

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Variance partitioning in spatio-temporal disease mapping models

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Franco Villoria, M., Ventrucci, M., Rue, H. (2022). Variance partitioning in spatio-temporal disease mapping models. STATISTICAL METHODS IN MEDICAL RESEARCH, 31(8 (August)), 1566-1578 [10.1177/09622802221099642].

Availability:

This version is available at: <https://hdl.handle.net/11585/900743> since: 2023-06-03

Published:

DOI: <http://doi.org/10.1177/09622802221099642>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

Variance partitioning in spatio-temporal disease mapping models

Maria Franco-Villoria
Department of Economics
University of Modena and Reggio Emilia

Massimo Ventrucchi
Department of Statistical Sciences
University of Bologna

Håvard Rue
CEMSE Division, King Abdullah University of Science and Technology,
Thuwal, Saudi Arabia

May 18, 2022

Abstract

Bayesian disease mapping, yet if undeniably useful to describe variation in risk over time and space, comes with the hurdle of prior elicitation on hard-to-interpret random effect precision parameters. We introduce a reparametrized version of the popular spatio-temporal interaction models, based on Kronecker product intrinsic Gaussian Markov Random Fields, that we name the variance partitioning (VP) model. The VP model includes a mixing parameter that balances the contribution of the main and interaction effects to the total (generalized) variance and enhances interpretability. The use of a penalized complexity prior on the mixing parameter aids in coding prior information in a intuitive way. We illustrate the advantages of the VP model using two case studies.

KEYWORDS: intrinsic Gaussian Markov Random Fields; intrinsic CAR; Kronecker product GMRF; penalized complexity prior; spatio-temporal smoothing

ADDRESS FOR CORRESPONDENCE: Maria Franco-Villoria, Department of Economics, University of Modena and Reggio Emilia. Email: maria.francovilloria@unimore.it

1 Introduction

The Covid-19 pandemic has put the world at stake. In Italy, the first two cases were confirmed on 31st January 2020 and on 9th March 2020 a national lockdown was put in place by the authorities to control and reduce the expansion of the virus. Data on newly infected people have been routinely collected since then to monitor the evolution of the disease. The study of the pandemic evolution can be tackled using disease mapping. Knowledge of how the infection has spread can help to evaluate the performance of containment measures. In particular, the quantification of the space-time interaction, which describes how the spatial patterns change over time, has been proposed as a way to deepen our understanding on the evolution of the disease (Diggle et al., 1995).

Disease mapping models (Martinez-Beneito and Botella-Rocamora, 2019; Lawson, 2018; MacNab, 2022; Wakefield, 2007) aim to describe the variation in risk of a particular disease over space and time. Data are usually available in the form of aggregated counts at some spatial level, such as counties, municipalities, etc. Additive time and space models have been long used to model disease rates (Knorr-Held and Besag, 1998). More recently, the availability of complex data has made

it possible to consider more complex models that include an interaction term to appropriately capture space-time relationships in the data (see for example Abellan et al. (2008), Knorr-Held (2000), Waller et al. (1997), Bernardinelli et al. (1995) to cite a few). Understanding the spatial distribution of disease risk or how it has evolved over time might be useful for public health authorities in planning resource allocation and identification of areas to be prioritized. In particular, the space-time interaction may reveal important information regarding the nature of the disease, for example suggesting whether a new disease is possibly infectious (Aldstadt, 2007) or the existence of additional causes in non-infectious cases (Robertson et al., 2010). Thus a model that is able to quantify the importance of this term is desirable from a practical point of view; in this paper we introduce a model parametrization that partitions the total variance into main and interaction effects so that the contribution of each of those can be quantified.

Crude disease rates are unreliable due to sampling variability so smoothing is used to borrow information across neighbouring areas and time points. For this reason, disease mapping has been developed mainly in a Bayesian hierarchical model formulation where the building blocks of a smooth in one and more dimensions can be modelled using intrinsic Gaussian Markov Random Fields (IGMRF) such as the first and second order random walk (Lindgren and Rue, 2008) or the ICAR (Besag, 1974) models. For modelling interactions statisticians have used tensor products smoothers, where, in a Bayesian framework, the penalty can be seen as a special type of GMRF called Kronecker product GMRFs (Rue and Held, 2005).

In Bayesian spatio-temporal disease mapping the precision parameter of the IGMRFs plays a role in controlling the degree of smoothing applied over time and space. A number of issues related to prior elicitation need to be addressed when dealing with intrinsic models. Firstly, the precision matrix is singular, which means that the total variance that we aim to partition is not finite. In order to define priors on the variance components we can rely upon the concept of generalized variance of an IGMRF; this has been defined by Sørbye and Rue (2014) as the geometric mean of the diagonal elements of the generalized inverse of the precision matrix of the IGMRF, and can only be computed upon linear constraints.

A second issue to bear in mind is that the generalized variance of an IGMRF depends on the structure matrix, and hence it changes depending on things like the temporal and spatial resolution or the size of the dataset at hand. This means that interpretation of the precision parameter becomes case-dependent, making prior elicitation and parameter interpretation difficult. To avoid this problem, Sørbye and Rue (2014) advise scaling the structure matrix so that the generalized variance is equal to 1; this way the precision parameter is automatically rescaled and the prior has the same meaning regardless of the graph structure (Freni-Sterrantino et al., 2018). Scaling becomes particularly relevant in the context of space-time models, as otherwise differences in the structure matrices of the spatial, temporal and spatio-temporal terms would have an impact on the priors for the corresponding precision parameters that we cannot control. By scaling the structure matrix of the temporal and spatial random effects, the structure matrix of the interaction, defined as a Kronecker IGMRF, is automatically scaled.

