus on this model.

4.4 | Car crashes rates

Figure 5 displays the posterior means of car accident rates, λ_{ij} , estimated using model (F) both for severe and slight crashes. The colours of the road segments were generated by dividing the

TABLE 4 Estimates of deviance information criterion (DIC) and Watanabe-Akaike information criterion (WAIC) values for the models described in Section 3. The columns *balanced accuracy—severe* and *balanced accuracy—slight* will be explained in Section 5

ID	DIC	WAIC	Balanced accuracy—severe	Balanced accuracy—slight
(A)	14462.56	14474.89	0.631	0.718
(B)	14408.06	14433.77	0.635	0.716
(C)	14269.79	14296.06	0.645	0.718
(D)	14154.09	14167.64	0.661	0.717
(E)	14122.46	14106.69	0.662	0.720
(F)	14103.44	14086.46	0.675	0.720

predicted values of each severity level into ten classes based on a set of quantiles ranging from red (highest quantile) to blue (lowest quantiles). In both cases, the highest values correspond to, approximately, one severe or ten slight car crashes every kilometre and every thousand of daily commuters (i.e. the two quantities that define the offset). Moreover, the inset maps highlight a few roads close to Leeds city centre (denoted by a black star).

The four maps show similar patterns, but some roads in the southern part of the city (especially M621) look more prone to severe car crashes. The city of Leeds appeared to be divided into several areas. The northern and north-eastern part of the city are associated with lower car accident rates compared to other suburbs. The areas located in north-west, east and south of the city centre seem to be associated with the highest levels of car crashes rates, especially severe ones. This is probably linked with some congested arterial thoroughfares reaching the centre. Finally, we note that the roads closer to the city centre are the safest part of the city network, as it is clearly shown by the two inset maps in Figure 5.

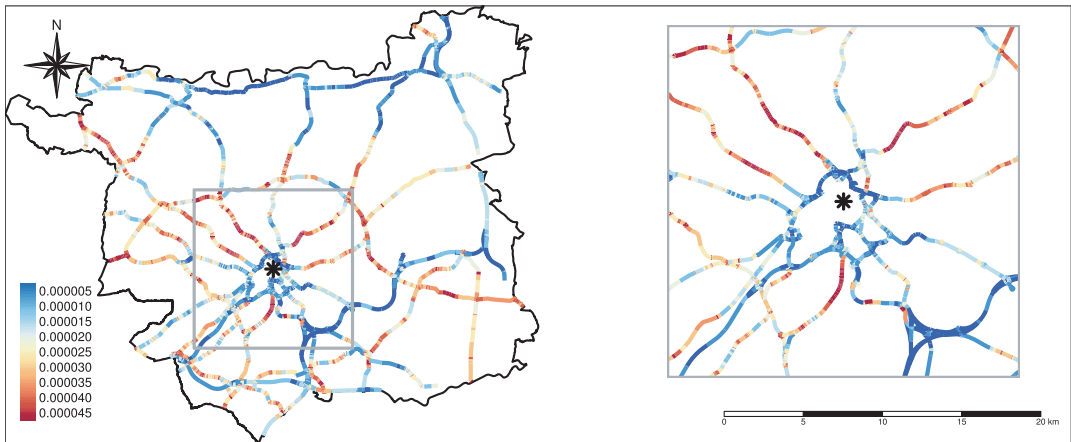
5 | MODEL CRITICISM AND SENSITIVITY ANALYSIS

Deviance information criterion and WAIC criteria were never intended to be absolute measures of model fit, and they cannot be used for *Model Criticism*. Hence, we tested the adequacy of model (F) using two strategies.


5.1 | First strategy for criticism

The classical criterion for criticism of a Bayesian hierarchical model is the probability integral transform (Held et al., 2010; Marshall & Spiegelhalter, 2003), typically adjusted in case of a discrete response variable (such as car crashes counts) using a continuity correction. Unfortunately, these adjustments do not seem to work appropriately when modelling sparse count data, such as severe crashes, since the correction is not adequate. We refer the interested reader to the supplementary material for more details.

Therefore, hereafter, we followed a different strategy. We binned the observed and predicted counts into two classes: *Zero* and *One or more* car crashes. Then, we built a confusion matrix and evaluated the model performance via some accuracy measures that are summarised in Table 5. A similar procedure for sparse count data was also presented in Ma et al. (2017). We decided to adopt



(a) Severe Car Crashes.



(b) Slight Car Crashes.

FIGURE 5 Maps representing the posterior means for severe and slight car crashes rates, estimated using model (F). The colours go from red (higher quantiles) to blue (lower quantiles). The black star indicates the city centre, while the inset maps highlight the road network in its proximity

one as a threshold to dichotomise the variables since more than 80% of road segments registered no severe car crash during 2011–2018. However, the algorithm proposed here can be extended to three or more classes, defined using a set of different thresholds (such as *Zero*, *One*, and *Two or more* road crashes).

