Review of spatio-temporal models
for disease mapping

Antonio López-Quílez, Facundo Muñoz

June 10, 2009

1   Introduction

The EUROHEIS2 project (http://www.euroheis.org) is aimed at improving the analysis, reporting and dissemination of environmental health information. The project will further develop the health and environment information system for health threat analysis (the Rapid Inquiry Facility – RIF) initiated in the previously funded EUROHEIS project. One of the specific objectives is to include spatio-temporal methods for disease mapping in the RIF.

Statistical techniques for disease mapping have become very popular in public health analysis. These methods enable the smoothing of ecological health indicators accounting for the geographical structure of the units under study. As a consequence, more reliable risk estimates in less populated areas are obtained due to the sharing of information between neighbouring regions, which are assumed to share common risk factors. In this way, it becomes possible to display the geographical distribution of risk even in small areas.

But disease risks are variable in space and time, and supporting risk management should ideally incorporate spatio-temporal analysis tools. Recently, several spatio-temporal disease mapping techniques have been proposed. However, the implementation of these methods is not always easy or adequate for a quick response tool. Furthermore, there is not a wide consensus on how to describe temporal and spatial evolution at the same time in a proper way. Therefore, a special effort is necessary to indentify which methods are suitable for inclusion in the RIF.

The objective of this report is to evaluate the most representative spatio-temporal methods for disease mapping found in the literature in order to promote the scientific discussion of their properties, and also to point out the key aspects that should be taken into account when considering their implementation in a RIF-like application.

In Section 2 we outline a general framework for spatio-temporal models, breaking them up into four stages: the probabilistic model for observations, the components of the linear predictor, the structures of the effects and the inference methodology. For each stage, the most commonly used alternatives in the literature are discussed. Then in Section 3 we do the actual review, classifying each model according to the structure of the temporal trends that may arise and discussing the relative advantages and disadvantages of the different approaches. In Section 4 we present the elements that we have identified as key aspects when it comes to deciding what kind of spatio-temporal model is more

suitable for an end-user, fast and general disease mapping application like the RIF. Finally, we summarize the conclusions in Section 5.

2 General framework and notation

Throughout this report, we will assume that the region under study is divided into I regions (counties, municipalities, etc.) indexed by $i = 1, \ldots, I$. The temporal dimension will be indexed by $j = 1, \ldots, J$ representing each period of time under study, typically years. Let n_{ij} denote the number of persons–time at risk in area i and period j, and y_{ij} the corresponding observed cases or deaths. In some cases, an additional categorization is used apart from region and period; for instance, when considering one or more covariates such as age, race, sex or risk factors. In these cases an additional subindex $k = 1, \ldots, K$ will identify each combination of existing categories. For example, if sex and race are used as covariates, with levels {male, female} and levels {white, non-white} respectively, then there would be four groups, namely: white males, white females, non-white males and non-white females.

2.1 Probabilistic model for observations

Conceptually, the observations y_{ijk} are assumed to be a conditionally independent random sample from a given probability distribution from the exponential family. Typically, a Poisson distribution will be preferred when the observed values in each region and period are expected to be low. With some non-rare diseases, a Binomial distribution could be more appropriate (Knorr-Held and Besag, 1998).

In any case, the observed data y_{ijk} depends, in first place, on the number of persons–time at risk n_{ijk}. For this reason, the expected value for y_{ijk} is factorized as $n_{ijk} \cdot r_{ijk}$, where r_{ijk} denotes the *risk* in region i, period j and group k.

Commonly, the number of persons–time at risk is standardized by age (external or internally); in these cases, for each combination of region, period and risk the *expected number of cases* E_{ijk} is computed, and the *relative risk* is modelled. We will denote the relative and absolute risk in the same way r_{ijk}, as it will always be clear from the context which one of the concepts is being used.

There is a lot of literature on whether to model data through a Poisson or Binomial distribution, when and how to standardize by age, and what the advantages and disadvantages of working with absolute or relative risks are. In addition, most of the time the decision is conditioned by data availability or data quality.

The modelling decisions taken at this level are completely independent of the spatio-temporal modelling structure, and consequently will not restrict or condition future decisions regarding spatio-temporal modelling, which is the focus of this report. Therefore, although the other authors have followed different approaches within the papers reviewed, we will not discuss this aspect any further.

2.2 Components of the linear predictor η_{ijk}

Depending on the model chosen for observations, the logarithm or the log-odds of the risk r_{ijk} (or rarely, some other *link* function of the risk) is called the *linear predictor* η_{ijk}, which is just a re-scaled version of the risk having nicer properties that make predictions on it more reliable.

The linear predictor is usually expressed additively as the sum of some components or *effects* that can be interpreted as individual and independent contributions to the risk in that region and period. Therefore, the linear predictor may have all of the following terms or, more commonly, a subset of them.

$$\eta_{ijk} = \text{Intercept} + \underbrace{C_k + S_i + T_j}_{\text{main effects}} + \underbrace{CS_{ik} + CT_{jk} + ST_{ij} + CST_{ijk}}_{\text{interaction terms}} + \varepsilon_{ijk} \quad (1)$$

Next, we briefly describe each one of the possible effects. Note that the contribution of a given term can increase or decrease the risk. In any case, we talk about the *additional* risk due to that effect, taking into account that it could be a *negative* or *positive* contribution.

Intercept: The intercept term gives a starting amount of risk that is shared by all regions, periods and groups. It is usually included in every model, but is sometimes included within the mean value of a random effect, in one of the following terms.