Further to the issues mentioned above, the choice of priors for variance parameters has received much attention in the literature (Gelman, 2006; Wakefield, 2007; Fong et al., 2010; Simpson et al., 2017). Part of the hassle in choosing a prior stems from the difficulty of interpreting variance parameters, especially for intrinsic processes, where the standard deviation is to be interpreted as a conditional one (Fong et al., 2010; Riebler et al., 2016). On top of that, in models with various terms, the tendency is to set priors independently for each precision parameter, while some authors are beginning to recognize that it might be more practical to think about total variability and how each term in the model contributes to that rather than to concentrate on single variance components separately (Wakefield, 2007; Riebler et al., 2016; Fuglstad et al., 2020; Ventrucci et al., 2020). In

the context of disease mapping, Wakefield (2007) proposes using an inverse Gamma prior on the total variability, along with a Beta prior that distributes the variance between a spatially correlated random field and a spatially unstructured effect (the so called BYM model Besag et al. (1991)). Using a similar parametrization, Riebler et al. (2016) present a prior that shrinks towards no spatial effect following the penalized complexity (PC) prior approach of Simpson et al. (2017). Outside the disease mapping literature, Ventrucci et al. (2020) develop a PC prior in one-factor mixed models for the relative contribution of group-specific variability. In a more general context, Fuglstad et al. (2020) introduce a framework for hierarchically distributing the variance in additive models, where, at each level of the total variance decomposition, ignorance or preference about the variance contribution of a term is expressed via a Dirichlet or a PC prior, respectively. We add to the literature by considering also the temporal dimension in disease mapping models. In particular, all the terms in the model (main effects and interaction) are assumed to follow intrinsic models. This differentiates our work from the literature mentioned above.

In this work, we revisit the spatio-temporal models proposed by Knorr-Held (2000), where the space-time interaction term can be one of four different types, depending on the degree of dependence assumed between time and space. These four types are characterized by different prior assumptions, expressed in terms of a Kronecker product. We propose an intuitive reparametrization that leads to partitioning the generalized variance between the main effects and interaction. The main and interaction effects are not independent, and hence using a joint prior on those terms is preferable. We do so by including a mixing parameter that 1) eases interpretation and 2) naturally leads to a prior that is intuitive to elicit. One of the advantages of the Bayesian framework is that whenever information on the disease process is available, it can be encoded into the prior (Robertson et al., 2010). Often, the epidemiologist might have an intuition on how important the interaction term is in explaining the spatio-temporal variation of a particular disease. However, translating this information in terms of a precision parameter is not trivial at all. We follow the penalized complexity prior (PC) framework of Simpson et al. (2017) to derive a prior for the mixing parameter that avoids overfitting by construction and allows the user to code any prior information easily. This way we alleviate both problems, by considering an interaction model that not only enhances interpretability but also permits a more intuitive construction of the prior. We call this reparametrized version the *variance partitioning (VP) model*. The proposed methodology is applicable to any of the four space-time interactions described in Knorr-Held (2000).

The rest of the paper is organized as follows. Section 2 covers spatio-temporal disease mapping models, with a particular emphasis on the space-time interaction framework by Knorr-Held (2000), followed by a brief discussion of priors for variance parameters with special attention to the PC prior approach. In Section 3 the VP model is described in detail and the PC prior for the mixing parameter is presented, while the technical details are relegated to the supplementary material. Section 4 illustrates the proposed model on two case studies, a well known example in the disease mapping literature and an Italian Covid-19 dataset. The paper closes with a discussion in Section 5.

2 Spatio-temporal disease mapping

Consider data on n_1 time points and n_2 non-overlapping areas, y_{ij} is the observed number of cases at time $i = 1, \dots, n_1$ and area $j = 1, \dots, n_2$. The most commonly used models for y_{ij} are the binomial and the Poisson; in either case, the model in the linear predictor scale can be written as $\eta_{ij} = \alpha + f_1(i) + f_2(j) + f_{12}(i, j)$, where $f_1(i)$ and $f_2(j)$ represent the main temporal and spatial effects respectively and the function $f_{12}(i, j)$ captures the space-time interaction. The model can be

parametrized with random effects as

$$\eta_{ij} = \alpha + \beta_{1i} + \beta_{2j} + \delta_{ij}, \quad (1)$$

where $\beta_1 = (\beta_{1,1}, \dots, \beta_{1,n_1})^T$ and $\beta_2 = (\beta_{2,1}, \dots, \beta_{2,n_2})^T$ are vectors of random effects describing the temporal and spatial main effect, respectively, and $\delta = \{\delta_{ij}\}, i = 1, \dots, n_1, j = 1, \dots, n_2$ is the vectorized spatio-temporal interaction term. The random effects β_1, β_2 and δ are typically assumed as smooth processes modelled using intrinsic Gaussian Markov Random Fields (IGMRF, Rue and Held (2005)), a special type of improper GMRF, defined below. Appropriate constraints (Goicoa et al., 2018) need to be imposed to ensure identifiability of the terms in (1). The constraints on the interaction term are summarized on Table 1, while on the temporal and spatial main effects it is enough to impose a sum to zero constraint. As usual, any available covariates can be included in model (1) as fixed effects.

Definition 1 (Improper GMRF). *Let \mathbf{Q} be an $n \times n$ symmetric positive semi-definite (SPSD) matrix with rank $n - p > 0$. Then $\mathbf{x} = (x_1, \dots, x_n)^T$ is an improper GMRF of rank $n - p$ with parameters $(\boldsymbol{\mu}, \mathbf{Q})$ if its density is*

$$\pi(\mathbf{x}) = (2\pi)^{-\frac{(n-p)}{2}} (|\mathbf{Q}|^*)^{1/2} \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{Q} (\mathbf{x} - \boldsymbol{\mu}) \right),$$

where $|\mathbf{Q}|^*$ is the generalized determinant of the precision matrix \mathbf{Q} . Improper GMRFs are used as smoothing priors in structured additive regression (STAR) models, a flexible class including generalized linear mixed models, temporally dynamic models, spatial varying coefficient models, etc; for an account of STAR models see Fahrmeir et al. (2013) and references therein.

Following Rue and Held (2005) we define an IGMRF of order 1 as an improper GMRF where $\mathbf{Q}\mathbf{1} = \mathbf{0}$, i.e., the precision matrix is singular with null space spanned by a column vector of ones, $\mathbf{1}_n$ of length n . Popular examples of an IGMRF of order 1 are the first order random walk (RW1), which is a possible option to model the temporal main effect β_1 , and the intrinsic conditional autoregressive (ICAR) model by Besag (1974), which is often assumed in disease mapping to model the spatial effect β_2 when smoothing across neighbouring regions is required.

An IGMRF of order 2 is an improper GMRF whose precision matrix is singular and its null space is spanned by a constant vector $\mathbf{1}_n$ and a linear vector $(1, \dots, n)^T$. A popular example is the second order random walk (RW2; Lindgren and Rue (2008)), popularly used for modelling smooth covariate effects in STAR models, and often implemented in spatio-temporal disease mapping for modelling the main temporal effect β_1 when smoothness in the disease risk over time is anticipated.

