

---

---

## STATISTICAL METHODS FOR DETECTING THE ONSET OF INFLUENZA OUTBREAKS: A REVIEW

---

---

- Authors: RUBÉN AMORÓS  
– Fundación para el Fomento de la Investigación Sanitaria y Biomédica de la Comunitat Valenciana, Spain  
ruben.amoros@uv.es
- DAVID CONESA  
– Depart. d'Estadística i Investigació Operativa, Univ. de València, Spain  
david.v.conesa@uv.es
- MIGUEL ANGEL MARTINEZ-BENEITO  
– Fundación para el Fomento de la Investigación Sanitaria y Biomédica de la Comunitat Valenciana,  
CIBER de Epidemiología y Salud Pública (CIBERESP), Spain  
martinez\_mig@gva.es
- ANTONIO LÓPEZ-QUÍLEZ  
– Depart. d'Estadística i Investigació Operativa, Univ. de València, Spain  
antonio.lopez@uv.es

Received: October 2014    Revised: November 2014    Accepted: December 2014

### Abstract:

- This paper reviews different approaches for determining the epidemic period from influenza surveillance data. In the first approach, the process of differenced incidence rates is modeled either with a first-order autoregressive process or with a Gaussian white noise process depending on whether the system is in an epidemic or a non-epidemic phase. The second approach allows us to directly model the process of the observed cases via a Bayesian hierarchical Poisson model with Gaussian incidence rates whose parameters are modeled differently, depending on the epidemic phase of the system. In both cases transitions between both phases are modeled with a hidden Markov switching model over the epidemic state. Bayesian inference is carried out and both models provide the probability of being in epidemic state at any given moment. A comparison of both methodologies with previous approaches in terms of sensitivity, specificity and timeliness is also performed. Finally, we also review a web-based client application which implements the first methodology for obtaining the posterior probability of being in an epidemic phase.

### Key-Words:

- *autoregressive modeling; Bayesian inference; influenza; hidden Markov models; public health; temporal surveillance.*

### AMS Subject Classification:

- 62F15, 62P10.



---

## 1. INTRODUCTION

---

Influenza is an infectious disease that affects the upper and/or lower parts of the respiratory tract and is caused by the influenza virus. Influenza spreads around the world in seasonal epidemics, resulting in severe infections and the deaths of hundreds of thousands worldwide annually, and millions in pandemic years (good descriptions about its impact can be found in Simonsen *et al.* [52], Fleming *et al.* [16], and Monto [36]). This propensity for causing large scale seasonal epidemics and pandemics has clearly turned influenza surveillance into a challenging issue in public health practice.

Prospective local and national influenza surveillance systems can provide important and timely information to health service providers on the circulation of the influenza virus in a population. The emergence of the A-H1N1 influenza virus in 2009 and its subsequent rapid global spread was a clear and real example of the need of having good surveillance infrastructures available. But at the same time it has also focused attention on surveillance capabilities worldwide (Lipsitch *et al.*, [31]). In particular, it has shown the need for sentinel surveillance systems that could strengthen a country's capacity for seasonal, novel, and pandemic influenza detection and prevention (Ortiz *et al.* [44]).

There are several surveillance information sources that make it possible to track influenza virus activity such as real-time internet surveys [58], queries [14], microblogging [30], over the counter sales [33], prescription pharmaceutical sales [45], absenteeism registers [13], syndromic/sentinel surveillance [21], laboratory test isolations [43], emergency room visit rates [10], hospital admissions [8], pneumonia and influenza mortality rates [48], etc. Nevertheless, there is no such thing as the "best" surveillance information source. As Cheng *et al.* [5] state, "each method only captures a portion of infections within the community with different timeliness and specificity". On one hand, laboratory test results are highly specific, but take days or even weeks and only capture a small fraction of the infected population. In contrary, data may be collected instantly from the internet, typically from searching engines or social networks, as Broniatowski *et al.* [2], Gesualdo *et al.* [18], Li and Cardie [30] or Grover and Aujla [22] do with Twitter data, or Google Flutrends does with Google queries (Ginsbert *et al.* [19]). A discussion about the role of internet data sources in disease surveillance can be found in Milinovich *et al.* [35]. However, these sources of information provide just indirect measures of the influenza incidence levels in the population so their accuracy may be sometimes poor. Nevertheless, the high amount and the immediate availability of the data from this non-conventional sources make its analysis particularly interesting.

No matter what kind of data the surveillance system uses, there is always the need for an algorithm that, applied to the data, could quickly detect meaningful increases in reported influenza incidence. This would make it possible health

services to get prepared for the incipient outbreak, which could have a great impact on the number of lives saved, outbreak management and the effective setting of prevention measures. This has made the statistical literature to pay considerable attention to early detection methods for influenza, or infectious diseases in general.


Methods based on historical limit are the most widely used for detecting the onset of influenza epidemics and with longer tradition in the epidemiologic literature. These methods are based on the model of Shewhart [50], where a warning is triggered when the difference between the current observation and a theoretical mean of the process surpasses a determined threshold, usually set using the estimated standard error. One way to determine this theoretical mean and threshold is to consider a window of observations of times  $t - m, \dots, t - 1$  from the present year and/or  $t - m, \dots, t + m$  times from previous years and compute some central estimator and standard error for the observations in these windows, as Stroup *et al.* [55] or Farrington *et al.* [15] do. Another option would be using all non epidemic data as training, fitting a regression model which includes time trend and Fourier periodical terms as proposed by Serfling's method [49], the approach used for influenza surveillance by the Center for Disease Control and Prevention (CDC) of the United States (Muscatello *et al.* [38]). This approach or modifications are also used in other works like Costagliola *et al.* [11], Simonsen *et al.* [52] and Boyle *et al.* [1].

These approaches have some drawbacks in practice (Rath *et al.* [47]). Firstly, a predefinition of epidemic and non-epidemic periods is needed for most of them in order to characterize the observations of the non-epidemic phase, when that division is precisely the final outcome that we want to draw. Secondly, time observations are treated as completely independent values, when we would expect that their temporal arrangement could induce some kind of dependence. Thirdly, the baseline (non-epidemic) period is often estimated with national data that maybe does not properly fit if we are mostly interested in a local influenza surveillance system. Finally, Goddard *et al.* [20] also point out as a fourth drawback that the use of temporally fixed threshold values to describe the levels of influenza activity can be misleading due to time trends in consultations for influenza. Specifically, they pointed out a decline in the number of influenza-related consultations in recent years that could reduce the sensitivity of these methods.

