



A spatio-temporal hierarchical Markov switching model for the early detection of influenza outbreaks

Rubén Amorós^{1,2} · David Conesa^{2,3} · Antonio López-Quílez^{2,3} · Miguel-Angel Martínez-Beneito^{4,5}

Published online: 11 February 2020
© The Author(s) 2020

Abstract

Rapidly detecting the beginning of influenza outbreaks helps health authorities to reduce their impact. Accounting for the spatial distribution of the data can greatly improve the performance of an outbreak detection method by promptly detecting the first foci of infection. The use of Hidden Markov chains in temporal models has shown to be great tools for classifying the epidemic or endemic state of influenza data, though their use in spatio-temporal models for outbreak detection is scarce. In this work, we present a spatio-temporal Bayesian Markov switching model over the differentiated incidence rates for the rapid detection of influenza outbreaks. This model focuses its attention on the incidence variations to better detect the higher increases of early epidemic rates even when the rates themselves are relatively low. The differentiated rates are modelled by a Gaussian distribution with different mean and variance according to the epidemic or endemic state. A temporal autoregressive term and a spatial conditional autoregressive model are added to capture the spatio-temporal structure of the epidemic mean. The proposed model has been tested over the USA Google Flu Trends database to assess the relevance of the whole structure.

Keywords Surveillance · Influenza · Outbreak detection · Bayesian · Spatio-temporal · Hidden Markov model · Markov switching model

Financial support from the Conselleria de Sanitat of the Generalitat Valenciana (the Valencian Regional Health Authority) is gratefully acknowledged. The authors would also like to acknowledge financial support from the Ministerio de Economía y Competitividad (the Spanish Ministry of Economy and Finance) for its support in the form of the research Grants TEC2016-81900-REDT and MTM2016-77501-P (jointly financed with the European Regional Development Fund, FEDER)

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s00477-020-01773-5>) contains supplementary material, which is available to authorized users.

✉ Rubén Amorós
ruben.amoros.salvador@gmail.com

David Conesa
david.v.conesa@uv.es

Antonio López-Quílez
antonio.lopez@uv.es

Miguel-Angel Martínez-Beneito
martinez_mig@gva.es

Extended author information available on the last page of the article

1 Introduction

Influenza is a disease that affects millions of people and causes hundreds of thousands of deaths each year (World Health Organization 2016). This disease is also the cause of large amounts of direct and indirect expenses due to health care costs, absenteeism and other effects of the epidemic (Gasparini et al. 2012). For these reasons, surveillance of this viral infection has considerable interest for health policymakers and, in particular, the early detection of outbreaks. Among other benefits, detecting the exact moment when the epidemic begins allows for a better use of resources. This decreases the morbidity and mortality while reducing the expense. Influenza surveillance is usually carried out by screening cases of Influenza-Like-Illness (ILI), which is generally defined as a set of symptoms, including fever and some upper-respiratory affection such as cough and/or sore throat. In this work we interchangeably use the expressions ‘influenza surveillance’ and ‘ILI surveillance’ bearing in mind that, in general, they refer to the same systems and processes.

Epidemics of influenza in temperate parts of the planet have a particular behaviour that shapes the statistical methods dedicated to their detection. Considering the temporal dimension, seasonal influenza occurs during the cold months of the year, but the onset and duration of each seasonal epidemic are variable and depend on multiple factors, such as the viral strain properties, the susceptibility and transmission conditions of the population, and the climate (Lofgren et al. 2007). Besides the seasonal outbreaks, other non-stationary influenza epidemics can happen at any time of the year, usually caused by strains of virus that jump species from animals to humans. The particular case of the swine influenza A(H1N1)pdm09 virus was monitored with special attention and led to the analysis and re-evaluation of several influenza surveillance strategies (Cook et al. 2011; De Lange et al. 2013; Gomez-Barroso et al. 2014; Ortiz et al. 2009).

For some time, the detection of influenza outbreaks has been only based on temporal methods applied to each region of interest separately. The list of statistical models proposed to do this detection is large, as reflected in several reviews (Amorós et al. 2015; Spreco and Timpka 2016; Unkel et al. 2012). Later on, in the same fashion as done in other fields such as species distribution (Martínez-Minaya et al. 2018), detection methods have started to combine spatial tools used in disease mapping and cluster detection with the temporal methodologies already in use in order to adapt to the intrinsic spatio-temporal behaviour of influenza spread. The epidemics usually start in one or several points and spread from them. Therefore, promptness and sensitivity of the methods can be improved by considering several time series of counts or rates associated with small areas instead of only one time series pooling all the areas; firstly, because the substantial increase in the incidence rates of some of these locations at the beginning of the epidemic will not be masked by the aggregation; and secondly, because the spatial correlation of the spread can be used to distinguish the epidemic behaviour.

Several approaches have been considered to build spatio-temporal methods for disease surveillance. Spatio-temporal CUSUM and EWMA models (Rogerson and Yamada 2004; Zhou and Lawson 2008) can raise an alarm when finding a mild but persistent shift of the mean rate, which is a common endemic behaviour of ILI upon the arrival of the cold months of the year, and this could cause false alarms. Scan statistics (Kulldorff 2001) usually require that the potentially detected clusters cover at most half of the locations considered, but it is usual for influenza epidemics to cover entire countries or even continents at some time of the season.

Several temporal methods have been proposed based on ARIMA models (Cowling et al. 2006; Rao and McCabe 2016). Proper and intrinsic conditional autoregressive

models (Besag et al. 1991) and other spatial structures used for disease mapping have been combined with these time series structures for disease surveillance under the frequentist paradigm. But it has been under the Bayesian paradigm that this kind of combination has been more prolific. The flexibility of the Bayesian hierarchical models allows for inference over complex models combining several different temporal and spatial structures with relative ease (Martínez-Bello et al. 2018). In that way, several methods have been proposed, usually based on a spatio-temporal Poisson model. Some of them locate regions with unusual temporal patterns of incidence (Li et al. 2012; Mugglin et al. 2002) or predict future incidence taking into account correlation among several diseases (Corberán-Vallet and Lawson 2014). Corberán-Vallet and Lawson (2011) model an endemic Poisson model and use the scaled surveillance conditional predictive ordinate (a measure of goodness of prediction) to trigger alarms, and Corberán-Vallet (2012) extended this to the spatio-temporal surveillance of several diseases by means of a shared component model. Rotejanaprasert and Lawson (2016) consider two components (endemic and epidemic) for the relative risk and use the Kullback-Leibler distance between the usual predictive distribution and the predictive distribution ignoring the epidemic component of the risk to set alarms.

A natural option for the detection of outbreaks by means of distinguishing an endemic and an epidemic behaviour of the data is the use of models including latent variables that allow them to classify observations. Considering that the endemic and epidemic incidence rates have different behaviours, two different models can be combined, allowing the method to select which model offers a better fit for each time t and each location i . This is achieved by associating a set of discrete latent variables Z_{it} to each of the locations and times, with two possible values, 0 or 1, indicating the non-epidemic or epidemic phase respectively. The values taken by these hidden variables determine which model (endemic or epidemic) is to be applied for each observation Y_{it} . One advantage of these latent variables for decision making is that the estimated probability of being in the epidemic phase for each time and location can be inferred from them, by means of their expected value in the Bayesian paradigm or by their estimated allocation probabilities in the frequentist paradigm (Grzegorzczak and Shafiee 2017). In this way, the answers offered by models using these latent variables are richer than a simple ‘yes’ or ‘no’. The intrinsic temporal structure of influenza data, where each epidemic week is usually followed by other epidemic weeks and each non-epidemic week is usually followed by more non-epidemic weeks, suggests linking the latent variables using temporal Markov chains. In this way, each latent variable is

conditionally dependent only on the previous latent variable in time. Hidden Markov models (HMMs) (Cappé et al. 2005) and Markov switching models (MSMs) (Douc et al. 2004) use hidden Markov chains and are usually seen in surveillance to model the observations and their associated latent variables. HMMs consider conditionally independent observations given the values of the hidden variables while MSMs also define a conditional distribution for consecutive observations which varies according to the epidemic state (the value of the hidden variable). The HMMs non-parametric modelling of temporal variability could show substantial advantage over some other spatio-temporal parametric proposals (Torres-Avilés and Martínez-Beneito 2015), which could have problems to describe epidemic peaks occurring at very different weeks for the different seasons of study.

HMMs have been applied in very diverse fields such as financial economics (Bhar and Hamori 2004), computer vision (Bunke and Caelli 2001), hydrology (Khadr 2016; Vasas et al. 2007) or biological sequence analysis (Yoon 2009). In the field of temporal surveillance of influenza and ILI, HMM have been applied both under the frequentist paradigm (Le Strat and Carrat 1999, Rafei et al. 2015) and the Bayesian paradigm (Lu et al. 2010; Madigan 2005; Rath et al. 2003; Sun and Cai 2009), while MSMs have been used mostly under the Bayesian paradigm (Conesa et al. 2015; Lu et al. 2010; Lytras et al. 2018; Martínez-Beneito et al. 2008). Spatio-temporal HMM models have been more scarce. The proposal of Knorr-Held and Richardson (2003) has been used for meningococcal infection surveillance (though not explicitly for outbreak detection). Another example of a model where the conditional dependence of the hidden variable Z_{it} is not only temporal but also spatial is found in the work of Banks et al. (2012). The authors present a theoretical multivariate Bayesian framework for syndromic surveillance where the parameter of a Poisson model is defined by an always present endemic component plus an epidemic component multiplied by the latent variable so that it is only added for the epidemic times and locations. The logarithm of each parameter is modelled by a linear regressor. Particular cases of this framework of model have been used for the detection of influenza outbreaks (Heaton et al. 2012; Zou et al. 2012), indicating that *the choices of hyperparameters represent vague prior information and ensure posterior propriety*. Nevertheless, it can be proven that a simplification of these models, where no spatial or spatio-temporal terms are taken into account, and where improper non-informative prior distributions are set for the parameters, has improper posterior distributions for certain parameters of the model (Amorós 2017). It can be also proven how an alternative switch model as seen in some other works (Li et al. 2012) has a proper posterior distribution.

Considering all this, in this work we present a spatio-temporal Bayesian Markov switching model over the differentiated rates. This model does not require strong prior information to estimate the transition to the epidemic state. It also does not require a predefinition of the epidemic periods in the historic data (as some other models do) and has no spatial coverage limitation. At the same time, this method benefits from considering the spatial structure of the data, which allows for a faster detection of the beginning of the outbreak and its location. Additionally, this model focuses the attention on the jumps of incidence rate from one week to the next. Paying attention to the behaviour of the increases and decreases instead of the size of the rates is key to improving the promptness of detection as, at the beginning of an epidemic, the rates are relatively low and quite similar to those of the endemic phase. It is the relative behaviour among contiguous rates which can characterize the moment of the outbreak. Therefore, it is fundamental to model different temporal and spatial dependencies of the differentiated data depending on whether each location is in the endemic or epidemic phase for each week. As stated before, the use of MSMs is perfectly suited for this, being able to detect the moments of transition from non-epidemic to epidemic periods for each of the locations.