All the IGMRFs described above have in common that their precision matrix can be written as $\mathbf{Q} = \tau \mathbf{R}$, where τ is a precision parameter and \mathbf{R} is a known structure matrix that encodes the dependence structure. In particular, for the RW1

$$R_{k,l} = \begin{cases} 1 & k = l \in \{1, n\} \\ 2 & k = l \in \{2, \dots, n-1\} \\ -1 & k \sim l \\ 0 & \text{otherwise,} \end{cases}$$

where notation $k \sim l$ indicates contiguous time points. For the ICAR, the structure matrix is given by

$$R_{k,l} = \begin{cases} m_k & k = l \\ -1 & k \sim l \\ 0 & \text{otherwise,} \end{cases}$$

where m_k is the number of neighbours for region k and notation $k \sim l$ indicates contiguous areas that share a common border. The structure matrix of a RW2 can be written as $\mathbf{R} = \mathbf{D}^T \mathbf{D}$ where \mathbf{D} is a second order difference matrix of dimension $(n - 2) \times n$.

It is common in the disease mapping literature to consider one or both main effects f_1 and f_2 as a sum of structured and unstructured effects, so that model (1) becomes

$$\eta_{ij} = \alpha + \beta_{1_i} + \epsilon_{1_i} + \beta_{2_j} + \epsilon_{2_j} + \delta_{ij}, \quad (2)$$

where $\epsilon_1 \sim N(\mathbf{0}, \tau_{\epsilon_1} \mathbf{I}_{n_1})$, $\epsilon_2 \sim N(\mathbf{0}, \tau_{\epsilon_2} \mathbf{I}_{n_2})$. Typically, a RW1 or RW2 model is assumed for the temporal effect $\beta_1 \sim N(\mathbf{0}, \tau_1^{-1} \mathbf{R}_1^-)$ and an ICAR is assumed for the spatial effect $\beta_2 \sim N(\mathbf{0}, \tau_2^{-1} \mathbf{R}_2^-)$, where notation \mathbf{M}^- indicates the generalized inverse of matrix \mathbf{M} . The combination of the structured and unstructured spatial terms $\beta_{2_j} + \epsilon_{2_j}$ is commonly known as the BYM model (Besag et al., 1991).

2.1 Modelling interactions via Kronecker product IGMRFs

We describe now the interaction term δ in Eq. (2). Smoothness is induced by assuming

$$\delta \sim N(\mathbf{0}, \tau_{12}^{-1} \mathbf{R}_I^-),$$

which is a Kronecker product IGMRF with precision $\mathbf{Q} = \tau_{12} \mathbf{R}_I$, i.e. an improper GMRF with precision given by the Kronecker product of two IGMRFs. These models are used for smoothing spatial and spatio-temporal data, and they are the Bayesian equivalent of tensor product spline models (Wahba, 1978). Knorr-Held (2000) envisions four different types of interactions, reported in Table 1. Interaction type I can be seen as unstructured variation due to unobserved covariates, while interaction types II and III allow for the temporal trend to change from location to location and the spatial trend to change over time, respectively, but in an independent manner. Interaction type IV is the most complex one, assuming that the temporal trend changes with location in a spatially dependent way, or equivalently, that the way in which the spatial trend changes over time is time-dependent.

Table 1: The four types of interactions in spatio-temporal smoothing according to Knorr-Held (2000). The IGMRF on the interaction parameter vector δ has structure \mathbf{R}_I given by a Kronecker product; $r_1 = 1$ or 2 depending on the order of the RW assumed for the time effect.

type	\mathbf{R}_I	$\text{rank}(\mathbf{R}_I)$	linear constraints on δ
I	$\mathbf{I}_{n_2} \otimes \mathbf{I}_{n_1}$	$n_1 n_2$	not needed
II	$\mathbf{I}_{n_2} \otimes \mathbf{R}_1$	$n_2(n_1 - r_1)$	$[\mathbf{I}_{n_2} \otimes \mathbf{1}_{n_1}]^T \delta = \mathbf{0}_{n_2}$
III	$\mathbf{R}_2 \otimes \mathbf{I}_{n_1}$	$(n_2 - 1)n_1$	$[\mathbf{1}_{n_2} \otimes \mathbf{I}_{n_1}]^T \delta = \mathbf{0}_{n_1}$
IV	$\mathbf{R}_2 \otimes \mathbf{R}_1$	$(n_2 - 1)(n_1 - r_1)$	$[\mathbf{I}_{n_2} \otimes \mathbf{1}_{n_1}]^T \delta = \mathbf{0}_{n_2}; \quad [\mathbf{1}_{n_2} \otimes \mathbf{I}_{n_1}]^T \delta = \mathbf{0}_{n_1}$

Model (1) includes different precision parameters τ_1 and τ_2 for smoothing over time and space and an additional one, τ_{12} , controlling the variance of the interaction term, which yields a model able to capture the smooth spatio-temporal structure underlying the data with high flexibility. However, these models have limitations in terms of interpretation of the results, as precision parameters are not informative about the total variance explained by the associated components and the priors are not easy to elicit (see Section 2.2). We propose an alternative parametrization to address these issues in Section 3.

2.2 Priors for the precision parameters

There are two main challenges in prior choice for the precision parameters in model (1). The first problem regards the so called scaling issue that affects IGMRFs in general; Sørbye and Rue (2014) proposed addressing this issue by scaling the precision structure \mathbf{R} so that the geometric mean of the diagonal elements in \mathbf{R}^{-1} is 1. In this way, the prior for τ will roughly encode the same degree of complexity across different types of structures and hence will have the same interpretation. Of particular interest is the spatial case where, after scaling the precision of the ICAR, the prior for the precision parameter becomes transferable across different applications using different graph structures.

The second challenge regards the structure of the Kronecker product IGMRF, which can be thought of as an extra layer of flexibility on top of the main effects model. The common practice is to set independent priors on each precision parameter, but this totally disregards the model structure. Popular choices are Gamma for τ , or half-t and uniform on the standard deviation $1/\sqrt{\tau}$ (Gelman, 2006). The Gamma prior has repeatedly been pointed out as a poor choice often made by convenience; among the reasons why it should be avoided is that it forces overfitting or underfitting depending on the choice of its parameters (Frühwirth-Schnatter and Wagner, 2010, 2011; Fong et al., 2010; Simpson et al., 2017; Ventrucchi and Rue, 2016).

In Section 3 we propose a novel modelling framework where the interaction is seen as a flexible extension of the main effects model, and the prior is set so that the interaction term shrinks to the main effects following the PC prior framework. Recently, PC priors have been proposed as a way to prevent overfitting, based on four simple principles, that we briefly summarize and illustrate below for the precision parameter τ of a Gaussian random effect. For further details the reader is referred to Simpson et al. (2017).