Le Strat and Carrat [28] pioneered the use of hidden Markov models to segment time series of influenza indicators into epidemic and non-epidemic phases. Hidden Markov models are a particular case of Markov switching models, which are stochastic models that consider a set of non observed variables  $Z_t$  (hidden states, usually  $Z_t = 0$  for the non epidemic state and  $Z_t = 1$  for the epidemic state) and a set of observed values  $y_t$  (observations), one for each time unit  $t \in \{1, \dots, T\}$ , so that  $\{Z_t\}$  is a Markov chain

$$(1.1) \quad P(Z_t|Z_1, \dots, Z_{t-1}, p_{ij}) = P(Z_t|Z_{t-1}, p_{ij})$$

where  $p_{ij}$  are the transition probabilities, and the observations  $y_t$  are dependent on previous observations, usually through an autoregressive process. The present state  $Z_t$  affects both the present observation  $y_t$  through the transition probabilities  $p(y_t|y_{t-1})$ . The conditional relationships in a Markov switching model are represented in Figure 1.



**Figure 1:** Diagrams of the conditional dependencies in a Markov switching model.

Hidden Markov models are usually formulated with an added restriction, so that the value of the observed variable at each time  $y_t$  is only dependent on the hidden state for that time, given the past observations and the present and past states

$$P(y_t|Z_1, \dots, Z_t, y_1, \dots, y_{t-1}, \boldsymbol{\theta}) = P(y_t|Z_t, \boldsymbol{\theta})$$

where  $\boldsymbol{\theta}$  are the remaining parameters of the model.

Le Strat and Carrat’s approach has two advantages, the first being that the method can be applied to historical data without the need to previously distinguish between epidemic and non-epidemic periods in the data. The second one is considering observations as dependent on the past observations and states or, at least, on the last epidemic state, whereas Serfling’s method assumes marginal independence of the data [47]. In subsequent papers, Rath *et al.* [47], Madigan [32] and Sun and Cai [56] further developed that modeling. Nevertheless, we find also convenient to mention some other contributions beyond Markov Switching models such Cowling *et al.* [12], Griffin *et al.* [21] and Boyle *et al.* [1] who use models based on statistical quality control and time series methods like CUSUM, EWMA or auto-regressive processes, Frisen *et al.* [17] who search for change points on the monotony of the process, Shmueli [51] who uses wavelet-based methods or Nuño and Pagano who adjust one or several Gaussian peaks in different locations for each year [42]. Three comprehensive reviews of statistical algorithms for the detection of infectious disease outbreaks can be found in Le Strat [27], Burkom [3] and Unkel *et al.* [57].

Bayesian methodology provides a unified theory for handling uncertainty in very different areas such as statistical inference, forecasting, decision-making

under uncertainty, analysis of expert systems, etc. This ability to deal with uncertainties is what makes Bayesian analysis a very advisable tool for many issues that arise in the decision-making process of a surveillance system. Specifically, Bayesian analyses enable to quantify whatever feature of interest of any variable in the model by means of its posterior distribution. In our setting, this makes the Bayesian methodology to be perfectly suited for quantifying the probability of being in an epidemic phase at any given moment. Bayesian studies are not new in surveillance literature, but in recent years there has been increasing interest in them (see for example Mugglin *et al.* [37], Cooper *et al.* [9], Sebastiani *et al.* [48], Niemi *et al.* [40], Zhou and Lawson [61], Charland *et al.* [4] and Neill and Cooper [39]).

Our main goal in this paper is to review two alternative approaches to influenza surveillance that avoid some of the above-mentioned disadvantages and take advantage of the ability to quantify epidemic probabilities of Bayesian methodology. The first proposal is to use a Markov switching model in order to determine the epidemic and non-epidemic periods from influenza surveillance data (Martínez-Beneito *et al.* [34]). This approach differs from those hidden Markov models previously mentioned in the sense that it models the series of differenced incidence rates rather than the series of incidence rates. The new differenced series is detrended, allowing us to take advantage of autoregressive (stationary) modeling to analyze the data. In particular, depending on whether the system is in an epidemic or a non-epidemic phase, the differenced series are, respectively, modeled either with a first-order autoregressive process or with a Gaussian white noise process. The transition between the phases of the disease is considered to follow a Markovian process. The Bayesian paradigm is used to estimate the probability of being in an epidemic phase at any given moment, which is the key to detecting influenza epidemics at their onset.

Two features of this model have proved to be very convenient in the influenza surveillance context. The first one is the use of Markov switching models to segment the time series of influenza into epidemic and non-epidemic phases. The second is the use of the variability of data as main tool to distinguish between both epidemic and non-epidemic phases. Thus, the underlying hypothesis of this model is that non-epidemic dynamics are characterized by small, time-independent random changes (since, supposedly, there is no underlying active process) meanwhile, in epidemic dynamics, changes are greater and possibly correlated.

Nevertheless, although the variability of data may enable to distinguish both dynamics, incorporating the magnitude of the incidence rates and not just their differences could also be very advantageous because this magnitude would also inform on the state of the illness (low incidence clearly meaning a non-epidemic phase). This could increase the capability of the method to distinguish between both epidemic and non-epidemic phases and so it could be easier to

determine the onset of epidemics. As a result, we also review here an enhanced version of that modeling (Conesa *et al.* [7]) that also incorporates the magnitude of the incidence rates. Moreover, this proposal directly models the weekly observed counts as a Poisson distribution depending on the incidence rates, thus these incidence rates are not considered deterministic quantities but, on the contrary, (random) variables in a Bayesian model. As a consequence, this proposal also takes into account the uncertainty in the available incidence rates.

The remainder of this paper is organized as follows. In Sections 2 and 3 we present, respectively, the model based on the differenced rates and on the observed cases. In Section 4 we describe the results obtained when applying the proposed methodologies in a particular setting related to sentinel surveillance data. That section also includes a validation of the performance of both models compared with other existing methodologies. Section 5 describes `fludetweb`, a web-based implementation of the first methodology for obtaining the posterior probability of being in an epidemic phase. Finally, in Section 6 we present some concluding remarks and some lines of development of the current methodology.

---

## 2. MODELING OF THE DIFFERENCED INCIDENCE RATES

---

We first review the modeling introduced by Martinez-Beneito *et al.* [34]. This model performs a segmentation of the differenced incidence rates series into epidemic and non-epidemic phases by using a Markov switching model. Specifically, let  $Y = \{Y_{i,j}, i = 1, \dots, I; j = 1, \dots, J\}$  denote the set of differences between the rates of weeks  $i + 1$  and  $i$  in year  $j$ . We consider a set of retrospective years so that the system has previous information about epidemic periods before observing it in the current year. The underlying idea of Markov switching models is to associate each  $Y_{i,j}$  with a random variable  $Z_{i,j}$  that determines the conditional distribution of  $Y_{i,j}$  given  $Z_{i,j}$ . In this case, each  $Z_{i,j}$  is an unobserved random variable that indicates which phase the system is in (1, epidemic; 0, non-epidemic). The unobserved sequence of  $Z_{i,j}$  follows a first order two-state Markov chain with transition probabilities:

$$P(Z_{i+1,j} = l | Z_{i,j} = k) = P_{k,l}$$

where  $k, l \in \{0, 1\}$ ,  $i \in \{1, \dots, I - 1\}$  and  $j \in \{1, \dots, J\}$ . This Markovian feature enables epidemic, respectively non-epidemic, weeks to be followed by epidemic, respectively non-epidemic, weeks with a high probability if the data required it. This performance could not be achieved with an independent modeling of the  $Z_{i,j}$ 's and it makes the epidemic/non-epidemic state to be more robust to sudden, although slight, changes in the differenced series.