The *accuracy* measure, usually adopted for evaluating the predictive performance of a model, is typically biased and overly optimistic in case of unbalanced classes (such as *Zero* and *One or more* severe car crashes per road segment), since, even in the worst case, it is as high as the percentage of observations in the more frequent class (Brodersen et al., 2010). The *balanced accuracy*, first introduced by Brodersen et al. (2010), is defined as the average of sensitivity and specificity, and it overcomes this drawback since it represents an average between the predictive performances on each class.

The output of a Bayesian hierarchical model is an estimate of the posterior distribution of predicted values, while the procedure reported in the previous paragraph can only be applied to binary data. For this reason, we simulated n Poisson random variables (one for each road segment)

TABLE 5 The first two rows sketch a confusion matrix between observed and predicted counts after binning them into two classes. The other rows define a few measures that can be used to evaluate the predictive fit

	Zero	One or more
Zero	A	B
One or more	C	D
Sensitivity	$= \frac{A}{A+B}$	
Specificity	$= \frac{D}{C+D}$	
Accuracy	$= \frac{A+D}{A+B+C+D}$	
Balanced accuracy	$= \frac{1}{2} \left(\frac{A}{A+B} + \frac{D}{C+D} \right)$	

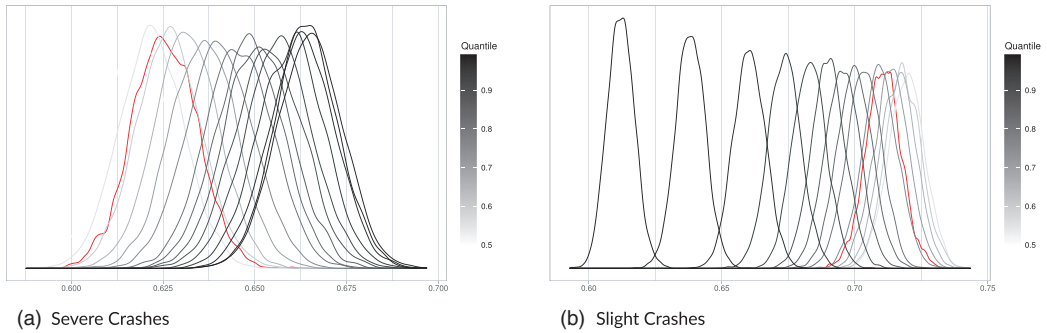


FIGURE 6 Distribution of balanced accuracy for severe crashes (left) and slight crashes (right), considering a binary classification using the posterior mean and a set of quantiles. The red curve represents the mean

with mean equal to the mean of each posterior distribution. Then, we binned the observed and sampled counts into two classes, that is Zero and one or more car crashes, and we compared the two values, obtaining a single estimate of balanced accuracy. The distribution of the balanced accuracy measure was finally approximated by repeating this procedure $N = 5000$ times.

Moreover, we calculated several quantiles of the posterior distribution of each predicted value, and we run the same steps as before, sampling from a Poisson distribution with mean equal to each of those quantiles. Lastly, being severe and slight car crashes potentially quite different processes, this algorithm was applied independently for the two severity levels. We reported in the supplementary material the pseudo-code for running this procedure, whereas results are displayed in Figure 6a (severe cashes) and Figure 6b (slight crashes). In both cases, the red curve represents the distribution of balance accuracy obtained by a binary classification based on the posterior means, whereas the other curves represent the same distribution obtained using the set of quantiles.

It looks like the optimal threshold for binary classification of severe car crashes is given by the 0.975-quantile, where the balanced accuracy distribution is concentrated around 0.675. The optimal threshold for binary classification of slight car crashes is given by the median, and the distribution of balanced accuracy is centred around 0.72. These plots remark the differences between the two severity levels in terms of sparsity, suggesting the adoption of a higher quantile for the prediction of the more sparse events. However, using the appropriate cut-off(s), Model (F) seems to

perform reasonably well in both cases with slightly better performance for slight car crashes. We also explored the goodness of fit of model (F) estimating the balanced accuracy measure separately for the three road classes and the two road carriageway types. We always obtained results similar to Figure 6. Moreover, we did not observe any relevant difference among the road typologies, highlighting that model (F) successfully predicts traffic collisions in all tested situations.

The third column in Table 4 summarises the estimates of Balanced Accuracy obtained for each model in Table 1 under the best scenario (i.e. the optimal quantile) for severe crashes, while the fourth column summarises the same quantities for slight crashes. We notice that the accuracy improves every time a new correlation term is included into a model, particularly for severe car accidents, which is the rarer car crash typology. In fact, considering that improving risk estimation of very rare events is one of the main reasons why one may want to adopt a multivariate model in the first place, our results seem to suggest that the approach proposed in this paper represents a reasonable way to investigate road collision dynamics.