Main effects C_k, S_i and T_j: These components represent the additional risk of belonging to group k, living in region i and period j respectively.

Second order interaction terms CS_{ik}, CT_{jk} and ST_{ij}: These components represent the contribution to the risk due to a combination of the effects that cannot be explained additively by the main effects. For example, suppose that for some reason there is a significant under-registration of cases in a region in a given year. This fact cannot be explained with the spatial effect, since the region has higher values at other periods of time. In the same way, it is not a temporal effect of that period. This is a typical example of a spatio-temporal effect. Abnormal values estimated for this effect can lead to the detection of hidden problems that are masked within the rest of the effects, and therefore are not easily identified from the data itself.

Third order interaction term CST_{ijk}: This effect represents an additional risk affecting one specific group of people k, living in a specific region i and in a specific period of time j. This component is rarely used because it greatly increases the complexity of the model.

Extra variability term ε_{ijk}: Frequently, an explicit unstructured extra variability term is included in order to capture the overall effect of other minor factors. It is implemented as a white noise random effect, and can introduce noise either globally or into a specific subspace.

$$\varepsilon_d | \lambda_\varepsilon \overset{\text{iid}}{\sim} \text{No}(0, \lambda_\varepsilon), \qquad d \in \{i, j, k, ij, ik, jk, ijk\} \quad (2)$$

Note that the correct interpretation of a given effect frequently depends on the presence of some other effects. While this section provides a general guide for interpretation, it should be handled carefully.

2.3 Structures of the effects

Covariates effect C_k: Usually the effect of the covariates is expressed as a linear model on them, or on a transformation of them. Depending on the type of covariate, it can also be stratified into several categories and included as a fixed effect or as a structured random effect.

A special case is the modeling of the age effect as a covariate, instead of standardizing by age. Usually age is stratified into age groups and included as a fixed effect.

Spatial effect S_i: When the region-specific data are scarce, the classical fixed effects model with maximum-likelihood estimation often leads to unsatisfactory estimates of the spatial effects in each area. In disease mapping this problem has been overcome by a Bayesian approach which models the spatial effects as *random effects*, through a prior distribution. Specifically, the spatial effects may give rise to a *spatially unstructured* variation (*heterogeneity*) or to a *spatially structured* variation (*clustering*).

In the model for *unstructured heterogeneity*, the spatial effects ϕ_i are assumed to be sampled from a normal distribution with mean 0 and precision λ_ϕ

$$S_i = \phi_i | \lambda_\phi \stackrel{\text{iid}}{\sim} \text{No}(0, \lambda_\phi) \qquad \text{(heterogeneity effect)} \qquad (3)$$

In the *clustering* model, the mean of the structured effect θ_i is allowed to depend on the neighbouring θ_js through the Gaussian Conditionally Autoregressive (CAR) distribution (see, for example, Besag and Kooperberg, 1995). Formally, the joint distribution of the vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_I)'$ is denoted

$$S_i = \boldsymbol{\theta} | \lambda_\theta \sim \text{CAR}(\lambda_\theta) \qquad \text{(clustering effect)} \qquad (4)$$

The choice between the *clustering* and the *heterogeneity* model depends upon our prior belief about the scope of dominant risk determinants. Risk determinants exceeding the limits of one or more regions leads to a *clustering* model since they induce similar risk values in neighbouring regions. On the contrary, when the scope of risk determinants is smaller than a region's size it leads to a *heterogeneity* model.

Among all the proposals for performing risk smoothing which have appeared in the literature, the one stated by Besag et al. (1991) has had a particular impact. Most of the models in this review base their spatial modelling on that approach, which we will refer to as *the BYM specification*. The risk associated with a region is broken down as the sum of a heterogeneity and a clustering effect

$$S_i = \phi_i + \theta_i \qquad \text{(BYM specification)} \qquad (5)$$

where vectors $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$ are distributed as in equations 3 and 4.

Another approach to modelling the spatial main effect is the use of (two-dimensional) splines. In particular, Penalized splines (Eilers and Marx,

1996) are well-established for smoothing Gaussian and non-Gaussian data in one or more dimensions (Currie et al., 2006). Recently, a combination of Penalized splines and CAR models were combined in order to separately model the spatial variation on a large and small scale respectively (Lee and Durbán, 2009).

Temporal effect T_j**:** In contrast to restrictive evolution models such as linear or polynomial parametric models, most of the time, a smooth and flexible evolution is preferred. So it is frequent to model this term as a structured random effect, ensuring that contiguous periods are likely to be similar, but allowing for flexible shapes in the evolution curve, specially when long periods of time are being considered. First or second order random walks (RW1, RW2), autoregressive processes (AR), or splines are the models that have been used within the reviewed papers.

Sometimes time has also been stratified into a few blocks of time and modelled as a fixed effect, thus estimating the effect of each block independently from the others.

Covariate interactions CS_{ik}**,** CT_{jk} **and** CST_{ijk}**:** These are unfrequent effects. It is more usual to assume independence between the covariates and the spatial and temporal effects.

An example of such an interaction can be found in the paper by Sun et al. (2000). They stratified age into four groups and assumed that each age-group could present a different evolution pattern in mortality rates. So they modelled this interaction as a linear parametric function of time, and found that mortality rates showed a decreasing trend for the youngest group but an increasing trend for the other three over the period under study. Note that this type of study is not possible when using age-standardized rates.