The next step is to model the behavior of the differenced series for both epidemic and non-epidemic periods. It seems reasonable to assume no underlying

process beyond Gaussian noise for the non-epidemic period since, supposedly, no underlying mechanism should be inducing dependence among the observations. On the other hand, the epidemic phase should show greater variability, and possibly dependent observations. Therefore, the conditional distribution of  $Y_{i,j}$  is modeled either as a Gaussian white noise process or as an autoregressive process of order 1, depending on whether the system is in, respectively, non-epidemic or epidemic phase, *i.e.*

$$\begin{aligned}
 & Y_{1,j}|(Z_{1,j} = 0) \sim N(0, \sigma_{0,j}^2) \\
 & Y_{1,j}|(Z_{1,j} = 1) \sim N(0, \sigma_{1,j}^2) \\
 (2.1) \quad & Y_{i,j}|(Z_{i,j} = 0) \sim N(0, \sigma_{0,j}^2) \quad i \in \{2, \dots, I\}, \quad j \in \{1, \dots, J\}, \\
 & Y_{i,j}|(Z_{i,j} = 1) \sim N(\rho Y_{i-1,j}, \sigma_{1,j}^2) \quad i \in \{2, \dots, I\}, \quad j \in \{1, \dots, J\},
 \end{aligned}$$

where the first subindex of the variance  $\sigma_{k,j}^2$  represents whether the system is in the epidemic phase ( $k = 1$ ) or not ( $k = 0$ ). This model assumes a different variance for each season in order to reflect that the variability in any of the phases is not necessarily the same in different years, as a consequence of differences in the shape of the corresponding epidemic waves. Note also that the conditional distribution of the first difference of rates cannot be modeled as an autoregressive process as there is no previous value to condition on.

Once the model is determined, the following step is to estimate its parameters. Martínez-Beneito *et al.* [34] propose using the following prior distributions for the parameters involved in the model:

$$\begin{aligned}
 & \rho \sim \text{Unif}(-1, 1) & \theta_{\text{low}} &= \lambda_{[1]} \\
 & P_{1,1} \sim \text{Beta}(0.5, 0.5) & \theta_{\text{mid1}} &= \lambda_{[2]} \\
 (2.2) \quad & P_{0,0} \sim \text{Beta}(0.5, 0.5) & \theta_{\text{mid2}} &= \lambda_{[3]} \\
 & \sigma_{0,j} \sim \text{Unif}(\theta_{\text{low}}, \theta_{\text{mid1}}) & \theta_{\text{sup}} &= \lambda_{[4]} \\
 & \sigma_{1,j} \sim \text{Unif}(\theta_{\text{mid2}}, \theta_{\text{sup}})
 \end{aligned}$$

where  $\{\lambda_{[1]}, \lambda_{[2]}, \lambda_{[3]}, \lambda_{[4]}\}$  corresponds to the ordered sequence of the variables  $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$  which follow as prior distribution:

$$(2.3) \quad \lambda_j \sim \text{Unif}(a, b) \quad j = 1, \dots, 4,$$

where  $a$  and  $b$  are hyperparameters to be fixed by the modeler, typically expressing vague prior knowledge.

Expressions (2.1) and (2.2) contain all the knowledge about the system but they do not yield analytical estimates. Therefore, Markov Chain Monte Carlo (MCMC) methods are necessary, WinBUGS [54] being an option for carrying out the inference. See Martínez-Beneito *et al.* [34] for more details on the specific implementation of this model.



---

### 3. DIRECT MODELING OF THE OBSERVED COUNTS

---

The key for developing a method that could be adapted to most kinds of surveillance data is to use a common feature to all of them. This feature is that most surveillance systems are usually fed of counting incidence data. In fact, most of the available surveillance data consist of time series of daily/weekly rates directly obtained by transforming time series of observed cases (such as the number of daily/weekly deaths due to influenza, the number of daily/weekly hospital admissions, etc.). Attending to this comment, Conesa *et al.* [7], model the weekly observed cases, which are subjected to sampling variation that should be taken into account in our proposal, instead of directly modeling the rates. In particular, if  $O_{i,j}$  denotes the number of observed cases of influenza during week  $i$  in season  $j$ , they model  $O_{i,j}$  by means of a Poisson distribution whose parameter is a function of the incidence rate  $r_{i,j}$  of the week  $i$  in season  $j$  via the following hierarchical structure:

$$(3.1) \quad \begin{aligned} O_{i,j} &\sim \text{Poisson}(\nu_{i,j}) \\ \nu_{i,j} &= f(r_{i,j}) \\ r_{i,j} &\sim \mathcal{N}(R_{i,j}(Z_{i,j}), \sigma_j^2(Z_{i,j})) . \end{aligned}$$

The function in the second line in (3.1) depends on the type of data we are working with. For instance, when working with sentinel data in which data are formed by the weekly percentage of patients with influenza, the function in the second line in (3.1) will be:

$$(3.2) \quad f(r_{i,j}) = \frac{N_{i,j} \cdot r_{i,j}}{100} ,$$

where  $N_{i,j}$  represents the total number of patients seen in the corresponding week.

The rates now are modeled again as a Normal distribution with both mean and variance depending on  $Z_{i,j}$ , an unobserved random variable that indicates which phase the system is in (1, epidemic; 0, non-epidemic). As in the previous model, this is the idea of a Markov switching model in which the unobserved sequence of  $Z_{i,j}$  follows a two-state Markov chain of first order with transition probabilities:

$$(3.3) \quad P(Z_{i+1,j} = l | Z_{i,j} = k) = P_{k,l} \quad k, l \in \{0, 1\} ,$$

with  $P_{0,1} = 1 - P_{0,0}$  and  $P_{1,0} = 1 - P_{1,1}$  for suitable probabilities  $P_{0,0}$  and  $P_{1,1}$ .

The model assumes constant but different variances for each phase of each season:  $\{\sigma_j^2(0), \sigma_j^2(1) : j = 1, \dots, J\}$ . Moreover, the variance of the epidemic phase will be assumed higher than that in the non-epidemic phase. As already mentioned, non-epidemic dynamics are characterized by small random changes while

in epidemic dynamics the changes are greater. This will help once again to separate the series in 2 different periods. Different variances are assumed for each season in order to reflect the different features of the different epidemic seasons since, for example, some years have higher and steeper incidence peaks in contrast to other years with flatter epidemic waves.