Let π_1 denote the density of a model component w with precision parameter τ . This model component can be seen as a flexible extension of a based model with density π_0 and $\tau = \infty$ (i.e. absence of random effects). The four principles are:

1. Parsimony: The prior for τ should give proper shrinkage to $\tau = \infty$ and decay with increasing complexity of π_1 , so that the simplest model is favoured unless there is evidence for a more flexible one.
2. The increased complexity of π_1 with respect to π_0 is measured using the Kullback-Leibler divergence (KLD, Kullback and Leibler, 1951),

$$\text{KLD}(\pi_1||\pi_0) = \int \pi_1(w) \log \left(\frac{\pi_1(w)}{\pi_0(w)} \right) dw.$$

For ease of interpretation, the KLD is transformed to a unidirectional distance measure

$$d(\tau) = d(\pi_1||\pi_0) = \sqrt{2\text{KLD}(\pi_1||\pi_0)}$$

that can be interpreted as the distance from the flexible model π_1 to the base model π_0 .

3. The PC prior is defined as an exponential distribution on the distance,

$$\pi(d(\tau)) = \lambda \exp(-\lambda d(\tau)),$$

with rate $\lambda > 0$. The PC prior for τ follows by a change of variable transformation, leading in this case to a type-2 Gumbel distribution with parameters $(1/2, \lambda)$:

$$\pi(\tau) = \frac{\lambda}{2} \tau^{-3/2} \exp(-\lambda \tau^{-1/2}), \quad \tau > 0, \lambda > 0. \quad (3)$$

4. The parameter λ in (3) can be selected by the user based on his prior knowledge of τ (or an interpretable transformation of it such as the standard deviation). This can be expressed in an intuitive way with a probability statement, e.g. setting U and a such that $\mathbb{P}(1/\sqrt{\tau} > U) = a$, so that $\lambda = -\log(a)/U$. Knowledge on the marginal standard deviation can aid in choosing a sensible value for U ; Simpson et al. (2017) provide a practical rule of thumb: once the precision τ is integrated out, the marginal standard deviation of the random effect for $a = 0.01$ is about $0.31U$.

3 Partitioning the variance between main and interaction

We present below the VP model assuming model (1), but everything applies straightforwardly to model (2) as well; details about the VP version of model (2) can be found in Section 4.

From model (1) it is hard to quantify the relative contribution of the main and interaction components to the total variance, because the involved precision parameters are not interpretable in terms of the variance explained by the associated components. Our proposal is to reparametrize model (1) as a weighted sum of two IGMRFs representing the main and interaction components by means of a mixing parameter $\gamma \in [0, 1]$. We include a further mixing parameter $\phi \in [0, 1]$ to distribute the variance between the temporal and spatial main effects. Assume model (1), the reparametrized version of the linear predictor is

$$\begin{aligned} \eta_{ij} &= \alpha + \sqrt{\tau^{-1}} \left[\sqrt{1-\gamma} (\sqrt{1-\phi}\beta_{1i} + \sqrt{\phi}\beta_{2j}) + \sqrt{\gamma}\delta_{ij} \right], \\ \beta_1 &\sim N(\mathbf{0}, \tilde{\mathbf{R}}_1^-), \quad \beta_2 \sim N(\mathbf{0}, \tilde{\mathbf{R}}_2^-), \quad \delta \sim N(\mathbf{0}, \tilde{\mathbf{R}}_I^-), \end{aligned} \quad (4)$$

where $\tau > 0$ is an overall precision parameter, $0 < \gamma < 1$, $0 < \phi < 1$. We consider a RW1 or a RW2 prior on the temporal main effect β_1 and an ICAR prior on the spatial main effect β_2 as specified in Section 2. Note that, differently from model (1), the precision structures $\tilde{\mathbf{R}}_1$, $\tilde{\mathbf{R}}_2$ have been scaled according to Sørbye and Rue (2014). The interaction term δ is modelled as a Kronecker product IGMRF; following Knorr-Held (2000) we consider interaction types I, II, III, and IV as described in Table 1.

Model (4) includes the same vectors of random effects as model (1), but in contrast to model (1), we now have very intuitive hyperparameters: τ is the total precision, i.e. τ^{-1} is the total generalized variance, and γ and ϕ are two interpretable mixing parameters. The value of γ can be interpreted as the proportion of total variance explained by the interaction δ . The variance explained by the main effects is therefore given by $\tau^{-1}(1-\gamma)$: $1-\phi$ quantifies the proportion of such variance which can be attributed to the temporal random effects β_1 , with ϕ being the proportion attributed to the spatial random effects β_2 .

We need to assign priors to the overall precision parameter τ and the mixing parameters γ and ϕ . In the next section we focus on the prior for γ , and leave prior choice for the remaining parameters to Section 4.

3.1 A Penalized Complexity prior for γ

Our choice of a PC prior for γ follows naturally from the model reparametrization in Eq. (4) and provides a way of eliciting the prior in a very intuitive way. Furthermore, it avoids overfitting by construction hence guaranteeing a parsimonious model. Our PC prior for γ (see Result 1 below) is based on the assumption that the interaction model in (4) shrinks to the main effects model ($\beta_1 + \beta_2$).

Result 1. *Let us assume a model of the form (4), for all types of interaction in Table 1:*

1. The distance from the base model is

$$d(\gamma) \simeq \sqrt{\gamma}, \quad 0 < \gamma < 1$$

2. The PC prior for γ with base model $\gamma = 0$ is

$$\pi(\gamma) = \frac{\theta \exp(-\theta \sqrt{\gamma})}{2\sqrt{\gamma}(1 - \exp(-\theta))} \quad 0 < \gamma < 1, \theta > 0. \quad (5)$$

The proof can be found in Supplemental material A.1-A.3.

