Bayesian approaches to epidemiological surveillance: a review and introduction for risk-assessors and decision-makers.

Vial F.[1*], Hedell R.[2], Hopp P.[3], Dórea, F.C.[4], Leblond A.[5], Gethmann J.[6], Andersson M.G.[7]

[1] Veterinary Public Health Institute, DCR-VPH, Vetsuisse Fakultät, Bern, Switzerland. Current affiliation: Animal and Plant Health Agency, Sand Hutton, UK

[2] Swedish National Forensic Centre, SE-581 94 Linköping, Sweden.

[3]Norwegian Veterinary Institute, Ås, Norway.

[4]Department of Disease Control and Epidemiology, National Veterinary Institute, Uppsala, Sweden.

[5] INRAE, VetAgro Sup, UMR EPIA, Marcy L'Etoile, France.

[6] Friedrich-Loeffler-Institut, Federal Research Institute for Animal Health, Institute of Epidemiology, Südufer 10, 17493 Greifswald – Insel Riems, Germany.

[7] Department of Chemistry, Environment and Feed Hygiene, National Veterinary Institute, Uppsala, Sweden.

[*] Corresponding author:

Email: flavie.vial@apha.gov.uk

Abstract

Quality decision making in public health and animal health surveillance relies on addressing the challenge of synthesizing health-related information from disparate sources into actionable information. In the case of early warning systems for impending outbreaks this challenge is compounded with the need for evidence generation in real-time, and timely decision-making. The analysts running and interpreting the output from the epidemiological surveillance algorithms must present those in a format that is appropriate to those who have responsibility for taking action. We argue that the Bayesian inference framework, which provides a posterior probability for a given disease state, can be easily combined with a decision theory framework to support decision-making for disease surveillance and control in a transparent way. We provide a simple introduction to Bayesian approaches to epidemiological surveillance, with a particular focus on syndromic surveillance (SyS), that covers:(i) full Bayes (hierarchical) models; (ii) empirical Bayes models; and (iii) semi-Bayes models that use Bayesian approaches to estimate model parameter distributions but that produce an output not intended for Bayesian inference. We illustrate the flexibility and robustness of applying Bayesian probabilistic reasoning with three working examples based on animal SyS data from France and Norway.

In more complex SyS scenarios, the main drawback of applying full Bayesian methods resides in the challenge of setting prior probabilities and the demanding computations, which may necessitate the use of approximate solutions. As an alternative approach, a framework for communicating SyS results based on the Bayes factor, i.e. the ratio between the posterior and prior odds that an outbreak is ongoing against an alternative hypothesis, is presented. Such explicit separation of prior information about a hypothesis and evidence from the data makes the framework useful for presenting results even when the modelling approach is not in itself Bayesian.

Keywords: Syndromic surveillance, health surveillance, Bayesian modeling, Time-series analysis, Hidden markov model, outbreak detection

Introduction

Health surveillance provides scientific and factual evidence to risk-assessors which are essential to inform decision-making, and to motivate timely and appropriate action. Increased availability of health-related data and methodological innovations have fostered new approaches to optimise surveillance systems for early-detection. One of them is syndromic surveillance (SyS). The public health sector has initiated the use of "health-related data that precede diagnosis and signal with sufficient probability of a case or an outbreak to warrant further public health response" (Fricker and Rolka, 2006) for surveillance at the turn of this century, and has been followed in the last decade by the veterinary health sector (Dórea and Vial, 2016; Dórea et al., 2011).

Risk-assessors are faced with an unprecedented amount of health-related data being passively collected on national, regional and even individual levels (from flu-related internet searches to antibiotic sales in the pig industry). Decision-makers must take an increasing number of routine decisions based on these data: which reports should be part of a formal investigation? Should they wait another day before acting? Answering such questions is not a trivial task given that passive surveillance data are associated with a higher degree of uncertainty compared to, for example, data on notifiable diseases (Onisko et al., 2006).

Observable data evaluated under different scientific hypotheses (e.g. the null hypothesis H^0: "no outbreak of the disease is currently occurring in this region" and the alternative hypothesis H: "an outbreak of the disease is currently occurring in this region") are typically modelled through probability distributions which depend on unknown quantities called parameters. There are two main approaches to the statistical inference of parameters and of hypothesis testing. These two classes are known as the *frequentist* and the *Bayesian* approach. Both classes have existed for centuries but in practice, frequentist methods have dominated, in large part due to the fact that they include a number of statistical tests that allows calculations to be performed by hand or using pre-calculated tables. Frequentist approaches assume that the data are a repeatable random sample (i.e. they can be associated with frequencies) from an infinite sampling scheme. The underlying parameters are treated as fixed at some unknown value that remains constant during this repeatable sampling process. In the Bayesian paradigm, the data are treated as observations from a realised sample (i.e. fixed) and parameters are described probabilistically, reflecting the uncertainty about their true value. Bayesian methods are very flexible but even moderately complex models will result in integrals that can only be solved by numerical methods such as Markov chain Monte Carlo (MCMC) methods. Consequently Bayesian methods did not gain popularity until computational power become readily available (Madigan, 2005).

Ideally, a decision should be made which maximises the expected benefit based on decision theory. We argue that Bayesian methods applied to health surveillance problems provide an output that is better apt to support decision-making than the corresponding frequentist approaches. While outputs may support decision making with transparency, the terminology and technical aspects of the Bayesian inference network may be daunting and difficult to grasp, hampering communication about the model approaches used in the surveillance system and the interpretation of statistical outputs. Decision-makers may, as a result, not fully understand what inference can be drawn from these outputs. The objectives of this paper are three-fold: 1) to provide a theoretical yet simple introduction to Bayesian methods commonly applied in health surveillance; 2) to discuss how Bayesian (inference) framework can be used as a general approach for presenting results from SyS to decision-makers; and 3) to illustrate, through three working examples, how a Bayesian framework can be applied to outbreak detection scenarios .

Theoretical overview

A literature search aimed at identifying Bayesian methods already used in the field of animal or public health surveillance was performed in Scopus using the following search string:

TITLE-ABS-KEY ((bayes AND surveillance) OR (bayes* AND syndromic))*

It was outside the scope of this work to provide a systematic literature review. The goal of this scoping exercise was to identify key references to build the theoretical framework presented, and find relevant examples to illustrate the use of this framework in practice. Title and abstract were screened for all 3225 resulting documents, and selected abstracts subjected to full-text evaluation. All reports identified which specifically report the use of Bayesian methods for outbreak detection were included and are cited in the relevant sections below.

A closer look at the articles retrieved from the first search revealed that the terminology used in literature is not consistent and that a search with fixed keywords would miss significant pieces of work. Thus, the initial search was followed up by a "snowball approach" in which we looked up the original works cited in books, reviews and research papers.

Working examples

To illustrate the different approaches to probabilistic reasoning in a Bayesian framework, three surveillance working examples were constructed:

1) Syndromic data on French horses presenting nervous symptoms and respiratory symptoms are evaluated each week with an empirical Bayesian network to detect incursions of West Nile virus (as in (Mats Gunnar Andersson et al., 2014)).

2) A dynamic empirical Bayesian network is applied to the same data for change point analysis.

3) A full Bayes approach to spatio-temporal SyS for bluetongue using on-farm mortality and late abortions data in Norwegian cattle. More details about this example are given in the Supplementary Material.

Although the examples we provide are based on animal health surveillance scenarios, the concepts they illustrate are very much transferable to public health surveillance systems.

An introduction to Bayesian inference

Important concepts in Bayesian statistics are the prior probability, which is the probability assigned to a hypothesis or event before the data were observed; and the posterior probability which is the probability for the same hypothesis given the prior probability and the relevant data (Bernardo and Smith, 1994; Christensen et al., 2011). The central part of Bayesian statistics is the Bayes' theorem (equation 1), which allows us to calculate the posterior probability of a hypothesis of interest H (e.g. disease present) given data regarding a chosen indicator event E (e.g. number of reported clinical cases):

$$P(H|E) = \frac{P(E|H)\,P(H)}{P(E)} = \frac{P(E|H)\,P(H)}{P(E|H)\,P(H) + P\left(E\middle|H^0\right)P(H^0)}, \qquad (1)$$

$P(H)$ is the prior probability of the disease within our population and $P(E|H)$ and $P(E)$ are the conditional probability of observing the symptoms in the presence of the disease, and of observing the symptoms regardless of the disease state of the population respectively. H^0 is the null hypothesis of no disease outbreak, while the alternative *H*, an ongoing outbreak of the disease.

A common form of Bayes´ theorem is the odds form (equation 2):

$$\frac{P(H|E)}{P(H^0|E)} = \frac{P(E|H)}{P(E|H^0)} \times \frac{P(H)}{P(H^0)} \qquad (2)$$

It can also be expressed as:

$$O_{post} = LR \times O_{prior} \quad (3)$$

In the form shown in equation (3) it is apparent that the ratio between posterior (O_{post}) and prior odds (O_{prior}) equals the likelihood ratio (LR) for the observed evidence under the two hypotheses. This form is extensively used when reporting results from the analysis of forensic evidence where LR is referred to as *value of evidence* (Aitken and Taroni, 2004). More generally the ratio of the posterior and the prior odds is known as the *Bayes factor*.

In a continuous sampling space, computation of $P(E|H)$ involves integration over the unknown model parameters θ:

$$P(E|H\ \ ) = \int \mathrm{P}(\mathrm{E}\,|\,\theta, \mathrm{H})\pi_y(\theta)d\theta \quad (4)$$

Where $\pi_y(\theta)$ is the probability distribution of θ based on training data or expert opinion y. Starting from a (possibly vague) prior distribution $\pi(\theta)$, an updated (posterior) distribution $\pi_y(\theta)$ given y may be found via the general form of Bayes' theorem, see e.g. (Christensen et al., 2011). In other cases, the Bayes factor may be assessed "directly" as the LR in equation (2) above, i.e. without any further averaging or integration over parameters or sub-hypotheses.

A natural extension of the Bayesian idea that the values of parameters arise from distributions is the use of models where parameters arise within hierarchies. In a SyS context, the probability that a disease case with a syndrome is observed would be a parameter of the distribution of reported cases, and when this parameter is unknown it may be modelled by a probability distribution. In Bayesian modelling, the parameters of a prior distribution are referred to as *hyperparameters* and their probability distribution referred to as *hyperprior distributions.* In the example above the hyperparameter "reporting probability" and its probability distribution have a real meaning but hyperparameters may also represent unknown statistical relationships. Using hyperprior distributions in addition to prior distributions is known as *hierarchical Bayes*. This approach is commonly used for multilevel modelling as it allows us to explicitly incorporate uncertainty from the multiple levels of the information.

For further reading about Bayesian inference see e.g. (Christensen et al., 2011).

Hierarchical Bayesian models and Bayesian Networks

Bayesian models for SyS will typically be based on multiple variables describing different stochastic events in which the output from one variable is the input for another. A chain of variables may for example describe (i) the distribution of infected animals; (ii) the number of symptomatic animals given (i); and (iii) the number of reported animals given (ii). Such models are usually referred to as hierarchical models or *Bayesian networks* (BN). Other names include Bayesian belief networks, probabilistic graphical models, or probabilistic independence networks. There is no clear difference between a hierarchical model and a BN and the latter may be seen as a graphical representation of the joint distribution of a set of variables in the model.

