Published in final edited form as:

Stat Med. 2006 June 15; 25(11): 1803–1825. doi:10.1002/sim.2566.

A Bayesian dynamic model for influenza surveillance

Paola Sebastiani^{1,*†}, Kenneth D. Mandl², Peter Szolovits³, Isaac S. Kohane², and Marco F. Ramoni²

¹Department of Biostatistics, Boston University, Boston, MA, U.S.A

²Children's Hospital Informatics Program, Harvard Medical School, Boston, MA, U.S.A

³Computer Science and Artificial Intelligence Laboratory, M.I.T., Cambridge, MA, U.S.A

SUMMARY

The severe acute respiratory syndrome (SARS) epidemic, the growing fear of an influenza pandemic and the recent shortage of flu vaccine highlight the need for surveillance systems able to provide early, quantitative predictions of epidemic events. We use dynamic Bayesian networks to discover the interplay among four data sources that are monitored for influenza surveillance. By integrating these different data sources into a dynamic model, we identify in children and infants presenting to the pediatric emergency department with respiratory syndromes an early indicator of impending influenza morbidity and mortality. Our findings show the importance of modelling the complex dynamics of data collected for influenza surveillance, and suggest that dynamic Bayesian networks could be suitable modelling tools for developing epidemic surveillance systems.

Keywords

dynamic Bayesian networks; influenza surveillance; syndromic data

1. INTRODUCTION

The 2001 Anthrax attacks, the spread of SARS during the Winter of 2002–2003, the shortage of flu vaccine in 2004, and the looming threat of an influenza pandemic caused by the H5N1 virus motivate the urgent need for surveillance systems able to provide early quantitative predictions of acute respiratory infections. Although less lethal than SARS, influenza is the seventh leading cause of death for people above 65 and below 4 years of age, and among the top 10 in almost all age groups [1]. Every year influenza infects up to 40 million Americans, causes the hospitalization of over 114 000 [2], kills approximately 36 000 [3] and costs the United States between two and five billion dollars in physician visits, lost productivity, and lost wages [4]. This cost could grow to 160 billion dollars in the event of an influenza pandemic: a sudden and widespread outbreak of a new strain of the influenza virus that has significant morbidity and mortality [5]. Historically, the periodicity of pandemics has ranged between 9 and 39 years, with the last one in 1968–1969. A new

pandemic in the next few years is considered today a looming if not inevitable threat [6], which is reinforced by the increasing number of outbreaks caused by the H5N1 bird-flu strain [7].

Because influenza viruses change continuously by antigenic drift and shift, nobody is immune and surveillance efforts are in place to detect changes in the viruses, to monitor their effects on hospitalization and death rates, and possibly to forecast the resurgence of a new influenza pandemic. The influenza surveillance system in the United States is managed by the Centers for Disease Control and Prevention (CDC) Influenza Branch, which collects and reports weekly information on influenza activity in the United States from September to May. Distribution of influenza activity and identification of viral types is ascertained through about 75 laboratories of the World Health Organization and 50 laboratories of the National Respiratory and Enteric Virus Surveillance System located throughout the United States. About 650 sentinel physicians report the weekly number of patients they have seen and the number of those patients with influenza-like illness (ILI) by age group that is defined as fever (temperature above 100°F) plus either a cough or a sore throat. The effect of influenza on mortality is monitored through weekly reports filed by the vital statistics offices of 122 cities that contain the total number of death certificates and the number of those for which pneumonia or influenza (P&I) were listed as a contributing cause of death.

The current influenza surveillance system monitors all these data streams individually, as shown in Plate 1. For example, ILI data are monitored to detect surges of influenza morbidity nationwide by comparing weekly data to a national baseline defined by a fixed threshold. It is acknowledged by CDC that the national baseline does not provide a useful threshold for local influenza surveillance so that the issue of local monitoring of influenza morbidity is still open [8]. The aggregate mortality data from the 122 cities are monitored to detect influenza epidemics. The detection consists of comparing the aggregated mortality data to an epidemic threshold that is calculated for each week using a baseline defined by a cyclic regression model [9]. To the best of our knowledge, no attempt is made to use and possibly integrate ILI and P&I data to forecast the course of an epidemic.

Furthermore, for both states and physicians, influenza activity reporting is voluntary and data do not become available to health officials until approximately 2 weeks after the reports have been filed. An emerging approach to more timely detection is syndromic surveillance [10], which identifies infected individuals early in the course of their disease, generally before a confirmed diagnosis is made. Patients are classified by syndrome, such as acute respiratory infection or gastrointestinal illness, based on a variety of data sources ranging from purchase of over the counter medications [11] to primary care physician and emergency department (ED) logs [12, 13]. The primary focus of these efforts has been to provide early detection of bioterrorism, following the Anthrax attacks of 2001, although this data could also provide early indication of naturally occurring disease outbreaks [14].

Despite the increasing amount of data that could potentially inform about outbreaks of disease, biosurveillance research to date has mainly focused on monitoring individual data streams [11, 13, 15]. When multiple signals have been considered, a system architecture based on a set of parallel, independent surveillance systems has been envisioned [16, 17].