The next step is to model the mean of the rates in both states. Note that  $R_{i,j}(0)$  and  $R_{i,j}(1)$  represent the level of magnitude of the incidence in case of being, respectively, at non-epidemic or epidemic phase. The model at every week decides which of them (jointly with the corresponding variance) fits better to the new observed data. The first and easiest way to model both means is to consider them as independent but constant. This modeling can be denoted as AR0-AR0, *i.e.* two order 0 (independent) autoregressive processes, these terms representing respectively the non-epidemic and epidemic phases. A second option would be to consider the mean of the rates as temporally dependent processes. In this setting, it could be convenient to model the mean of the rates as a first order autoregressive process. As a result, three more modelings could be considered combining the AR0 and AR1 proposals, that is: AR0-AR1, AR1-AR0 and AR1-AR1. In a similar way, we could think that the rates are related to rates from two or more previous weeks. Then a suitable option would clearly be to consider the mean of the rates as an autoregressive process of higher order. When dealing with second order autoregressive processes, five more modelings could be considered, that is AR2-AR0, AR2-AR1, AR2-AR2, AR1-AR2 and AR0-AR2. For simplicity, here we just present the AR1-AR1 model, with both means (those corresponding to non-epidemic and epidemic settings) being first order autoregressive processes, *i.e.*:

$$(3.4) \quad \begin{aligned} R_{i,j}(0)|r_{1,j}, \dots, r_{(i-1),j} &= \mu_0 + \rho_0 \cdot (r_{i-1,j} - \mu_0) \\ R_{i,j}(1)|r_{1,j}, \dots, r_{(i-1),j} &= \mu_1 + \rho_1 \cdot (r_{i-1,j} - \mu_1) \end{aligned}$$

with  $\mu_0 < \mu_1$  in order to set the epidemic period as that having a higher expected rate.

The specification proposed by Conesa *et al.* [7] for the prior distributions of each of the parameters involved in this model is the following:

$$\begin{aligned} P_{0,0} &\sim \text{Beta}(0.5, 0.5) & \theta_{\text{low}} &= \lambda_{[1]} \\ P_{1,1} &\sim \text{Beta}(0.5, 0.5) & \theta_{\text{mid1}} &= \lambda_{[2]} \\ \sigma_j(0) &\sim \text{Unif}(\theta_{\text{low}}, \theta_{\text{mid1}}) & \theta_{\text{mid2}} &= \lambda_{[3]} \\ \sigma_j(1) &\sim \text{Unif}(\theta_{\text{mid2}}, \theta_{\text{sup}}) & \theta_{\text{sup}} &= \lambda_{[4]} \end{aligned}$$

where  $\{\lambda_{[1]}, \lambda_{[2]}, \lambda_{[3]}, \lambda_{[4]}\}$  corresponds to the ordered sequence of the variables  $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$  which follow the non-informative prior distribution:

$$(3.5) \quad \lambda_j \sim \text{Unif}(0, c) \quad j = 1, \dots, 4,$$

where  $c$  is usually high enough so that it does not condition the posterior distribution of  $\lambda_1, \dots, \lambda_4$ .

With respect to the parameters involved in the modeling of the mean of the rates, that is,  $\mu_0$  and  $\mu_1$ , they propose a slightly different structure than the one used for the variance:

$$(3.6) \quad \mu_0 = \theta_{[1]} , \quad \mu_1 = \theta_{[2]} ,$$

where  $\{\theta_{[1]}, \theta_{[2]}\}$  corresponds to the ordered sequence of the variables  $\{\theta_1, \theta_2\}$  which follow the non-informative prior distribution:

$$(3.7) \quad \theta_j \sim \text{Unif}(0, d) , \quad j = 1, 2 ,$$

where  $d$  is again a hyperparameter chosen in a way that makes vague the former distribution.

Finally, they also propose considering flat prior distributions for any of the parameters of the autoregressive processes. Specifically, they choose uniform distributions in the region where the processes are stationary. As an example, the selection in the AR1-AR1 modeling would be:

$$(3.8) \quad \rho_0, \rho_1 \sim \text{Unif}(-1, 1) .$$

Again, expressions from (3.1) to (3.8) contain all the knowledge of the system but these expressions do not yield analytical expressions for the posterior distribution of the parameters. MCMC methods and WinBUGS are again a good option for carrying out the inference. See Conesa *et al.* [7] for more details on the specific implementation of this model.

---

#### 4. COMPARING METHODOLOGIES ON SENTINEL SURVEILLANCE DATA

---

One of the most popular kinds of influenza surveillance data comes from sentinel systems. In spite of their limitations, these data have proved to be very useful to follow up influenza during the last decade, rapid information transmission being one of their main advantages. Basically, sentinel systems are formed by volunteer practitioners that, depending on the system, report in a weekly basis the percentage of patients with Influenza-like illness (ILI), usually defined as fever plus acute respiratory symptoms such as cough and/or sore throat, from the total number of patients seen, or just the number of consultations with patients reporting ILI symptoms. Data are collected at least during the periods of non-negligible influenza activity. Thus in Western countries in the temperate climate

zone, data are typically collected in seasons (lasting around 30–35 weeks) that extend over two consecutive years (as the epidemic activity usually extends across both of them), while in other places data are collected throughout the year.

In what follows we review a case study presented in Conesa *et al.* [7] in order to check how these two methodologies can be applied in real settings. In particular data analyzed in this example are retrieved from the North Carolina Influenza Sentinel Surveillance Program. The dataset is formed by the weekly percentage of ILI patients from the total number of seen patients. Information about this system can be found in its website [41], which allows to export raw data in standard format. In this system, sentinels also collect representative samples for virus strain identification. As we mention later, this laboratory data on influenza isolates can be very useful as a gold standard for validating the performance of a method.

It is worth noting that the methodology introduced in Section 3 requires the selection of the model that fits each phase better from the different options introduced in the previous section. As a result, the first step must be to use a model selection criterion. One option is to use DIC, a Bayesian model comparison criterion introduced by Spiegelhalter *et al.* [53], based on the trade-off between the fit of the data and the corresponding complexity of the model. In our experience, models with different structures in both phases tend to give more weight to the phase with more structure. As a result we usually consider just models with both phases having the same structure, in order to avoid decompensation between phases. In this case, we have only analyzed those three options with the same structure in both phases (AR0-AR0, AR1-AR1 and AR2-AR2). For this data, results indicate that models AR1-AR1 and AR2-AR2 outperform (as they have lower DIC, 2097.6 and 2095.9 respectively) the AR0-AR0 model (with DIC 2125.8). As a result the AR0-AR0 model can be discarded, the next step being to choose between the remaining two models.

The assessment of models can be done using the measures proposed by Cowling *et al.* [12] and Kleinman and Abrams [25]. These measures summarize the information on sensitivity, specificity and timeliness for a detection method and a particular data set in a unique value between 0 and 1, achieving 1 for a perfect performance. In particular, the measures used here are **AUWROC1** and **VUTROS1**, which weigh the information given by the ROC curve by the timeliness achieved for each method, building a 3D ROC curve which height is a function of the timeliness of detection and measuring the area left behind this curve, and **VUTROS3** and **VUTROCS**, which construct several reference ROC curves restricted by different maximum delay of detection and integrate the information from them all.