5.2 | Second strategy for criticism

Following the results illustrated before, we estimated the 0.975-quantile of λ_{i1} (severe crashes) and the median of λ_{i2} (slight crashes), and we multiplied them by the corresponding offset values, that is E_i . Then, we created a sequence of histograms of predicted values, grouped by the observed counts categorised in four levels: 0, 1, 2 and 3 or more. The results are summarised in Figure 7. Both graphs show a good agreement between predicted and observed number of crashes, since the distributions corresponding to higher observed counts progressively move more and more to the right. Moreover, Figure 7a shows the importance of our previous analysis and the pitfalls of predicting severe car crashes counts using the posterior means.

5.3 | Sensitivity analysis and the modifiable areal unit problem

Finally, we performed a sensitivity analysis evaluating the robustness of model (F) under different specifications for (a) the hyperprior distributions, (b) the adjacency matrix and (c) the definition of the segments in the road network.

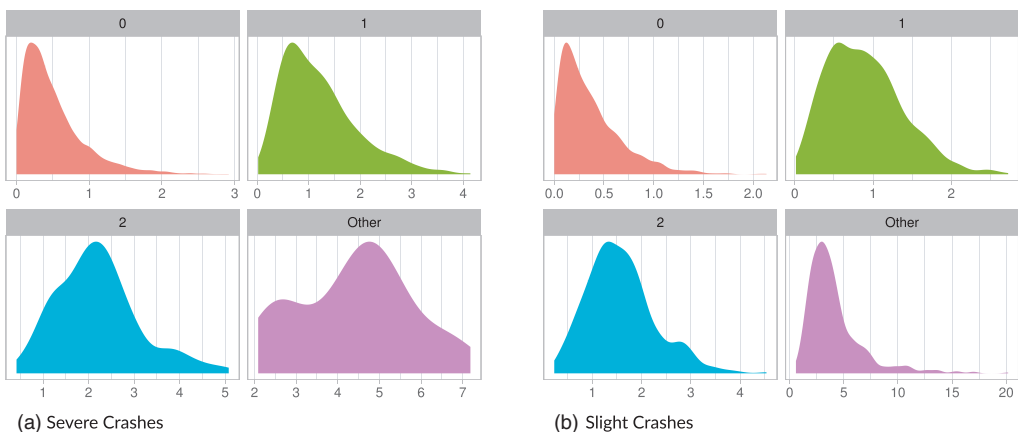


FIGURE 7 Histogram of posterior 0.975-quantile (left) and posterior median (right), grouped by the corresponding observed counts. Other means ‘Three or more’

The models described in Section 3 considered a Wishart hyperprior for the precision matrix $\mathbf{\Omega}$ with rank equal to 2 and scale matrix equal to \mathbf{I}_2 . We repeated the analysis using more vague and more informative Wishart distributions, setting the scale matrix equal to $\text{diag}(2, 2)$ and $\text{diag}(0.5, 0.5)$. We did not find any noticeable differences amongst alternative specifications. Hence results are not reported hereafter, but we refer the interested readers to the supplementary material.

We compared different definitions for the adjacency matrix \mathbf{W} , testing second and third order neighbours and distance-based spatial neighbours, considering the first order neighbours in case two adjacent road segments were longer than the given threshold. Also in this case, we did not find any noticeable differences as far as the estimation of the fixed effects is concerned, whereas only small differences were found in the posterior distributions of the random effects (especially for $\sigma_{\phi_1}^2$ and $\sigma_{\phi_2}^2$ when we considered a spatial adjacency matrix with a threshold equal to 500m). However, worse DIC and WAIC values were found for models using alternative definitions of \mathbf{W} matrix, and we refer to the supplementary material for more details. Similar findings are also reported by Aguero-Valverde and Jovanis (2008), Wang et al. (2016) and Alarifi et al. (2018).

Finally, we explored the influence of a particular configuration of the network segments on our results. In fact, the location of the vertices (and, hence, the edges) in a road network created with OS data is essentially arbitrary (although some minimal consistency requirements must be satisfied, see Karduni et al. (2016) and Gilardi et al. (2020)), which implies that there is no unique and unambiguous way of defining the lengths and relative positions of the road segments. We, therefore, considered an alternative network configuration reshaping and contracting the road network using an algorithm implemented in Padgham (2019). This algorithm manipulates a network by excluding all *redundant* vertices, that is those vertices that connect two contiguous segments without any other intersection (Padgham, 2019).

A toy example representing the ideas behind the contraction of a road network is sketched in Figure 8. The red dots in Figure 8a represent redundant vertices since they can be removed without tampering the shape or the routability of the network, meaning that excluding those vertices does not add any new cluster to the graph structure. The goal is to remove all redundant vertices and merge the corresponding edges, creating a graph which looks identical to the original one but with fewer edges. Figure 8b shows the results of the contraction operations applied to the toy network sketched in Figure 8a. We can see that the redundant vertices were removed, combining the road segments that touched them.

Applying this algorithm produced a contracted road network with, approximately, 2700 segments (instead of the original 3661). Following the same procedures detailed in Sections 2 and 4, we calculated the number of severe and slight car crashes that occurred in each road segment, the traffic volumes (which are used as an offset), the edge betweenness centrality measures, and the socio-economic covariates. The road type and the dual-carriageway dummy variables were automatically determined since the algorithm did not merge two road segments with different classifications. Finally, we estimated model (F) and reported a summary of means and standard errors of the posterior distributions of fixed and random effects in Tables 6 and 7, respectively.