The intuition underlying our approach is that the integration of different data streams should provide complementary information about a disease and improve our ability for early detection of outbreaks. This integration, however, requires a coherent modelling framework to integrate data into a global model.

In this paper, we introduce a novel approach to integrate different data streams into a multivariate model for influenza surveillance. The novelty of our approach is in both the modelling framework that we use and in the type of data. We build a dynamic Bayesian networks that relates pediatric and adult syndromic data in two EDs to the traditional measures of influenza morbidity and mortality, and we show how to use this model for 'active' influenza surveillance by forecasting the course of influenza epidemics. Our analysis shows that the use of a dynamic Bayesian network for influenza surveillance has several advantages. By directly modelling measures of influenza morbidity and mortality, our model can be used to forecast the beginning of epidemics, as well as peaks of epidemics. Furthermore, the joint modelling of the four data streams shows that pediatric patients are infected with respiratory viruses well before the general population. In particular, the number of respiratory syndrome cases in a pediatric ED predicts influenza morbidity in the general population as early as 2 weeks in advance and influenza mortality as early as 3 weeks in advance. These findings suggest that children with respiratory syndromes seen at the ED act as sentinels for surges in influenza morbidity and mortality and that active surveillance of pediatric populations could become an important component of the influenza surveillance effort.

The next section describes the data that we used to build our dynamic models. Section 3 gives a brief introduction to dynamic Bayesian networks and Section 4 describes how we built the network with the data available and how we use it to forecast the course of an influenza epidemic. Conclusions are reported in Section 5.

2. DATA

We use variations in the weekly frequency of patients with respiratory syndromes presenting to two urban, tertiary care teaching hospitals in Boston, Massachusetts, to predict ILI data for the influenza seasons 1997–1998–2001–2002, and to deaths for pneumonia and influenza (P&I) in New England from January 1998 to March 2003 published by CDC.[‡] Both ILI and P&I data includes all age groups. The hospitals, one pediatric (CH) and the other adult (AH), share the same catchment area and each emergency department has an annual census of approximately 50 000 visits, with an average age of approximately 6 years for CH and 48 years for AH. Syndromic grouping was based on ED chief complaints as described in Reference [18]. Chief complaints were used to select those ED visits that were related to respiratory illness. At AH, chief complaints are entered as free-text during the triage process. We used two procedures to classify complaints: a text-string search, and a publicly available naïve Bayesian classification program [16]. Default probabilities from an ED data set, supplied with the software, were applied to the AH data. At CH, chief complaint codes were chosen during the triage process, from a pre-defined on-line list of

[‡]<http://wonder.cdc.gov/mmwr/mmwmort.asp>

181 choices. A previously validated subset of the constrained chief complaint set was chosen *a priori* for inclusion in the respiratory syndromic grouping. The ED data set dates back to June 1992 at CH and to June 1998 at AH. We will refer to the number of respiratory syndrome cases in either hospitals as CH and AH data.

The four time series are displayed in Figure 1. The plot shows the apparent regularity of the four time series that are characterized by evident winter peaks during the influenza seasons, and more modest peaks in the fall seasons. Typically, the fall peaks correspond to increased activity of respiratory infections that follow the opening of schools. An intriguing feature of these time series is that the peaks do not occur simultaneously but there is an evident order with CH data (···) peaking first, followed by AH (– · –) and ILI data (—) and then P&I mortality data (– –).

To ascertain the temporal order of the four time series, we analysed their pair-wise cross-correlations over 52 weeks. We use the notation $\rho_x(A, B)$ to denote the cross-correlation between two time series A and B , with B shifted back by x weeks when $x > 0$, and A shifted back by x weeks when $x < 0$. The cross-correlation plot in the left panel of Figure 2 suggests that CH data are the earliest indicator of influenza-like illness (—: $\rho_x(\text{CH}, \text{ILI}) > 0.5$ for $1 \leq x \leq 5$) and deaths in the community (– – –: $\rho_x(\text{CH}, \text{P\&I}) > 0.5$ for $2 \leq x \leq 8$). The large correlation between CH and ILI data for a 3 week lag suggests that the first effects of influenza emerge in CH data well before they become evident in the general population. CH data lead AH data by 1–5 weeks (···: $\rho_x(\text{CH}, \text{AH}) > 0.5$ for $1 \leq x \leq 5$) and the large correlations $\rho_1(\text{CH}, \text{AH}) = 0.64$ for a 1 week lag, and $\rho_{2,3}(\text{CH}, \text{AH}) = 0.63$ for a 2–3 week lag show that CH data can be used to predict a substantial proportion of the variability in AH data.

AH data are also able to predict influenza mortality but with a shorter lead time compared to CH data and they appear to provide only slightly earlier information than ILI data currently monitored through federal surveillance programs. This finding is further supported by the cross-correlation between AH data and ILI and P&I mortality data displayed in the right panel of Figure 2: AH data lead P&I data by 1–3 weeks (– – –: $\rho_x(\text{AH}, \text{P\&I}) > 0.5$ for $1 \leq x \leq 3$). Similarly, ILI data lead P&I mortality data by 1 week (· · ·: $\rho_x(\text{ILI}, \text{P\&I}) > 0.5$ for $0 \leq x \leq 3$) thus confirming that AH and ILI data contain similar information to predict P&I mortality. Compared to AH data, CH data show a smaller correlation with P&I deaths, possibly because the elderly are the most at risk of dying from influenza and pneumonia. Nonetheless, this analysis suggests that pediatric patients are the first group to show symptoms of respiratory infections, and that these early effects are seen in the ED.