In order to perform this comparison we need to know approximately the real epidemic periods in all the seasons. The approach used here to obtain the gold-standard is that presented in Cowling *et al.* [12]. In particular, using laboratory

data on influenza isolates from the North Carolina Sentinel Network, the period between the first and last week in which the proportion of positive isolations for influenza surpassed 30% of the maximum seasonal level has been taken as the gold standard epidemic phase.

This criteria may also be used beyond the comparison of the AR1-AR1 and AR2-AR2 models to compare these proposals with alternative models in the literature. In particular, here we present the results of applying those metrics to the on-line results obtained with both models and those obtained using four alternative methods for the automatic detection of influenza, again using information only from past weeks. The first approach is the widely-used method proposed by Serfling [49], while the second one uses the `depmix` package [59] of R [46] to reproduce Le Strat and Carrat's [28] hidden Markov model. The third approach is a simple regression method which is a slight modification of Stroup *et al.* [55] and the last one is the model based on the differences of the rates (MBDR from now on) presented in Section 2.


It is worth noting that this validation has been done using an on-line version of all the models instead of the results obtained by applying them to the whole data set. That is, by applying all models sequentially, starting from the fourth season (in order to use at least three seasons as training data) and then predicting, for each week, the probability of being in an epidemic phase by only taking into account the information from past weeks (within the current season and all the weeks from previous seasons). The reason for this is that the on-line version reproduces the sequential arrival of data that is habitual in real situations.

Table 1 shows the values of the metrics AUWROC1, VUTROS1, VUTROS3 and VUTROCS obtained when applying the selected models (AR1-AR1 and AR2-AR2) and the four alternative methods to the North Carolina data set. As can be appreciated, results indicate that both models perform better than the other alternative methods, since the values of the four metrics are greater and closer to 1. Moreover, the AR2-AR2 behaves substantially better and so it will be our selection from now on for analyzing the North Carolina data set.

**Table 1:** Comparison of metrics AUWROC1, VUTROS1, VUTROS3 and VUTROCS when applying different early warning systems (higher being better) to the North Carolina data set.

	AUWROC1	VUTROS1	VUTROS3	VUTROCS
Serfling	0.612	0.553	0.349	0.698
depmix	0.608	0.556	0.341	0.682
Stroup	0.540	0.517	0.404	0.807
MBDR	0.601	0.544	0.356	0.713
AR1-AR1	0.676	0.595	0.412	0.824
AR2-AR2	0.726	0.649	0.420	0.840

The results of applying the on-line version of the AR2-AR2 model can be seen in Figure 2, which shows the weekly ILI incidence rates, the white, grey and black dots indicating those weeks where the posterior probability of being in an epidemic phase exceeds the values 0.25, 0.5 and 0.75, respectively. As mentioned above, these probabilities have been sequentially obtained, starting from the fourth season and only taking into account the information from past weeks (within the current season and the all weeks from previous seasons). This is the kind of graph that Health Authorities could use to raise the alarm at those precise moments in which there is a high probability of being in an epidemic phase. In particular, values exceeding 0.5 indicate that we are observing for that week a higher probability of being in an epidemic phase than of being in a non-epidemic one, and so an alarm could be triggered if considered convenient.



**Figure 2:** On-line weekly results for the North Carolina data set from seasons 4 to 8. Graphs show the influenza incidence rates in which the white, grey and black dots indicate those weeks where the posterior probability of being in an epidemic phase exceeds the values 0.25, 0.5 and 0.75 respectively.

---

## 5. FLUDETWEB

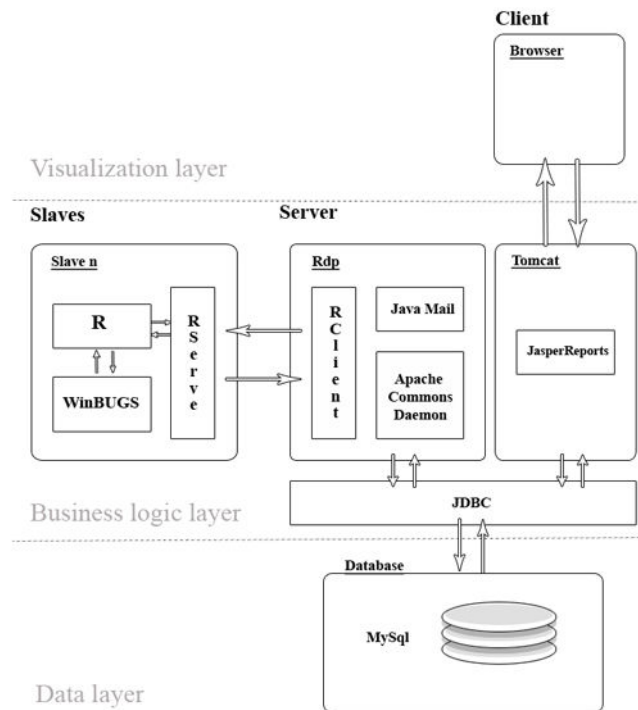
---

The complexity of disease surveillance methods has been progressively increasing. In fact, most of the methods mentioned in the Introduction are not easy to implement. On the contrary, most of them and, in general, the most advanced surveillance systems require skilled personnel to implement, fine-tune and maintain them. These requirements have kept these new developments from common usage. In order to resolve this issue, there has been a recent interest in enhancing existing disease surveillance methodologies by using tools for presenting data and information to users. Hauenstein *et al.* [23] describe in detail the processes and tools (such as system architecture, web-based applications, etc.) needed to do so.

In this section, we review *fludetweb* (Conesa *et al.* [6]), an enhanced web implementation of the MBDR, the surveillance methodology described in Section 2. This implementation is available on-line at: <http://www.geeitema.org/meviepi/fludetweb/>.

Fludetweb's implementation has been done using a thin client application design for ease of user interaction with the program, that is through a web application that could be accessed by any network-enabled device (PC's, tablets, cell phones, etc.) with a web browser. But moreover, the computational requirements of the detection algorithm, which could need several minutes to return the results, also motivated the use of a master-slave intranet architecture to take advantage of other secondary available computers.

Figure 3 shows the internal architecture of the server and its connections with the slaves and clients. The system has been implemented as a three-tier architecture by separating its functions into three separate layers. The top tier corresponds to the presentation layer and is responsible for interaction between the user and the system through data and personal information querying, visualization of results, etc. The second tier is the business logic tier, which is the core of the system as it controls the running of the influenza surveillance algorithm. The final layer is the data tier and it is responsible for data storage, not only of the influenza rates but also of the user's personal information, the availability and state of slaves, IP addresses, assigned tasks, etc.



**Figure 3:** Internal architecture of the fludetweb implementation, including its internal connections with slaves and clients.