Spatio-temporal interaction ST_{ij}**:** This is the key aspect of spatio-temporal models, and possibly the most difficult one, because of the many possibilities and the lack of an accepted standard that functions well.

Knorr-Held (2000) established a classification of four possible types of interaction between spatial and temporal random effects. Here we present a generalization, using the same abstract types of interaction but not restricted to any specific form of spatial and temporal effects. In this way we are able to classify all of the reviewed papers into one of these types, according to conceptual aspects –instead of technical– of the spatio-temporal modelling.

Type I interaction This can be thought of as independent unobserved covariates for each combination of region and period (i, j), thus without any structure. However, note that if spatial and temporal main effects are present in the model, then this interaction effect only implies independence in the *deviations* from them. Contribution to risk in neighbouring regions or in consecutive periods of time can still be highly correlated, due to the main effects. This is a global space-time heterogeneity effect, and is usually modelled as white noise.

This interaction can represent all kinds of –non persistent– circumstances that can cause a slight increase or decrease in the rates in a specific region-period, allowing for random –independent– oscillations around the expected rates. Moreover, it is a simple way of implementing a spatio-temporal interaction, allowing data to show if there is anything worthy of further investigation.

Type II interaction Here each region has a specific evolution structure that is independent of that in the neighbouring regions. The evolution structure for each region may have as many forms as the temporal main effect itself.

In the same way as before, this does not mean that each region has an evolution independent of the neighbouring ones, since they may share a common temporal main effect. Independence only affects the *deviation* from the global trend.

This is suited to modelling factors affecting specific regions and inducing deviations from the global trend. It was the type of interaction selected in the example studied in Knorr-Held (2000) because it provided a good balance between fit and complexity.

Type III interaction Analogously, the interaction can be assumed to have a spatial structure for each period, independent of adjacent periods (its neighbours in time). Like the spatial clustering effect, this is typically modelled with a CAR distribution for each period. Note that the inclusion of an additional heterogeneity spatial term (such as in the BYM specification) would produce a Type I interaction term. Implicitly, here it is assumed that each specific region may have a slight deviation from the global trend, but that this deviation is likely to be similar to that in the neighbouring regions while, at the same time, independent of that in the previous or subsequent period of time.

Such an interaction could represent situations where an unobserved regional factor is affecting an area containing two or more adjacent regions, but not persistent in time. For example, a one-off waste spill into a river could affect those regions in its way, temporally deviating observed rates from what is expected for a set of neighbouring regions.

Type IV interaction From a theoretical point of view, the most interesting form of interaction arises when deviations from the global trends are assumed to be correlated with their neighbours both in space and time.

This can model hidden factors whose effects exceed the limits of one or more regions and also persist for more than one period of time. It is also the most efficient way of extracting information from data, specially in the case of rare-diseases or less populated regions, since the risk estimation for a region-period is performed not only based on the locally observed data but also on that in neighbouring regions and periods. This is called *borrowing strength* from neighbours. Note, however, that such a dense structure would not be likely to fit an unexpected outbreak of cases.

Regarding the implementation of this type of interaction, the most simple approach is to use a parametric function of time (typically

polynomials of first or second degree), where the parameters are spatially dependent. In this way, neighbouring regions have similar trends.

Another approach is to model the interaction term as a random effect, which allows for the implementation of any of the interaction types as explained next.

Finally, a non parametric smoothing function (such as splines) might be specified for each region, incorporating spatial correlation into the distribution of its coefficients.

Regarding the technical implementation of the interaction effect, one remarkable approach is that suggested by Clayton (1996), in which they are modelled as a random effect with a precision matrix λK, where λ is an unknown scalar to be estimated from the data and K is a *structure matrix* computed as the Kronecker product of the structure matrices of those main effects which are assumed to interact. This is the approach followed by Knorr-Held (2000), and in all posterior proposals that are based on random effects.

Note that the four types of interaction are suited to modelling different kinds of phenomena, mainly depending on their scope. It would be nice to be able to include one term for each type of interaction in order to capture all kinds of possible situations. However, this would cause an overmodelling of the data. In practice, it is necessary to find out what kind of phenomenon is dominating in order to capture it and explain it with a suitable model. This is why many authors try a few alternatives for a given dataset, and finally analyze the model that best explains it. In some cases, however, authors include two types of interaction in the same model. Usually a Type II and a Type IV are combined, in a setting analogous to the clustering and heterogeneity spatial effects in the BYM specification. Globally, the model is considered to have a Type IV interaction since there is spatial dependence between neighbouring trends.

2.4 Inference methodology

Inference may be performed from a frequentist or Bayesian approach. In the former, the fixed effects and nonlinear functional effects of covariates are considered as deterministic, while in the latter they are interpreted as realizations of random variables or random functions.

Spatio-temporal models are commonly within the class of structured additive regression (STAR) models. In them, and from the Bayesian approach, all unknown functions and parameters can be treated within a unified general framework by assigning appropriate prior distributions with the same general structure but different forms and degrees of smoothness.

Moreover, additional structure might be put on the hyperparameters. For instance, Nobre et al. (2005) present models with Type I or Type III interaction terms, in the form of unstructured or spatially dependent random effects respectively. The variance parameters of these random effects modulates the degree of dependence between neighbouring pixels (i, j). It may be thought of as a parameter controlling the smoothing degree. While it is common to assume

that this smoothing degree is constant, the authors allowed it to vary smoothly over time. In summary, they assumed the variance hyperparameter to be another random variable varying over time with a first order random walk prior distribution.