The structure of a BN is made up of variables (called *nodes*) which are connected probabilistically (through *arcs*). The arcs may indicate a direct, functional relationship or an observed statistical relationship where the cause for the correlation may be outside the model. If the arcs embody direct causal relationships, then the model is called a *causal BN.* Each node X_i, represents a function which takes the value of the parental nodes as input to calculate the probability for each *state* of the node (value of that variable): $P(X_i|\mathrm{parents}(X_i))$. If X_i has no parent, the prior probabilities of its states are specified. Nodes in BN must respect the *Markov condition*, i.e. a node must be independent of its non-descendants, given the state of its parents. This Markov condition allows us to factor the complete joint distribution of the variables in the model as following equation (5):

$$P(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} P(X_i|\text{parents}(X_i)) \quad (5)$$

From equation (5) it then becomes possible to derive the probability of any subset of nodes conditioned on the state of another subset of nodes.

Figure 1 illustrates how an empirical Bayesian model, where syndromic data are evaluated for each week independently, as in (Mats Gunnar Andersson et al., 2014), can be represented as a BN. We used data on French horses presenting nervous symptoms (NeurSy) and respiratory symptoms (RespSy) to detect incursions of an exotic disease, West Nile virus. The probability distributions of neurological symptoms (NeurSy) and respiratory symptoms (RespSy) would typically be estimated by dynamic regression. The numbers shown in figure 1 are hypothetical.

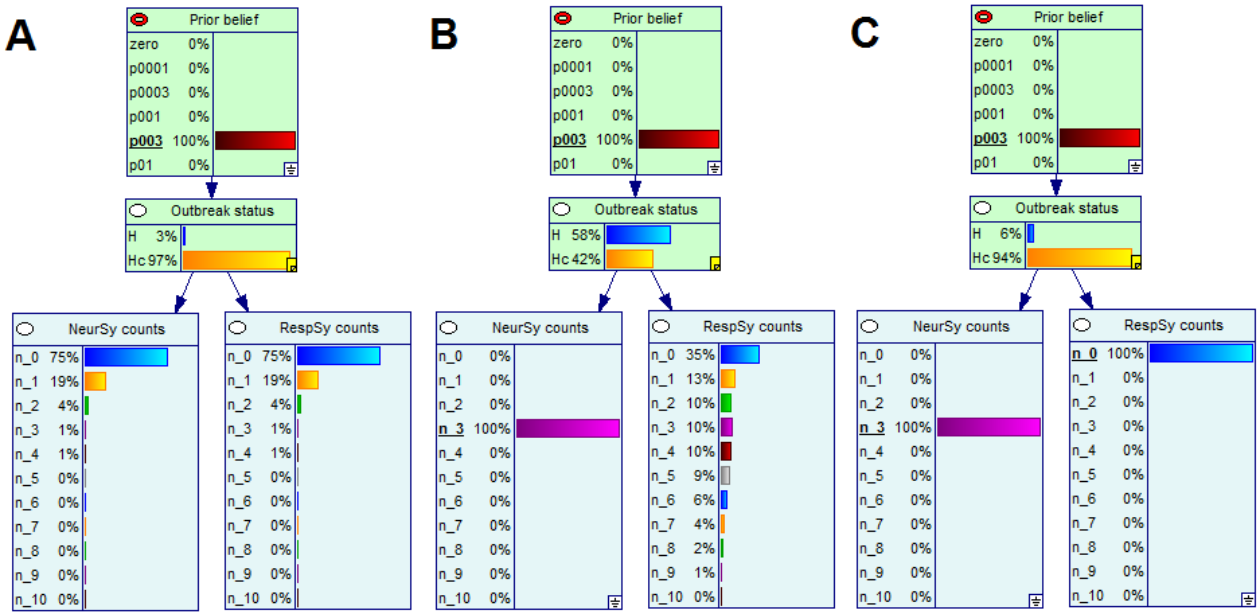

Figure 1: In this empirical Bayesian model, nodes with a red circle (e.g. "prior Belief") are decision nodes with values set by the user, whereas nodes with a white filled circle are chance nodes. When a chance node shows several bars they represent the estimated probability of the node being in that state (e.g. observing 1, 2, 3,...n syndromic cases). When the chance node has a single bar, 100%, the value of the node is known. (e.g. Fig 1b, NeurSy counts). In this case the value of the node is used to recalculate the probability of the states of the other nodes in the model.

BN are quite robust to imperfect prior knowledge and probabilities need not be exact to be useful. This is an interesting trait of BN as causal conditional properties are often easier to estimate than the reverse. For example, clinicians would find it easier to estimate $P(\text{NeuroSy}|\text{WNV})$ than $P(\text{WNV}|\text{NeuroSy})$.

Bayesian inference and decision theory

Bayes' theorem can easily be applied *a posteriori* to derive the probability of an outbreak given a statistical alarm (signal above threshold: yes/no) derived from frequentist methods (equation 6):

$$p(outbr.|alarm) = \frac{p(alarm|outbr.) * p(\text{outbr})}{p(alarm|outbr.) * p(\text{outbr.}) + p(alarm|no\ outbr.) * p(\text{no outbr.})} \quad (6)$$

Risk-assessors will usually have some idea about the prior probability of an outbreak $P(\text{outbr.})$ and the false alarm rate of the system $p(\text{alarm}|\text{no outbr.})$. Sensitivity $p(\text{alarm}|\text{outbr.})$ is much harder to quantify as it will depend on the shape and magnitude of the outbreak. However, it is possible to compute upper and lower bounds for the posterior probability $p(\text{outbr.}|\text{alarm})$ by assuming that the sensitivity is, at worst equal to the false alert rate, and at best equal to 1. In practice, the probability $p(\text{outbr.}|\text{alarm})$ would depend on whether the counts are near or very much above a pre-defined threshold (Grossi, 2008).

Within the Bayesian framework, it is not necessary to define a threshold for the signal (e.g. number of clinical cases observed). We may instead use the probability densities for the observed signal given an outbreak $f(\text{signal}|\text{outbr.})$ or no outbreak $f(\text{signal}|\text{no outbr.})$:

$$p(outbr.|signal) = \frac{f(signal|outbr.)*p\ (\text{outbr.})}{f\ (signal|outbr.)*p(\text{outbr.})+f(signal|no\ outbr.)*\ p(\text{no outbr.})} \quad (7)$$

and let the threshold be defined for a given posterior probability or expected utility of action (Mats Gunnar Andersson et al., 2014). The latter is the average amount of clinical cases, or loss of animals, that we expect to see. Since an unmanaged outbreak as well as actions will result in costs, the expected utility will always be zero or negative.

The Bayesian approach, centered on sequential inference, constitutes a transparent support to risk-assessors (Heaton et al., 2012) but in some situations, deriving the posterior probability of an outgoing outbreak may not be enough to take informed mitigation measures (e.g. vaccination, quarantine). The costs and benefits of the possible actions must be considered in a way that is adaptive, i.e. relies on the latest collected information. This is particularly important when the surveillance goal of the system is early detection and decision-makers require an understanding of the explicit trade-offs between waiting another day for more data and acting today based on the information collected so far.

Decision theory is a framework for making optimal decisions given values and uncertainties and is closely related to game theory. Since Bayesian models for SyS will present as output a posterior probability for each state, given prior knowledge and evidence, they are easily combined with a decision theory framework as discussed in (Mats Gunnar Andersson et al., 2014; Onisko et al., 2006) to design a surveillance system that is optimal, accounting for costs of actions and consequences of outbreaks. Many software for BN, including the GeNIe modeling environment ("BayesFusion, LLC," n.d.; Druzdzel, 1999), allows the user to incorporate functions for calculating utility as extra nodes in BN models. BNs with utility nodes are referred to as influence diagrams (Figure 2). Decision support models can be built in which a Bayesian model is combined with a decision theory framework to evaluate the best decision given the evidence or the added value of "waiting" for more data; or to analyse alternative ways of epidemic control under imperfect information (Lin and Ludkovski, 2014).

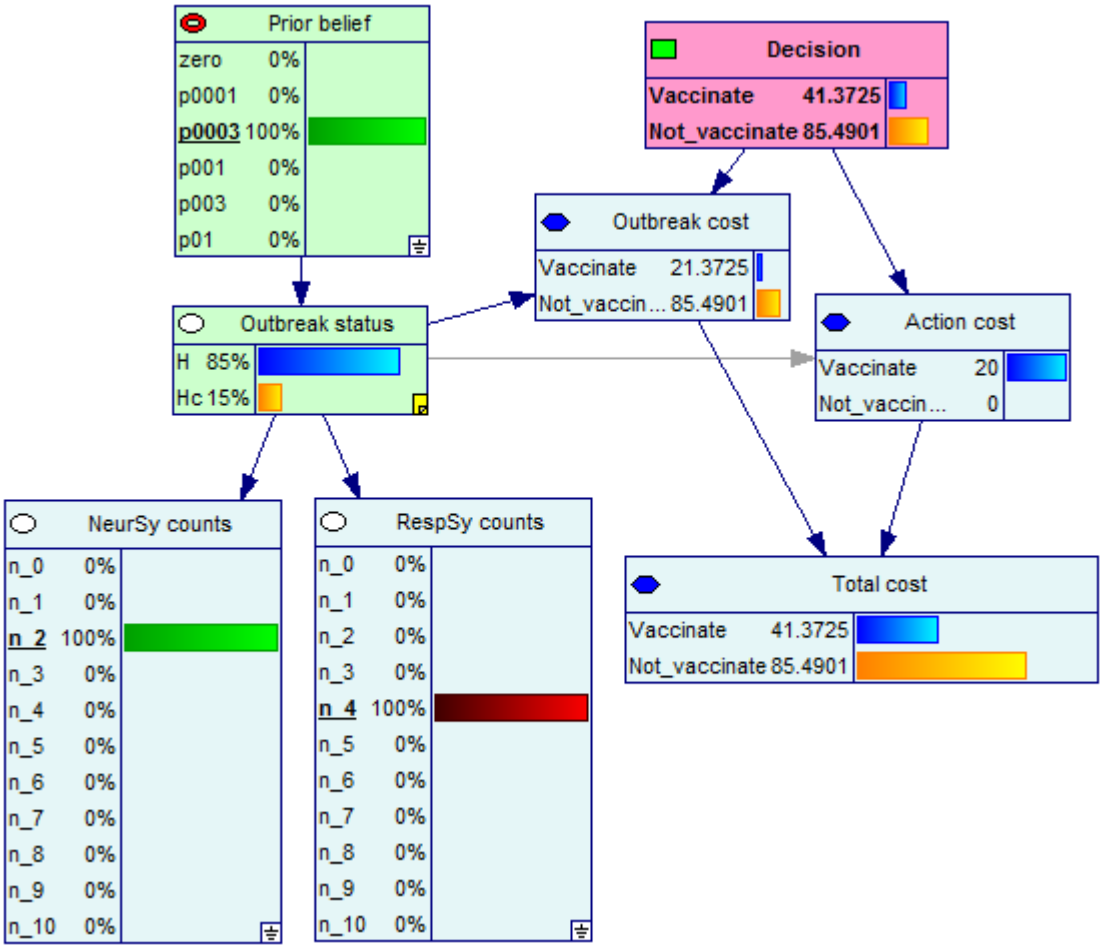

Figure 2: Example of a simple Influence Diagram where the posterior probability of an outbreak is used to decide on whether to vaccinate. The evaluation nodes (blue hexagon) takes different values for each combination of states of the parental decision and chance nodes. The values shown in the example are hypothetical.