In practice, users can introduce and edit their own data consisting of a series of weekly influenza incidence rates. Users may also obtain estimates of

the probability of being in an epidemic phase for the weeks of interest. The estimation process is not immediate, so the system has been designed to respond to requests from a multi-user environment on a first-come, first-served basis. After completion of the process, the system returns the probability of being in an epidemic phase jointly with a forecast of the probability of an increase in the incidence rate during the following week. `Fludetweb` also provides two further graphs. The first one shows the weekly rates of the last two seasons, indicating whether the posterior probability of being in an epidemic phase in the analyzed week is greater than 0.5 or not. The second graph shows all the weekly rates with flags only for requested weeks. In particular, flags indicate whether the posterior probability of being in an epidemic phase is greater than 0.5 or not. Its ease of use and on-line availability should make `fludetweb` a valuable tool for public health practitioners.

---

## 6. FURTHER LINES OF DEVELOPMENT

---

Our interest in this paper has been to review two possible methodologies for detecting influenza epidemics at the very moment of their onset. Both modelings can be used for any kind of surveillance data that show either epidemic periods, describing an increase and a decrease of the epidemic curve, or non-epidemic periods, where the data follows a random noise process. Both methodologies provide the posterior probability of being in an epidemic phase, and so they can be very useful for health authorities who could use them to raise the alarm at those precise moments in which there is a high probability of being in, or even better starting, an epidemic phase. We have also reviewed a web-based implementation of the first methodology for obtaining the posterior probability of being in an epidemic phase.

With respect to possible extensions of these proposals, a first improvement could be to include a spatial component that would help us to raise geographically-referenced alarms. Now, we would have as many time series as areal units in the region of study and the goal would be to induce spatial dependence on these series in order to take advantage of the information of neighboring series. As an example Zou *et al.* [62] and Heaton *et al.* [24] propose to use, for the epidemic phase, spatio-temporally correlated random effects, with every random effect being conditional dependent to its temporal and spatial neighbors. On the other hand, they consider either a white noise or a spatially correlated process for the non-epidemic phase. Finally, the transition probability between states for every region is considered to depend on the number of neighbors in epidemic state. Knorr-Held and Richardson [26], Watkins *et al.* [60], Li *et al.* [29] or Li and Cardie [30] are just some other examples of studies addressing the incorporation of a spatial component in outbreak detection studies.



Spatio-temporal models improve traditional temporal outbreak detection models by using a second source of information, that provided by neighboring regions. A second alternative would be to incorporate complementary information sources, not necessarily linked to other geographical locations. Thus, Nunes *et al.* [43] consider a bivariate information source (for every week) instead of the traditional univariate rates for distinguishing between epidemic and non-epidemic periods. Thus, on one hand, they use for each week the number of laboratory-confirmed cases of the previous week and, on the other hand, the number of reported ILI cases of the current week. That is, we have two information sources in this system one of higher quality (confirmed cases in the previous week) and a second one corresponding to more recent data. This kind of analysis would be more sensible than traditional univariate studies as they are based on a larger amount of information in order to decide the current epidemic/non-epidemic state.

A particular case of multivariate data use is that of Twitter. Twitter is a microblogging social network in which users publicly post short messages of less than 140 characters that may also be geolocated. Several words or sets of words related to influenza can be used to predict the onset, spread and decay of the epidemic. Works like those of Li and Cardie [30] or Grover and Aujla [22] deal with the necessary preprocessing of the raw data and the use of Markov chains to model several phases like rising, stationarity and declining of the epidemic.

Flexibility is one of the main advantages of Bayesian hierarchical models. These models can be easily adapted to any specific feature of any dataset, what has made them particularly common in outbreak detection studies. Nevertheless, the main drawback of Bayesian hierarchical models is that they usually resort to (frequently slow) MCMC simulation to carry inference out. Although WinBUGS usually makes relatively easy the inference process, once a new weekly observation arrives, we are forced to repeat the whole simulation with the new (and all the previous) observation(s). This makes simulations to become progressively slower as new observations are incorporated into the system. Sequential MCMC methods, such as particle filters, would be a solution to this, avoiding to run again the model each week with all historical and, obviously, the new data. The incorporation of this kind of inference tools to influenza outbreak detection problems could be very advantageous for a problem where the computing time is a limitation if results want to be used for on-line epidemiological surveillance.

Finally, we would like to mention a general drawback of this kind of models. Model selection is a delicate issue in outbreak detection studies. Most model selection criteria are based on fit or predictive properties of models. Nevertheless, in our particular case we would be mostly interested in some other particular criteria such as sensitivity, specificity or timeliness in the detection of outbreaks. As mentioned in this paper, several measures have been proposed in the literature (Cowling *et al.* [12] and Kleinman and Abrams [25]) paying specific attention to these aspects and, therefore, particularly suited for outbreak detection.

Regretfully these methods depend on the availability of a gold standard, that make it possible to assess the goodness of any particular method. Acceptable gold standards are not generally available excepting maybe for very few publicly accessible datasets. This makes model selection in this area a particularly cumbersome issue where more research (or more publicly accessible datasets for this purpose) would be very welcome.

---

## ACKNOWLEDGMENTS

---

This paper constitutes part of the work developed by the authors during their visit as Research Fellows to the Statistical and Applied Mathematical Sciences Institute in North Carolina in the working group “Fundamentals of Spatial Modeling” of the program “Space-time analysis for environmental mapping, epidemiology and climate change”. 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 Educación y Ciencia (the Spanish Ministry of Education and Science) via research grants MTM2010-19528, MTM2013-42323-P (jointly financed with the European Regional Development Fund) and FUT-C2-0047 (as part of the INGENIO-MATHEMATICA research project) and from the Generalitat Valenciana via the research grant EVES-015/2008.

---

## REFERENCES

---

- [1] BOYLE, J.R.; SPARKS, R.S.; KEIJZERS, G.B.; CRILLY, J.L.; LIND, J.F. and RYAN, L.M. (2011). Prediction and surveillance of influenza epidemics, *The Medical Journal of Australia*, **194**, 4, S28–33.
- [2] BRONIATOWSKI, D.A.; PAUL, M.J. and DREDZE, M. (2013). National and local influenza surveillance through Twitter: an analysis of the 2012-2013 influenza epidemic, *PloS ONE*, **8**, 12, e83672.
- [3] BURKOM, H.S. (2007). *Alerting algorithms for biosurveillance*. In “Disease Surveillance, a Public Health Informatics Approach” (J.S. Lombardo and D.L. Buckeridge, Eds.), John Wiley & Sons, Ltd, 143–192.
- [4] CHARLAND, K.M.L.; BUCKERIDGE, D.L.; STURTEVANT, J.L.; MELTON, F.; REIS, B.Y.; MANDL, K.D. and BROWNSTEIN, J.S. (2009). Effect of environmental factors on the spatio-temporal patterns of influenza spread, *Epidemiology and Infection*, **137**, 10, 1377–1387.