The scaling of the PC prior for γ , i.e. the choice of θ in Eq. (5), is done by defining the probability of a tail event on γ . The parameter θ controls the strength of penalisation for deviating from the base model; the higher the θ the greater the penalty. We suggest setting U and a such that $\mathbb{P}(\gamma < U) = a$; this way θ is obtained by numerically solving:

$$\frac{1 - \exp(-\theta \sqrt{U})}{1 - \exp(-\theta)} = a, \quad a > \sqrt{U}.$$

Note that it is not possible to assign equal weight to the main and interaction terms in the model, i.e. $U = a = 0.5$ because of the constraint $a > \sqrt{U}$. However, we can always encode a fair amount of uncertainty into the prior by choosing a close to 1 and large values of U . In the left panel of Figure 1, θ is obtained using $a = 0.99$ and three different values for U . A large U allows for more flexibility as the corresponding density curve decreases steadily towards zero as γ increases, while for a small value of U the density curve drops towards zero quite sharply, strongly penalizing any deviation from the base model. For comparison, the right panel in Figure 1 shows the prior on γ that corresponds to using a Gamma prior on all three precision parameters in model (1) for three different parameter choices. The figure illustrates how the resulting prior on γ depends strongly on the chosen values for the Gamma parameters, going from one extreme to the other in terms of prior weight on the base model.

Results from a simulation study reported in Supplemental material B indicate that the posterior mean estimates of γ are reasonably close to the true value under different scenarios. We have observed stable results for several choices of U , unless one defines on purpose an unflexible prior, where most of the probability mass is placed near the base model (e.g. when adopting $a = 0.99$ and a small $U = 0.05$). Results are comparable to those obtained using a Uniform prior on γ unless there is no interaction (i.e. $\gamma = 0$), in which case the uniform leads to greater bias when the population at risk is small.

4 Examples

As introduced in Section 2, model (2) is more common in practice and indeed it is the model adopted in this section for both case-studies. In the case of structured and unstructured main effects, another set of parameters ψ_1 and ψ_2 can be included to further distribute the variance, so that model (4) becomes:

$$\eta_{ij} = \alpha + \sqrt{\tau^{-1}} \left(\sqrt{1 - \gamma} \left(\sqrt{1 - \phi} \left(\sqrt{1 - \psi_1} \beta_{1i} + \sqrt{\psi_1} \epsilon_{1i} \right) + \sqrt{\phi} \left(\sqrt{1 - \psi_2} \beta_{2j} + \sqrt{\psi_2} \epsilon_{2j} \right) \right) + \sqrt{\gamma} \delta_{ij} \right), \quad (6)$$




Figure 1: Left panel: PC prior $\pi(\gamma)$ using $a = 0.99$ and three different values for U . Right panel: implied prior on γ when a Gamma prior is used on all three precision parameters $\tau_1, \tau_2, \tau_{12}$.

where $\tau > 0, 0 < \gamma < 1, 0 < \phi < 1, 0 < \psi_1 < 1, 0 < \psi_2 < 1, \epsilon_1 \sim N(\mathbf{0}, \mathbf{I}_{n_1}), \epsilon_2 \sim N(\mathbf{0}, \mathbf{I}_{n_2})$ and β_1, β_2 and δ as in model (4). Result 1 about the PC for γ still holds; see Supplemental material A.3.

Note that the parameters in model (6) are identifiable as the model is just a reparametrized version of the classic space-time interaction model (2), where each random effect has its corresponding precision parameter. The number of parameters is exactly the same; in fact, it can be shown that there is a one-to-one mapping between the parameters of both versions of the model. As in model (1), appropriate constraints need to be imposed to ensure identifiability of the terms in (6). The constraints on the interaction term are summarized on Table 1, while on the temporal and spatial structured main effects it is enough to impose a sum to zero constraint.

In the next two examples, we use the PC prior in Eq. (3) for τ and the PC prior in Eq. (5) for γ . Regarding ϕ, ψ_1 and ψ_2 , we simply choose a uniform on $(0,1)$ as a prior for each of them, but other choices are possible. In fact, a PC prior could also be used for ϕ following the work by Fuglstad et al. (2020), who also consider the use of a Dirichlet prior where the base model attributes equal weights to each component, thus expressing ignorance about how the variance is distributed. Similarly, one could use a PC prior on each ψ_1 and ψ_2 as in Riebler et al. (2016), considering as base model the absence of structured effects.

All the VP models presented in the next two examples were run using R-INLA (Rue et al., 2009), see code in Supplemental material D.

4.1 Ohio lung cancer

We illustrate our model using the Ohio lung cancer data (Knorr-Held, 2000; Knorr-Held and Besag, 1998; Waller et al., 1997) which is available at <http://www.biostat.umn.edu/~brad/data2.html>. These data report yearly counts of lung cancer deaths for white males from 1968 to 1988, in the 88 counties of Ohio. Figure 2 left panel displays the time series of mortality rate for all counties. Our aim is not to find the best model for this data, but to show what our approach can add

in terms of interpretability of the results compared to a classical analysis as performed in Knorr-Held (2000).

4.1.1 Model

Let y_{ij} be the number of deaths at time $i = 1, \dots, 21$ in county $j = 1, \dots, 88$ and pop_j be the population at risk in county j , we consider the model proposed in Knorr-Held (2000) assuming structured and unstructured effects for both space and time main effects, plus a space-time interaction term. The classical parameterization in Knorr-Held (2000) follows,

$$\begin{aligned} y_{ij} &\sim \text{Bin}(\text{pop}_j, \exp(\eta_{ij}) / \exp(1 + \eta_{ij})), \\ \eta_{ij} &= \alpha + \underbrace{\beta_{1_i} + \epsilon_{1_i} + \beta_{2_j} + \epsilon_{2_j}}_{\text{main}} + \underbrace{\delta_{ij}}_{\text{int}}, \end{aligned} \quad (7)$$

where the main effects are modelled as:

$$\begin{aligned} \epsilon_1 &\sim N(\mathbf{0}, \tau_{\epsilon_1}^{-1} \mathbf{I}_{n_1}); & \epsilon_2 &\sim N(\mathbf{0}, \tau_{\epsilon_2}^{-1} \mathbf{I}_{n_2}); \\ \beta_1 &\sim N(\mathbf{0}, \tau_1^{-1} \mathbf{R}_1^-); & \beta_2 &\sim N(\mathbf{0}, \tau_2^{-1} \mathbf{R}_2^-). \end{aligned}$$

where \mathbf{R}_1 and \mathbf{R}_2 are the unscaled structure matrices of a RW1 (for time) and an ICAR (for space). The space-time interaction is modelled by a Kronecker product IGMRF built on the precision matrices of the structured components β_1 and β_2 . All the interaction types in Table 1 are considered in the following analysis. This model would require priors for the precision hyperparameters $\tau_{\epsilon_1}, \tau_{\epsilon_2}, \tau_1, \tau_2, \tau_{12}$.