Dynamic Bayesian Networks (DBN) and Hidden Markov Models HMM (HMM)

If a BN is used to model time-series data, i.e. the arcs flow forward in time, it is known as a *dynamic Bayesian network* (DBN) or sometimes as temporal BN or two-time slice BN. The DBN represents graphically conditional independencies (arcs) between a set of time instances (nodes) with probabilities. The simplest DBN for a sequence of observations $\{Y_1, \dots, Y_t\}$ is the first-order Markov model which only uses Y_{t-1} to derive the value of Y_t. If the observation Y_t is generated by some variable which state S_t is discrete and hidden from the observer, the DBN will form a hidden Markov model (HMM). In HMM, both the sequence of states and observations follow a first-order Markov order – that is, a given state S_t is independent of all the states prior to $t-1$; and given S_t, Y_t is independent of the states and observations at all other time indices.

A belief propagation algorithm (Pearl, 1988) is used to update the probabilities of all the nodes in the network to incorporate new evidence (i.e. new observations). The objective of HMM is to compute the optimal estimate of the hidden state (and its uncertainty) given the observed data - the posterior probability distribution (or density) $P(S_t|Y)$ - which can be derived as a recursive form of Bayes' rule. This can be computed through a forward filtering and backward smoothing approach (Scott, 2002). Parameter

estimation is then performed using methods such as Gibbs sampling (Carter and Kohn, 1994), the expectation-maximisation algorithm or Markov chain Monte Carlo (MCMC) sampling (Ryden, 2008).

Markov switching models & State-space models

Markov switching models (MSM) are an extension of HMM which include lagged observations. The observable random variables in the MSM depend on their historical values as well as the hidden state variables.

State-space models (SMM) (also known as linear Gaussian state-space models and Kalman filters) are also an extension of HMM in which the latent variable is continuous and normally distributed (as opposed to discrete and following a multinomial distribution). A good accessible introduction to HMM, DBN and their numerous extensions can be found in (Ghahramani, 2001).

Not one but several Bayesian approaches

The general term Bayesian may be applied to "any method that seeks to approximate the posterior (probability) distribution for some variable(s) or parameter(s) of interest" (definition attributed to (Bernardo and Smith, 1994)). Generally, three types of Bayesian models can be considered:

I. Models in which data are used to obtain posterior probability distributions for parameters which are then used for inference. Such hierarchical models, that include hyperparameters and hyperpriors, are referred to as ***full Bayes*** (Lawson and Kleinman, 2005). Some publications use the term "full likelihood method" for these models (Frisén and Andersson, 2009). A full Bayes approach entails formulating subjective prior probabilities to express pre-existing information; carefully modelling the data structure; checking and allowing for uncertainty in model assumptions; and possibly formulating a set of possible decisions and a utility function for the value or cost for correct and incorrect decisions.

II. Methods that seek to estimate the posterior probability of a variable of interest, but do not build full Bayes models are collectively named ***empirical Bayes*** (Lawson and Kleinman, 2005). In this case, point estimates for (some) input parameters are used (e.g. based on maximum likelihood or linear regression) rather than applying a parameter distribution.

III: Models that use Bayesian approaches to estimate model parameter distributions but that produce an output in a format not intended for Bayesian inference. These can be referred to as ***semi-Bayes***.

From a decision- maker's perspective the semi-Bayes approach (III) is indistinguishable from the traditional frequentist approach. Bayesian models may be used to estimate parameters of a regression model for time series analysis, as shown in Jung et al.(Jung et al., 2006), but the resulting model is used to compute confidence intervals for the syndromes of interest and define action thresholds just as in a traditional SyS. These were not often reported in the SyS literature and will not be discussed further.

Our literature informed discussion presented below will focus on the first two cases, full Bayes and empirical Bayes models. Within each of these classes, the models developed may differ significantly in their mathematical representation and technical implementation. In many cases, the naming of the methods and the technical description (e.g. using discrete distributions for small counts but approximating data with continuous distributions for large counts) may mask the fact that the models are, at a more fundamental level, very similar.

Bayesian framework as applied to SyS

1. One region-one time unit

The simplest situation is a surveillance system that, as in working example 1, is implemented for a region as a whole, looking at data for one time unit at a time. Observation Y (e.g. number of clinical cases for a syndrome) at time t is generated by some process (health/disease status of the population under surveillance) whose state S_t is hidden from the observer. The standard frequentist solution to this scenario is the Shewhart method which compares a daily sample statistic against an upper control limit typically set to be a multiple of the standard deviation of the mean (Shewhart, 1939). While most efficient in detecting spike-type outbreaks, the Shewhart method does not perform well when the increase in the mean of the process is slow.

I. Full Bayes

Schmidt and Pereira (Schmidt and Pereira, 2011) reviewed generalized dynamic models commonly used for modeling time series of count data, and demonstrated how the parameters of interest could be estimated using a full Bayes approach. In particular, the authors demonstrated how a Bayesian framework could be used to estimate the probability of disease (in their case dengue fever) given the absence of reported positive cases. The "one region – one time unit" is the simplest biosurveillance scenario we can consider, and most examples found in the literature using a full Bayesian approach addressed the more complex scenarios of multiple time points and/or multi-dimensional surveillance data, as presented in the next sections.

II. Empirical Bayes

Andersson et al. (Mats Gunnar Andersson et al., 2014) developed an Empirical Bayes model for syndromic surveillance of WNV and Equine Influenza using syndromic data as described in the working example 1. In this approach, where only one time unit is considered, the distributions of reported cases are modelled using linear regression models fitted by maximum likelihood. This allows calculation of ratio of the likelihood of observing *n* cases given an ongoing outbreak, over the likelihood of observing *n* cases given no outbreak. This ratio, mentioned above as the value of evidence (V) can then be multiplied by the prior odds of an ongoing outbreak to obtain the posterior odds.

The approach may be extended to handle multiple data streams by assuming that they are conditionally independent given any of the hypotheses the likelihood ratios from the two data streams can be combined by multiplication:

$$LR = \frac{P(E_1,E_2|H_1)}{P(E_1,E_2|H_2)} = \frac{P(E_1|H_1)}{P(E_1|H_2)} \times \frac{P(E_2|H_1)}{P(E_2|H_2)} \quad (8)$$

where E_1 is the evidence from syndrome 1 and E_2 is the evidence from syndrome 2. This equation is equivalent to the simple BN in working example 1.

2. One region-multiple time units

A more advanced surveillance scenario would be to look at data from one region accumulated over several time units at a time. This is exemplified by extending working example 1 into working example 2.

We use the term *change point analysis* (CPA) for methods that seek to detect subtle change(s) in incidence and to characterise the direction of the change in a time series between change points. Knowing the time at which process parameters have started to shift, the so-called change point, makes it easier to initiate a search to identify and eliminate the source of variation. Frequentist statistical process control charts are

used to detect shifts in a process parameter by distinguishing between assignable cause variation (i.e. new emergent properties) and common causes of the process variation (i.e. variation that is predictable probabilistically). They may incorporate the cumulative sums of the deviations of the sample values from the target values (CUSUM) (Lucas, 1985) or may use a weighted average of the sample statistics with exponentially decaying weights (termed exponentially weighted moving average chart or EWMA) (Holt, 2004; Winters, 1960). These methods and others are reviewed in (Unkel et al., 2012), and more recently in (Yuan et al., 2019).

Application of the Bayesian framework to the change point estimation problem allows us to draw inferences based on posterior distributions for the time and the magnitude of a change (Barry and Hartigan, 1993). Kass-Hout et al. (Kass-Hout et al., 2012), for example, applied the CUSUM CPA method to detect changes in emergency admission trends that can indicate influenza illness in USA. The authors showed that the use of CPA, in comparison to single-point detection, allowed decision-makers to make a more informed decision on which alarms warranted response, prioritizing time-series changes depending on whether they represented decreasing, stable or increasing trends. However, the authors brought attention to the inherent assumption of a normal distribution for the time-series data, an issue that Texier et al. (2016) (Texier et al., 2016) later suggested could be the reason for a more accurate and less biased performance of frequentist methods in their particular evaluation using simulated data.

I. Full Bayes

Full Bayes approaches for CPA are often based on HMM. Le Strat and Carrat (1999) (Le Strat and Carrat, 1999) pioneered the use of HMM in biosurveillance which many biostatisticians have built upon. The recursive nature of HMM allows us to easily run them in real time as only the present observation(s) and the previously estimated state and uncertainty matrix are required, i.e. no additional past data are needed. Unkel et al. (Unkel et al., 2012) provided a comprehensive review of statistical methods used for prospective detection of infectious disease outbreaks, and pointed HMM among the class of methods that can explicitly account for the correlation structure among observations in a time-series.

HMM and MSM can be used in a purely temporal setting, where the transition from non-outbreak to outbreak scenarios is seen as a Markovian process. That is, disease outbreak states are modeled as hidden state variables which control the observed time series. The process was evaluated for the detection of simulated anthrax outbreaks in a time-series or clinic visits collected from a metropolitan area (Lu et al., 2010, 2008). The authors showed higher sensitivity compared to traditional deterministic SyS methods. They also pointed out that the method was less sensitive to extreme values than traditional approaches, as a jump component could be introduced to absorb sporadic extreme values. This is expected to reduce the number of false alarms and increase robustness to variations in the data, which is the result reported by other authors applying HMM to determine the epidemic and non-epidemic periods from influenza surveillance data (Conesa et al., 2015; Martínez-Beneito et al., 2008; Rath et al., 2003). Amorós et al. (Amorós et al., 2020) later extended the temporal model presented by Martinez-Beneito et al. (2008)(Martínez-Beneito et al., 2008) to model specifically the differentiated incidence rates between equally spaced time points, improving detection in earlier stages of the epidemic, when incidence rates are low.

Influenza monitoring was also used as the test case for a framework that used Bayesian networks both to estimate the probability of an individual patient having the specific disease, based on its electronic medical records, and then the probability of an outbreak in the population (Cooper et al., 2015). The population component is based on a SEIR compartmental spread disease model (Susceptible-Exposed-Infected-Recovered). In this framework, which is unique in its integrated approach to combining patient and population outbreak detection and characterization, counts are not the only source of evidence. The population probabilities of an ongoing outbreak also become more informative with more and better information about the individual clinical cases.

Dawson et al. (2015)(Dawson et al., 2015) used a similar Bayesian network methodology, informed by individual medical records and a spread model, but improved by incorporating a particle filter. In the words of the authors, "[t]his inclusion allows the system to track the fraction of the population sick as a continuous parameters, rather than as a few coarse discrete states, which is especially important when the number of cases are small". The gain in timeliness and specificity comes at the cost of a high parametrisation burden, which can be particularly disadvantageous for detection of unknown diseases or rare diseases used as biological weapons.

Watkins et al. (Watkins et al., 2009) used reported case counts of hepatitis A, superimposed with simulated outbreaks, to evaluate the suitability of HMM as a surveillance methods for use in sparse small area count data, with limited availability of baseline data. At false alarm rates around 0.05, the Bayesian method did not outperform traditional CUSUM methods, but at 0.01 false alarm rates the HMM had both greater sensitivity and shorter timeliness for detection.