- [5] CHENG, C.K.Y.; LAU, E.H.Y.; IP, D.K.M.; YEUNG, A.S.Y.; HO, L.M. and COWLING, B.J. (2009). A profile of the online dissemination of national influenza surveillance data, *BMC Public Health*, **9**, 339.
- [6] CONESA, D.; LÓPEZ-QUÍLEZ, A.; MARTINEZ-BENEITO, M.A.; MIRALLES, M.T. and VERDEJO, F. (2009). FluDetWeb: an interactive web-based system for the early detection of the onset of influenza epidemics, *BMC Medical Informatics and Decision Making*, **9**, 36.
- [7] CONESA, D.; MARTINEZ-BENEITO, M. A.; AMORÓS, R. and LÓPEZ-QUÍLEZ, A. (2015). Bayesian hierarchical Poisson models with a hidden Markov structure for the detection of influenza epidemic outbreaks, *Statistical Methods in Medical Research*, online first as doi: 10.1177/0962280211414853.
- [8] CONGDON, P. (2005). *Bayesian Models for Categorical Data*, John Wiley & Sons, Inc.
- [9] COOPER, G.F.; DASH, D.H.; LEVANDER, J. D.; WONG, W.K.; HOGAN, W.R. and WAGNER, M.M. (2004). *Bayesian biosurveillance of disease outbreaks*. In “UAI’04: Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence” (J.S. Lombardo and D.L. Buckeridge, Eds.), Arlington, Virginia, AUAI Press, 94–103.
- [10] CORBERÁN-VALLET, A. and LAWSON, A.B. (2014). Prospective analysis of infectious disease surveillance data using syndromic information, *Statistical Methods in Medical Research*, **23**, 6, 572–590.
- [11] COSTAGLIOLA, D.; FLAHAULT, A.; GALINEC, D.; GARNERIN, P.; MENARES, J.; VALLERON, A.J. and GAMENN, P. (1991). A routine tool for detection and assessment of epidemics of influenza-like syndromes in France, *American Journal of Public Health*, **81**, 11, 97–99.
- [12] COWLING, B.J.; WONG, I.O.L.; HO, L.M.; RILEY, S. and LEUNG, G.M. (2006). Methods for monitoring influenza surveillance data, *International Journal of Epidemiology*, **35**, 5, 1314–1321.
- [13] CRAWFORD, G.B.; MCKELVEY, S.; CROOKS, J.; SISK, K.; RUSSO, K. and CHAN, J. (2011). Influenza and school-based influenza-like illness surveillance: a pilot initiative in Maryland, *Public Health Reports*, **126**, 4, 591–596.
- [14] DUKIC, V.; LOPES, H.F. and POLSON, N.G. (2012). Tracking epidemics with Google Flu Trends data and a state-space SEIR model, *Journal of the American Statistical Association*, **107**, 500, 1410–1426.
- [15] FARRINGTON, C.P.; ANDREWS, N.J.; BEALE, A.D. and CATCHPOLE, M.A. (1996). A statistical algorithm for the early detection of outbreaks of infectious disease, *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **159**, 3, 547–563.
- [16] FLEMING, D.M.; ZAMBON, M.; BARTELD, A.I. and DE JONG, J.C. (1999). The duration and magnitude of influenza epidemics: a study of surveillance data from sentinel general practices in England, Wales and the Netherlands, *European Journal of Epidemiology*, **15**, 5, 467–473.
- [17] FRISÉN, M.; ANDERSSON, E. and SCHIÖLER, L. (2009). Robust outbreak surveillance of epidemics in Sweden, *Statistics in Medicine*, **28**, 3, 476–493.
- [18] GESUALDO, F.; STILO, G.; AGRICOLA, E.; GONFIANTINI, M.V.; PANDOLFI, E.; VELARDI, P. and TOZZI, A.E. (2013). Influenza-like illness surveillance on Twitter through automated learning of naïve language, *PloS ONE*, **8**, 12, e82489.

- [19] GINSBERG, J.; MOHEBBI, M.H.; PATEL, R.S.; BRAMMER, L.; SMOLINSKI, M.S. and BRILLIANT, L. (2009). Detecting influenza epidemics using search engine query data, *Nature*, **457**, 1012–1014.
- [20] GODDARD, N.L.; KYNCL, J. and WATSON, J.M. (2003). Appropriateness of thresholds currently used to describe influenza activity in England, *Communicable Disease and Public Health / PHLIS*, **6**, 3, 238–245.
- [21] GRIFFIN, B.A.; JAIN, A.K.; DAVIES-COLE, J.; GLYMPH, C.; LUM, G.; WASHINGTON, S.C. and STOTO, M.A. (2009). Early detection of influenza outbreaks using the DC Department of Health’s syndromic surveillance system, *BMC Public Health*, **9**, 483.
- [22] GROVER, S. and AUJLA, G.S. (2014). Prediction model for influenza epidemic based on Twitter data, *International Journal of Advanced Research in Computer and Communication Engineering*, **3**, 7, 7541–7545.
- [23] HAUSTEIN, L.; WOJCIK, R.; LOSCHEN, W.; ASHAR, R.; SNIEGOSKI, C. and TABERNEIRO, N. (2007). *Putting it together: the biosurveillance information system*. In “Disease Surveillance: A Public Health Informatics Approach” (J.S. Lombardo and D.L. Buckeridge, Eds.), John Wiley & Sons, Ltd, 193–261.
- [24] HEATON, M.J.; BANKS, D.L.; ZOU, J.; KARR, A.F.; DATTA, G.; LYNCH, J. and VERA, F. (2012). A spatio-temporal absorbing state model for disease and syndromic surveillance, *Statistics in Medicine*, **31**, 19, 2123–2136.
- [25] KLEINMAN, K.P. and ABRAMS, A.M. (2006). Assessing surveillance using sensitivity, specificity and timeliness, *Statistical Methods in Medical Research*, **15**, 5, 445–464.
- [26] KNORR-HELD, L. and RICHARDSON, S. (2003). A hierarchical model for space-time surveillance data on meningococcal disease incidence, *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **52**, 2, 169–183.
- [27] LE STRAT, Y. (2005). *Overview of temporal surveillance*. In “Spatial and Syndromic Surveillance for Public Health” (A.B. Lawson and K. Kleinman, Eds.), John Wiley & Sons, Ltd, 13–29.
- [28] LE STRAT, Y. and CARRAT, F. (1999). Monitoring epidemiologic surveillance data using hidden Markov models, *Statistics in Medicine*, **18**, 24, 3463–3478.
- [29] LI, G.; BEST, N.; HANSELL, A.L.; AHMED, I. and RICHARDSON, S. (2012). BaySTDetect: detecting unusual temporal patterns in small area data via Bayesian model choice, *Biostatistics*, **13**, 4, 695–710.
- [30] LI, J. and CARDIE, C. (2013). Early Stage Influenza Detection from Twitter. Working paper.
- [31] LIPSITCH, M.; RILEY, S.; CAUCHEMEZ, S.; GHANI, A.C. and FERGUSON, N.M. (2009). Managing and reducing uncertainty in an emerging influenza pandemic, *The New England Journal of Medicine*, **361**, 2, 112–115.
- [32] MADIGAN, D. (2011). *Bayesian data mining for health surveillance*. In “Spatial and Syndromic Surveillance for Public Health” (A.B. Lawson and K. Kleinman, Eds.), John Wiley & Sons, Ltd, 203–221.
- [33] MAGRUDER, S.F. (2003). Evaluation of over-the-counter pharmaceutical sales as a possible early warning indicator of human disease, *Johns Hopkins Applied Physics Laboratory Technical Digest*, **24**, 4, 349–353.