As far as the fixed effects are concerned, the new network configuration influences results quite mildly, since the estimates of the coefficients did not change in sign, order of magnitude or significance. As one could expect, the impact of the reshaping is slightly more pronounced for the random effects, in particular for $\sigma_{\phi_1}^2$ and $\sigma_{\phi_2}^2$, two of the four hyperparameters in the PMCAR prior. The model trained on the contracted network presents a greater spatial uncertainty than




FIGURE 8 Sketching of the algorithm used for contracting the road network. Red points on the left represent redundant vertices

TABLE 6 Means and standard deviations for the posterior distributions of the fixed effects in the model estimated after contracting the road network

	β_0	Betw	Motorways	Prim. Roads	Ratio Empl.	Pop. Dens.	Dual Carr.
Severe:	-15.187	0.015	-0.886	0.397	-0.015	0.218	-0.055
	(0.303)	(0.058)	(0.186)	(0.143)	(0.070)	(0.064)	(0.120)
Slight:	-13.521	0.029	-0.201	0.533	0.014	0.276	-0.321
	(0.306)	(0.038)	(0.130)	(0.114)	(0.056)	(0.050)	(0.085)

TABLE 7 Means and standard deviations for the posterior distributions of the hyperparameters in the model estimated after contracting the road network

	$\sigma_{\theta_1}^2$	$\sigma_{\theta_2}^2$	ρ_θ	ρ	$\sigma_{\phi_1}^2$	$\sigma_{\phi_2}^2$	ρ_ϕ
mean:	0.331	0.500	0.374	0.997	3.543	3.978	0.938
sd:	(0.108)	(0.114)	(0.032)	(0.002)	(0.645)	(0.645)	(0.021)

model (F), but similar posterior distributions for car crashes rates. We refer the interested readers to the supplementary material, where we also reported the maps of car crashes rates using the new network configuration.

Network reshaping and contraction is a network readaptation of the classical modifiable area unit problem (MAUP), only recently explored by a handful of authors in the literature of road safety models for areal data (see, for example, Xu et al. (2018) for an introduction, and Briz-Redón et al. (2019b) for an extensive application). The main conclusion of these papers is that MAUP severely impacts car crashes models, affecting the magnitude and significance of the estimates for both fixed and random effects, hence it should never be ignored for a reliable road safety analysis.

Our results tell a somewhat different story. The statistical analysis is found quite robust to MAUP when carried out on a network lattice, possibly because the road network has a physical geometrical meaning and, hence, a lower degree of arbitrariness than administrative boundaries.

Hence, we suggest not to ignore the network structure of the data whenever it is available when analysing car crashes data or other phenomena that naturally occur on a network. The only paper that performs a descriptive analysis of the influence of road segment configurations on car crashes counts is Thomas (1996), and, to the best of our knowledge, this is the first attempt to explore the robustness of a statistical model for lattice network data to MAUP.

6 | DISCUSSION AND CONCLUSIONS

This paper investigated the spatial distribution of road crashes in a major city using Bayesian methods for road network analysis. The relationships between crashes of different severity levels, either *slight* or *severe*, were modelled using a range of multivariate models to explore their spatial dynamics. Key to the approach was constraining crash locations to the city's road network, a one-dimensional linear network composed of segments representing a spatial lattice. We tested a range of multivariate hierarchical models with different random effects, and we found that the best model (according to DIC, WAIC and *balanced accuracy* criteria) includes a multivariate spatially unstructured random effect and a multivariate spatially structured PMCAR random effect.

The physical and social environment of the street segments were considered in the model specification. Our results, summarised in Tables 2 and 3, suggest that population density is positively correlated with severe and slight crashes, although the effect is stronger for the latter case. On the other hand, the employment rate does not seem related with car crash occurrences, highlighting the importance of testing alternative proxies for local economic levels. As far as street category is concerned, Primary roads have been found safer than A roads both for severe and slight car crashes, whereas Motorways are significantly less prone to severe car accidents than A roads. No significant effect is found for slight accidents. The betweenness centrality does not seem to influence slight or severe accidents; this is an unexpected result which may deserve further investigation, possibly with alternative proxies to measure the VMT. Finally, dual carriageway roads have been found significantly less prone to slight car accidents, whereas no impact has been found for severe car crashes. Other potentially relevant physical characteristics might be the traffic speed, the line width, the presence of speed limits or junctions, the slope and the curvature of the road segments. Unfortunately, these variables were not available to us at the time when this paper was written. However, given the detailed network adopted in this paper, which is composed of thousands of short segments, we believe that some of those variables may be of limited relevance due to the short length of many segments. These unmeasured effects are accounted, to some extent, by the second order parameters of the random effects included in our models.