Höhle et al. (Höhle and An Der Heiden, 2014) used a hierarchical Bayesian model to account for reporting delay during an outbreak of Shiga toxin-producing *Escherichia coli*. The authors explicitly modelled the delay distribution (discrete time survival regression) in parallel to the epidemic curve, allowing for changes in the reporting delay as intervention measures are implemented, and an effective "nowcasting" of the epidemic burden in order to inform control strategies. The approach was also used for surveillance of foodborne disease in China (Wang et al., 2018), and dengue fever in Brazil (Bastos et al., 2009) and Thailand (Rotejanaprasert et al., 2020). A Bayesian nowcasting model which accounts for reporting delay while explicitly modeling the temporal relationship between cases – to accurately model reporting delay accounting for the fact that future cases are intrinsically linked to past reported cases – was recently introduced and made available as the R package {nobBS} (McGough et al., 2020). Liu et al. (Liu et al., 2020) used a Bayesian model to learn from individuals reporting behaviour in online participatory health surveillance systems, and estimating their probability of reporting every week. The model was applied to ILI prevalence estimation in Australia to demonstrate how the framework can be used to correct inconsistent participation and sampling bias in prevalence estimations.

II. Empirical Bayes

Working example 2 illustrates a simple empirical DBN (Figure 3). Such absorbing state model is only appropriate when the goal of the surveillance system is the early detection of emerging diseases for example, but it is not suitable for the monitoring of recurrent seasonal diseases such as influenza.

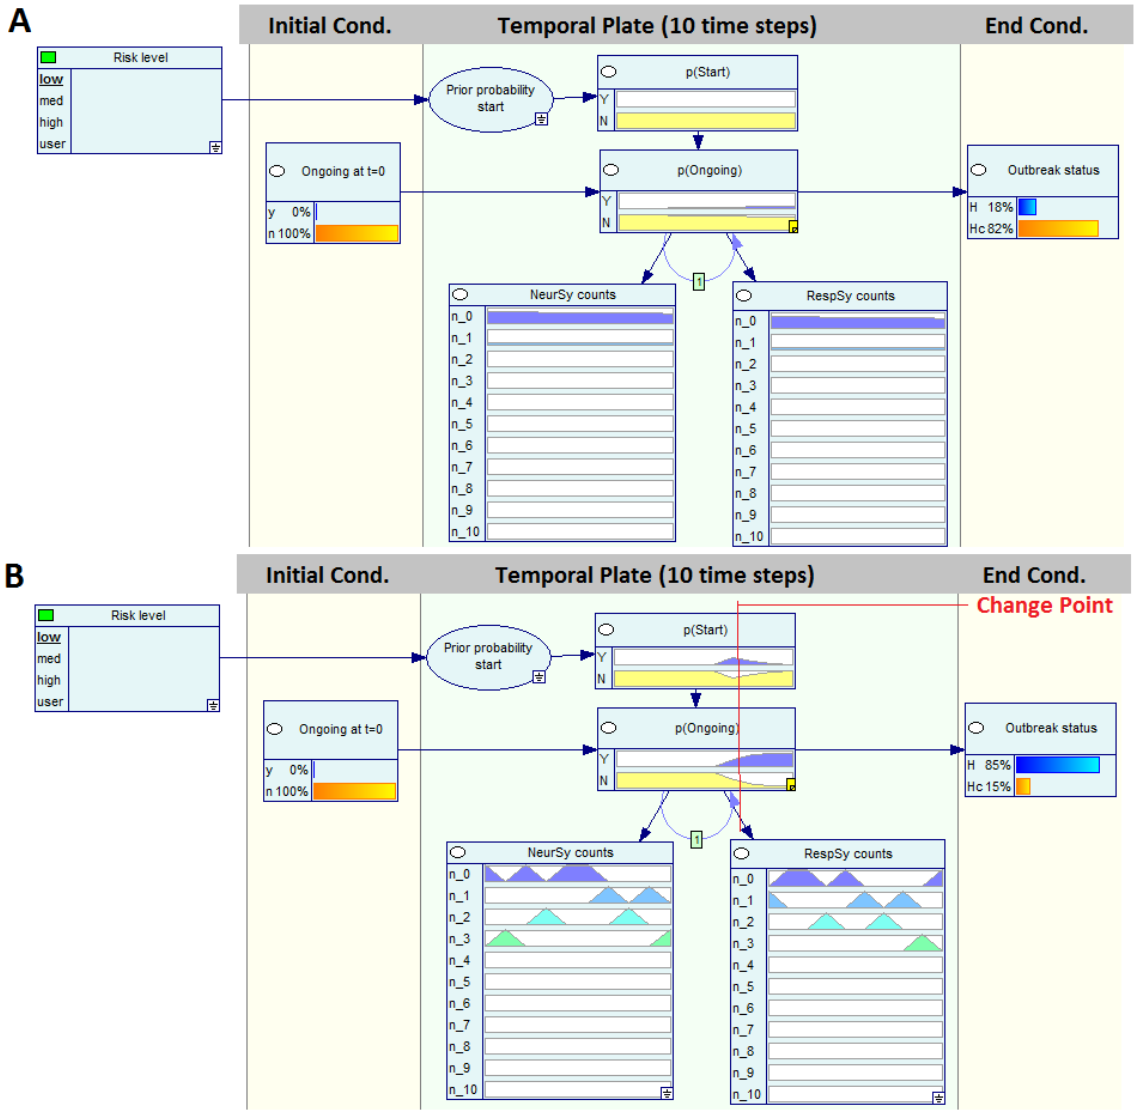

Figure 3: In this Dynamic Bayesian Network, the probability distributions for respiratory and neurological syndromes are the same as in Figure 1 and assumed to be dependent only on the state of node p(Ongoing). The nodes in the "dynamic" part of the network, however, take one value for each time-step. The node P(ongoing) takes the values "yes" or "no". The state may change from "no" to "yes" with a probability defined in node P(onset). Once the node is in state "yes" (outbreak started), it cannot switch back. Such models are referred to as "absorbing state models" (Heaton et al., 2012).

The DBN in Figure 3 is a first order Markov model since the state of each time-instance is only dependent on the state of the previous time-instance (represented by the circular arc at node "Ongoing"). A DBN may be extended to allow higher order (>1) interactions between time instances such as an n^{th} order DBN allows arcs from $\{Y_{T-n}, \dots, Y_{T-1}\}$ to Y_T.

While the approach above uses the Bayesian decision framework to weight in the likelihood of being in an epidemic versus non epidemic state based on the number of observed cases, García et al. (2015) (García et al., 2015) proposed a method based on the shape of the distribution for the number of cases, which they tested for detection of influenza-like illness (ILI). Their method is based on the rationale that the number of reported cases, which follows an autoregressive dynamic in the absence of an outbreak, will change to

exponential growth during the early phase of the outbreak. The authors noted that the method, which could be implemented in a straightforward algorithm, relies only on training based on historical data, without the need to tune free parameters. Furthermore, it allows quantitative estimations of epidemic parameters. Polyakov et al. (Polyakov and Breban, 2016) also explored the idea of a breakpoint change, but in this case the authors looked for a statistically significant change in the epidemic's basic reproduction number, R_0.

Brooks et al. (Brooks et al., 2015) proposed a framework for monitoring ILI seasonal epidemic in which the prior for an upcoming seasonal curve is calculated considering sets of transformations of past seasons' curves. As the season progresses and data becomes available, the likelihood of being in any of these curves are weighted based on actual observed data. These weights are used to generate distributions – rather than point values – for forecasts of specific epidemic targets, such as peak and duration.

Izadi et al. (Izadi et al., 2009) presented innovative work using BN to evaluate the performance of SyS algorithms. The authors evaluated algorithms both intended for the one-time unit evaluation problem, as well as options that consider multiple time-points. Their work is mentioned in this section because as for other work that we have classified as empirical Bayes, detection is based on frequentist statistical approaches, but a Bayesian framework is used to deal with uncertainty in the process, and aid decision making – in this case decision about algorithm choice an parameterization. Ebel et al. (Ebel et al., 2017) used an empirical Bayes approach to account for the variability and uncertainty associated with reporting of foodborne illness cases, and estimate the actual power of outbreak detection through surveillance.

3. Multi-dimensional surveillance data

Multivariate SyS systems, which concurrently monitor several health-related data streams, have greater sensitivity and are more reliable than univariate systems (Rolka et al., 2007). This is because no single data source captures data from all the individuals involved in an outbreak. Some diseases will cause a wide variety of clinical symptoms in different people or animals (e.g. diarrhoea in some, fever in others) and/or will affect different strata of the population (e.g. different age or production groups). Since there is generally different information contained in observations from different data sources, SyS systems should seek to simultaneously evaluate various combinations of multiple data sets using multivariate approaches – overviews are provided in (Frisén, 2010; Sonesson and Frisén, 2005).

Many information systems will also record some sort of spatial information related to a syndromic case (e.g. postal code of patient, geographic coordinates of a farm). Including this extra layer of information in the analytical methods can allow detection of localised outbreaks of a disease or identify variations in regional patterns. Spatial and spatio-temporal frequentist aberration-detection algorithms have been developed, ranging from spatial CUSUM (Dassanayaka, 2015) to space-time scan statistics (Kulldorff, 2001) and spatio-temporal regression methods (Kleinman et al., 2004). For a comprehensive review of methods for space-time disease surveillance, readers are referred to (Robertson et al., 2010).

The application of Bayesian alternatives to multi-dimensional surveillance in general, and spatial-temporal monitoring in particular, are reviewed below with examples.

I. Full Bayes

Dynamic Bayesian Networks can be used to discover the interplay among multiple data sources monitored for health surveillance. Such method was applied by Sebastiani et al. (Sebastiani et al., 2006) to jointly monitor four data sources employed for influenza surveillance. The joint model can be used to forecast the beginning of epidemics, as well as the peaks of epidemics, and in their work showed that paediatric patients were infected with respiratory viruses before the rest of the population.

Later coining this as "Bayesian Information Fusion Networks", a series of papers demonstrated the enhancement of disease surveillance systems by this method's advantages of combining multiple data

sources and providing Bayesian decision support capabilities. Mnatsakanyan et al. (2009) (Mnatsakanyan et al., 2009) described a system for detection of influenza-like events combining chief complaints from emergency department (ED) visits, International Classification of Diseases Revision 9 (ICD-9) codes from records of outpatient visits to civilian and military facilities, and influenza surveillance data from state health departments. The system showed results high sensitivity and specificity for timely detection compared to confirmed laboratory cases. Burkom et al. (Burkom et al., 2011) "fused" environmental data with public health data, including water quality data in the surveillance and early detection of waterborne disease outbreaks. The fusion networks method was later refined to improve sensitivity while reducing false alarm rate (Burkom et al., 2013). The refined method considers the inclusion of many different sources of data, which are then tested individually for their inclusion in a process of hierarchical training of the Bayesian networks. Hierarchical model selection was also used by (Ertem et al., 2018) to enable combination of multiple predictors into the previously referenced model developed in (Brooks et al., 2015).