The remainder of this paper is organized as follows. After this introduction, Sect. 2 proposes the spatio-temporal Bayesian MSM. Section 3 shows the application of the model on the USA Google Flu Trends data, and Sect. 4 shows a comparison with several simplifications and variations of the model, in order to stress the necessity of several elements of the hierarchical structure. Conclusions are presented in Sect. 5.

2 The spatio-temporal Markov switching model for influenza outbreaks detection

In this section we present a novel proposal for the spatio-temporal detection of influenza outbreaks which can deal with the behaviour of the spread of influenza outbreaks while avoiding several drawbacks observed in several methods in the literature. This method is an extension of a temporal model presented by Martínez-Beneito et al. (2008) that can use any data coming from any surveillance system, provided that the influenza incidence rates are observed for a discrete and fixed set of locations at equally spaced times.

Each surveillance system obtains data from a different source and there is usually a trade-off between how fast data is obtained from these sources and how specific the system is (Cheng et al. 2009). Hence, data obtained from an Internet source such as Twitter can be almost

immediate, but its specificity would be very low. This is even more evident if we compare it with a mortality registry, a very specific system where data can take several weeks or months to be available. The nature of the data also varies with the source from which it has been retrieved. In this way, spatial locations of the data can be expressed in several ways: individual cases may be associated with their home address; individual or aggregated cases may be assigned to a health care facility—hospitals, emergency rooms, pharmacies, etc.—or they can be associated to an administrative region. All the possibilities can, in general, be translated to aggregated rates in administrative regions, which results in spatially discrete support of the information structured as a lattice, where a neighbouring rule can be established. This last format may in some cases be coarser, but ensures that almost every spatial data of any surveillance system can be translated to it. Therefore, with the intention of making our method as versatile as possible, we model it to deal with lattice data.

The usual way of reporting data in surveillance systems is to associate it with discrete and equally spaced times (days or weeks, usually). If the data is presented for each time on a lattice support, the result is a set of time series spatially related. In this line, our proposed dataset for exemplifying the performance of our proposal is an illustrative example of this kind of data (as shown in Fig. 1, dataset described in detail in Sect. 3.1). But, it is worth noting that all the methods here presented could be used in many other similar datasets with the same characteristic.

The remainder of this section has been organized in subsections reflecting the hierarchical structure of the model proposed. Moreover, priors have been set up at every step in order to facilitate comprehension.

2.1 Modelling the differentiated rates

We consider our data to be the first order differentiated rates of influenza observed on a finite set of locations on a given neighbourhood graph (lattice data). Usually the edges are defined by adjacency, where two regions share a border. Equidistant time points are considered for the temporal dimension of the data (for example, weekly rates), which is the usual way of reporting data to surveillance systems.

In the same line as done in the proposal of Martinez-Beneito et al. (2008), modelling the differentiated rates instead of the raw rates themselves allow us to divert the focus of attention from the magnitude of the rates. This can be an advantage because, in some cases, using the magnitude of the rates can make difficult the detection of the beginning of an outbreak, as most start when influenza rates are still low. Instead of that, the size and behaviour of the increments of the rates are the distinctive features

which allow to classify each location and time as being in the epidemic or endemic phase.

For simplicity in the notation, given a location i , we name y_{it} to the rate in time t (r_{it}) minus the rate in time $t - 1$ (r_{it-1}), that is

$$y_{it} = r_{it} - r_{it-1}, \quad t = 1, \dots, T$$

where T represents the total number of periods and the time index for the rates starts at 0 to have the index for the differences starting at 1. We also denote Z_{it} the variable that indicates the latent epidemic state for each time and location, with a value of 1 for the epidemic state and a value of 0 for the non-epidemic state.

As observed in the lower graph in Fig. 1, the differentiated rates can be seen to show two behaviours, one with smaller jumps around the zero value and another with sharper increases or decreases which tend to be followed by values of the same sign. Therefore, we consider that the differentiated rates follow a distribution with mean and variance depending on the epidemic state. Indeed, as in Martinez-Beneito et al. (2008), we consider it to be normal:

$$y_{it} \sim N(R_{itZ_{it}}, \sigma_{Z_{it}}^2). \quad (1)$$

Finally, it is worth mentioning that another option could have been to model the logarithm of the rates with a Gaussian distribution, a common practice in disease surveillance and disease mapping. Nevertheless, this would have made difficult the detection of the onset of epidemics. If we were to compute the logarithms of the rates and then differentiated them, we would have that our data would be $\log(r_{it+1}) - \log(r_{it}) = \log\left(\frac{r_{it+1}}{r_{it}}\right)$. In this situation, the model would only consider the percentage of the growth and not its magnitude. Doubling the rate would result in the same value for the response variable, whether the jump on incidence was from a rate of 20 cases per 100 000 inhabitants to 40 (reasonable jump during the non-epidemic phase) or it was from 2000 to 4000 (more likely during the epidemic phase).

2.2 The hidden Markov structure for the epidemic phase Z_{it}

In order to model the two possible phases in which each location and time can be classified, we consider the set of binary latent variables Z_{it} in Eq. (1) which indicate the epidemic state of the process (equal to 1 to indicate epidemic phase and equal to 0 to indicate non-epidemic phase) as a hidden Markov chain for each location. The distribution of the latent variable Z_{it} for each location i conditioned to Z_{it-1} follows a Bernoulli distribution with transition probabilities common for all times and locations:

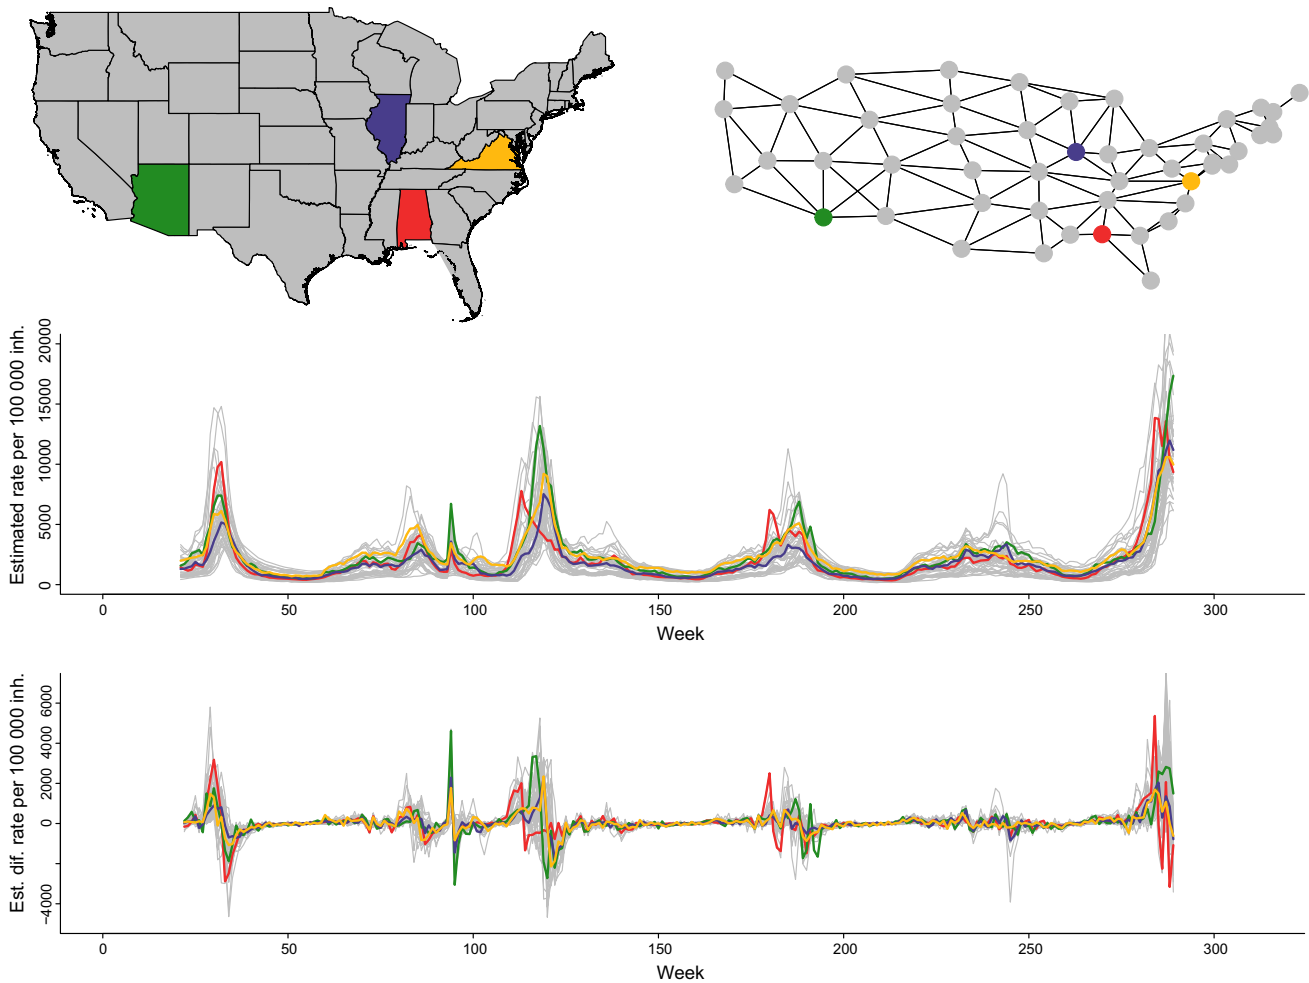


Fig. 1 Map of USA (leaving out Alaska and Hawaii), graph of neighbourhood defined by adjacency, estimated influenza rates by Google Flu Trends between 2007 and 2013 and the corresponding

differentiated rates. Highlighted regions: Alabama in red, Arizona in green, Illinois in blue and Virginia in yellow

$$Z_{it} \sim \text{Ber}(p_{Z_{it-1}}), \tag{2}$$

where

$$\begin{aligned} p_{kl} &= P(Z_{it} = l | Z_{it-1} = k), \\ p_k &= P(Z_{i1} = k), \end{aligned} \tag{3}$$

with p_k being the probabilities for time 1, that is, $Z_{i1} \sim \text{Ber}(p_1)$.

Jeffreys non-informative prior densities for Bernoulli trials are set for the transition probabilities and for the initial probabilities:

$$p_{00}, p_{11}, p_0 \sim \text{Beta}\left(\frac{1}{2}, \frac{1}{2}\right), \tag{4}$$

and the rest of probabilities p_{10} , p_{01} and p_1 are obtained by complementarity. This Markovian structure in time for Z_{it} could also be induced in space (Green and Richardson 2002). Nevertheless, this would make much more complex our proposal which is already complex enough. Hence, we

have preferred to induce temporal dependence in a less complex form as we will see later.

As said before, the variables Z_{it} distinguish the latent epidemic and non-epidemic states, and their posterior expected values are estimations of the probabilities of being in the epidemic state, which are used for declaring epidemic alarms. This response is not dichotomous and can be used by the public health authorities to make better-informed decisions which actually take into account the uncertainty of the estimation.