Concerning this last point, we underline that the unstructured random effects are supposed to reflect the uncertainty due to scarce sampling information (which may occur in small spatial domains like road segments) as well as the differences among the segments, such as faults in the pavements, crosses or junctions. These characteristics act as local shocks on the car crash occurrences and, although unrelated in space, they act both on severe and slight car accidents. Hence, the corresponding random effects are expected to be correlated and the parameter ρ_θ is meant to measure this correlation.

Similarly, spatially structured random effects are expected to account for the impact of unmeasured variables that have regularities in space (i.e. meteorological conditions or structural characteristics of the roads not available in our analysis). These variables act both on severe and slight car accidents and parameter ρ_ϕ is expected to measure the correlation between spatially

structured random effects due to these components. Both correlations are found relevant in our data, although the correlation between spatially structured effects is found stronger than the other. Hence, we argue they both should be considered when modelling car crashes data. The relevant interactions between the two severity levels allow one level to borrow strength from the other and improve the estimation of the risk associated with each piece of the city road network, especially for the rarer event.

We evaluated the sensitivity of our modelling approach to different hyperprior specifications and adjacency configurations of the components of the lattice network, showing that the statistical model presents substantial robustness in this respect. We finally considered the impact of MAUP when modelling data collected on a spatial network. An algorithm was proposed in the paper to assess the magnitude of MAUP effect in the estimates and model predictions. As mentioned above, differently from several previous studies that considered the MAUP for various areal partitions of the spatial region of interest, we found that our results are quite robust under an alternative configuration of the road network. This can be related to the fact that road networks have a physical meaning, hence they are expected to suffer MAUP less. Nevertheless, further research, possibly in different fields, is definitely necessary to better understand the impact of MAUP on network lattice data.

Finally, we remark that using areal units as spatial support ignores the fact that car crashes cannot occur outside the road network, and ascribes an estimated risk to all the streets in that polygon, whether or not they are actually exposed to it. Differently, using a lattice based on a network structure draws attention to limited and specific parts of the spatial support. Hence, adopting a more appropriate spatial disaggregation can be fundamental for local authorities to plan actions (such as to install a new traffic light, add or remove roundabouts or enforce police control) to mitigate this risk where it is found too high.

The ideas presented in this paper could be extended in several directions. A first step forward could be focused towards the development of a spatiotemporal extension of model (F), following the suggestions in Miaou and Song (2005), Wang et al. (2011), Boulieri et al. (2017) and Ma et al. (2017). We point out, however, that this is not straightforward (and, to the best of our knowledge, it was pursued only by Ma et al. (2017) using a single road divided into a few segments) given the extreme sparse spatiotemporal nature of severe car crashes on a metropolitan road network. Indeed, for the data set at hand, more than 95% of all car crashes registered no fatal or serious car accident for any given year, something that could require a different methodological approach. The procedure for MAUP detection could also be improved by developing new routines for testing alternative algorithms for network reshaping and contraction, which were first developed for areal data in the field of geography (see, e.g. Xu et al. (2018) and references therein).

An additional improvement to the approach could involve the development of spatial or spatio-temporal theoretical point pattern models for car crashes on networks (Baddeley et al., 2021). This approach represents a flexible and powerful way to investigate the spatial dynamics of random events on a graph support and may provide a tool that largely circumvents the MAUP problem. However, whereas it is moderately straightforward to include covariates in the bivariate lattice models suggested in this paper, the theory of bivariate point pattern models on networks and the inclusion of spatial covariates in this framework would require substantial methodological development. In addition, most of the work in point pattern modelling is non-parametric or semiparametric in nature, whereas the approach adopted in the present paper and, to some extent, in the analysis spatial lattice data, is grounded in the Bayesian framework. In this sense, the two approaches are complementary when modelling spatial or spatio-temporal dynamics of crash data. A bridge between the two would be the log-Gaussian Cox processes, where a

stochastic component is included in the intensity function to deal with the unexplained spatial variation (Møller et al., 1998). Dealing with an intensity function governed by a spatial random field, however, raises challenging methodological problems since defining a proper covariance function on graph spatial support is not straightforward. Considering this issue is beyond the scope of the present paper; hence, we do not discuss it further and leave it for future research.

It is clear that more research is needed to evaluate the full range of possible models for identifying crash ‘hot spots’ and highlight segments on the network. The approach presented in this paper demonstrates the potential of network-based approaches to work at city scales for flexible and robust estimates of crash rates, down to the road segment level, providing a foundation for further work in the field.

ACKNOWLEDGEMENTS

Contains OS data © Crown copyright [and database right] 2020. J. Mateu is partially supported by funds PID2019-107392RB-I00 and AICO/2019/198. We greatly acknowledge the DEMS Data Science Lab for supporting this work by providing computational resources. We thank the Editor, Associate Editor and all the reviewers for their comments, which we believe have substantially improved our work. Open Access Funding provided by Università degli Studi di Milano-Bicocca within the CRUI-CARE Agreement. [Correction added on 24 May 2022, after first online publication: CRUI funding statement has been added.]

CODE AND DATA AVAILABILITY STATEMENT

The data and the code that were used to estimate models from (A) to (G) and the model on the contracted network can be obtained at the following link: <https://github.com/agila5/multivariate-analysis-car-crashes>.