Instead of working with the above model, we assume the VP model in Eq. (6), with scaled structure matrices, with the priors for $\tau, \gamma, \phi, \psi_1, \psi_2$ stated in Section 4. For τ 's PC prior we set $a = 0.01$ and $U = 1/0.31$ following the rule of thumb described in Section 2.2. The PC prior for γ is scaled by imposing $U = 0.5, a = 0.99$; results (not shown here) were stable for varying $U = \{0.05, 0.5, 0.95\}$.

4.1.2 Results

Table 2 reports various model selection criteria for the VP model, for interaction types I, II, III and IV, namely DIC (Spiegelhalter et al., 2002), WAIC (Watanabe, 2013), leave-one-out log score (LOOLS), computed as $-\sum_{i=1}^n \log \pi(y_i | y_{-i})$, and the log-marginal likelihood (logMLIK), $\pi(y | \mathcal{M})$, which quantifies the likelihood of the data y under a given model \mathcal{M} . PC priors enhance the marginal likelihood as a simple and effective tool for *fair* model comparison, when the compared models have similar structure and only differ on a particular component (Sørbye and Rue, 2018; Ventrucci et al., 2020). Assume \mathcal{M}_1 and \mathcal{M}_2 are the interaction type I and II models, respectively: these models are the same except for a different type of interaction. The Bayes factor (Kass and Raftery, 1995) is defined as

$$K = \frac{\pi(y | \mathcal{M}_1)}{\pi(y | \mathcal{M}_2)} = \frac{\pi(\mathcal{M}_1 | y) \pi(\mathcal{M}_2)}{\pi(\mathcal{M}_2 | y) \pi(\mathcal{M}_1)}. \quad (8)$$

The scale parameter θ of the PC prior for γ , which controls the decay rate from the base model (the model with no interaction), has to be chosen for both \mathcal{M}_1 and \mathcal{M}_2 . We can handle this choice conveniently by setting the same θ for \mathcal{M}_1 and \mathcal{M}_2 , which implies that $\pi(\mathcal{M}_2)/\pi(\mathcal{M}_1)$ in Eq.(8) cancels out and the Bayes factor turns out to be the ratio of the posterior odds. For all the four interaction models in Table 2 we follow this strategy and set the same decay rate for the PC prior on γ . The advantage is that the *a priori* contribution of the interaction to the total (generalized) variance

is the same, no matter what interaction type is assumed; this is a desirable feature when having to choose among different types of interaction models the one that best fits the data. Therefore, we suggest comparing the logMLIK values for model choice purposes here; furthermore, DIC is known to favour complexity in models with many random effects (Riebler and Held, 2010).

From Table 2 we see that DIC, WAIC and LOOLS point to type II (followed by type IV) as the best model; a similar conclusion based on DIC was found in Knorr-Held (2000). Interestingly, logMLIK is largest for type I which indicates that the model with main effects plus an individual-level random effects capturing unstructured variation may be a better description of the Ohio data.

Table 2: Model comparison criteria (computed using R-INLA) for the VP model, under the four interaction types.

interaction type	logMLIK	DIC (deviance; p_D)	WAIC	LOOLS
I	-5623.52	10945.23 (10722.68; 222.55)	10957.75	5489.75
II	-6759.84	10916.00 (10739.09; 176.91)	10931.3	5469.43
III	-6098.89	10957.86 (10792.68; 165.18)	10980.99	5496.28
IV	-7200.13	10919.23 (10755.06; 164.17)	10934.82	5470.89

In order to show now the gain of using our approach compared to a classical analysis we start by discussing some plots obtained for type I interaction about the main effects. The top right panel in Figure 2 displays the estimated main temporal effect, in the scale of the linear predictor, decomposed into its structured and unstructured (iid) components. The unstructured effects looks very flat compared to the structured ones which is probably responsible for most of the temporal variation in the relative risk. The relative risk increases roughly linearly in time, with a less steep increase towards the end of the time window. The bottom panels in Figure 2 display the estimated structured (left) and iid (right) spatial effects in the scale of the linear predictor. Here the unstructured effect shows larger variability than the structured one which shows a very smooth spatial gradient from north-west to south-east. A visual inspection of this sort, also possible when the classical model is used, gives useful insights into the spatial and temporal patterns in the data. However, it does not allow proper quantification of the variance attributable to the various sources (main, interaction, spatial and temporal effects, etc), while this quantification is readily available from our VP model.

Table 3 reports the mean (with 2.5 and 97.5 quantiles between brackets) of the posterior distribution of the mixing parameters $\gamma, \phi, \psi_1, \psi_2$, from which we can understand quantitatively the contribution of the main sources of variation. In particular, rows 1 and 2 report the total variance partitioned into interaction versus main effects; rows 3 and 4 quantify how the variance attributable to the main effects is partitioned into space and time; rows 5 to 6 and 7 to 8 give the variance partitioning for the structured versus iid for space and time, respectively. The main findings about the variation of the spatio-temporal mortality risk pattern are as follows. First, the estimated contribution of the interaction is about 4.8%, which means that the main effects are responsible for most of the variability in mortality risk with the interaction playing a minor role in describing this data. This is reasonable for non-infectious diseases such a cancer (Abellan et al., 2008). Second, space is responsible for about 87.5% of the variation in risk explained by the main effects, which is hard to grasp from only looking at the plot of the main temporal and spatial effects in Figure 2. This result highlights the fact that lung cancer in Ohio had, in the period of time considered, larger variability over space than time which could be informative for policy makers and epidemiologists and may contribute to generate hypothesis on the role played by possible environmental risk factors in the region. Third, within the main spatial and temporal effects, the structured component is predominant for time, while the iid component is predominant for space. However, these estimated contributions, and in particular the latter, are affected by greater uncertainty than the previous estimates, indicating

that the data are less informative about the posterior for ψ_1 and ψ_2 than they are for γ and ϕ . These findings are stable across the different types of interactions (see Supplemental material C).

Table 3: Variance partitioning table for Ohio lung cancer, type I interaction. The column named *contribution* reports the posterior mean of the hyper-parameters displayed in the column named *estimator*, with 0.025 and 0.975 posterior quantiles between brackets. All values are in a $(0, 1)$ interval and indicate the proportional contribution of the model component *level 2* to the variance explained by the model component *level 1*.