Morrison et al. (2016) (Morrison et al., 2016) applied the hierarchical Bayesian framework to improve environmental health models, which usually focus only on monitoring health outcomes, with the explicit modeling of environmental exposure (risk factors) as a latent process. The authors implemented the computation efficiently by using integrated nested Laplace approximations, and demonstrated the superiority of the method compared to univariate models.

A causal Bayesian network to model an entire population of people (not just those seeking treatment) was introduced by Cooper et al. (Cooper et al., 2004) to monitor emergency department chief complaint data. The so called Population-wide Anomaly Detection and Assessment (PANDA) algorithm was later extended to simultaneously monitor two data sources of different granularity: aggregated regional counts for OTC sales and multivariate ED records for individual patients (W.-K. Wong et al., 2005). The latter work used the extended PANDA for detection of anthrax release, but the authors pointed out that the algorithm could be used even to model the effects of noncontagious disease outbreaks. Later, the population-wide Bayesian network model idea was applied again to emergency department data for detecting both specific and non-specific disease outbreaks (Shen and Cooper, 2009). This hybrid approach can jointly model known diseases (e.g., influenza and anthrax) by using informative prior probabilities, and unknown diseases (e.g., a new, highly contagious respiratory virus that has never been seen before) by using relatively non-informative prior probabilities.

Multi-dimensional surveillance data also emerges when the spatial dimension is explicitly taken into account. The Bayesian-network-based spatial scan statistic (BNetScan), similarly to the temporal approaches described above, is an entity-based BN that models the underlying state and observable variables for each individual in a population. This network is then used to determine the posterior probability of each sub-region containing a cluster. It has been applied to simulated outbreaks of influenza and cryptosporidiosis injected into Emergency Department data (Jiang et al., 2010).

Spatio-temporal data are also the most common use case for HMMs. The transition probabilities of the Markov chain can be allowed to vary over space and time, in the fashion of conditional autoregressive modeling (CAR) models in which spatial correlation in the disease data is modelled by a set of random effects. Several extensions to HMM (such as beyond normality (Rolka et al., 2007), beyond two hidden states, multivariate extensions and random observation time (Pearl, 1988)) are available. More recent advances go further. The spatio-temporal conditional autoregressive HMM with an absorbing state proposed in (Heaton et al., 2012) combines good sensitivity and specificity, use of covariate information, inclusion of spatio-temporal dynamics, and transparent support to decision-makers.

Working example 3 illustrates a full Bayes approach applied to a spatio-temporal SyS for bluetongue disease using two data-streams: increased mortality, and late abortions in the Norwegian cattle population. A HMM model with an absorbing state was implemented, similar to Heaton et al. (2012) (Heaton et al., 2012) but with some changes. For example, our model includes time delays between infection and resulting observed

data. An overview is given in Figure 4, and a detailed description of the model is given in the Supplementary Material. The model for the baseline rates (i.e. number of deaths and late abortions without an outbreak) was fitted to actual historical data from 2006. Data on the individual level were aggregated on a grid of 25 km x 25 km cells, with a total of 42 grid points covering the south-west region of Norway from which the data originate. A model for the bluetongue outbreak signal was fitted using simulated outbreak data (Szmaragd et al., 2009), and this signal was then added to the observed baseline cases. Another simulated outbreak of bluetongue in the same region was used for the SyS with infections status unknown for the whole county. The main parameter of interest is the infections status (infected or not infected with the disease) of each grid point for each week. The first simulated infection occurs five weeks after the start of the syndromic surveillance, but the increase of deaths and late abortions does not pick up speed until week 14. Detection is set to occur when the estimated probability of infection is above 50%, which occurs at week 15. Approximate posterior distributions of model parameters were obtained using MCMC simulation using OpenBUGS version 3.2.3 (Thomas et al., 2006). A similar framework was applied to the monitoring of nervous symptoms in horses, showing how spatio-temporal monitoring can be applied simultaneously for detection of a specific disease (in this case West Nile Virus) and a more general class of diseases related to the syndromic manifestation (Hedell et al., 2019).

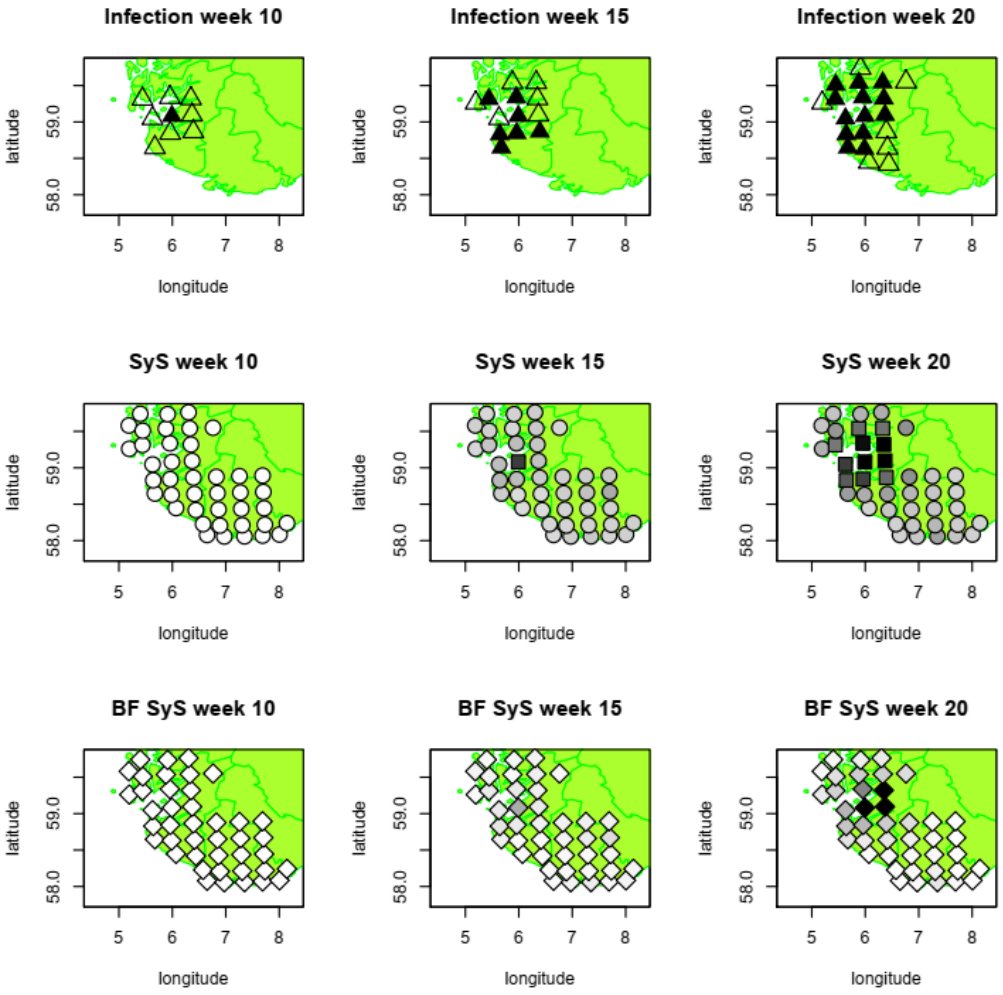

Figure 4: A full Bayes approach to spatio-temporal surveillance for bluetongue in Norwegian cattle. The first row shows the grid points infected at various times in the simulation, with open triangles indicating infection but not yet any additional deaths or late abortions. The second and third rows illustrate the ability of the SyS model to predict the infection status. The second row shows the posterior probability that each region is infected, while the third row shows the Bayes factor.

(Zamiri et al., 2015) modelled influenza spread within populations using SIR compartment models (Susceptible-Infected-Recovered). Individual sets of SIR parameters are used to model the epidemic dynamics within each area, and spatial spread and information are then explicitly modelled by adding a set of transition probabilities between every pair of monitored geographical areas. The optimal Bayesian predictor for the unknown number of ongoing epidemics is an extension of previous formulations of nonlinear multitarget filtering to account explicitly for spatial spread of disease, and be able to process multiple syndromic data streams representing the reports from multiple geographical areas. While the full Bayesian optimal solution is computationally intractable, the authors implemented an estimation algorithm based on a probability hypothesis density (PHD) filter with particle systems (known as particle-PHD or sequential Monte Carlo (SMC) PHD filter in tracking literature (Jégat et al., 2008). While the framework was shown to be useful in providing timely prediction of the epidemic peak and duration, the authors highlight that this is rather a conceptual solution, and implementation in practice requires further research.

(Zou et al., 2018) also used a hybrid hierarchical Bayesian framework in which a spatial explicit model is coupled with a particle filter, which the authors highlight allows for online updating of streaming data."

A Bayesian model that outperformed SaTScan was introduced by Li et al. (Li et al., 2012). BaySTDetect is a mixture of a component that describes the background effect of the disease in a study region as a whole, accounting for spatial and temporal autocorrelations, and a second component that estimates the time trend for each area. Boulieri et al. (Boulieri et al., 2020) later extended the model by addressing important limitations, such as allowing for the mixing parameter which designates areas as following a usual or unusual trend to vary in time. It is important to note that the framework has only been applied in the surveillance of non-communicable, non-infectious conditions. These have not been extensively explored in this article. For a recent review of spatiotemporal models for non-communicable disease surveillance readers are referred to Blangiardo et al. (Blangiardo et al., 2020)

II. Empirical Bayes

The main challenge to multi-dimensional monitoring is handling the complex correlation structure among the data sources monitored, which cannot be easily addressed without the help of Bayesian networks, as described in the full Bayes approaches listed above. While less common, empirical Bayes frameworks have also been constructed to handle the spatial dimension or the combination of multiple sources of data.

One example is the Bayesian spatial scan statistic proposed by Neil et al. (Neil et al., 2006). The (original) Kulldorff scan statistic (Kulldorff, 1997), which is based on the concept of likelihood ratios, "preludes" Bayesian models but lacks a Bayesian interpretation. Since it is possible to estimate, by simulation, the probability of obtaining a particular value of a scan spatial or spatiotemporal statistic under baseline and outbreak conditions (Lawson and Kleinman, 2005), it is also possible to apply the continuous form of Bayes Theorem to estimate the Bayes ratio between H and HA and thus make inference on the posterior probability of H.

The framework for syndromic surveillance based on the value of evidence presented in examples 1 and 2 (Struchen et al., 2017) has been extended to monitor multiple syndromes (Faverjon et al., 2016), and account for reporting delays (Struchen et al., 2017).

Manitz et al. (2013) (Manitz and Höhle, 2013) extended the widely used Farrington aberration detection based on generalized linear models (GLM) algorithm (Farrington et al., 1996; Noufaily et al., 2013), to a Bayesian generalized additive model (GAM). This extends the original algorithm by allowing adjusting for concurrent processes influencing the case counts. The authors demonstrated this by incorporating the influence of absolute humidity when modeling weekly reports of campylobacteriosis cases in Germany. Fast and efficient integrated nested Laplace approximations allowed the method to be made available through an

easily accessible implementation in the R (statistical programming environment R (R Core Team and Team, 2016)) package {surveillance} (Höhle, 2007; Salmon et al., n.d.). Salmon et al. (Salmon et al., 2015) later built further into this framework by including an adjustment to account for reporting delays. Vial et al. (F Vial et al., 2016) applied the framework to multivariate time series surveillance in animal health. While the authors highlighted challenges such as defining the expected covariance structure among series in the presence and absence of an outbreak, they also demonstrated such as the ability to include known covariates.