ORCID

Andrea Gilardi  <https://orcid.org/0000-0002-9424-7439>

Jorge Mateu  <https://orcid.org/0000-0002-2868-7604>

Riccardo Borgoni  <https://orcid.org/0000-0002-2520-3512>

Robin Lovelace  <https://orcid.org/0000-0001-5679-6536>

REFERENCES

- Abdel-Aty, M., Lee, J., Siddiqui, C. & Choi, K. (2013) Geographical unit based analysis in the context of transportation safety planning. *Transportation Research Part A: Policy and Practice*, 49, 62–75.
- Aguero-Valverde, J. & Jovanis, P.P. (2006) Spatial analysis of fatal and injury crashes in Pennsylvania. *Accident Analysis & Prevention*, 38, 618–625.
- Aguero-Valverde, J. & Jovanis, P.P. (2008) Analysis of road crash frequency with spatial models. *Transportation Research Record: Journal of the Transportation Research Board*, 2061(1), 55–63. Available from: <https://doi.org/10.3141/2061-07>
- Alarifi, S.A., Abdel-Aty, M.A., Lee, J. & Wang, X. (2018) Exploring the effect of different neighboring structures on spatial hierarchical joint crash frequency models. *Transportation Research Record*, 2672, 210–222.
- Baddeley, A., Nair, G., Rakshit, S., McSwiggan, G. & Davies, T.M. (2021) Analysing point patterns on networks — a review. *Spatial Statistics*, 42, 100435. Available from: <https://doi.org/10.1016/j.spasta.2020.100435>
- Barua, S., El-Basyouny, K. & Islam, M.T. (2014) A full Bayesian multivariate count data model of collision severity with spatial correlation. *Analytic Methods in Accident Research*, 3, 28–43.
- Besag, J. (1974) Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36, 192–225.
- Besag, J. & Kooperberg, C. (1995) On conditional and intrinsic autoregressions. *Biometrika*, 82, 733–746.

- Borgoni, R., Gilardi, A. & Zappa, D. (2021) Assessing the risk of car crashes in road networks. *Social Indicators Research*, 156(2–3), 429–447. Available from: <https://doi.org/10.1007/s11205-020-02295-x>
- Botella-Rocamora, P., Lopez-Quilez, A. & Martinez-Beneito, M. (2013) Spatial moving average risk smoothing. *Statistics in Medicine*, 32, 2595–2612.
- Boulieri, A., Liverani, S., de Hoogh, K. & Blangiardo, M. (2017) A space–time multivariate Bayesian model to analyse road traffic accidents by severity. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180, 119–139.
- Braunholtz, D. & Elliott, D. (2019) Estimating and adjusting for changes in the method of severity reporting for road accidents and casualty data. Department for Transport. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/820588/severity-reporting-methodology-final-report.odt
- Briz-Redón, Á., Martínez-Ruiz, F. & Montes, F. (2019a) Estimating the occurrence of traffic accidents near school locations: a case study from Valencia (Spain) including several approaches. *Accident Analysis & Prevention*, 132, 105237.
- Briz-Redón, Á., Martínez-Ruiz, F. & Montes, F. (2019b) Investigation of the consequences of the modifiable areal unit problem in macroscopic traffic safety analysis: a case study accounting for scale and zoning. *Accident Analysis & Prevention*, 132, 105276.
- Brodersen, K.H., Ong, C.S., Stephan, K.E. & Buhmann, J.M. (2010) The balanced accuracy and its posterior distribution. In *2010 20th International Conference on Pattern Recognition*, 3121–3124.
- Carlin, B.P., Banerjee, S. (2003) Hierarchical multivariate CAR models for spatio-temporally correlated survival data. *Bayesian Statistics*, 7, 45–63.
- Castle, J., Lynam, D., TRL, J.M., Lawson, S.J. & Klassen, N. (2008) Star rating roads for safety: UK trials 2006–07. Available from: <https://web.archive.org/web/20080227050304/>, http://www.eurorap.org/library/pdfs/20071130_UK_RPS_report.pdf
- Cressie, N.A. (1993) *Statistics for spatial data*. Hoboken: John Wiley & Sons, Ltd.
- Cronie, O., Moradi, M. & Mateu, J. (2020) Inhomogeneous higher-order summary statistics for point processes on linear networks. *Statistics and Computing*, 30, 1221–1239.
- Department for Transport. (2011) Instructions for the Completion of Road Accident Reports from non-CRASH Sources. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/995423/stats20-2011.pdf. Accessed: 15/09/2021.
- Department for Transport. (2020) Reported road casualties in Great Britain: 2019 annual report. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/922717/reported-road-casualties-annual-report-2019.pdf
- Dubin, R.A. (1988) Estimation of regression coefficients in the presence of spatially autocorrelated error terms. *The Review of Economics and Statistics*, 70, 466–474. Available from: <http://www.jstor.org/stable/1926785>
- El-Basyouny, K. & Sayed, T. (2009) Urban arterial accident prediction models with spatial effects. *Transportation Research Record*, 2102, 27–33.
- Freni-Sterrantino, A., Ventrucci, M. & Rue, H. (2018) A note on intrinsic conditional autoregressive models for disconnected graphs. *Spatial and Spatio-Temporal Epidemiology*, 26, 25–34. Available from: <http://www.sciencedirect.com/science/article/pii/S1877584517301600>
- Gelfand, A.E. & Vounatsou, P. (2003) Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics*, 4, 11–15.
- Gelman, A., Hwang, J. & Vehtari, A. (2014) Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24, 997–1016.
- Gilardi, A., Lovelace, R. & Padgham, M. (2020) Data structures and methods for reproducible street network analysis: overview and implementations in R. osf.io/78yub.
- Gómez-Rubio, V. (2020) *Bayesian inference with INLA*. Boca Raton: CRC Press.
- Held, L., Schrödle, B. & Rue, H. (2010) Posterior and cross-validators predictive checks: a comparison of MCMC and INLA. In: *Statistical modelling and regression structures*, 91–110. Springer.
- Hodges, J.S., Carlin, B.P. & Fan, Q. (2003) On the precision of the conditionally autoregressive prior in spatial models. *Biometrics*, 59, 317–322.
- Huang, H., Zhou, H., Wang, J., Chang, F. & Ma, M. (2017) A multivariate spatial model of crash frequency by transportation modes for urban intersections. *Analytic Methods in Accident Research*, 14, 10–21.