2.3 Modelling the mean of the differentiated rates $R_{it}Z_{it}$

In Eq. 1, the differentiated rates were modelled to follow a normal distribution with mean and variance depending on the epidemic state. After defining the latent variables Z_{it} which indicate this state, let us define the distribution of the means according to the epidemic state.

The non-epidemic phase is characterized for jumps on the rates close to zero. A first approach could be just leaving the mean of the differentiated endemic rates to be 0, but observation of empirical data shows that there are non-epidemic weeks where small growths or decreases of the rates happen for all locations at the same time. This may be caused by the stationarity of influenza and of some other respiratory diseases, which raises the expectancy of notifiers of an imminent outbreak, regardless if an outbreak will occur or not. In that way, it would be common for some diseases like the cold to be reported as ILI by practitioners or to cause internet users to perform more influenza-related queries or post more influenza-related tweets during autumn and winter. For this reason, we consider each week t to have a mean value μ_{t0} which is common to all locations at that week, but different from the mean of other weeks for fitting the mentioned stational pattern. No further spatial or temporal structure is considered for the endemic phase since given μ_{t0} we would expect the differentiated rates belonging to the endemic phase to be distributed as random noise around its mean.

In order to characterize the epidemic phases, we used three features of the expected value of the differentiated rates, namely its expected spatial and temporal correlation, and the higher structured variability for that state.

Firstly, the epidemic phase is modelled with a more complex temporal structure than that of the endemic phase as we expect to find that behaviour in data. It is expected for a region in the epidemic state to show several positive jumps (differentiated rates) until reaching the peak of the epidemic, and then the jumps become negative whereas the rates decrease again to the endemic level. Therefore, we would expect the differentiated rates for each location in the epidemic state to be temporally dependent.

The second expected behaviour caused by the contagious nature of influenza is that if one region has epidemic growths on its incidence rates, people of neighbouring regions will become infected and those regions will be expected to show similar growths. As a consequence, we model the mean of the differentiated rates by a term μ_{t1} , common for all the locations (the overall epidemic rate for that week) but different for each time, plus a temporal autoregressive structure of order 1 on the observations for each location with parameter ρ , plus a spatial intrinsic conditional autoregressive (ICAR) model (Besag et al. 1991) for each time. In this way, those means reflect the temporal and spatial dependence that we expected them to show.

The third feature which helps to characterize the epidemic phase is the higher structured variability of the data, as the jumps shown by the epidemic phase tend to be bigger (in the positive or negative direction) than those in the endemic phase. This higher structured variability comes from the higher variability of the epidemic mean of the

normal distribution in Eq. 1 given by the common terms for each time unit, the temporal autoregressive structure and the spatial ICAR term.

As a result, the mathematical expressions for the mean of both endemic and epidemic periods are:

$$\begin{aligned} R_{it0} &= \mu_{t0}, \\ R_{it1} &= \mu_{t1} + \rho y_{it-1} + \psi_{it}, \end{aligned} \quad (5)$$

with ψ_{it} being the ICAR term. Let us describe each of these components of both states in more depth:

2.3.1 The common term for each time unit μ_{t0} and μ_{t1} .

As stated before, the common term of the endemic phase μ_{t0} can capture the mild common seasonality of non-epidemic ILI. In the case of the epidemic season, μ_{t1} models the common rise or fall of the rates during the epidemic. In both cases, these terms model the weekly consensus along all the states in the same epidemic phase, which in general is positive in ascending phases (before the epidemic peak) and negative in descending phases (after the epidemic peak). We consider μ_{t0} and μ_{t1} as two random effects over time, with larger variability for μ_{t1} , as common rises or falls of the epidemic phase are expected to be larger than those of the endemic phase. In order to avoid identifiability problems with the main unstructured variability $\sigma_{Z_{it}}^2$ of the main distribution, we set the standard deviations of the two random effects to be proportional to those of the main distribution. We can express the modelling of the common temporal terms as follows:

$$\begin{aligned} \mu_{t0} &\sim N(0, \sigma_{\mu_{t0}}^2), \\ \mu_{t1} &\sim N(0, \sigma_{\mu_{t1}}^2), \end{aligned} \quad (6)$$

where λ is the estimated proportion factor and a is a hyperparameter to be set by the modeller expressing a vague prior knowledge. For example, $a = 100$ would be a possible choice, as it is highly unlikely that this part of the structured variability is so much larger than the unstructured one.

2.3.2 The autoregressive structure in the epidemic mean.

During the epidemic state of influenza, it is common to observe a rise in the rates during several weeks and, after reaching the peak, a decline in the rates that also lasts several weeks. The autoregressive structure on the growths is able to fit this behaviour.

To ensure the stationarity of the autoregressive process we must bound the parameter ρ to the interval $[-1, 1]$. Anyhow, considering that the correlation between subsequent growths in the epidemic phase is expected to be positive (differentiated rates will be positive in general

during several weeks until reaching a peak and then they will be mostly negative for some more weeks), in practice we can restrict the interval to $[0, 1]$:

$$\rho \sim \text{Unif}(0, 1). \tag{7}$$

To ensure that the variance of y_{i1} is equal to the stationary variance of the series $\{y_{it}\}_{t=1}^{\infty}$, the Expression (1) for the first week is modified as follows:

$$y_{i1} \sim N\left(R_{i1Z_{i1}}, \frac{\sigma_{Z_{i1}}^2}{1 - Z_{i1}\rho^2}\right). \tag{8}$$

This type of correction has been introduced, for example, in some works in the context of spatio-temporal disease mapping (Martinez-Beneito et al. 2008).

2.3.3 The ICAR structure in the epidemic mean ψ_{it} .

We have seen how the autoregressive structure models similar jumps for consecutive times for each location during the epidemic phase. In an analogous way, the ICAR structure induces spatial dependence in the epidemic phase. Clusters of locations where influenza is spreading show growths on rates that have that spatial dependence, while neighbouring regions where the disease has not extended yet adapt better to the absence of any spatial structure (endemic state). This behaviour is theoretically expected due to the infectious nature of influenza, but it has also been empirically suggested (Fox and Dunson 2015). In their work, Fox and Dunson (2015) propose a Bayesian non-parametric covariance regression, which allows the variance matrix in a multivariate regression model to vary with the predictors among time. The model is exemplified with an application on USA Google Flu Trends (2017) data between 2003 and 2009 (more details about this data source in Sect. 3.1). The application does not assume any prior spatial structure but the estimated posterior correlation matrices among regions for each week show it. In addition, the estimated correlation values are higher during the epidemic weeks and lower on non-epidemic weeks, confirming the appropriateness of our proposal.

Our proposal is that the terms ψ_{it} for a given week t follow a joint ICAR distribution as follows:

$$\psi_{it} | \psi_{-it} \sim N\left(\frac{\sum_{j \sim i} \psi_{jt}}{n_i}, \frac{\sigma_{\psi}^2}{n_i}\right), \tag{9}$$

with $j \sim i$ meaning locations i and j are neighbours, $\psi_{-it} = \{\psi_{jt} : j \neq i\}$ and n_i being the number of neighbours of location i . The ICAR joint distribution is improper and an additional restriction of sum to zero, $\sum_i \psi_{it} = 0$, is usually added in order to be able to do inference.

We set a non-informative prior for the common conditional standard deviation of ψ_{it} :

$$\sigma_{\psi} \sim \text{Unif}(0, b), \tag{10}$$

with b a hyperparameter to be fixed by the modeller expressing a vague prior knowledge, for example, any value above the highest differentiated rate in absolute value.

Note that the ψ_{it} do not really reproduce spatio-temporal dependence. Instead, they are a collection of temporally independent spatial patterns. Pure spatio-temporal dependence patterns could be used for defining ψ (Knorr-Held 2000; Martinez-Beneito et al. 2008; Adín et al. 2017). Nevertheless, as for Z_{it} , this would introduce further complexity in our already complex proposal so we have preferred to avoid that possibility.

2.4 Modelling the variance of the differentiated rates $\sigma_{Z_{it}}^2$

We have seen that one important aspect for the discrimination between epidemic and non-epidemic phases in our proposal is the behaviour of the expected value of the differentiated rates, its spatial and temporal correlation and the higher structured variability for the epidemic state. Another feature that helps characterize the epidemic phase will be the unstructured variability.

In general, differentiated rates are centred in zero, therefore, though the behaviour of the mean of the response variable gives important information to distinguish between epidemic and endemic phases, the behaviour of the variance of the data is critical for this task. The jumps of the rates on the non-epidemic state are relatively small in absolute value, while the growths or decreases of the rates on the epidemic state are usually larger. A wider variance on the differentiated rates is therefore expected for the epidemic state.

Characterizing the epidemic state by a higher variance has been successfully done on the differentiated rates (Martinez-Beneito et al. 2008) and on the raw rates (Nunes et al. 2013). In some spatio-temporal proposals the extra variance of the data on the epidemic phase is modelled by the extra variance of the added spatio-temporal structure (Heaton et al. 2012; Zou et al. 2012). In our case, besides the structured noise that is fitted by the spatio-temporal term for the mean, we also consider an unstructured noise in the epidemic phase, getting a combination of structured and unstructured noise in the fashion of the popular Besag, York and Mollie’s model (Besag et al. 1991). We also ensure that the unstructured noise of the epidemic phase is higher than that of the endemic phase,



emphasizing the importance of the variance as a tool for classification and avoiding the interchangeability of these terms.

Therefore, we model the unstructured variability of the two phases by obtaining both standard deviations from a uniform prior and ensuring that they are ordered:

$$\begin{aligned}\sigma_0 &= \min(\theta_1, \theta_2), \\ \sigma_1 &= \max(\theta_1, \theta_2),\end{aligned}\quad (11)$$

with c a hyperparameter to be set by the modeller expressing vague prior knowledge. Setting c as a value above the largest differentiated rate in absolute value is a sensible choice, as suggested before for the ICAR structure.

2.5 The complete model

For clarity in the understanding of the proposed method, we present again all the previously introduced components which make up the hierarchical model:

$$\begin{aligned}y_{it} &\sim N(R_{it}Z_{it}, \sigma_{Z_{it}}^2), \\ R_{it0} &= \mu_{t0}, \\ R_{it1} &= \mu_{t1} + \rho y_{it-1} + \psi_{it}, \\ \rho &\sim \text{Unif}(0, 1), \\ \sigma_0 &= \min(\theta_1, \theta_2), \\ \sigma_1 &= \max(\theta_1, \theta_2), \\ \sigma_{\mu0} &= \lambda\sigma_0, \\ \sigma_{\mu1} &= \lambda\sigma_1,\end{aligned}$$

3 Application of the model

This section describes the USA Google Flu Trends database and shows an application of the proposed spatio-temporal model on this database, stressing the difference between the retrospective and real-time (online) application of the method and showing a comparison of the method with several simplifications and variations and with the temporal model by Martinez-Beneito et al. (2008).