This kind of modelling can rarely be done outside the Bayesian paradigm. Bayesian Hierarchical Models enable the greatest flexibility and the most exhaustive posterior information on every parameter of the model. Posterior mean or median values, posterior credible intervals or threshold exceedance probabilities can be computed straightforward.

However, great care is needed in the specification of priors, especially when stepping away from standard models which have well-studied priors. Besides, inference is usually done with Markov Chain Monte Carlo methods (MCMC) which are far from fast, especially regarding spatio-temporal models. Simulations may take hours or days, and technical considerations about mixing, convergence, etc. have to be taken into account. In summary, MCMC methods are not suitable for a fast consulting service oriented to non-technical users.

Recently, there have been some significant advances in Bayesian approximate inference. Rue and Martino (2007) and Rue et al. (2009) developed a library (GRMFLib) with basic functions for approximate inference, and a program (INLA) using this library that allows the formulation of models and inference to be performed on them, as well as an R package (INLA) that interfaces the program. The method, together with the computational tools, provides an extremely fast and straightforward environment for Bayesian inference. However, its scope is restricted to a subset of structured additive regression (STAR) models called *latent Gaussian Markov Mandom Field* (GMRF) models. The aforementioned paper by Nobre et al. (2005) is an example that does not meet the conditions. Still, many of the models reviewed here are or can be written as GMRF models, meaning that they can be fitted in a matter of minutes or even seconds. To our knowledge, there is still no published work on spatio-temporal modelling using this approach; however, we are aware of some work in progress.

Within the frequentist approach, inference can be performed with empirical-Bayes (EB) techniques. Breslow and Clayton (1993) popularized the use of penalized quasi-likelihood (PQL) methods for inference in generalized linear mixed models (GLMM). Moreover, STAR models can be represented by GLMMs after appropriate reparameterization, thus allowing for EB inference based on GLMM methodology (Lin and Zhang, 1999; Fahrmeir et al., 2004).

The variance or smoothing parameters are considered to be unknown constants and can be estimated by using (approximate) restricted maximum likelihood (REML). For given or estimated smoothing parameters, covariate effects and unknown functions can be obtained by maximizing their posterior density.

However, this approach may not be straightforward for complex models, since it involves a reparameterization and numerical iterative algorithms for efficient REML estimation.

In addition, although EB estimation typically yields consistent and nearly unbiased point estimates of the relative risks, the variability in these estimates is often understated, since the uncertainty that arises from estimating the hyperparameters of the random effects distribution is not taken into account. MacNab et al. (2004) have proposed the use of bootstrap methodology to address this issue in the context of (spatial) disease mapping.

We would like to emphasize that all inference methods involve some kind of approximation. While MCMC methods rely on simulation, empirical Bayes approaches make use of point estimates of certain parameters in the model, and INLA, as its name states, is based on Laplace approximations. Practitioners need to be conscious of these approximations and their limitations, and expertise is needed to overcome possible spurious effects due to the methodology.

In summary, Hierarchical Models are powerful and flexible enough to represent almost any possible relationship. In addition, MCMC methods provide the means of performing Bayesian inference on these models. However, great technical expertise is needed to work with complex models, and even in the simpler ones simulation usually takes many hours. MCMC methods seem not to be suitable for fast and automatic tools. Restricted to GMRF models, Bayesian inference can still be done in a matter of seconds or minutes with the INLA methods and tools. Empirical Bayes inference is feasible in many situations. Some effort is needed in order to parameterize the model in a proper way, and to develop efficient estimation algorithms and strategies. There is much work in progress in this field.

3 Spatio-temporal models in the literature

Next we provide a brief review of the most prominent papers in the spatio-temporal modelling of diseases. We keep the focus on the spatio-temporal modelling approach, keeping to one side other technical aspects like the centering of the covariates, or whether a Poisson or a Binomial distribution is preferred for the observations, or the prior specification for hyperparameters, etc. Notation is homogenized to that established in the previous section.

Models are classified into three categories according to the structure of the temporal evolution of the estimated risk for each region. Namely, *parametric models* have a predefined shape (linear, quadratic, ...), while on the contrary, *temporally independent models* estimate the risks for each period independently of those from previous periods, and finally *smooth temporal evolution models* allow for structured trends without restricting to any predefined shape (see Figure 1).

Note that the temporal evolution for a region is determined by the sum of the temporal main effect with and the potential interaction terms involving time. Therefore, parametric models imply spatio-temporal interactions of Type

(a)   (b)   (c)

Figure 1: Representation of (a) parametric, (b) temporally independent and (c) smooth temporal evolution models

II or IV and temporally independent models imply interactions of Type I or III. However, smooth temporal evolution models can be constructed with any type of spatio-temporal interaction provided that it is combined with a smooth temporal main effect.

3.1 Parametric models

SUMMARY: The log risk in a given region is assumed to be a linear or quadratic function of time. The coefficients in the function are region-specific and are spatially structured so that neighbouring regions have similar evolutions.

PROS: Information is shared in both space and time. The parametric formulation is straightforward.

CONS: The parametric evolution in time seems to be inappropriate for long periods of time, as it is too restrictive.

3.1.1 Bernardinelli *et al.* (1995) and Assunção, Reis and Oliveira (2001)

$$\eta_{ij} = \text{Intercept} + S_i + T_j + ST_{ij}$$

where S_i is modelled as a heterogeneity (3) or a clustering (4) effect depending on the problem, $T_j = \beta t_j$ and $ST_{ij} = \delta_i t_j$ where δ_i is again a heterogeneity or a clustering effect.