- [34] MARTINEZ-BENEITO, M.A.; CONESA, D.; LÓPEZ-QUÍLEZ, A. and LÓPEZ-MASIDE, A. (2008). Bayesian Markov switching models for the early detection of influenza epidemics, *Statistics in Medicine*, **27**, 22, 4455–4468.
- [35] MILINOVICH, G.J.; WILLIAMS, G.M.; CLEMENTS, A.C.A. and HU, W. (2014). Internet-based surveillance systems for monitoring emerging infectious diseases, *The Lancet Infectious diseases*, **14**, 2, 160–168.
- [36] MONTO, A.S. (1999). Individual and community impact of influenza, *Pharmacoeconomics*, **16**, Suppl 1, 1–6.
- [37] MUGGLIN, A.S.; CRESSIE, N. and GEMMELL, I. (2002). Hierarchical statistical modelling of influenza epidemic dynamics in space and time, *Statistics in Medicine*, **21**, 18, 2703–2721.
- [38] MUSCATELLO, D.J.; MORTON, P.M.; EVANS, I. and GILMOUR, R. (2008). Prospective surveillance of excess mortality due to influenza in New South Wales: feasibility and statistical approach, *Communicable Diseases Intelligence*, **32**, 4, 435–442.
- [39] NEILL, D.B. and COOPER, G.F. (2010). A multivariate Bayesian scan statistic for early event detection and characterization, *Machine Learning*, **79**, 3, 261–282.
- [40] NIEMI, J.B.; SMITH, M. and BANKS, D.L. (2008). *Test power for drug abuse surveillance*. In “Bio-surveillance and Biosecurity”, Springer, Ltd, 131–142.
- [41] NORTH CAROLINA INFLUENZA SENTINEL SURVEILLANCE PROGRAM. (2014). Communicable disease: North Carolina Influenza update. Available in <http://www.flu.nc.gov/data/>. Data of access: 2014 October 22.
- [42] NUÑO, M. and PAGANO, M. (2007). *A model for characterizing annual flu cases*. In “Intelligence and Security Informatics: Biosurveillance”, Springer Berlin Heidelberg, 37–46.
- [43] NUNES, B.; NATÁRIO, I. and CARVALHO, M.L. (2013). Nowcasting influenza epidemics using non-homogeneous hidden Markov models, *Statistics in Medicine*, **32**, 15, 2643–2660.
- [44] ORTIZ, J.R.; SOTOMAYOR, V.; UEZ, O.C.; OLIVA, O.; BETTELS, D.; MCCARRON, M.; BRESEE, J.S. and MOUNTS, A.W. (2009). Strategy to enhance influenza surveillance worldwide, *Emerging Infectious Diseases*, **15**, 8, 1271–1278.
- [45] PATWARDHAN, A. and BILKOVSKI, R. (2012). Comparison: Flu prescription sales data from a retail pharmacy in the US with Google Flu trends and US ILINet (CDC) data as flu activity indicator, *PloS ONE*, **7**, 8, e43611.
- [46] R CORE TEAM. (2014). *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- [47] RATH, T.M.; CARRERAS, M. and SEBASTIANI, P. (2003). *Automated detection of influenza epidemics with hidden Markov models*. In “Advances in Intelligent Data Analysis V” (M.R. Berthold, H.-J. Lenz, E. Bradley, R. Kruse, and C. Borgelt, Eds.), Springer, Berlin, 521–532.
- [48] SEBASTIANI, P.; MANDL, K.D.; SZOLOVITS, P.; KOHANE, I.S. and RAMONI, M.F. (2006). A Bayesian dynamic model for influenza surveillance, *Statistics in Medicine*, **25**, 11, 1803–1825.
- [49] SERFLIN, R.E. (1963). Methods for current statistical analysis of excess pneumonia-influenza deaths, *Public Health Reports*, **78**, 6, 494–506.

- [50] SHEWHART, W.A. (1931). *Economic Control of Quality of Manufactured Product*, Princeton: Van Nostrand Reinhold.
- [51] SHMUELI, G. (2005). Wavelet-based monitoring for modern biosurveillance. Technical report, Robert H. Smith School of Business, University of Maryland, College Park.
- [52] SIMONSEN, L.; CLARKE, M.J.; WILLIAMSON, G.D.; STROUP, D.F.; ARDEN, N.H. and SCHONBERGER, L.B. (1997). The impact of influenza epidemics on mortality: introducing a severity index, *American Journal of Public Health*, **87**, 12, 1944–1950.
- [53] SPIEGELHALTER, D.J.; BEST, N.G.; CARLIN, B.P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**, 4, 583–639.
- [54] SPIEGELHALTER, D.J.; THOMAS, A.; BEST, N. and LUNN, D.J. (2003). *WinBUGS User Manual, Version 1.4.*, MRC Biostatistics Unit.
- [55] STROUP, D.F.; WILLIAMSON, G.D.; HERNDON, J.L. and KARON, J.M. (1989). Detection of aberrations in the occurrence of notifiable diseases surveillance data, *Statistics in Medicine*, **8**, 3, 323–332.
- [56] SUN, W. and CAI, T.T. (2009). Large-scale multiple testing under dependence, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **71**, 2, 393–424.
- [57] UNKEL, S.; FARRINGTON, C.P.; GARTHWAITE, P.H.; ROBERTSON, C. and ANDREWS, N. (2012). Statistical methods for the prospective detection of infectious disease outbreaks: a review, *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **175**, 1, 49–82.
- [58] VANDENDIJK, Y.; FAES, C. and HENS, N. (2013). Eight years of the Great Influenza Survey to monitor influenza-like illness in Flanders, *PloS ONE*, **8**, 5, e64156.
- [59] VISSER, I. (2007). depmix: An R-package for fitting mixture models on mixed multivariate data with Markov dependencies. R-package manual.
- [60] WATKINS, R.E.; EAGLESON, S.; VEENENDAAL, B.; WRIGHT, G. and PLANT, A.J. (2009). Disease surveillance using a hidden Markov model, *BMC Medical Informatics and Decision Making*, **9**, 39.
- [61] ZHOU, H. and LAWSON, A.B. (2008). EWMA smoothing and Bayesian spatial modeling for health surveillance, *Statistics in Medicine*, **27**, 28, 5907–5928.
- [62] ZOU, J.; KARR, A.F.; BANKS, D.; HEATON, M.J.; DATTA, G.; LYNCH, J. and VERA, F. (2012). Bayesian methodology for the analysis of spatial-temporal surveillance data, *Statistical Analysis and Data Mining*, **5**, 3, 194–204.