3.1 The USA Google Flu Trends data

The automated collection of data from search engines provides almost immediate large amounts of data from vast populations and territories without the purposeful collaboration of the individuals. This is the case of Google Flu Trends, a web tool which provides estimates of influenza incidence through the application of an algorithm on Google queries (Ginsberg et al. 2009). In particular, the data we are using to illustrate our proposal is the USA Google Flu Trends (USA GFT), consisting of weekly

estimates of influenza incidence for the 48 spatially connected states of the USA plus Washington, D.C., between 2007-12-02 (48th week of the year) and 2013-01-20 (3rd week of the year). These estimated rates are a proxy of the reported influenza rates by the CDC. While publicly available data by the CDC was aggregated in 10 regions, the USA GFT offers data segregated by states. Raw and differentiated incidence rates (our proposal directly models the differentiated rates, see Sect. 2.1) of the USA GFT dataset and the vicinity structure considered are shown in Fig. 1. Four of the 49 chains are highlighted in different colours as examples.

There has been some debate about the quality of the data given by Google Flu Trends. After the A(H1N1)pdm09 epidemic on 2009, it became clear that the behaviour of the search engine users is not stable and the algorithm needed to be re-evaluated and adjusted (Cook et al. 2011). In fact, the algorithm for the USA was reassessed in 2009, 2013 and 2014 (Olson et al. 2013). Currently, data is available at the website but the algorithm does not offer new estimations since August 19th, 2015. In any case, though the data might not be narrowly accurate for some weeks in comparison to some of the CDC published incidence rates, this dataset is realistic and perfectly suited for illustrative purposes.

3.2 Retrospective estimates of the epidemic phase in space and time

We can distinguish two different ways of applying a detection method over a certain dataset according to whether the data is considered to be available all at once or not. The first case is a retrospective application, where the model is applied once using all the data. The second case is a prospective or ‘online’ application, where the model is applied several times, consecutively adding each time the data from one week, as it would happen in a real surveillance system, where each week the method would be reapplied adding the most recently collected data. The real use of a detection method is prospective, but the retrospective application is computationally cheap and can offer a proxy of its performance which can be used for its evaluation.

As a first approach to evaluate the performance of the spatio-temporal model on the USA GFT data described above, we show the results of a retrospective application of our model. This and all subsequent inference processes in this work were carried out using WinBUGS (Lunn et al. 2000), running a process with 2 chains with 5 000 iterations of burning and 5000 subsequent iterations for each chain. After thinning, 1000 iterations were kept, 500 from each chain. Convergence was checked by observing the

effective sample size, the \hat{R} statistic (Gelman et al. 2013) and visual check of the chains of simulations. The consistent convergence of the model has been further checked by running the model with a considerably larger amount of iterations (2 chains with 50 000 burning iterations and 100 000 subsequent iterations for each chain), observing only small differences in the posterior means and variances that can be attributed to the MCMC error. The code of the model can be found in Appendix 1 in Supplemental Material.

In Fig. 2 one can observe the estimated posterior probability of the epidemic phase (posterior expected values of the latent variables Z_{it}) for four randomly chosen states of the 49 possible ones which we use as an example from now on (Alabama, Arizona, Illinois and Virginia). In Appendix 2 in Supplemental Material, all estimates for the 49 states are displayed. The retrospective detection of the epidemic phase shows a sensible behaviour that appears to give a high probability of epidemic to those weeks with steep growth or decay around the peak of the epidemic for each season.

In Table 1 the posterior mean and 95% credible interval for the most relevant parameters of the model are shown; the histograms and density plots for the simulated samples from the posterior distributions for these parameters are shown in Fig. 3. The high values of p_{00} and p_{11} indicate, as is expected, that weeks in a certain phase tend to be followed by weeks of the same phase. The standard deviations associated with the endemic phase, σ_0 and σ_{μ_0} , take notably lower values than those of the epidemic phase, σ_1 , σ_{μ_1} and σ_{ψ} . This indicates that the variance of the differentiated rates is a key point for the classification of the two states of the Markov chain. It is also worth noting how a large part of the epidemic variability is captured by σ_{ψ} , which indicates the importance of the spatial structure to capture the spatial behaviour of the data. All the credible intervals are quite narrow, which is an indicator of the identifiability of all the parameters of the model.

3.3 Online application of the model

In a realistic application of the spatio-temporal proposal in a surveillance system, the estimates of the epidemic phase would have been obtained week by week. Therefore, for each week data from only that and previous weeks (but not following ones) would be used to estimate the epidemic states. To reproduce this realistic behaviour one must apply the model in an online basis, that is, running the model once per week excluding posterior data to do the inference. Figure 4 compares the posterior probability of epidemic for

the last season of the USA GFT data for the four states used as an example previously mentioned.

The results of both ways of applying the model are similar, though the online lines of estimates tend to be slightly less smooth. This behaviour is not surprising, as the retrospective way of applying the method can use information from posterior weeks when estimating the phase, while the online application can not.

Figure 5 shows maps of the estimated posterior probability of being in epidemic phase for some of the weeks of the last season when applying the model on an online basis. One can appreciate how the model makes a sensible estimation of the spread of the outbreak, which starts from the south and east of the United States and ends up covering the whole country.

4 Comparison with alternative proposals

In order to evaluate the relevance of the proposed spatio-temporal model, in this section we show a comparison of the new proposal with two simplifications of that proposal, one variation and the temporal model of Martinez-Beneito et al. (2008) (from now on, M-B 2008) on which the above proposal is based.

4.1 Simplifications and variations of our proposals

Here we describe the two simplifications and the variation we are going to compare our spatio-temporal proposal with. We also explain how we apply the M-B 2008 model over the GFT spatio-temporal data.

Removing the spatial component The first proposed simplification consists in removing the structured spatial random effect, the intrinsic CAR component, from the modelling of the epidemic phase. To do so, the lower equation in Expression (5) loses the ψ_{it} term. The modified expression is as follows:

$$\begin{aligned} R_{it0} &= \mu_{t0} \\ R_{it1} &= \mu_{t1} + \rho y_{it-1}. \end{aligned} \quad (12)$$

Thus, the prior structure of this model would not induce spatial dependence at all. By comparing this simplification with the original novel proposal one can assess how the assumption of similar growths of incidence among neighbours in epidemic phase affects the fitting and the phase classification.

Removing μ_{t0} . As it is indicated in Sect. 2.3.3, the term μ_{t1} for the epidemic phase is needed for fitting the evident overall time trend of this phase. Hence the term μ_{t0} is added to the model to capture the mild common jumps on

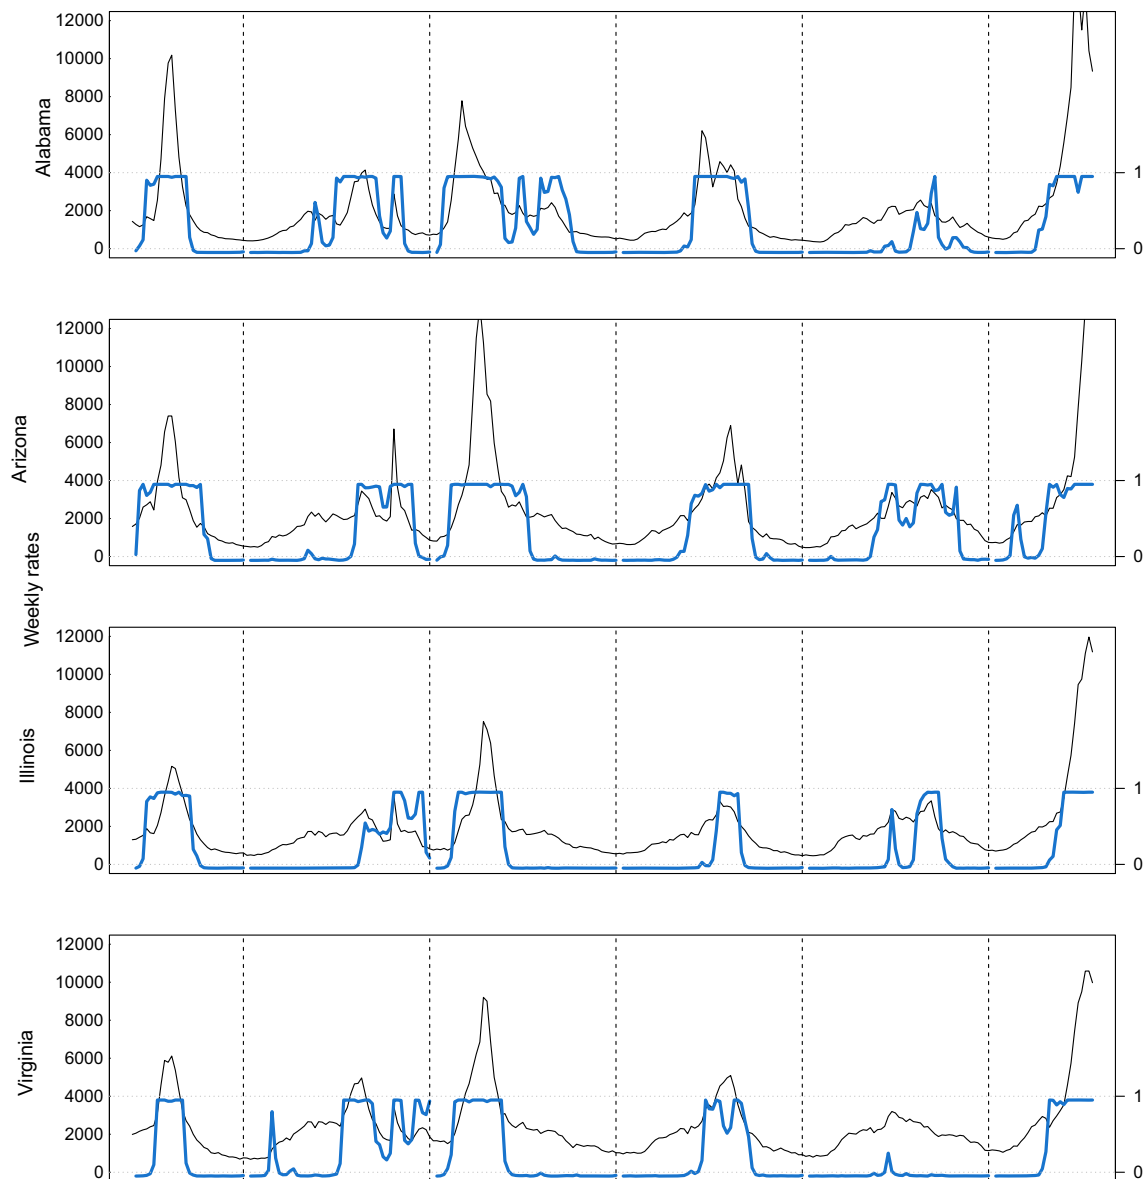


Fig. 2 Retrospective estimated probability of being in epidemic phase by the spatio-temporal model on USA GFT data for 4 states. In black: weekly estimated influenza incidence per 100 000 inhabitants during seasons from 2007–2008 to 2012–2013

the endemic phase, but also thinking that it brings balance of complexity among the regressors of the epidemic and endemic phases, preventing an excessive sensitivity of the model. To check that this choice is appropriate, the second simplification we consider consists in fixing $\mu_{t0} = 0$ for all endemic observations. Expression (5) for this modification becomes now as follows:

$$\begin{aligned} R_{it0} &= 0 \\ R_{it1} &= \mu_{t1} + \rho y_{it-1} + \psi_{it}. \end{aligned} \quad (13)$$

Also, as $\sigma_{\mu0}$ is no longer present in the model, λ is no longer needed for ensuring identifiability of the parameters. Therefore, Expression (6) is substituted by the following,

so that μ_{t1} has a variance that is conditionally independent of σ_0 and σ_1 :

$$\mu_{t1} \sim N(0, \sigma_{\mu1}^2) \quad (14)$$

The hyperparameter d is set to a value above the largest differentiated rate (in absolute value), as suggested in Sects. 2.3 and 2.4.