In this way, the log risk is a linear function of time, with region-specific intercepts and slopes. When δ_i is a heterogeneity effect the model has an interaction of Type II, but of Type IV if δ_i is a clustering effect.

The authors also propose an extension that allows for the *a priori* correlation between the spatial and spatio-temporal random effects, using an additional level in the hierarchical model.

They perform Bayesian inference with MCMC methods.

Assunção et al. (2001) follow the same line, incorporating an additional quadratic term in time that allows for curved trends with convex or concave shapes. Random effects are modelled as clustering effects (4), thus producing spatio-temporal interactions of Type IV.

3.1.2 Sun *et al.* (2000)

$$\eta_{ijk} = C_k + S_i + CT_{jk} + ST_{ij} + \varepsilon_{ijk}$$

Sun et al. (2000) incorporate the age group as a fixed-effect covariate C_k, and model the spatial effect S_i with a clustering effect (4).

The covariate-time and space-time interaction terms are modelled as linear functions of time, with a slope depending on the age group and region respectively: $CT_{jk} = \text{age}_k t_j$, and $ST_{ik} = \theta_i t_j$. Furthermore, while age_k is a fixed effect, θ_i is a spatial clustering effect. Therefore, spatio-temporal interaction is of Type IV.

Finally, an overall heterogeneity term (2) is added.

3.2 Temporally independent spatial models

SUMMARY: These can be seen simply as a set of spatial models, one for each period of time, with almost no relation between them, except possibly for some restrictions in their precision parameters.

PROS: Temporal evolution is not restricted to any specific shape. Information is shared in space.

CONS: For each period of time one spatial model is estimated, therefore not sharing information in time. These are high dimensional models, with negative consequences on identifiability and overmodelling.

3.2.1 Waller *et al.* (1997) and Xia and Carlin (1998)

Waller et al. (1997) discuss the inclusion of spatial and temporal main effects together with sex and race covariates.

$$\eta_{ijk} = C_k + ST_{ij}$$

In the first place they assume additivity, meaning that sex and race effects are not affected by region and year, therefore there are no CS_{ik} or CT_{jk} interaction terms. The covariate component C_k is that of a linear model with race and sex fixed effects plus a sex-race interaction effect.

The spatio-temporal interaction term has the form $ST_{ij} = \phi_i^{(j)} + \theta_i^{(j)}$, where for each period of time j, the vector $\boldsymbol{\phi}^{(j)} + \boldsymbol{\theta}^{(j)}$ follows the BYM specification (5), with different precision parameters $\lambda_\phi^{(j)}$ and $\lambda_\theta^{(j)}$ for each period of time, therefore being an interaction of Type III.

This approach results in a spatio-temporal model where the spatial dimension is *nested* within time, meaning that for each period of time, a spatial model is fitted.

This spatial model is not tied in any way to its temporal neighbours, allowing for free evolution, but not sharing information in time.

Xia and Carlin (1998) extend the model by Waller et al. (1997) introducing another covariate (smoking prevalence) in a more sophisticated way, namely introducing a way to account for errors in covariates. However, there is nothing new with respect to the spatio-temporal modelling.

3.2.2 Nobre, Schmidt and Lopes (2005)

Nobre et al. (2005) follow the same approach, fitting a spatial model nested within time with some modifications.

$$\eta_{ijk} = CT_{jk} + ST_{ij}$$

They generalize the covariate effect, allowing its coefficient to evolve over time through a RW1 prior.

Regarding the spatio-temporal modelling, $ST_{ij} = \delta_i^{(j)}$ where, for each period j, $\delta_i^{(j)}$ is a heterogeneity effect (3) or a clustering effect (4), but not both, therefore representing spatio-temporal interactions of Types II or IV respectively.

Furthermore, the corresponding precision parameters $\lambda_\delta^{(j)}$ can be either independent or such that their logarithms follow a RW1.

In this way, variability in the degree of spatial dependence over time is allowed while neighbouring years have similar values.

In fact they run the four different models arising from these two choices, and based on the DIC, their best choice is a model with a spatio-temporal heterogeneity term whose log-variance follows a Gaussian random walk, showing no spatial dependence.

3.3 Smooth temporal evolution models

SUMMARY: The evolution of the estimated risk in each region is a smooth function of time.

PROS: Temporal evolution is not restricted to any predefined shape. Information is shared in time.

CONS: Many possible alternatives for interaction term, so there is a need for model selection criteria.

3.3.1 Knorr-Held (2000)

$$\eta_{ijk} = \text{Intercept} + S_i + T_j + ST_{ij}$$

The spatial main effect S_i follows the BYM specification (5), and analogously, the temporal main effect T_j is specified as the sum of a RW1 structured effect and an unstructured random effect.

Knorr-Held (2000) tries five different alternatives for the spatio-temporal interaction ST_{ij}, the first of which is not having any interaction term at all. The interaction terms of the other four alternatives are modelled as random effects with precision matrices λK, where λ is an unknown scalar to be estimated from the data and K is a *precision matrix* computed as the Kronecker product of the structure matrices of either the structured or unstructured spatial effect with either the structured or unstructured temporal effect, following a rationale by Clayton (1996). In this way, if both unstructured effects are combined then a Type I interaction results, while on the contrary the combination of both structured effects produces an interaction of Type IV. On the other hand, if one of the structured effects is combined with the unstructured component of the other effect, then a Type II or III arises, depending on the specific combination.