- Karduni, A., Kermanshah, A. & Derrible, S. (2016) A protocol to convert spatial polyline data to network formats and applications to world urban road networks. *Scientific Data*, 3, 1–7.
- Kirk, D.S., Cavalli, N. & Brazil, N. (2020) The implications of Ridehailing for risky driving and road accident injuries and fatalities. *Social Science & Medicine*, 250, 112793.
- Kolaczyk, E.D. & Csárdi, G. (2014) *Statistical analysis of network data with R*, vol.65. Berlin: Springer.
- Lindgren, F. & Rue, H. (2015) Bayesian spatial modelling with R-INLA. *Journal of Statistical Software*, 63, 1–25. Available from: <https://www.jstatsoft.org/v63/i19/>. Accessed: 08/02/2021.
- Lord, D. & Mannering, F. (2010) The statistical analysis of crash-frequency data: a review and assessment of methodological alternatives. *Transportation Research Part A: Policy and Practice*, 44, 291–305.
- Lovelace, R. & Hama, L. (2020) pct: Propensity to Cycle Tool. Available from: <https://CRAN.R-project.org/package=pct>. R package version 0.4.0
- Lovelace, R., Goodman, A., Aldred, R., Berkoff, N., Abbas, A. & Woodcock, J. (2017) The propensity to cycle tool: an open source online system for sustainable transport planning. *Journal of Transport and Land Use*, 10. <https://www.jtlu.org/index.php/jtlu/article/view/862>
- Lovelace, R., Morgan, M., Hama, L., Padgham, M., Ranzolin, D. & Sparks, A. (2019) stats 19: a package for working with open road crash data. *The Journal of Open Source Software*, 4, 1181.
- Møller, J., Syversveen, A.R. & Waagepetersen, R.P. (1998) Log Gaussian cox processes. *Scandinavian Journal of Statistics*, 25, 451–482. Available from: <https://www.jstor.org/stable/4616515>. Accessed: 03/06/2021.
- Ma, X., Chen, S. & Chen, F. (2017) Multivariate space-time modeling of crash frequencies by injury severity levels. *Analytic Methods in Accident Research*, 15, 29–40.
- MacKay, M. (1972) Traffic accidents - a modern epidemic. *International Journal of Environmental Studies*, 3, 223–227.
- Mardia, K. (1988) Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *Journal of Multivariate Analysis*, 24, 265–284.
- Marshall, E. & Spiegelhalter, D. (2003) Approximate cross-validators predictive checks in disease mapping models. *Statistics in Medicine*, 22, 1649–1660.
- Marshall, S., Gil, J., Kropf, K., Tomko, M. & Figueiredo, L. (2018) Street network studies: from networks to models and their representations. *Networks and Spatial Economics*, 18, 735–749.
- Martínez-Beneito, M.A. & Botella-Rocamora, P. (2019) *Disease mapping: from foundations to multidimensional modeling*. Boca Raton: CRC Press.
- Miaou, S.-P. (1994) The relationship between truck accidents and geometric design of road sections: poisson versus negative binomial regressions. *Accident Analysis & Prevention*, 26, 471–482.
- Miaou, S.-P. & Lum, H. (1993) Modeling vehicle accidents and highway geometric design relationships. *Accident Analysis & Prevention*, 25, 689–709.
- Miaou, S.-P. & Song, J.J. (2005) Bayesian ranking of sites for engineering safety improvements: decision parameter, treatability concept, statistical criterion, and spatial dependence. *Accident Analysis & Prevention*, 37, 699–720.
- Miaou, S.-P., Song, J.J. & Mallick, B.K. (2003) Roadway traffic crash mapping: a space-time modeling approach. *Journal of Transportation and Statistics*, 6, 33–58.
- Nantulya, V.M. & Reich, M.R. (2002) The neglected epidemic: road traffic injuries in developing countries. *BMJ: British Medical Journal*, 324, 1139–1141.
- Noland, R.B. & Quddus, M.A. (2004) A spatially disaggregate analysis of road casualties in England. *Accident Analysis & Prevention*, 36, 973–984.
- Openshaw, S. & Taylor, P.J. (1981) The modifiable areal unit problem. In: Wrigley, N. & Bennett, R.J. (Eds.) *Quantitative geography: a British view*. London: Routledge and Kegan Paul, pp. 60–70.
- Ordnance Survey. (2020) Ordnance survey. Available from: <https://www.ordnancesurvey.co.uk/>. Accessed: 21/05/2020.
- FACTS. (2020) Roads policing and its contribution to road safety. *Tech. rep.*, Parliamentary Advisory Council for Transport Safety.
- Padgham, M. (2019) dodgr: An R package for network flow aggregation. *Transport Findings*.
- Palmi-Perales, F., Gómez-Rubio, V. & Martínez-Beneito, M. A. (2021) Bayesian multivariate spatial models for lattice data with INLA. *Journal of Statistical Software*, 98, 1–29.
- Porta, S., Crucitti, P. & Latora, V. (2006) The network analysis of urban streets: a dual approach. *Physica A: Statistical Mechanics and its Applications*, 369, 853–866.