Leroux. The variation of the model we compare the new proposal with here consists of the substitution of the ICAR term on the epidemic regressor by the spatial structure proposed by Leroux et al. (2000). Expression (5) remains the same, but in this case, the definition of the conditional distribution of the terms $\psi_{it} | \psi_{-it}$ is set as a Normal

Table 1 Posterior mean and 95% credible interval for the parameters of the spatio-temporal model applied on the USA GFT data in a retrospective basis

Parameter	Posterior mean	95% Credible interval
p_{00}	0.95	[0.95, 0.96]
p_{11}	0.89	[0.87, 0.90]
σ_0	85.00	[82.99, 87.01]
σ_1	430.89	[404.70, 458.70]
$\sigma_{\mu 0}$	127.61	[116.60, 138.40]
$\sigma_{\mu 1}$	646.69	[583.28, 701.00]
σ_{ψ}	770.10	[716.90, 822.70]
ρ	0.35	[0.31, 0.38]
λ	1.50	[1.38, 1.62]

distribution with conditional mean and variance as those expressed in the work of Leroux et al. (2000). The conditional distribution of the Leroux effect is described as follows:

$$\psi_{it} | \psi_{-it} \sim N \left(\frac{\phi}{1 - \phi + \phi n_i} \sum_{j \sim i} \psi_{jt}, \frac{\sigma_{\psi}^2}{1 - \phi + \phi n_i} \right) \quad (15)$$

$$\phi \sim \text{Unif}(0, 1).$$

Taking into account that ϕ is bounded between 0 and 1, one may notice the following: the mean of the conditional distribution monotonically increases with ϕ , while the variance decreases as ϕ gets bigger (except when $n_i = 1$, where the variance is constantly σ_{ψ}^2). In the limit cases, the distribution is that of an unstructured random noise with 0 mean and variance equal to σ_{ψ}^2 when $\phi = 0$ and an ICAR distribution, with mean equal to $\frac{1}{n_i} \sum_{j \sim i} \psi_{jt}$ and variance equal to $\frac{\sigma_{\psi}^2}{n_i}$, when $\phi = 1$. All intermediate values of ϕ result in a set of Gaussian conditional distributions with mean and variance values in between. In that way, the parameter ϕ distributes the variability of the Leroux term between spatially structured and unstructured variability. We compare our proposal with this variation to check if the ICAR term forces a too strong spatial relation which the Leroux model could be able to soften.

M-B 2008 This purely temporal model is independently run for each of the 49 locations of the USA GFT data set, so no information (whether spatial or otherwise) is shared among them. The comparison with this model indicates the relevance of the spatio-temporal modelling against the purely temporal modelling which does not share information among different geographic units at all.

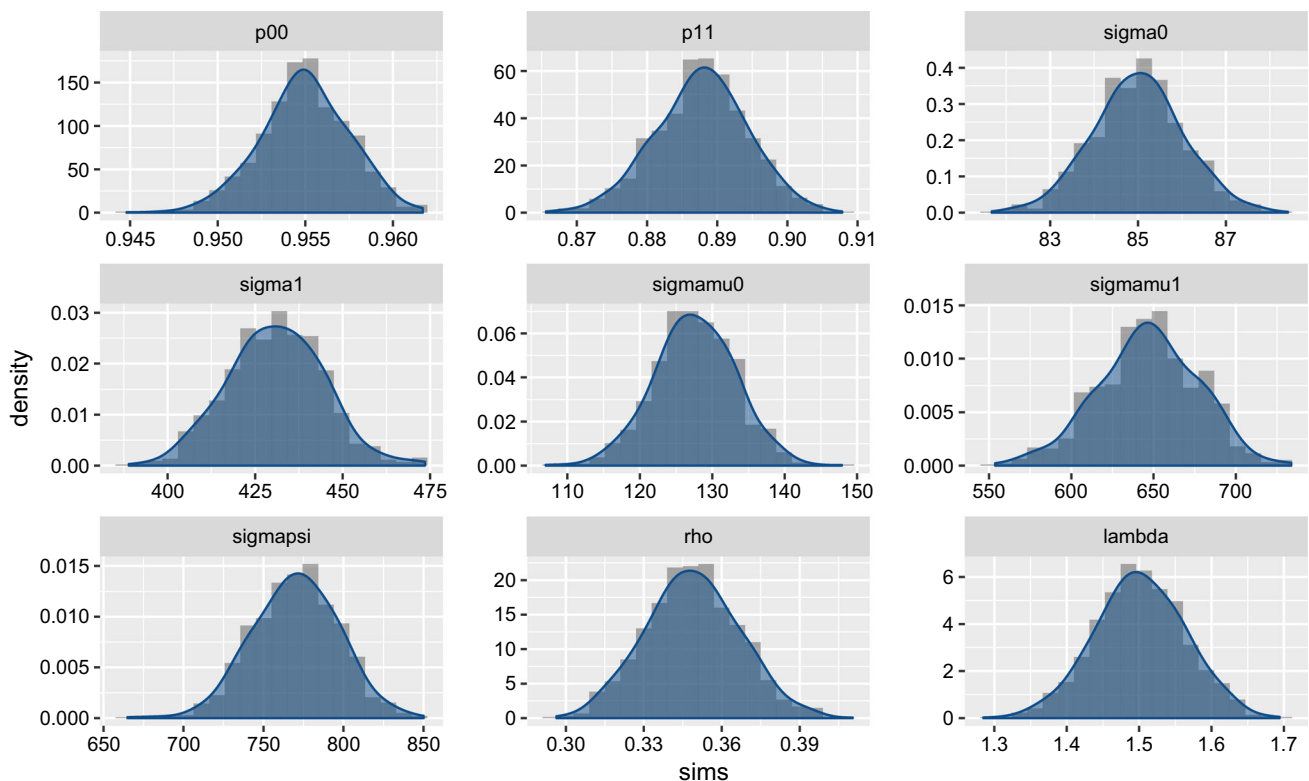


Fig. 3 Posterior histograms and density plots for the parameters in the model, applied retrospectively, for the USA GFT data

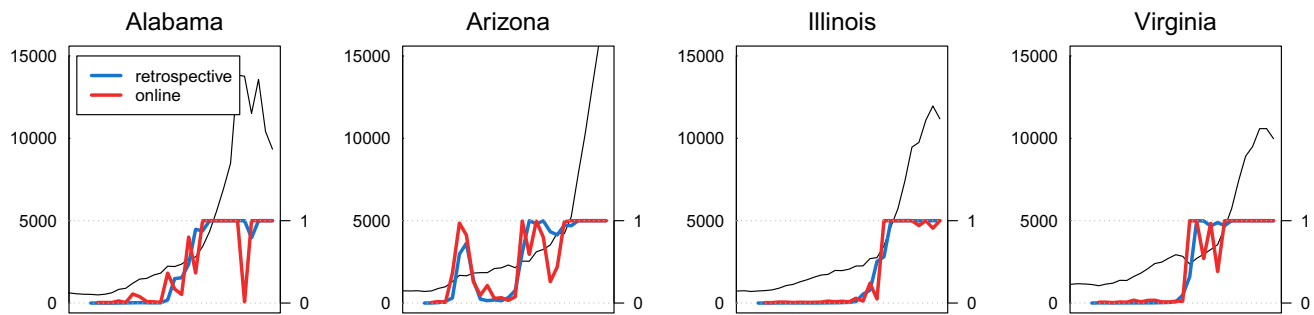


Fig. 4 Comparison of the online and retrospective estimated probability of being in epidemic phase by the spatio-temporal model on USA GFT data for four states. In black: weekly estimated influenza incidence per 100 000 inhabitants during season 2012–2013

4.2 Comparison of estimated parameters and epidemic phases

In order to give an insight into how the different versions considered alter the detection of the epidemic phase, retrospective inference on the USA GFT dataset was carried out using each model. Calculations were performed in an Intel® Core™ CPU I7–3770 with 4 cores at 3.40GHz and 8Gb of RAM, with OS Windows 7 Professional 64 bits. Table 2 shows computation time in minutes for the spatio-temporal proposal and its variations. The greatest difference in terms of computational time is observed between the new proposal and the M-B 2008 model, which runs over three times faster. All observed computational costs are affordable if the model is to be run once a week in a real surveillance system.

Figure 6 displays the estimated posterior probability of being in the epidemic phase for the retrospective application of the spatio-temporal proposal, its variations and the M-B 2008 model for four of the locations. The same graph for all the locations is depicted in Appendix 3 in Supplemental Materials. The estimated posterior mean for the principal parameters of the models are shown in Table 3.

Taking a look at the yellow line in Fig. 6, one can observe how the detection of the simplification without μ_{t0} is consistently higher than the rest. The higher posterior estimated mean for p_{11} , shown in Table 3, also indicates the higher sensibility and lower specificity of this simplification. As it was conjectured, excluding μ_{t0} gives an advantage to the epidemic structure over the non-epidemic structure to be able to adapt to the data, and this raises the number of weeks classified as epidemic.

The red line represents the detection when removing the spatial dependence, that is, when not considering an ICAR component for the epidemic phase. The layout of this red line is similar to that of the blue line, which represents the novel spatio-temporal proposal, but with lower values in certain weeks. Contrary to the model without spatial component, the model with spatial component can use the information about the epidemic state of the neighbours to

raise the probability of epidemic on these weeks. It seems that the addition of the spatially structured component helps with the correct classification of these weeks as epidemic ones thanks to the sharing of the information with neighbours. Observing the estimates for the means of the parameters in Table 3, one can see how the variability that σ_{ψ} can not capture (because it is absent) is split among the other two variance terms of the epidemic phase; σ_1 and $\sigma_{\mu 1}$, increasing the error term of the epidemic phase.