In their application to Ohio Lung Cancer data, they find that any of the extended models outperforms the basic model. For comparison between the four types of model with an interaction term, they compute a few indicators for goodness of fit and model complexity based on the posterior saturated deviance. For this specific dataset, they find that Type II interaction gives the best fit with moderate posterior deviance variation.

3.3.2 Lagazio, Biggeri and Dreassi (2003)

Lagazio et al. (2003) extend the Knorr-Held (2000) model by turning it into an Age-Period-Cohort model.

While keeping the BYM specification for the spatial effect S_i, the temporal effect T_j is split into three random effects corresponding to age group, calendar period and birth cohort, all of them with first order random walk priors. Regarding the interaction term, they considered interactions between the spatial clustering effect with both period and cohort effects. Hence, two possible interactions of Type IV arise.

They perform model comparison based on Expected Predictive Deviance, and also by measuring the differences based on the Kullback-Leibler divergence. Finally, they choose the model with the three temporal main effects plus space-cohort interaction, which seems to play a very important role.

3.3.3 Schmid and Held (2004)

Schmid and Held (2004) present the same model as Lagazio et al. (2003), with slight modifications. Namely, the observations are assumed to be drawn from a binomial distribution, the spatial effect is modelled as a clustering effect (4), and a heterogeneity effect scoping age, period and space dimensions is added.

They investigate whether a period-space interaction, a cohort-space interaction, or no interaction at all is more appropriate for their data. In either case of spatio-temporal interaction, they also determine whether a Type II or a Type IV interaction produces the best results, and finally whether a first or second order random walk prior is more adequate for the temporal main effects.

Computations are based on simulation, applying a couple of *tricks* described in the literature for faster computation and optimization of the process.

Finally, they show how to make future predictions from these models, although specific sophisticated algorithms are needed for achieving this.

They compare by using the DIC the 10 models that arise from the alternative specifications and find almost no differences between RW1 and RW2 priors for the temporal main effects. They also find Type II interaction giving the best result.

3.3.4 Macnab and Dean (2001, 2002)

MacNab and Dean (2001) and MacNab and Dean (2002) follow an empirical Bayes approach, using generalized additive mixed models (GAMM). They develop methods of inference based on PQL (Breslow and Clayton, 1993) for simultaneous modelling of spatial effects, temporal effects, and spatiotemporal interaction using low-order fixed knot (1-knot cubic) B-splines.

Specifically, they model the linear predictor as having both temporal and spatial main effects and an interaction term, and a second model incorporating the effect of age, without any interaction with other terms. The brackets indicates that the term is present in the second model only.

$$\eta_{ijk} = \text{Intercept} + [C_k] + S_i + T_j + ST_{ij}$$

The temporal main effect is modelled with a cubic B-spline, allowing for a smooth and quite flexible evolution. For the spatial main effect they use a clustering effect (4).

They explore two alternatives for the spatio-temporal interaction term. The most general one is fitting a cubic spline for each region, allowing for region-specific evolution. This option produces an interaction of Type II. Although

they declare it could be possible to add spatial dependence onto these splines, they choose to keep them independent.

The second alternative is to use a parametric interaction term as in Bernardinelli et al. (1995). Although the coefficient could be a clustering spatial effect, they find that no spatial correlation is evident, therefore, $ST_{ij} = \phi_i t - j$ where ϕ is a heterogeneity spatial effect. The interaction is then of Type II. Note that this does not mean that the temporal evolution of disease rates for each region is linear, since the temporal main effect is not, yet the region-specific departure from global evolution is modelled as linear.

Inference was done using penalized quasi-likelihood (PQL) methods, although the MCMC approach can also be used.

3.3.5   Richardson, Abellan and Best (2006)

Most of the work reviewed here focus on a single disease of interest. However, when two or more phenomena (diseases) are related, better results typically arise from a joint study than from an individual study. Richardson et al. (2006) have treated male and female lung cancer as two related diseases, and developed a methodology to study their spatio-temporal behaviour jointly.

Instead of considering sex as a covariate, they consider male and female lung cancer as two different but correlated diseases. For each disease they fit a spatio-temporal model, having one common component and one specific component that calibrates the differential between the two diseases.

However, the spatio-temporal behavior of each disease follows a rather simple model.

$$\eta_{ijk} = \text{Intercept} + S_i + T_j + ST_{ij}$$

Where S_i is a clustering effect, T_j is a RW1, and ST_{ij} is a global heterogeneity effect, thus producing an interaction of Type I.

3.3.6   Martínez-Beneito, López-Quílez and Botella-Rocamora (2008)

Martínez-Beneito et al. (2008) present a completely new approach called the Autoregressive Linking of spatial patterns.

$$\eta_{ijk} = \text{Intercept} + T_j + ST_{ij}$$

The temporal main effect follows a RW1 prior specification. However, there is no spatial main effect. Instead, the estimated differential risk for a region-period pixel (i, j) is specified as a multivariate autoregressive temporal model with an unknown parameter ρ and a spatially structured vector of errors:

$$ST_{ij} = \rho ST_{i(j-1)} + \delta_{ij}$$

where the vector $\boldsymbol{\delta}_j$ follows a BYM specification with parameters λ_ϕ and λ_θ, $\forall j$.

In this way, the BYM specification is not used to model the spatial distribution of the differential rates, but their increment from period to period. This approach produces a spatio-temporal interaction of Type IV.