Model component level 1	level 2	Variance Partitioning	
		estimator	contribution
main+int	main	$1 - \hat{\gamma}$	0.952 (0.913, 0.979)
	int	$\hat{\gamma}$	0.048 (0.021, 0.087)
main	space	$\hat{\phi}$	0.875 (0.765, 0.946)
	time	$1 - \hat{\phi}$	0.125 (0.054, 0.235)
time	iid	$\hat{\psi}_1$	0.069 (0.010, 0.229)
	str	$1 - \hat{\psi}_1$	0.931 (0.771, 0.990)
space	iid	$\hat{\psi}_2$	0.658 (0.273, 0.925)
	str	$1 - \hat{\psi}_2$	0.342 (0.075, 0.727)

4.2 Covid-19 in Italy

We use the VP model to study Covid-19 incidence variations across space and time in Italy. Data cover all of the 107 Italian provinces and span a period of time that goes from the onset of the pandemic on 24th February 2020 to late July 2021 for a total of 70 weeks; the full dataset is made available by the Italian National Institute of Health through the website <https://github.com/pcm-dpc/COVID-19>. Data are originally available on a daily basis, but we aggregate them by week to smooth out artefactual patterns mainly due to delays in reporting new cases. The final dataset consists of weekly counts of new Covid-19 cases y_{ij} , for week $i = 1, \dots, 70$ and province $j = 1, \dots, 107$, and the population at risk for each province pop_j .

Our goal is to analyze the sources of variation in Covid-19 incidence rates in a scale between 0 and 1, which is easy to interpret and visualize and provides a clear idea of the contribution of each source. We follow the ideas in Picado et al. (2007) in considering the interaction term as a measure of local heterogeneity, which can be seen as an indirect measure of how effective the control measures are. Hence a primary interest is to quantify the contribution of the interaction to the total variability, i.e. the posterior estimate for γ . Our second interest is to investigate changes in the estimated local heterogeneity across geographical macro-regions and time windows. We run two analysis: in the first one we fit the VP to the full dataset, in the second one we run the same VP model to separate subsets of the data which are constructed using combinations of geographical area, with levels north (N), centre (C) and south (S), and pandemic wave, with levels W1 and W2. The first wave (W1) covers the first 18 weeks and roughly indicates the national lock down period, while the second wave (W2) covers the rest of the time frame and indicates the period where restriction measures were set at a regional level. Data are displayed in Figure 3.

4.2.1 Model

We consider the binomial model in Eq. (7), where structured and unstructured random effects are specified for both space and time as main effects. We model the temporally structured effects as




Figure 2: Top left panel: time series of lung cancer (white males) disease rates per 10000 population at risk, for the 88 counties in the Ohio dataset. Top right panel: temporally structured and temporally unstructured components for type I interaction model, in the scale of the linear predictor. Bottom left and right panels show, respectively, the spatially structured and unstructured components for type I interaction model, in the scale of the linear predictor.

a RW1 (as we do not anticipate smoothness) and the spatially structured effects as an ICAR, and assume a type IV space-time interaction to capture potential complex space-time patterns which are not explained by the main space and time components. In this particular example, the spatial main effect may reflect differences on the public health policy strategies adopted in each area (for example different testing rates across provinces). Again, we avoid the classic parametrization and take advantage of the VP approach described in (6). Doing so, we can elicit the prior easily and describe the various sources of variability in the data in an intuitive way in terms of the mixing parameters $\gamma, \phi, \psi_1, \psi_2$.

Available information on the nature of the disease can be used to aid in parameter choice for the PC priors on γ and τ . Since we know that we are dealing with a contagious disease which evolves over time possibly in a different manner across provinces we anticipate a relevant contribution of the interaction term. Thus we choose θ in Eq. (5) by setting $U = 0.95, a = 0.99$, which implies a large probability that $\gamma < 0.95$. In choosing the scale parameter λ of the PC prior for τ in Eq. (3) we consider the scale of the logit transformed incidence rates and use the rule of thumb described in Simpson et al. (2017), imposing a marginal standard deviation equal to 2 for the incidence rates in the linear predictor (logit) scale.

4.2.2 Results

The left panel in Figure 4 reports the variance partitioning plot for the full Covid-19 dataset; this plot is just a graphical version of the variance partitioning table that was presented in Table 3 for the Ohio lung cancer data. This plot resembles the graphs in Gelman (2005) that summarize anova results in terms of estimated standard deviation for each bunch of random effects in the model. Our variance partitioning plot follows the same idea but represents the contribution of each source in a scale $(0, 1)$. The main effects account for the greatest proportion of the total variation. Within the main effects, the variability in incidence rates is mostly driven by the spatial component, in particular by the unstructured part of it (although the corresponding posterior estimates are highly uncertain), while for the temporal part is the structured component that explains most of the variability.

The middle and right panels in Figure 4 report the variance partitioning plot for the models fitted to different subsets of the full dataset to investigate whether the spatio-temporal pattern in Covid-19 cases is consistent or not across geographical areas (N, C, S) and pandemic waves (W1, W2). It is interesting to see that the impact of the interaction term is greater in the second wave than in the first one for all three areas, suggesting greater local heterogeneity during the second wave. This could reflect the fact that restriction measures went from being national in the first wave to being regional in the second one, so we expect greater heterogeneity over space during the latter. Within the first wave, the main effects are responsible for a greater proportion of variation in all three areas, but that attributable to the interaction is slightly greater in the South, followed by the North and then the Centre.

5 Discussion

In this paper, we revisit spatio-temporal disease mapping, with particular attention to the interaction models discussed in Knorr-Held (2000), and propose a new model that allows variance partitioning among the main effects and the space-time interaction. When defining priors on the hyperparameters that control complexity of each intrinsic GMRF component, it is important to bear in mind that the main effects belong to the null space of the interaction term. This means that the interaction can naturally be regarded as an extension of the model including the main effects alone. This idea leads




Figure 3: Weekly Covid-19 incidence rates in the North (left panel), Centre (central panel) and South (left panel) of Italy. The vertical dashed line marks the separation between the first (W1) and second (W2) wave.




Figure 4: Variance partitioning plot for Covid-19 full dataset (left panel), first wave (middle panel) and second wave (right panel). The middle and right panels allow comparison across northern (black), central (green) and southern (red) areas in Italy.

to a model reparametrization where a mixing parameter γ balances out the contribution of the main and interaction effects to the total variance. The proposed approach implicitly defines a joint prior on the precision parameters of the various terms in the classic parametrization of the model.