- R Core Team. (2020) *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Available from: <https://www.R-project.org/>
- Rakshit, S., Davies, T., Moradi, M.M., McSwiggan, G., Nair, G., Mateu, J. & Baddeley, A. (2019) Fast kernel smoothing of point patterns on a large network using two-dimensional convolution. *International Statistical Review*, 87, 531–556.
- Rue, H., Riebler, A.I., Sørbye, S.H., Illian, J.B., Simpson, D.P. & Lindgren, F.K. (2017) Bayesian computing with INLA: a review. *Annual Reviews of Statistics and Its Applications*, 4, 395–421. Available from: <http://arxiv.org/abs/1604.00860>
- Savolainen, P.T., Mannering, F.L., Lord, D. & Quddus, M.A. (2011) The statistical analysis of highway crash-injury severities: a review and assessment of methodological alternatives. *Accident Analysis & Prevention*, 43, 1666–1676.
- Shankar, V., Mannering, F. & Barfield, W. (1995) Effect of roadway geometrics and environmental factors on rural freeway accident frequencies. *Accident Analysis & Prevention*, 27, 371–389.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P. & Van Der Linde, A. (2002) Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64, 583–639.
- Thomas, I. (1996) Spatial data aggregation: exploratory analysis of road accidents. *Accident Analysis & Prevention*, 28, 251–264.
- UK Data Service Census Support. (2014) Census support: Flow data. Available from: <http://wicid.ukdataservice.ac.uk/>
- Ukkusuri, S., Miranda-Moreno, L.F., Ramadurai, G. & Isa-Tavarez, J. (2012) The role of built environment on pedestrian crash frequency. *Safety Science*, 50, 1141–1151.
- Wall, M.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.
- Wang, C., Quddus, M.A. & Ison, S.G. (2009) Impact of traffic congestion on road accidents: a spatial analysis of the M25 motorway in England. *Accident Analysis & Prevention*, 41, 798–808.
- Wang, C., Quddus, M.A. & Ison, S.G. (2011) Predicting accident frequency at their severity levels and its application in site ranking using a two-stage mixed multivariate model. *Accident Analysis & Prevention*, 43, 1979–1990.
- Wang, X., Yang, J., Lee, C., Ji, Z. & You, S. (2016) Macro-level safety analysis of pedestrian crashes in Shanghai, China. *Accident Analysis & Prevention*, 96, 12–21.
- Watanabe, S. (2010) Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11, 3571–3594.
- World Health Organization. (2018) *Global status report on road safety 2018*. Geneva: World Health Organization.
- Xu, P., Huang, H. & Dong, N. (2018) The modifiable areal unit problem in traffic safety: basic issue, potential solutions and future research. *Journal of Traffic and Transportation Engineering (English edition)*, 5, 73–82.
- Zhai, X., Huang, H., Xu, P. & Sze, N.N. (2019) The influence of zonal configurations on macro-level crash modeling. *Transportmetrica A: Transport Science*, 15(2), 417–434. Available from: <https://doi.org/10.1080/23249935.2018.1493550>
- Ziakopoulos, A. & Yannis, G. (2020) A review of spatial approaches in road safety. *Accident Analysis & Prevention*, 135, 105323.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Gilardi, A., Mateu, J., Borgoni, R. & Lovelace, R. (2022) Multivariate hierarchical analysis of car crashes data considering a spatial network lattice. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 1–28. Available from: <https://doi.org/10.1111/rssa.12823>