4 Discussion

4.1 Additional information

Many of the models reviewed exhibit their strength when used with convenient covariates. A model that fits the data well may perform worse than others if not fed with these covariates.

In general, as more additional information is available, the spatio-temporal structure becomes less significant, thus models with a simpler spatio-temporal structure perform better. Conversely, when additional information is scarce, more structure is needed in order to capture the spatio-temporal effects produced by the unobserved factors.

In the Euroheis project a general model is required for all types of data, with few specific covariates. This means that almost everything has to be explained in terms of space, time and space-time interaction. Therefore, from this point of view, we should encourage the most flexible models in terms of space-time modelling.

4.1.1 To standardize or not to standardize

Age is one of the most obvious factors affecting risk for almost all diseases. It is well-known that two populations may show differences in incidence or mortality for a given disease as a consequence of differences in age distributions alone.

This is commonly addressed by working with age-standardized rates, which allows age and cohort effects to be ignored in the analysis. However, the uncertainty in the corresponding estimates is neglected in such a two-step analysis; a further practical problem is the choice of the reference rates for standardization (Schmid and Held, 2004).

On the other hand, age/cohort effects can be introduced into the model in order to produce estimates for age/cohort-specific risk rates, and also allowing these effects to interact with others, such as spatial terms.

The latter option is applicable only when age-specific count data is available. (Zhang et al., 2006) provides an interesting and sophisticated implementation of age-gender-spatio-temporal modelling within the Bayesian approach.

However, this choice is completely independent of the spatio-temporal structure, and it should be decided based on other criteria.

4.2 Interactions

While it always seems preferable to use a model with maximum-flexibility, which allows for any possible kind of interaction, it is well-known that this is far from being true. Super-models have lots of parameters that have to be estimated, and there is often not enough information stored in the available data to do it. Besides, many of these parameters may conflict with each other, trying to explain the same thing, and finally adding noise to the model, which produces greater uncertainties. Thus, interaction terms should be added only when there is evidence to support it. A first exploratory model should be quite simple, although it is good to have the possibility of adding new interaction terms if a well-driven analysis suggests doing so.

4.3 Model selection

In almost every single paper reviewed, a basic modelling idea is developed producing from four to ten possible variations of the model, depending on what specific subset of the possible terms are to be included, and how. An important part of the articles is dedicated to model selection, after fitting all of the possibilities to a specific dataset. Model selection is one of the most difficult aspects of statistical modelling, and it is far from being solved in general. This is obviously not desirable for a tool oriented to the final user and for producing fast analyses.

4.4 Implementation

When considering the implementation of a general purpose, user-oriented application, there are some key aspects to take into account:

Specificity of the model The most sophisticated models perform very well in the specific datasets and situations for which they were designed. However, they are typically not suitable for other diseases or populations. Conversely, simple models perform moderately well in many situations, and provide a starting point for further analysis.

Complexity and scope of the inference method All of the methodologies have some technical difficulties that often require the intervention of a specialist. Sophisticated models that account for multiple effects and interactions will certainly require some fine tuning for almost every dataset, in order to avoid practical problems like identifiability or overmodelling. Besides, the structure of the model determines the inference methods that can be used. For simple models there are more methods to choose from.

The fastest and most straightforward inference method that we are aware of is the Integrated Nested Laplace Approximation (INLA) approach (Rue et al., 2009). However, it can *only* perform inference in a subset of structured additive regression (STAR) models called *latent Gaussian Markov Mandom Field* (GMRF) models. Still, many of the models reviewed here are, or can be written as, GMRF models.

Empirical Bayes (EB) techniques may provide a simple and fast inference approach for some models. Theoretically, STAR models can be represented by generalized linear mixed models (GLMM) after appropriate reparameterization, thus allowing for EB inference based on GLMM methodology (Fahrmeir et al., 2004). However, for complex models this representation and posterior EB inference may not be straightforward. Convergence problems may arise and *ad-hoc* programming may be needed.

Bayesian inference can be performed in STAR models with MCMC techniques. There are software tools like WinBUGS (Lunn et al., 2000) that prevent the need for programming specific algorithms. However, MCMC inference is not exempt from technical difficulties like convergence and mixing problems, specially for sophisticated models.

Speed This is a direct consequence of the choices on the model and the inference method. Models based on simulation usually take hours or days

to fit; therefore, they are unsuitable for fast analyses. Simple and general models are more likely to be successfully fitted with other approaches such as INLA or EB techniques within an acceptable time.

Accuracy EB estimation typically yields consistent and nearly unbiased point estimates of the relative risks but it often understates their variability, as a consequence of using point estimates of hyperparameters.

MCMC estimation requires a careful diagnosis of the simulation outcomes in order to ensure suitable estimations of the posterior densities.

5 Conclusions

Among the specific objectives of the EUROHEIS 2 project is the inclusion of spatio-temporal methods for disease mapping in the Rapid Inquiry Facility (RIF). However, there is not a wide consensus on how to describe temporal and spatial evolution at the same time in a proper way. Although several spatio-temporal disease mapping techniques have been proposed recently, the implementation of these methods is not always easy or adequate for a quick response tool.

We have outlined a general framework for spatio-temporal models, breaking them up into four stages: the probabilistic model for observations, the components of the linear predictor, the structures of the effects and the inference methodology. We established a standard notation for all the possible components in the linear predictor, and discussed the most common alternatives for their modelling, paying special attention to those related to the spatial and temporal dimensions. In particular, we generalized the classification by Knorr-Held (2000) for the spatio-temporal interaction.