The advantages of this reparametrization are twofold; on the one hand, prior choice can be made in an intuitive manner using a PC prior, avoiding the issue of eliciting priors on hard-to-interpret precision parameters. In space-time disease mapping, the nature of the disease can provide useful information to elicit the prior; for example, for non-infectious diseases such as the one considered in the first case study most of the variation is expected to be explained by the main effects (Abellan et al., 2008). This knowledge can be easily passed onto the PC prior for the mixing parameter γ , while coding this information into a precision parameter in the classic parametrization would be far from easy. On the other hand, the posterior for γ becomes a useful tool to investigate variations in disease risk on a very practical scale and can provide useful insights into epidemiological interpretations. We have illustrated the use of the VP model in two examples; the variance partitioning tables and plots summarize the contribution of the different sources of variation in terms of proportion of explained (generalized) variance.

In a broader perspective, our work falls within the framework of variance distributing models as introduced by Fuglstad et al. (2020), and adds to the literature in considering intrinsic GMRF models. The variance partitioning approach proposed here may be adopted in all those applications where intrinsic GMRFs are meant as tools to perform smoothing in more than one dimension; for instance in the analysis of grid-data such as those arising from agricultural field trials or spatio-temporal data from environmental studies and ecological surveys.

References

- Abellan, J., Richardson, S., and Best, N. (2008). Use of space–time models to investigate the stability of patterns of disease. *Environmental Health Perspectives*, 116(8):1111–1119.
- Aldstadt, J. (2007). An incremental Knox test for the determination of the serial interval between successive cases of an infectious disease. *Stochastic Environmental Research and Risk Assessment*, 21:487–500.
- Bernardinelli, L., Clayton, D., Pascutto, C., Montomoli, C., M., G., and Songini, M. (1995). Bayesian analysis of space-time variation in disease risk. *Statistics in Medicine*, 14.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society Series B*, 36(2):192–225.
- Besag, J., York, J., and Mollie, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43:1–21.
- Diggle, P., Chetwynd, A., Häggkvist, R., and Morris, S. (1995). Second-order analysis of space-time clustering. *Statistical Methods in Medical Research*, 4(2):124–136. PMID: 7582201.
- Fahrmeir, L., Kneib, T., Lang, S., and Marx, B. (2013). *Regression: models, methods and applications*. Springer-Verlag, Berlin.
- Fong, Y., Rue, H., and Wakefield, J. (2010). Bayesian inference for generalized linear mixed models. *Biostatistics*, 11(3):397–412.

- Freni-Sterrantino, A., Ventrucci, M., and Rue, H. (2018). A note on intrinsic conditional autoregressive models for disconnected graphs. *Spatial and Spatio-Temporal Epidemiology*, 26:25–34.
- Frühwirth-Schnatter, S. and Wagner, H. (2010). Stochastic model specification search for Gaussian and partial non-Gaussian state space models. *Journal of Econometrics*, 154(1):85–100.
- Frühwirth-Schnatter, S. and Wagner, H. (2011). Bayesian variable selection for random intercept modeling of Gaussian and non-Gaussian data. In *J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith and M. West (Eds.)*, pages 165–200. Bayesian Statistics 9, Oxford.
- Fuglstad, G.-A., Hem, I. G., Knight, A., Rue, H., and Riebler, A. (2020). Intuitive joint priors for variance parameters. *Bayesian Analysis*, 15(4):1109–1137.
- Gelman, A. (2005). Analysis of variance-why it is more important than ever. *The Annals of Statistics*, 33(1):1 – 53.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3):515–534.
- Goicoa, T., Adin, A., Ugarte, M.D., and Hodges, J.S. (2018). In spatio-temporal disease mapping models, identifiability constraints affect PQL and INLA results. *Stochastic Environmental Research and Risk Assessment*, 32: 749–770.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.
- Knorr-Held, L. (2000). Bayesian modelling of inseparable space-time variation in disease risk. *Statistics in Medicine*, 19(17-18):2555–2567.
- Knorr-Held, L. and Besag, J. (1998). Modelling risk from a disease in time and space. *Statistics in Medicine*, 17(18):2045–2060.
- Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22:79–86.
- Lawson, A. (2018). *Bayesian Disease Mapping: Hierarchical Modeling in Spatial Epidemiology (3rd ed.)*. Chapman and Hall/CRC.
- Lindgren, F. and Rue, H. (2008). On the second-order random walk model for irregular locations. *Scandinavian Journal of Statistics*, 35(4):691–700.
- MacNab, Y.C. (2022). Bayesian disease mapping: Past, present, and future. *Spatial Statistics*, <https://doi.org/10.1016/j.spasta.2022.100593>
- Martinez-Beneito, M. and Botella-Rocamora, P. (2019). *Disease Mapping: From Foundations to Multidimensional Modeling (1st ed.)*. Chapman and Hall/CRC.
- Picado, A., Guitian, J., and Pfeiffer, D. (2007). Space-time interaction as an indicator of local spread during the 2001 FMD outbreak in the UK. *Preventive Veterinary Medicine*, 79:3–19.
- Riebler, A. and Held, L. (2010). The analysis of heterogeneous time trends in multivariate age-period-cohort models. *Biostatistics*, 11(1):57–69.

- Riebler, A., Sørbye, S. H., Simpson, D., and Rue, H. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, 25(4):1145–1165. PMID: 27566770.
- Robertson, C., Nelson, T., MacNab, Y., and Lawson, A. (2010). Review of methods for space-time disease surveillance. *Spatial and Spatio-temporal Epidemiology*, 1:105–116.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields*. Chapman and Hall/CRC.
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*, 71(2):319–392.
- Simpson, D., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H. (2017). Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, 32(1):1–28.
- Sørbye, S. and Rue, H. (2014). Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spatial Statistics*, 8:39–51.
- Sørbye, S. and Rue, H. (2018). Fractional gaussian noise: Prior specification and model comparison. *Environmetrics*, 29(5-6):e2457.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B*, 64:583–639(57).
- Ventrucci, M., Cocchi, D., Burgazzi, G., and Laini, A. (2020). Pc priors for residual correlation parameters in one-factor mixed models. *Statistical Methods & Applications*, 29(4):745–765.
- Ventrucci, M. and Rue, H. (2016). Penalized complexity priors for degrees of freedom in bayesian p-splines. *Statistical Modelling*, 16(6):429–453.
- Wahba, G. (1978). Improper priors, spline smoothing and the problem of guarding against model errors in regression. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40(3):364–372.
- Wakefield, J. (2007). Disease mapping and spatial regression with count data. *Biostatistics*, 8(2):158–183.
- Waller, L. A., Carlin, B. P., Xia, H., and Gelfand, A. E. (1997). Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association*, 92(438):607–617.
- Watanabe, S. (2013). A Widely Applicable Bayesian Information Criterion. *Journal of Machine Learning Research*, 14:867–897, 2013.