The detection performance of the Leroux model is almost the same as that of the original proposal, as it is shown by the almost complete superposition of the blue and grey lines. The estimated mean of the parameters is quite similar, with the notable exception of the variances for the epidemic phase. As can be appreciated in Table 3, part of the non-spatial variability modelled in σ_1 in the spatio-temporal proposal moves to σ_{ψ} in the Leroux model, which is a parameter that captures both spatially structured and unstructured variability. The parameter ϕ indicates that spatial dependence in the data could be milder than that induced by the ICAR distribution. One slight advantage of the Leroux model is that it offers more freedom to decide the variance of the random effects for the mean $\sigma_{\mu 0}$ and $\sigma_{\mu 1}$. This is so because the unstructured variability is split among σ_1 , which directly affects the value of $\sigma_{\mu 1}$, and σ_{ψ} , which does not. Thus, overall, we consider the spatial modelling with a spatial Leroux's proposal a competitive option, although for $\phi \approx 0$ some problems of identifiability may arise between the mean and the error term of the epidemic phase. So, some more work is due on the use of this proposal.

The M-B 2008 model, represented in green, also shows important differences in the classification of epidemic states as compared to the rest of spatio-temporal models. This shows the evident benefit of sharing information among various regions, even if the neighbouring structure is not taken into account (as is the case of the model without the ICAR component).

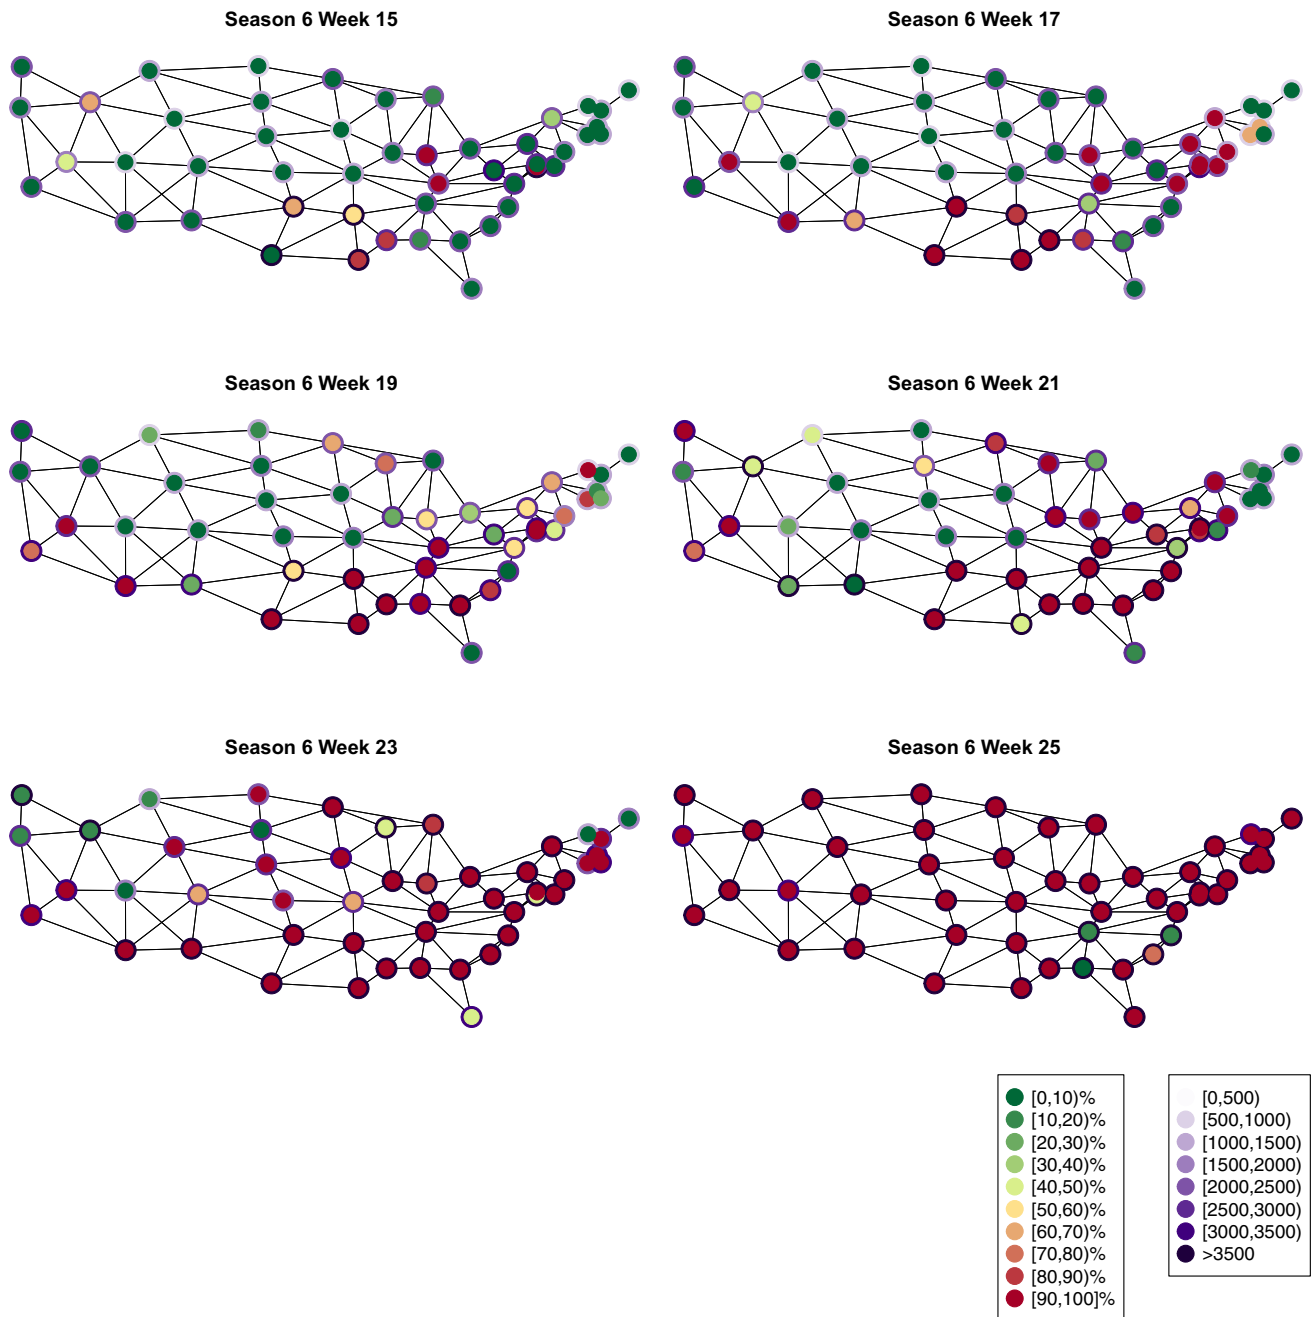


Fig. 5 Online estimated probability of being in epidemic phase by the spatio-temporal model on USA GFT data (green-red scale). Weeks 15, 17, 19, 21, 23 and 25. Contours in purples: GFT reported influenza incidence per 100 000 inhabitants

Table 2 Computational cost in minutes for the spatio-temporal model applied on the USA GFT data in a retrospective basis (reference), its variations and the M-B 2008 model

Reference	Removing ICAR	Removing μ_0	Leroux	M-B 2008
64	40	38	55	19

4.3 Comparison of the predictive performance

After validating the relevance of all the components of our proposal, we move now to validate it in terms of its predictive behaviour against the other variations of our detection method. When comparing detection methods, a first approach one can consider is assessing the sensitivity, specificity and timeliness of the models, which requires a gold standard. Obtaining a reliable gold standard is not

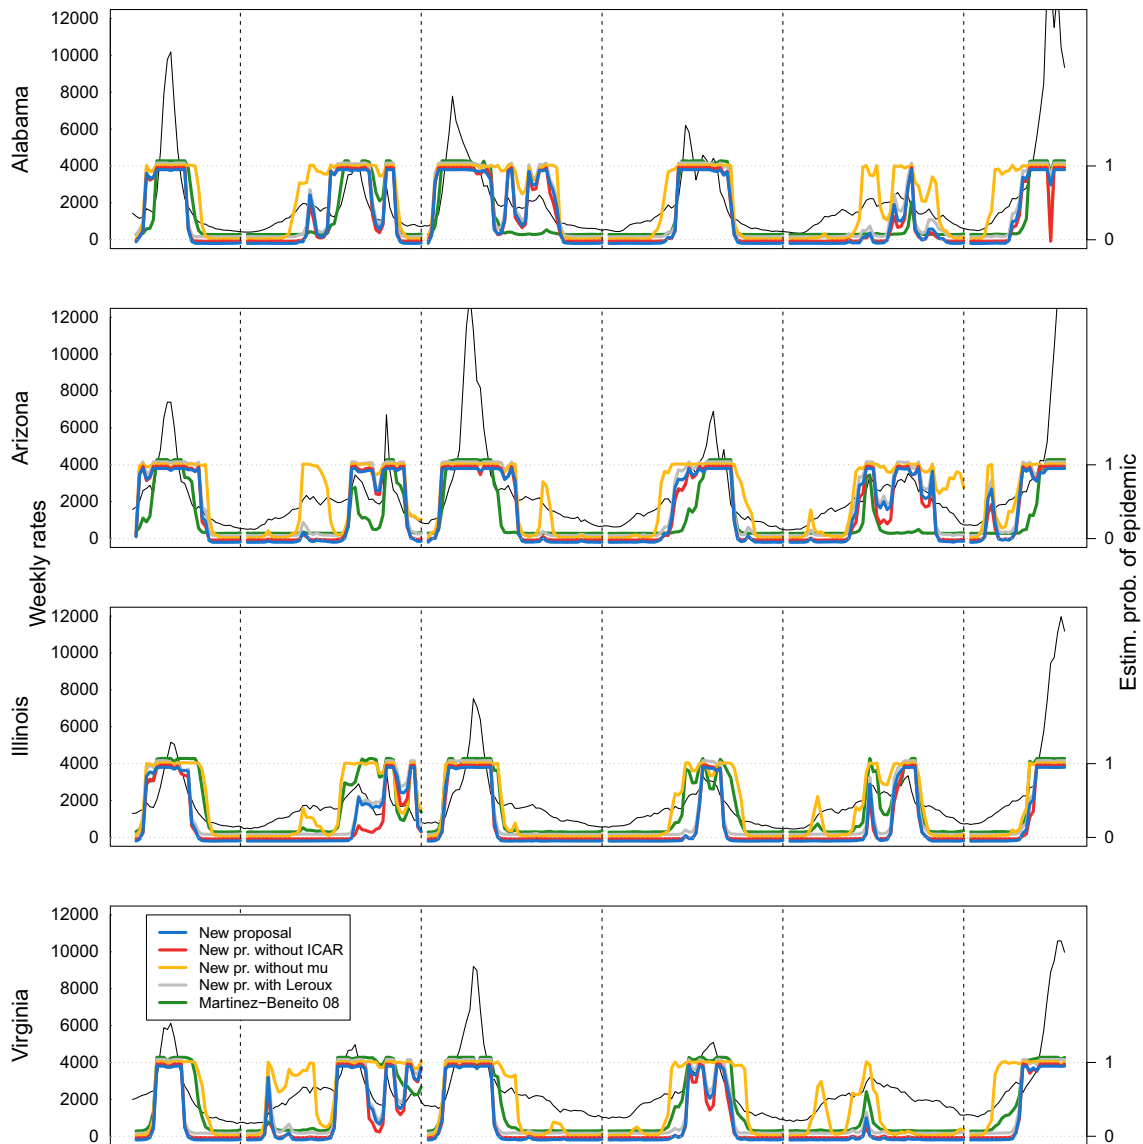


Fig. 6 Estimated probability of being in epidemic phase by the spatio-temporal model, its variations and the M-B 2008 model on USA GFT data for 4 states. In black: weekly estimated influenza incidence per 100 000 inhabitants during seasons from 2007–2008 to 2012–2013