We classified the models proposed in the literature into parametric, temporally independent and smooth temporal evolution models, reflecting the structure of the temporal trends that may arise. Parametric models provide a straightforward and simple formulation, allowing for spatio-temporal interaction and the sharing of information both in space and time. On the other hand, it may be too restrictive for long periods of time. Temporally independent models do not restrict the temporal trend, but at the cost of not sharing information in time. Smooth temporal evolution models provide the greatest flexibility and adaptability, but with increased complexity.

The final discussion summarizes what we have identified as the most relevant aspects to be considered in the selection of a methodology for spatio-temporal disease mapping to be included in a RIF-like application. Namely, the specificity of the model, the complexity and scope of the inference method, the processing speed and the accuracy of the results. We hope this provides all the necessary elements to take a well-based decision.

References

Assunção, R. M., I. A. Reis, and C. D. Oliveira (2001). Diffusion and prediction of leishmaniasis in a large metropolitan area in brazil with a bayesian space-time model. *Statistics in Medicine 20*(15), 2319–2335.

Bernardinelli, L., D. G. Clayton, C. Pascutto, C. Montomoli, M. Ghislandi, and M. Songini (1995). Bayesian analysis of space-time variation in disease risk. *Statistics in Medicine 14*(21-22), 2433–2443.

Besag, J. and C. Kooperberg (1995, December). On conditional and intrinsic autoregressions. *Biometrika 82*(4), 733–746.

Besag, J., J. York, and A. Mollié (1991, March). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics 43*(1), 1–20.

Breslow, N. E. and D. G. Clayton (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association 88*(421), 9–25.

Clayton, D. G. (1996, December). Generalized linear mixed models. In W. R. Gilks, S. Richardson, and D. J. Spiegelhalter (Eds.), *Markov Chain Monte Carlo in Practice*, pp. 275–301. London: Chapman & Hall/CRC.

Currie, I. D., M. Durban, and P. H. Eilers (2006, April). Generalized linear array models with applications to multidimensional smoothing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68*(2), 259–280.

Eilers, P. H. C. and B. D. Marx (1996). Flexible smoothing with b-splines and penalties. *Statistical Science 11*(2), 89–121.

Fahrmeir, L., T. Kneib, and S. Lang (2004, July). Penalized structured additive regression for space-time data: a bayesian perspective. *Statistica Sinica 14*(3), 731–761.

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 J. Besag (1998). Modelling risk from a disease in time and space. *Statistics in Medicine 17*(18), 2045–2060.

Lagazio, C., A. Biggeri, and E. Dreassi (2003). Age-period-cohort models and disease mapping. *Environmetrics 14*(5), 475–490.

Lee, D.-J. and M. Durbán (2009, June). Smooth-car mixed models for spatial count data. *Computational Statistics & Data Analysis 53*(8), 2968–2979.

Lin, X. and D. Zhang (1999, April). Inference in generalized additive mixed modelsby using smoothing splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 381–400.

Lunn, D. J., A. Thomas, N. Best, and D. Spiegelhalter (2000, October). Winbugs — a bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing 10*(4), 325–337.

MacNab, Y. C. and C. B. Dean (2001). Autoregressive spatial smoothing and temporal spline smoothing for mapping rates. *Biometrics 57*(3), 949–956.

MacNab, Y. C. and C. B. Dean (2002). Spatio-temporal modelling of rates for the construction of disease maps. *Statistics in Medicine 21*(3), 347–358.

MacNab, Y. C., P. J. Farrell, P. Gustafson, and S. Wen (2004, December). Estimation in bayesian disease mapping. *Biometrics 60*(4), 865–873.

Martínez-Beneito, M. A., A. López-Quilez, and P. Botella-Rocamora (2008). An autoregressive approach to spatio-temporal disease mapping. *Statistics in Medicine 27*(15), 2874–2889.

Nobre, A. A., A. M. Schmidt, and H. F. Lopes (2005). Spatio-temporal models for mapping the incidence of malaria in pará. *Environmetrics 16*(3), 291–304.

Richardson, S., J. J. Abellan, and N. Best (2006, August). Bayesian spatio-temporal analysis of joint patterns of male and female lung cancer risks in yorkshire (uk). *Statistical Methods in Medical Research 15*(4), 385–407.

Rue, H. and S. Martino (2007, October). Approximate bayesian inference for hierarchical gaussian markov random field models. *Journal of Statistical Planning and Inference 137*(10), 3177–3192.

Rue, H., S. Martino, and N. Chopin (2009). Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of hte Royal Statistical Society 71*, 1–35.

Schmid, V. and L. Held (2004). Bayesian extrapolation of space-time trends in cancer registry data. *Biometrics 60*(4), 1034+.

Sun, D., R. K. Tsutakawa, H. Kim, and Z. He (2000). Spatio-temporal interaction with disease mapping. *Statistics in Medicine 19*(15), 2015–2035.

Waller, L. A., B. P. Carlin, H. Xia, and A. E. Gelfand (1997). Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association 92*, 607–617.

Xia, H. and B. P. Carlin (1998). Spatio-temporal models with errors in covariates: mapping ohio lung cancer mortality. *Statistics in Medicine 17*(18), 2025–2043.

Zhang, S., D. Sun, C. Z. He, and M. Schootman (2006). A bayesian semi-parametric model for colorectal cancer incidences. *Statistics in Medicine 25*(2), 285–309.