Discussion

As the number of health data sources grows in volume and complexity, so does the number of approaches developed to continuously monitoring these data and provide decision-makers with information to support surveillance. Making this information actionable requires being able to incorporate the outputs of the data analyses into surveillance practice in a transparent way. In this scoping review we have focused specifically on the surveillance goal of early disease detection, and reviewed the use of Bayesian frameworks in syndromic surveillance systems.

We set claim already in the introduction that Bayesian inference can support decision-making in transparent ways by providing a posterior probability of being in specific disease occurrence states, and doing so by incorporating many sources of data, as well as prior knowledge of disease dynamics. As Moss et al. (2016) (Moss et al., 2016) pointed out, the ongoing challenge of integrating data analytics into surveillance practice "requires close collaboration between modellers, epidemiologists, and public health staff". Many of the challenges lifted by the authors – such as underestimation of the true number of cases and inherent biases in reporting-based data – were addressed with Bayesian frameworks in the examples presented and discussed here. These examples further highlighted the flexibility and robustness of the Bayesian framework, while also pointing our challenges such as complex implementation.

The Bayesian approach can easily incorporate different sources of data (W. Wong et al., 2005), while accounting for both the uncertainties in the estimations, and the stochasticity of the model (Salmon et al., 2015). Bayesian models are flexible enough to deal with trends, seasonality and other covariates, and different distributions (e.g. Poisson, Gamma). The ability to incorporate prior knowledge about the distribution of parameters can also help hone inference (e.g. through the narrowing of credible intervals) for highly-dimensional models which may result from the scaling of BN to millions of nodes (e.g. modelling of an entire population) for real-time surveillance applications. Traditionally, a strength of Bayesian methods has been the ability to monitor for specific diseases, based on their known characteristics; while frequentist methods, being non-specific, were better apt at monitoring for unknown diseases. However, recent methodological developments have proved successful in applying BN to detect known diseases by using informative prior probabilities, and unknown diseases by using relatively non-informative prior probabilities (Shen and Cooper, 2007). Examples reviewed – both full and empirical Bayes – demonstrated that BN are suited to model the underlying epidemiological process, but can also be applied to models that rely on the observed data alone to both predict and detect epidemic curves (Brooks et al., 2015).

A full Bayesian model where all parameters retain their probability density distribution is robust to overfitting and may handle complex data-streams with correlated data. Such approach is considered by some to be the "gold standard" and fulfil important optimality criteria. This is, at least in theory, an advantage over the empirical approaches that include shortcuts (e.g. using maximum likelihood or expectation maximisation to estimate parameters) and simplified assumptions regarding dependency which may both serve to propagate errors. Under particular scenarios (in particular when sampling size is low), Bayesian credible interval estimates obtained from MCMC, such as those generated through Gibbs sampling, will be narrower than confidence intervals calculated on the basis of large sample approximations (Salameh et al., 2014).

Despite the flexibility and robustness gained with the inclusion of prior probabilities into Bayesian based monitoring models, their specification can be challenging. This is particularly the case when dealing with multiple sources of evidence. While Bayesian models allow the incorporation of many sources of data, and provide robustness in the estimation of their probability distributions (Morrison et al., 2016), deciding on the structure of covariance among data streams, both in the absence and presence of an outbreak, is not a trivial task (Flavie Vial et al., 2016). A combination of expert opinion and previous data are often required to find suitable informative prior distributions. In many cases, weakly informative priors are used; weakly in the sense that the final results (posterior distributions) are mainly influenced by data rather than the priors. When the amount of data is scarce the choice of prior distribution may have a large influence on the final results. Therefore, investigating the sensitivity of the results for different priors may be important. LeStrat and Carrat (Le Strat and Carrat, 1999) proposed to detect outbreak and non-outbreak phases of influenza with HMMs using Gaussian distributed priors while Rath et al. (Rath et al., 2003) later demonstrated better detection accuracy using a mixture of Gaussian and exponential distributions. However, it is important to remember that in frequentist SyS, implicit priors are used. The sensitivity and specificity of the system are calculated based on outbreaks simulated from a probability distribution hence making the performance parameters conditional on these prior assumptions.

Irrespectively whether the analysis is based on frequentist or Bayesian approaches, the existence of multiple hypotheses is a major challenge for interpretation and communication. In the examples above the hypotheses of main concern (H) would be that that there is an outbreak of the disease (or disease group) of interest and the null hypothesis (H0) that everything is normal. Bayes rule is used to estimate the posterior probability that an outbreak is ongoing under the assumption that the set of hypotheses is exhaustive, that is P(H) + P(H0)=1. When we deal with a finite number of hypotheses (H1, …, Hn) we may assign a prior probability to each hypothesis and apply Bayes rule. In many cases, it is impossible to include all hypotheses in the model. There may be other possible explanations for a peak in a data-stream which is not covered by any model hypothesis and in such instances, it is not possible to calculate the posterior probability of H. The same problem applies when, for example, estimating the specificity of a detection algorithm. In the Bayesian decision framework, however, we can still use likelihood ratios to calculate the odds ratio between two hypotheses included in the decision framework, for instance outbreak presence over non-outbreak (Taroni et al., 2010).

Thus, presenting the result as a posterior probability of a specific disease may be dangerous. One option is to make the set of hypotheses exhaustive by explicitly defining H as “an outbreak of disease X or another disease displaying similar symptoms” and assuming all other scenarios explaining the results are covered by H0. Another possibility is to present the results as the likelihood ratio or the Bayes factor between each pair of hypotheses that are part of the model.

The main drawback of applying full Bayesian approaches to real-life SyS scenarios resides in the resulting computations, which can be quite complex and necessitate the use of approximate solutions. For example, models based on both continuous and discrete distributions may be handled in a MCMC framework (typically Gibbs sampling as in WinBUGS/OpenBUGS/JAGS). In some cases, it is possible to find a solution in “closed form”, i.e. the integrals have an algebraic solution, circumventing the need for finding distributions for parameters by simulation and thereby speeding up calculations. To overcome the long computational time associated with most MCMC-based models, approximate solutions to BN, including integrated nested Laplace approximations (Schrödle et al., 2011) have been developed. In the most complex cases, implementation of the full Bayes method may be theoretically too demanding and computationally too difficult and thus, in practice, the simpler empirical methods will still have a part to play (Lawson and Kleinman, 2005).

Conclusion

One of the main advantages of Bayesian approaches to epidemiological surveillance is their conceptual simplicity and the fact that their fundamental principles are based on relatively few concepts. Studies have linked the easiness of explanation of medical decision-support systems to user perception of the system and the accuracy of decision-making) (Suermondt and Cooper, 1992). The Bayesian framework where results are presented as posterior probabilities (and strength of evidence as LR's or Bayes factors) has been found to be analogous to intuitive human reasoning and thus useful for presenting and interpreting results. They have been adopted, for example, as the golden standard for presenting forensic evidence in many countries (Aitken and Taroni, 2004). Nonetheless, the posterior probability for a given node will depend on several factors that may need explaining to the decision-makers: the evidence (itself possibly arising from multiple data streams), the BN structure (nodes and arcs), and the BN parameters (local conditional probabilities). Thankfully, methods (e.g. the hierarchical explanation method) do exist for selecting and organising information to explain BN inference in the context of outbreak detection (Madigan et al., 1997; Šutovský and Cooper, 2008). The Bayesian approach fits nicely in a decision theoretic framework. If utility functions and the posterior probabilities of the hypotheses are provided, it is possible to find the best action for the decision-making problem at hand.

A common objection to Bayesian methods is that the posterior probabilities can be strongly influenced by the priors. As an alternative approach (M G Andersson et al., 2014), it is proposed a framework for communicating SyS based on explicit separation of prior information about an hypothesis and evidence from data. In this framework the results from SyS would be presented as the Bayes factor, i.e. the ratio between the posterior and prior odds that an outbreak is ongoing against an alternative hypothesis. A specific advantage of this approach is that is it logical for communicating results also when the set of hypotheses is not exhaustive. Furthermore, it is analogous with scale of evidence adopted at many forensic institutes (Nordgaard et al., 2012). From that perspective the proposed framework is useful for presenting results even when the modelling approach is not in itself Bayesian.

List of abbreviations

BN: Bayesian network

CPA: change point analysis

CUSUM: cumulative sum control chart

DBN: dynamic Bayesian network

EWMA: exponentially weighted moving average chart

HMM: hidden Markov models

ID: influence diagrams

LR: likelihood ratio

MCMC: Markov chain Monte Carlo

SMM: state-space models

SyS: syndromic surveillance

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

MGA, FV, PH, AL and JG conceived the concept and framework developed in this manuscript. MGA and RH analysed the data presented in the working examples. FV, FCD GMA and RH contributed to the literature review and wrote the paper. PH, AL and JG commented on various versions of the manuscript. All authors read and approved the final manuscript.

Acknowledgements

This research was carried within the VICE project funded under the FP7 EMIDA ERA-NET initiative. National funding within ERA-Net was provided by the Research Council of Norway (Project nr. 219101/O10), the Swedish research council Formas (Contract nr. 221-2011-2214). RH was partly funded by a grant from the Swedish Civil Contingencies Agency (MSB) (MSB-Säkprov Kvalitetssäkring av provtagning för farliga ämnen I foder). FCD was partly funded by the Swedish government research council for sustainable development (FORMAS) with grant number 2017-00779.

The authors thank Petter Mostad for comments on working example 3; and Yves Van der Stede for insightful comments on various versions of the manuscript. Data on on-farm mortality and late abortions was gathered from the Norwegian Dairy Herd Recording System (NDHRS, Hamar, Norway) and the Norwegian Cattle Health Services (NCHS, Ås, Norway) in accord with agreement number 004.2012.

References

Aitken, C., Taroni, F., 2004. Statistics and the Evaluation of Evidence for Forensic Scientists, 2nd ed. John Wiley & Sons Ltd, Chichester, UK.

Amorós, R., Conesa, D., López-Quílez, A., Martinez-Beneito, M.A., 2020. A spatio-temporal hierarchical Markov switching model for the early detection of influenza outbreaks. Stoch. Environ. Res. Risk Assess. 34, 275–292. https://doi.org/10.1007/s00477-020-01773-5

Andersson, M G, Faverjon, C., Vial, F., Legrand, L., Leblond, A., 2014. Using Bayes' rule to define the value of evidence from syndromic surveillance. PLoS One 9, e111335. https://doi.org/10.1371/journal.pone.0111335

Andersson, Mats Gunnar, Faverjon, C., Vial, F., Legrand, L., Leblond Agnès ED - Gubbins, S., 2014. Using Bayes' rule to define the value of evidence from syndromic surveillance. PLoS One 9, e111335–e111335.

Barry, D., Hartigan, J.A., 1993. A Bayesian Analysis for Change Point Problems. J. Am. Stat. Assoc. 88, 309. https://doi.org/10.2307/2290726

Bastos, A.D.S., Arnot, L.F., Jacquier, M.D., Maree, S., 2009. A host species-informative internal control for molecular assessment of African swine fever virus infection rates in the African sylvatic cycle Ornithodoros vector. Med. Vet. Entomol. 23, 399–409. https://doi.org/10.1111/j.1365-2915.2009.00828.x

BayesFusion, LLC [WWW Document], n.d. URL http://www.bayesfusion.com/ (accessed 5.21.21).