Table 3 Comparison of the estimated mean for the parameters of the spatio-temporal model applied on the USA GFT data in a retrospective basis and its variations, and the average value of the mean for the parameters of the M-B 2008 model applied separately on the 49 temporal strains of data of each state. Reference denotes the novel model proposed in Sect. 2

Param.	Reference	Without ICAR	Without μ_{t0}	Leroux	M-B 2008
p_{00}	0.95	0.96	0.94	0.95	0.95
p_{11}	0.89	0.88	0.92	0.87	0.89
σ_0	85.00	88.27	92.75	84.45	126.78
σ_1	430.89	712.33	390.94	259.51	935.40
$\sigma_{\mu 0}$	127.61	118.72	–	153.39	–
$\sigma_{\mu 1}$	646.69	958.09	418.72	469.81	–
σ_{ψ}	770.10	–	582.51	930.34	–
ρ	0.35	0.39	0.36	0.36	0.51
λ	1.50	1.34	–	1.82	–
ϕ	–	–	–	0.59	–

trivial and different approaches present different problems. Constructing a simulated data set would require a generating spatio-temporal model capable of reproducing the influenza behaviour. Due to the particular spatio-temporal behaviour of influenza, constructing such a generating model is far from trivial, and different proposals can result in very different outcomes of the validation. Another approach would be to use an arbitrary threshold over the laboratory virus isolations (Cowling et al. 2006), but different thresholds can also give inconsistent validity outcomes, and this kind of data is not always available. When a gold standard is not available, as is the case here, other approaches should be adopted. An interesting option used in some works (Boyle et al. 2011) is the comparison of the predictive power of the models. As all models in the comparison have the latent variables Z_{it} in their formulation, the posterior probability for these latent variables in all the models is the estimated probability of being in the epidemic phase. It seems sensible to assume that a model that gives better prediction for the differentiated rates will also give better estimates for the latent variable, from which they directly depend. Therefore the predictive performance can be sensibly used as an indirect measure of the quality of the state classification for any observation. This method has the virtue of not depending on any arbitrary definition of what is epidemic and what is not.

Approximate cross-validated predictive assessment (Marshall and Spiegelhalter 2003) was performed to evaluate the predictive power of the methods. This is a computationally cheaper alternative to the full leave-one-out cross-validated assessment where the model is run only once instead of once per removed observation. Prediction for each region i was calculated ignoring the estimate of the parameter ψ_{it} but taking into account the estimated parameters of the neighbours ψ_{jt} (with $j \sim i$). This process was done for each week on an online basis, doing prediction using only data from the same or previous weeks. In order to do the approximate cross-validated predictive assessment, a measure of the discrepancy between predictions and observed values is needed. As all the models in the comparison are defined under the Bayesian paradigm, predictions are expressed as probability distributions and not as punctual estimates. For this reason, in order to evaluate the goodness of the predictions in comparison to the observations, the Continuous Rank Probability Score (CRPS) was used (Gneiting and Raftery 2007). CRPS is a measure of discrepancy between a probability distribution and any observed value which considers not only the posterior expected value of a distribution but also its precision and shape to calculate the distance between the prediction and the observation, offering lower scores for better predictions.

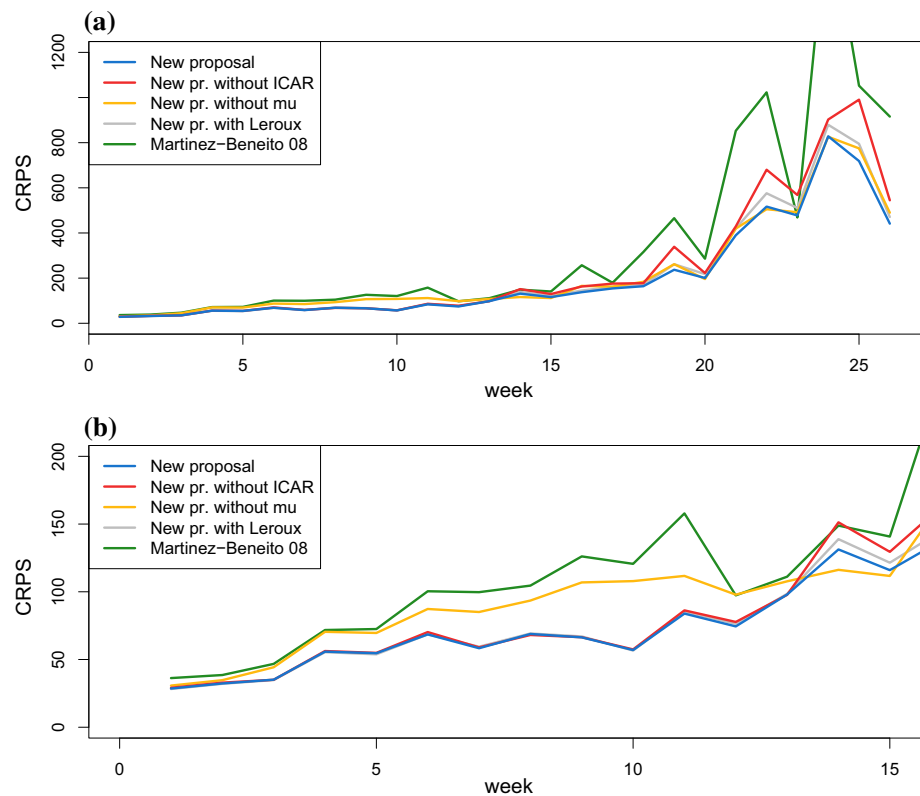
Figure 7 shows the average CRPS of the cross-validated predictive assessment for all 49 locations and each week of season 2012–2013 for the new spatio-temporal proposal, its variations and the M-B 2008 model. This cross-validated predictive assessment has been performed applying each one of the models on an online basis for all the weeks of this last season. The graph at the bottom of that same figure shows a detail of the first 15 weeks, where most locations are classified as endemic. We can observe that the new proposal, depicted in blue, offers the best (lowest) scores of CRPS both in the first weeks, where the majority of the states are in the endemic phase and in the last weeks, where the majority of locations are in the epidemic phase. The red line, representing the model without spatial structure, shows almost equivalent scores in the first weeks, as should be expected; as for the non-epidemic observations the spatial component is not present. The suppression of the ICAR term in the epidemic linear regressor, though, results in worse scores in the last weeks. The same phenomenon, but much milder, happens with the model that substitutes the ICAR structure for a Leroux structure, represented by the grey line. The quality of the prediction is close to that of the new proposal, but somewhat worse for the epidemic weeks. The model represented by the yellow line differs from the original novel proposal only in the non-epidemic regressor, which is modelled without the μ_{t0} term. For this reason, the most visible differences are in the first weeks, as their behaviour is not well captured by the simplified model. The M-B 2008 model, shown in green, gives the worst values among all the compared models both in the first and last weeks. This indicates that sharing information among locations is important to correctly model both endemic and epidemic weeks in a spatio-temporal context.

5 Conclusions

In this work we have presented a spatio-temporal hierarchical Markov switching model for the early detection of influenza outbreaks, shown its application and checked its relevance by comparing it with some variations and with the purely temporal model of Martinez-Beneito et al. (2008).

In the comparison, the purely temporal model for outbreaks detection of Martinez-Beneito et al. (2008) showed a worse performance than our proposal or any of its variations. This demonstrates the superiority of using a spatio-temporal approach over the separate analysis of several time series for a certain set of regions. Even the modification of our proposal without a structured spatial component performed notably better than the purely temporal approach. This suggests that detection for all the regions is

Fig. 7 Cross-validators predictive assessment using CRPS as measure of discrepancy for the new proposal, its variations and M-B 2008. (a) Average CRPS for the prediction of the 49 states is shown for all the weeks of the season 2012–2013. (b) Panel at the bottom shows a zoom for the first 15 weeks of the season so the differences may be better appreciated



benefited by sharing information about the temporal parameters of each time series and about the overall behaviour of all regions over time.

Furthermore, the way influenza epidemics spread suggests that people in neighbouring regions tend to infect each other. Indeed, the use of spatially structured terms (like the ICAR or the Leroux models) in combination with temporal global and local structures improves the performance of the detection models. Our modelling proposal for the detection of influenza outbreaks has therefore proven to be valid by demonstrating the relevance of its components to adequately capture the different aspects of the spatio-temporal evolution of influenza data.

Other advantages of this new proposal come from how these spatio-temporal structures are embedded in a MSM. The use of latent variables for the classification of the observations as endemic or epidemic allows for the new proposal to detect epidemics that start at several foci at different times and extend to all the locations considered. Also, no assumption of the epidemic or endemic state of the training data is required, and the output of the model is directly interpretable as the probability of being in the epidemic state for each time and location.

To summarize, our interest in this study has been to describe a spatio-temporal methodology for the detection of influenza outbreaks at the very moment of their onset. Results show the relevance of using a method which

takes into account both temporal and spatial structures of the data, as the one described does, for a better performance of surveillance systems.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adín A, Martínez-Beneito MA, Botella-Rocamora P, Goicoa T, Ugarte MD (2017) Smoothing and high risk areas detection in space-time disease mapping: a comparison of P-splines, autoregressive and moving average models. *Stoch Env Res Risk Assess* 31:403–415. <https://doi.org/10.1007/s00477-016-1269-8>
- Amorós R (2017) Bayesian temporal and spatio-temporal Markov switching models for the detection of influenza outbreaks. Ph.D. thesis. Department of Statistics and Operational Research, Universitat de València. <http://roderic.uv.es/handle/10550/59265>

- Amorós R, Conesa D, Martínez-beneito M, López-Quílez A (2015) Statistical methods for detecting the onset of influenza outbreaks: a review. *REVSTAT* 13:41–62
- Banks D, Datta G, Karr A, Lynch J, Niemi J, Vera F (2012) Bayesian CAR models for syndromic surveillance on multiple data streams: theory and practice. *Inf Fusion* 13:105–116
- Besag J, York J, Mollié A (1991) Bayesian image restoration, with two applications in spatial statistics. *Ann Inst Stat Math* 43:1–20
- Bhar R, Hamori S (2004) Hidden markov models: applications to financial economics, vol 40. Kluwer Academic Publishers, Boston
- Boyle J, Sparks R, Keijzers G, Crilly J, Lind J, Ryan L (2011) Prediction and surveillance of influenza epidemics. *Med J Aust* 194:S28–33
- Bunke H, Caelli T (2001) Hidden Markov models: applications in computer vision. World Scientific, Singapore
- Cappé O, Moulines E, Rydén T (2005) Inference in hidden Markov models. Springer, New York
- Cheng CK, Lau E, Ip D, Yeung A, Ho L, Cowling B (2009) A profile of the online dissemination of national influenza surveillance data. *BMC Public Health* 9:339
- Conesa D, Martínez-Beneito M, Amorós R, López-Quílez A (2015) Bayesian hierarchical Poisson models with a hidden Markov structure for the detection of influenza epidemic outbreaks. *Stat Methods Med Res* 24:206–223
- Cook S, Conrad C, Fowlkes A, Mohebbi M (2011) Assessing Google Flu Trends performance in the United States during the 2009 influenza virus A (H1N1) pandemic. *PLoS ONE* 6:e23610
- Corberán-Vallet (2012) Prospective surveillance of multivariate spatial disease data. *Stat Methods Med Res* 21:457–477
- Corberán-Vallet A, Lawson A (2011) Conditional predictive inference for online surveillance of spatial disease incidence. *Stat Med* 30:3095–3116
- Corberán-Vallet A, Lawson A (2014) Prospective analysis of infectious disease surveillance data using syndromic information. *Stat Methods Med Res* 23:572–590
- Cowling B, Wong I, Ho L, Riley S, Leung G (2006) Methods for monitoring influenza surveillance data. *Int J Epidemiol* 35:1314–1321
- De Lange M, Meijer A, Friesema I, Donker G, Koppeschaar C, Hooiveld M, Ruigrok N, van der Hoek W (2013) Comparison of five influenza surveillance systems during the 2009 pandemic and their association with media attention. *BMC Public Health* 13:881
- Douc R, Moulines É, Rydén T (2004) Asymptotic properties of the maximum likelihood estimator in autoregressive models with Markov regime. *Ann Stat* 32:2254–2304
- Fox E, Dunson D (2015) Bayesian nonparametric covariance regression. *J Mach Learn Res* 16:2501–2542
- Gasparini R, Amicizia D, Lai P, Panatto D (2012) Clinical and socioeconomic impact of seasonal and pandemic influenza in adults and the elderly. *Hum. Vaccines Immunother* 8:21–28
- Gelman A, Carlin J, Stern H, Dunson D, Vehtari A, Rubin D (2013) Bayesian data analysis, 3rd edn. Chapman and Hall/CRC, Boca Raton
- Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, Brilliant L (2009) Detecting influenza epidemics using search engine query data. *Nature* 457:1012–1014
- Gneiting T, Raftery A (2007) Strictly proper scoring rules, prediction, and estimation. *J Am Stat Assoc* 102(477):359–378
- Gomez-Barroso D, Martínez-Beneito M, Flores V, Amorós R, Delgado C, Botella P, Zurriaga O, Larrauri A (2014) Geographical spread of influenza incidence in Spain during the 2009 A(H1N1) pandemic wave and the two succeeding influenza seasons. *Epidemiol Infect* 142:2629–2641
- Google: Google Flu Trends. <http://www.google.org/flutrends/about/>. Accessed 19 Feb 2017
- Green P, Richardson S (2002) Hidden Markov models and disease mapping. *J Am Stat Assoc* 97:1055–1070
- Grzegorzczak M, Shafiee M (2017) Comparative evaluation of various frequentist and Bayesian non-homogeneous Poisson counting models. *Comput Stat* 32:1–33
- Heaton M, Banks D, Zou J, Karr A, Datta G, Lynch J, Vera F (2012) A spatio-temporal absorbing state model for disease and syndromic surveillance. *Stat Med* 31:2123–2136
- Khadr M (2016) Forecasting of meteorological drought using Hidden Markov Model (case study: The upper Blue Nile river basin, Ethiopia). *Ain Shams Eng J* 7(1):47–56
- Knorr-Held L (2000) Bayesian modelling of inseparable space-time variation in disease risk. *Stat Med* 19:2555–2567
- Knorr-Held L, Richardson S (2003) A hierarchical model for space-time surveillance data on meningococcal disease incidence. *J R Stat Soc: Ser C* 52:169–183
- Kulldorff M (2001) Prospective time periodic geographical disease surveillance using a scan statistic. *J Royal Stat Soc: Ser A* 164:61–72
- Le Strat Y, Carrat F (1999) Monitoring epidemiologic surveillance data using hidden Markov models. *Stat Med* 18:3463–3478
- Leroux B, Lei X, Breslow N (2000) Estimation of disease rates in small areas: a new mixed model for spatial dependence. In: Halloran ME, Berry D (eds) *Statistical models in epidemiology, the environment, and clinical trials*. Springer, New York, pp 179–191
- Li G, Best N, Hansell A, Ahmed I, Richardson S (2012) BaySTDetect: detecting unusual temporal patterns in small area data via Bayesian model choice. *Biostatistics* 13:695–710
- Lofgren E, Fefferman N, Naumov Y, Gorski J, Naumova E (2007) Influenza seasonality: underlying causes and modeling theories. *J Virol* 81:5429–5436
- Lu H, Zeng D, Chen H (2010) Prospective infectious disease outbreak detection using Markov switching models. *IEEE T Knowl Data En* 22(4):565–577
- Lunn D, Thomas A, Best N, Spiegelhalter D (2000) WinBUGS: a Bayesian modelling framework: concepts, structure, and extensibility. *Stat Comput* 10:325–337
- Lytras T, Gkolfinopoulou K, Bonovas S, Nunes B (2018) FluHMM: a simple and flexible Bayesian algorithm for sentinel influenza surveillance and outbreak detection. *Stat Methods Med Res*
- Madigan D (2005) Bayesian data mining for health surveillance. In: Lawson A, Kleinman K (eds.) *Spatial and syndromic surveillance for public health*, chap. 12, pp. 203–221. Wiley, Chichester
- Marshall E, Spiegelhalter D (2003) Approximate cross-validatory predictive checks in disease mapping models. *Stat Med* 22:1649–1660
- Martínez-Bello D, López-Quílez A, Torres Prieto A (2018) Spatiotemporal modeling of relative risk of dengue disease in Colombia. *Stoch Env Res Risk Assess* 32:1587–1601
- Martínez-Beneito M, Conesa D, López-Quílez A, López-Maside A (2008) Bayesian Markov switching models for the early detection of influenza epidemics. *Stat Med* 27:4455–4468
- Martínez-Beneito M, López-Quílez A, Botella-Rocamora P (2008) An autoregressive approach to spatio-temporal disease mapping. *Stat Med* 27:2874–2889
- Martínez-Minaya J, Cameletti M, Conesa D, Pennino MG (2018) Species distribution modeling: a statistical review with focus in spatio-temporal issues. *Stoch Env Res Risk Assess* 32:3227–3244. <https://doi.org/10.1007/s00477-018-1548-7>
- Mugglin A, Cressie N, Gemell I (2002) Hierarchical statistical modelling of influenza epidemic dynamics in space and time. *Stat Med* 21:2703–2721

- Nunes B, Natário I, Carvalho M (2013) Nowcasting influenza epidemics using non-homogeneous hidden Markov models. *Stat Med* 32:2643–2660
- Olson D, Konty K, Paladini M, Viboud C, Simonsen L (2013) Reassessing Google Flu Trends data for detection of seasonal and pandemic influenza: a comparative epidemiological study at three geographic scales. *PLoS Comput Biol* 9:e1003256
- Ortiz J, Sotomayor V, Uez O, Oliva O, Bettels D, McCarron M, Bresee J, Mounts A (2009) Strategy to enhance influenza surveillance worldwide. *Emerg Infect Dis* 15:1271–1278
- Rafei A, Pasha E, Jamshidi Orak R (2015) A warning threshold for monitoring tuberculosis surveillance data: an alternative to hidden Markov model. *Trop Med Int Health* 20:919–929
- Rao Y, McCabe B (2016) Real-time surveillance for abnormal events: the case of influenza outbreaks. *Stat Med* 35:2206–2220
- Rath T, Carreras M, Sebastiani P (2003) Automated detection of influenza epidemics with hidden Markov models. In: Berthold MR, Lenz H, Bradley E, Kruse R, Borgelt C (eds) *Advances in intelligent data analysis V*. Springer, Berlin, pp 521–532
- Rogerson P, Yamada I (2004) Monitoring change in spatial patterns of disease: comparing univariate and multivariate cumulative sum approaches. *Stat Med* 23:2195–2214
- Rotejanaprasert C, Lawson A (2016) Bayesian prospective detection of small area health anomalies using Kullback-Leibler divergence. *Stat Methods Med Res* 27:1076–1087
- Spreco A, Timpka T (2016) Algorithms for detecting and predicting influenza outbreaks: Metanarrative review of prospective evaluations. *BMJ Open* 6:e010683
- Sun W, Cai T (2009) Large-scale multiple testing under dependence. *J Royal Stat Soc: Ser B* 71:393–424
- Torres-Avilés F, Martínez-Beneito MA (2015) STANOVA: a smooth-ANOVA-based model for spatio-temporal disease mapping. *Stoch Env Res Risk Assess* 29(1):131–141. <https://doi.org/10.1007/s00477-014-0888-1>
- Unkel S, Farrington C, Garthwaite P, Robertson C, Andrews N (2012) Statistical methods for the prospective detection of infectious disease outbreaks: a review. *J Royal Stat Soc: Ser A* 175:49–82
- Vasas K, Elek P, Márkus L (2007) A two-state regime switching autoregressive model with an application to river flow analysis. *J Stat Plann Inference* 137(10):3113–3126
- World Health Organization: Influenza (seasonal). Fact sheet n.211. <http://www.who.int/mediacentre/factsheets/fs211/en/>. Accessed 12 Feb 2016
- Yoon BJ (2009) Hidden Markov models and their applications in biological sequence analysis. *Curr Genomics* 10(6):402–415
- Zhou H, Lawson A (2008) EWMA smoothing and Bayesian spatial modeling for health surveillance. *Stat Med* 27:5907–5928
- Zou J, Karr A, Banks D, Heaton M, Datta G, Lynch J, Vera F (2012) Bayesian methodology for the analysis of spatial-temporal surveillance data. *Stat Anal Data Min* 5:194–204

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Rubén Amorós^{1,2}  · David Conesa^{2,3} · Antonio López-Quílez^{2,3} · Miguel-Angel Martínez-Beneito^{4,5}

¹ School of Mathematics, University of Edinburgh, The King's Buildings, Peter Guthrie Tait Road, Edinburgh EH9 3FD, UK

² Valencia Bayesian Research Group, Valencia, Spain

³ Departament d'Estadística i Investigació Operativa, Universitat de València, C/ Dr. Moliner, 50, 46100 Burjassot, Spain

⁴ Fundació per al Foment de la Investigació Sanitària i Biomèdica de la Comunitat Valenciana, Avda. de Catalunya, 21, 46020 Valencia, Spain

⁵ CIBER Epidemiología y Salud Pública (CIBERESP), Madrid, Spain

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.