Bernardo, J.M., Smith, A.F.M., 1994. Bayesian Theory. Wiley, Chichester, UK.

Blangiardo, M., Boulieri, A., Diggle, P., Piel, F.B., Shaddick, G., Elliott, P., 2020. Advances in spatiotemporal models for non-communicable disease surveillance. Int. J. Epidemiol. 49, I26–I37. https://doi.org/10.1093/ije/dyz181

Boulieri, A., Bennett, J.E., Blangiardo, M., 2020. A Bayesian mixture modeling approach for public health surveillance. Biostatistics 21, 369–383. https://doi.org/10.1093/biostatistics/kxy038

Brooks, L.C., Farrow, D.C., Hyun, S., Tibshirani, R.J., Rosenfeld, R., 2015. Flexible Modeling of Epidemics with an Empirical Bayes Framework. PLoS Comput. Biol. 11, 1–18. https://doi.org/10.1371/journal.pcbi.1004382

Burkom, H., Elbert, Y., Ramac-Thomas, L., Cuellar, C., Hung, V., 2013. Refinement of a Population-Based Bayesian Network for Fusion of Health Surveillance Data. Online J. Public Health Inform. 5, e6.

Burkom, H.S., Ramac-Thomas, L., Babin, S., Holtry, R., Mnatsakanyan, Z., Yund, C., 2011. An integrated approach for fusion of environmental and human health data for disease surveillance. Stat. Med. 30, 470–479. https://doi.org/10.1002/sim.3976

Carter, C.K., Kohn, R., 1994. On Gibbs Sampling for State Space Models. Biometrika 81, 541–553. https://doi.org/10.2307/2337125

Christensen, R., Johnson, W., Branscum, A., Hanson, T.E., 2011. Bayesian Ideas and Data Analysis – An Introduction for Sciensists and Statisticians. CRC Press, Boca Raton, USA.

Conesa, D., Martínez-Beneito, M.A., 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. https://doi.org/10.1177/0962280211414853

Cooper, G., Dash, D., Levander, J., Wong, W., Hogan, W., Wagner, M., 2004. Bayesian biosurveillance of disease outbreaks | Department of Biomedical Informatics [WWW Document]. Proc. Conf. Uncertain. Artif. Intell. URL https://www.dbmi.pitt.edu/node/2074 (accessed 4.16.21).

Cooper, G.F., Villamarin, R., Tsui, F.C., Millett, N., Espino, J.U., Wagner, M.M., 2015. A method for detecting and characterizing outbreaks of infectious disease from clinical reports. J. Biomed. Inform. 53, 15–26. https://doi.org/10.1016/j.jbi.2014.08.011

Dassanayaka, S.K., 2015. Spatial CUSUM Chart Based Method for Rapid Detection of Outbreaks. Online J. Public Health Inform. 7, e122. https://doi.org/10.5210/ojphi.v7i1.5788

Dawson, P., Gailis, R., Meehan, A., 2015. Detecting disease outbreaks using a combined Bayesian network and particle filter approach. J. Theor. Biol. 370, 171–183. https://doi.org/10.1016/j.jtbi.2015.01.023

Dórea, F.C., Vial, F., 2016. Animal health syndromic surveillance: a systematic literature review of the progress in the last 5 years (2011--2016). Vet. Med. Reports 7, 157–169.

Dórea, F.C.F.C., Sanchez, J., Revie, C.W.C.W., 2011. Veterinary syndromic surveillance: Current initiatives and potential for development. Prev. Vet. Med. 101, 1–17. https://doi.org/10.1016/j.prevetmed.2011.05.004

Druzdzel, M.J., 1999. GeNIe: A Development Environment for Graphical Decision-Analytic Models. Proc. AMIA Annu. Symp.

Ebel, E.D., Williams, M.S., Schlosser, W.D., 2017. Estimating the Type II error of detecting changes in foodborne illnesses via public health surveillance. Microb. Risk Anal. 7, 1–7. https://doi.org/10.1016/j.mran.2017.10.001

Ertem, Z., Raymond, D., Meyers, L.A., 2018. Optimal multi-source forecasting of seasonal influenza. PLoS Comput. Biol. 14, 1–16. https://doi.org/10.1371/journal.pcbi.1006236

Farrington, C.P., j. Andrews, N., a. d. Beale, Catchpole, M.A., 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. J. R. Stat. Soc. A 159, 547–563.

Faverjon, C., Andersson, M.G., Decors, A., Tapprest, J., Tritz, P., Sandoz, A., Kutasi, O., Sala, C., Leblond, A., 2016. Evaluation of a Multivariate Syndromic Surveillance System for West Nile Virus. Vector Borne Zoonotic Dis. 16, 382–90. https://doi.org/10.1089/vbz.2015.1883

Fricker, R.D., Rolka, H., 2006. Protecting against Biological Terrorism: Statistical Issues in Electronic Biosurveillance. CHANCE 19, 4–14. https://doi.org/10.1080/09332480.2006.10722809

Frisén, M., 2010. Principles for multivariate surveillance, in: Frontiers in Statistical Quality Control 9. Physica-Verlag, pp. 133–144. https://doi.org/10.1007/978-3-7908-2380-6_9

Frisén, M., Andersson, E., 2009. Semiparametric Surveillance of Monotonic Changes. Seq. Anal. 28, 434–454. https://doi.org/10.1080/07474940903238029

García, Y.E., Christen, J.A., Capistrán, M.A., 2015. A Bayesian outbreak detection method for influenza-like illness. Biomed Res. Int. 2015. https://doi.org/10.1155/2015/751738

Ghahramani, Z., 2001. An introduction to hidden Markov models and Bayesian networks. Int. J. Pattern Recognit. Artif. Intell. 15, 9–42. https://doi.org/10.1142/s0218001401000836

Grossi, E., 2008. The single individual in medicine: how to escape from the probability theory trap. Cases J. 1, 58. https://doi.org/10.1186/1757-1626-1-58

Heaton, M.J., Banks, D.L., Zou, J., Karr, A.F., Datta, G., Lynch, J., Vera, F., 2012. A spatio-temporal absorbing state model for disease and syndromic surveillance. Stat. Med. 31, 2123–2136. https://doi.org/10.1002/sim.5350

Hedell, R., Andersson, M.G., Faverjon, C., Marcillaud-Pitel, C., Leblond, A., Mostad, P., 2019. Surveillance of animal diseases through implementation of a Bayesian spatio-temporal model: A simulation example with neurological syndromes in horses and West Nile Virus. Prev. Vet. Med. 162, 95–106. https://doi.org/10.1016/j.prevetmed.2018.11.010

Höhle, M., 2007. surveillance: An R package for the surveillance of infectious diseases. Comput. Stat. 22, 571–582.

Höhle, M., An Der Heiden, M., 2014. Bayesian nowcasting during the STEC O104: H4 outbreak in Germany, 2011. Biometrics 70, 993–1002. https://doi.org/10.1111/biom.12194

Holt, C.C., 2004. Forecasting seasonals and trends by exponentially weighted moving averages. Int. J. Forecast. 20, 5–10. https://doi.org/https://doi.org/10.1016/j.ijforecast.2003.09.015

Izadi, M., Buckeridge, D., Okhmatovskaia, A., Tu, S.W., O'Connor, M.J., Nyulas, C., Musen, M.A., O$\blacksquare$Connor, M.J., Nyulas, C., Musen, M.A., 2009. A Bayesian Network Model for Analysis of Detection Performance in Surveillance Systems, in: AMIA 2009 Symposium Proceedings. pp. 276–280.

814 Jégat, C., Carrat, F., Lajaunie, C., Wackernagel, H., 2008. Early Detection and Assessment of Epidemics by
815 Particle Filtering, in: Soares, A., Pereira, M.J., Dimitrakopoulos, R. (Eds.), GeoENV VI – Geostatistics for
816 Environmental Applications. C Springer Science+Business Media B.V.

817 Jiang, X., Neill, D.B., Cooper, G.F., 2010. A Bayesian network model for spatial event surveillance. Int. J.
818 Approx. Reason. 51, 224–239. https://doi.org/https://doi.org/10.1016/j.ijar.2009.01.001

819 Jung, R., Kukuk, M., Liesenfeld, R., 2006. Time series of count data: modeling, estimation and diagnostics.
820 Comput. Stat. & Data Anal. 51, 2350–2364.

821 Kass-Hout, T.A., Xu, Z., McMurray, P., Park, S., Buckeridge, D.L., Brownstein, J.S., Finelli, L., Groseclose, S.L.,
822 2012. Application of change point analysis to daily influenza-like illness emergency department visits. J
823 Am Med Inf. Assoc 19, 1075–1081. https://doi.org/10.1136/amiajnl-2011-000793

824 Kleinman, K., Lazarus, R., Platt, R., 2004. A Generalized Linear Mixed Models Approach for Detecting Incident
825 Clusters of Disease in Small Areas, with an Application to Biological Terrorism. Am. J. Epidemiol. 159,
826 217–224. https://doi.org/10.1093/aje/kwh029

827 Kulldorff, M., 2001. Prospective time periodic geographical disease surveillance using a scan statistic. J. R.
828 Stat. Soc. Ser. A (Statistics Soc. 164, 61–72. https://doi.org/10.1111/1467-985x.00186

829 Kulldorff, M., 1997. A spatial scan statistics. Commun. Stat. - Theory Methods 26, 1481–1496.

830 Lawson, A.B., Kleinman, K., 2005. Spatial and Spatio-temporal Disease Analysis, in: Spatial and Syndromic
831 Surveillance for Public Health. Chichester, John Wiley and Sons Ltd., Chichester.

832 Le Strat, Y., Carrat, F., 1999. Monitoring epidemiologic surveillance data using hidden Markov models. Stat.
833 Med. 18, 3463–3478.

834 Li, G., Best, N., Hansell, A.L., Ahmed, I., Richardson, S., 2012. BaySTDetect: Detecting unusual temporal
835 patterns in small area data via Bayesian model choice. Biostatistics 13, 695–710.
836 https://doi.org/10.1093/biostatistics/kxs005

837 Lin, J., Ludkovski, M., 2014. Sequential Bayesian inference in hidden Markov stochastic kinetic models with
838 application to detection and response to seasonal epidemics. Stat. Comput. 24, 1047–1062.
839 https://doi.org/10.1007/s11222-013-9419-z

840 Liu, D., Mitchell, L., Cope, R.C., Carlson, S.J., Ross, J. V., 2020. Elucidating user behaviours in a digital health
841 surveillance system to correct prevalence estimates. Epidemics 33.
842 https://doi.org/10.1016/j.epidem.2020.100404

843 Lu, H., Zeng, D., Chen, H., 2010. Prospective Infectious Disease Outbreak Detection Using Markov Switching
844 Models. IEEE Trans. Knowl. Data Eng. 22, 565–577. https://doi.org/10.1109/TKDE.2009.115

845 Lu, H., Zeng, D., Chen, H., 2008. Bioterrorism event detection based on the Markov switching model: A
846 simulated anthrax outbreak study, in: 2008 IEEE International Conference on Intelligence and Security
847 Informatics. pp. 76–81. https://doi.org/10.1109/ISI.2008.4565033

848 Lucas, J.M., 1985. Counted Data {CUSUM}'s. Technometrics 27, 129–144.

849 Madigan, D., 2005. Bayesian Data Mining for Health Surveillance, in: Lawson, A.B., Kleinman, K. (Eds.), Spatial
850 and Syndromic Surveillance. John Wiley and Sons Ltd, Chichester.

851 Madigan, D., Mosurski, K., Almond, R.G., 1997. Graphical Explanation in Belief Networks. J. Comput. Graph.
852 Stat. 6, 160–181. https://doi.org/10.1080/10618600.1997.10474735

Manitz, J., Höhle, M., 2013. Bayesian outbreak detection algorithm for monitoring reported cases of campylobacteriosis in Germany. Biometrical J. 55, 509–526. https://doi.org/10.1002/bimj.201200141

Martínez-Beneito, M.A., 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. https://doi.org/10.1002/sim.3320

McGough, S.F., Johansson, M.A., Lipsitch, M., Menzies, N.A., 2020. Nowcasting by Bayesian smoothing: A flexible, generalizable model for real-time epidemic tracking. PLoS Comput. Biol. 16. https://doi.org/10.1371/journal.pcbi.1007735

Mnatsakanyan, Z.R., Burkom, H.S., Coberly, J.S., Lombardo, J.S., 2009. Bayesian Information Fusion Networks for Biosurveillance Applications. J. Am. Med. Informatics Assoc. 16:6, 855–863.

Morrison, K.T., Shaddick, G., Henderson, S.B., Buckeridge, D.L., 2016. A latent process model for forecasting multiple time series in environmental public health surveillance. Stat. Med. 35, 3085–3100. https://doi.org/10.1002/sim.6904

Moss, R., Zarebski, A.E., Dawson, P., Franklin, L.J., Birrell, F.A., Mccaw, J.M., 2016. Communicable Diseases Intelligence 2019 - Anatomy of a seasonal influenza epidemic forecast.

Neil, D.B., Moore, A.W., Cooper, G.F., 2006. A Bayesian Scan Statistic for Spatial Cluster Detection. Adv. Dis. Surveill. 1, 55.

Nordgaard, A., R, A., W, D., L, J., 2012. Scale of conclusions for the value of evidence. Law Probab. Risk 11, 1–24. https://doi.org/10.1093/lpr/mgr020

Noufaily, A., Enki, D.G., Farrington, P., Garthwaite, P., Andrews, N., Charlett, A., 2013. An improved algorithm for outbreak detection in multiple surveillance systems. Stat. Med. 32, 1206–1222. https://doi.org/10.1002/sim.5595

Onisko, A., Wallstrom, G., Wagner, M.M., 2006. Decision analysis, in: Wagner MM Aryel RM., M.A.W. (Ed.), Handbook of Biosurveillance. Elsevier Academic Press.

Pearl, J., 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, 2nd ed. Morgan Kaufmann, San Francisco, USA.

Polyakov, P., Breban, R., 2016. Bayesian monitoring of emerging infectious diseases. PLoS One 11. https://doi.org/10.1371/journal.pone.0152629

R Core Team, Team, R.C., 2016. R: A Language and Environment for Statistical Computing.

Rath, T.M., Carreras, M., Sebastiani, P., 2003. Automated Detection of Influenza Epidemics with Hidden Markov Models, in: Lenz, H.-J., Bradley, E., Kruse, R., Borgelt, C. (Eds.), Advances in Intelligent Data Analysis V. Springer Berlin Heidelberg, Berlin, Heidelberg LB - 10.1007/978-3-540-45231-7_48, pp. 521–532.

Robertson, C., Nelson, T.A., MacNab, Y.C., Lawson, A.B., 2010. Review of methods for space-time disease surveillance. Spat. Spatiotemporal. Epidemiol. 1, 105–116. https://doi.org/10.1016/j.sste.2009.12.001

Rolka, H., Burkom, H., Cooper, G.F., Kulldorff, M., Madigan, D., Wong, W., 2007. Issues in applied statistics for public health bioterrorism surveillance using multiple data streams : Research needs. Stat. Med. 26, 1834–1856. https://doi.org/10.1002/sim

Rotejanaprasert, C., Ekapirat, N., Areechokchai, D., Maude, R.J., 2020. Bayesian spatiotemporal modeling

with sliding windows to correct reporting delays for real-time dengue surveillance in Thailand. Int. J. Health Geogr. 19, 1–13. https://doi.org/10.1186/s12942-020-00199-0

Ryden, T., 2008. EM versus Markov chain Monte Carlo for estimation of hidden Markov models: a computational perspective. Bayesian Anal. 3, 659–688. https://doi.org/10.1214/08-BA326

Salameh, P., Waked, M., Khayat, G., Dramaix, M., 2014. Bayesian and Frequentist Comparison for Epidemiologists: A NonMathematical Application on Logistic Regressions. Open Epidemiol. J. 7, 17–26.

Salmon, M., Schumacher, D., Höhle, M., n.d. Monitoring Count Time Series in R: Aberration Detection in Public Health Surveillance. 2016 70, 35. https://doi.org/10.18637/jss.v070.i10

Salmon, M., Schumacher, D., Stark, K., Höhle, M., 2015. Bayesian outbreak detection in the presence of reporting delays. Biometrical J. 57, 1051–1067. https://doi.org/10.1002/bimj.201400159

Schmidt, A.M., Pereira, J.B.M., 2011. Modelling Time Series of Counts in Epidemiology. Int. Stat. Rev. / Rev. Int. Stat. 79, 48–69.

Schrödle, B., Held, L., Riebler, A., Danuser, J., 2011. Using integrated nested Laplace approximations for the evaluation of veterinary surveillance data from Switzerland: a case-study. J. R. Stat. Soc. Ser. C (Applied Stat. 60, 261–279. https://doi.org/10.1111/j.1467-9876.2010.00740.x

Scott, S.L., 2002. Bayesian Methods for Hidden Markov Models. J. Am. Stat. Assoc. 97, 337–351. https://doi.org/10.1198/016214502753479464

Sebastiani, P., Mandl, K.D., Szolovits, P., Kohane, I.S., Ramoni, M.F., 2006. A Bayesian dynamic model for influenza surveillance. Stat. Med. 25, 1803–1825. https://doi.org/10.1002/sim.2566

Shen, Y., Cooper, G., 2007. A Bayesian Biosurveillance Method That Models Unknown Outbreak Diseases. https://doi.org/10.1007/978-3-540-72608-1_21

Shen, Y., Cooper, G.F., 2009. Bayesian modeling of unknown diseases for biosurveillance. AMIA ... Annu. Symp. proceedings. AMIA Symp. 2009, 589–593.

Shewhart, W.A., 1939. Statistical Method from the Viewpoint of Quality Control. Lancaster Press, New York, USA.

Sonesson, C., Frisén, M., 2005. Multivariate surveillance, in: Lawson, A., Kleinman, K. (Eds.), Spatial Surveillance for Public Health. Wiley, Chichester, UK, pp. 169–186.

Struchen, R., Vial, F., Andersson, M.G., 2017. Value of evidence from syndromic surveillance with cumulative evidence from multiple data streams with delayed reporting. Sci. Rep. 7. https://doi.org/10.1038/s41598-017-01259-5

Suermondt, H.J., Cooper, G.F., 1992. An evaluation of explanations of probabilistic inference. Proceedings. Symp. Comput. Appl. Med. Care 579–585.

Šutovský, P., Cooper, G.F., 2008. Hierarchical explanation of inference in Bayesian networks that represent a population of independent agents, in: Conference on ECAI 2008: 18th European Conference on Artificial Intelligence. IOS Press, Amsterdam, The Netherlands, pp. 214–218.

Szmaragd, C., Wilson, A.J., Carpenter, S., Wood, J.L.N., Mellor, P.S., Gubbins, S., 2009. A Modeling Framework to Describe the Transmission of Bluetongue Virus within and between Farms in Great Britain. PLoS One 4, e7741. https://doi.org/10.1371/journal.pone.0007741

Taroni, F., Bozza, S., Biedermann, A., Garbolino, P., Aitken, C., 2010. Data Analysis in Forensic Science: A Bayesian Decision Perspective. John Wiley & Sons Ltd, Chichester, UK.

Texier, G., Farouh, M., Pellegrin, L., Jackson, M.L., Meynard, J.B., Deparis, X., Chaudet, H., 2016. Outbreak definition by change point analysis: A tool for public health decision? BMC Med. Inform. Decis. Mak. 16. https://doi.org/10.1186/s12911-016-0271-x

Thomas, A., O'Hara, R., Ligges, U., Sturtz, S., 2006. Making BUGS open. R News 6, 12–17.

Unkel, S., Farrington, C.P., Garthwaite, P.H., Robertson, C., Andrews, N., 2012. Statistical methods for the prospective detection of infectious disease outbreaks: a review. J. R. Stat. Soc. Ser. A (Statistics Soc. 175, 49–82. https://doi.org/10.1111/j.1467-985X.2011.00714.x

Vial, F, Wei, W., Held, L., 2016. Methodological challenges to multivariate syndromic surveillance: A case study using Swiss animal health data. BMC Vet. Res. 12. https://doi.org/10.1186/s12917-016-0914-2

Vial, Flavie, Wei, W., Held, L., 2016. Methodological challenges to multivariate syndromic surveillance: a case study using Swiss animal health data. BMC Vet. Res. 12. https://doi.org/10.1186/s12917-016-0914-2

Wang, X., Zhou, M., Jia, J., Geng, Z., Xiao, G., 2018. A Bayesian approach to real-time monitoring and forecasting of Chinese foodborne diseases. Int. J. Environ. Res. Public Health 15, 1–13. https://doi.org/10.3390/ijerph15081740

Watkins, R.E., Eagleson, S., Veenendaal, B., Wright, G., Plant, A.J., 2009. Disease surveillance using a hidden Markov model. BMC Med. Inform. Decis. Mak. 9. https://doi.org/10.1186/1472-6947-9-39

Winters, P.R., 1960. Forecasting Sales by Exponentially Weighted Moving Averages. Manage. Sci. 6, 324–342.

Wong, W.-K., Cooper, G., Dash, D., Levander, J., Dowling, J., Hogan, W., Wagner, M., 2005. Use of multiple data streams to conduct Bayesian biologic surveillance. Morb. Mortal. Wkly. Rep. 54 Suppl, 63–69.

Wong, W., Cooper, G., Dash, D., Levander, J., Dowling, J., Hogan, W., Wagner, M., 2005. Use of multiple data streams to conduct Bayesian biologic surveillance. MMWR Suppl. 54, 63–69.

Yuan, M., Boston-Fisher, N., Luo, Y., Verma, A., Buckeridge, D.L., 2019. A systematic review of aberration detection algorithms used in public health surveillance. J. Biomed. Inform. https://doi.org/10.1016/j.jbi.2019.103181

Zamiri, A., Yazdi, H.S., Goli, S.A., 2015. Temporal and spatial monitoring and prediction of epidemic outbreaks. IEEE J. Biomed. Heal. Informatics 19, 735–744. https://doi.org/10.1109/JBHI.2014.2338213

Zou, J., Zhang, Z., Yan, H., 2018. A hybrid hierarchical Bayesian model for spatiotemporal surveillance data. Stat. Med. https://doi.org/10.1002/sim.7909