3. DYNAMIC BAYESIAN NETWORKS

By examining a pair of time series at a time, cross-correlation analysis may overlook the overall temporal dependency structure by introducing spurious or masking important dependencies. To better identify the dynamic structure among the four time series, we built a model based on a dynamic Bayesian network [19].

A dynamic Bayesian network is described by a directed acyclic graph in which nodes represent stochastic variables and arrows represent temporal dependencies that are quantified by probability distributions. Following the direction of the arrows, a node Y_1 with

an incoming arrow from a node Y_2 is called a child of Y_2 , and Y_2 is called a parent of Y_1 . We assume that the probability distributions of the temporal dependencies are time invariant, so that the directed acyclic graph of a dynamic Bayesian network represents only the time transitions that are necessary and sufficient to reconstruct the overall temporal process. Figure 3 shows the directed acyclic graph of a dynamic Bayesian network with three variables. The subscript of each node denotes the time lag, so that the arrows from the nodes $Y_{2(t-1)}$ and $Y_{1(t-1)}$ to the node $Y_{1(t)}$ describe the dependency of the probability distribution of the variable Y_1 at time t on the value of Y_1 and Y_2 at time $t - 1$. Similarly, the directed acyclic graph shows that the probability distribution of the variable Y_2 at time t is a function of the value of Y_1 and Y_2 at time $t - 1$. It is worth noting that a dynamic Bayesian network is not restricted to represent temporal dependency of order 1. For example the probability distribution of the variable Y_3 at time t depends on the value of the variable at time $t - 1$ as well as the value of the variable Y_2 at time $t - 2$.

The topology of the directed acyclic graph describes Markovian properties of the variables that allow us to decompose the network into related modules. The local Markov property asserts that a node Y is independent of its nondescendant nodes, given the parent nodes [20], and leads to a direct factorization of the transition distribution of the network variables $Y_1, \dots, Y_k$ into the product of the conditional distribution of each variable given its parents Π_i :

$$p(y_{1(t)}, \dots, y_{k(t)} | h_{t-1}) = \prod_i p(y_{i(t)} | \pi_i)$$

where h_{t-1} is the history of the process through time $t - 1$, $y_{1(t)}, \dots, y_{k(t)}$ denotes the value of the network variables at time t , and π_i are the parents of Y_i . For example, we can fully describe the transition distribution of the three variables in Figure 3 given the past history

$$h_{t-1} = y_{1(t-1)}, \dots, y_{1(0)}, y_{2(t-1)}, \dots, y_{2(0)}, y_{3(t-1)}, \dots, y_{3(0)}$$

by using only the three transition distributions:

$$\begin{aligned} p(y_{1(t)} | h_{t-1}) &= p(y_{1(t)} | y_{1(t-1)}, y_{2(t-1)}) \\ p(y_{2(t)} | h_{t-1}) &= p(y_{2(t)} | y_{1(t-1)}, y_{2(t-1)}) \\ p(y_{3(t)} | h_{t-1}) &= p(y_{3(t)} | y_{3(t-1)}, y_{2(t-2)}) \end{aligned}$$

Assuming that these probability distributions are time invariant, they are sufficient to compute the probability that a process that starts from known values $y_{1(0)}$, $y_{2(0)}$, and $y_{3(0)}$ evolves into $y_{1(T)}$, $y_{2(T)}$, $y_{3(T)}$, or that a process with values $y_{1(T)}$, $y_{2(T)}$, $y_{3(T)}$ at time T started from the initial states $y_{1(0)}$, $y_{2(0)}$, $y_{3(0)}$. Exact algorithms exist to perform this inference when the network variables are all discrete, all continuous and modelled with Gaussian distributions, or the network topology is constrained to particular structures [21–23]. For general network topologies and nonstandard distributions, we need to resort to stochastic simulation [24]. Among the several stochastic simulation methods currently available, Gibbs sampling [25, 26] is particularly appropriate for Bayesian network reasoning because of its

ability to leverage on the decomposition of joint multivariate distributions induced by the directed acyclic graph to improve computational efficiency.

4. MODEL BUILDING AND VALIDATION

We leveraged on the local Markov property to build our dynamic Bayesian network by modules. For each variable, we identified the best multiple regression model by selecting the significant predictors in the set defined by the variable itself observed at $t - 1, \dots, t - 10$ week lags, the other variables observed over the same temporal range, and two auxiliary variables that model the seasonal pattern of the time series. We followed suggestions in Reference [27] and limited attention to additive models. Because of their significant skewness, we modelled the number of P&I deaths by a Poisson log-linear model and the number of patients with ILI symptoms with a log-normal distribution, while we modelled AH and CH data with normal distributions. We used backward step-wise regression to avoid overlooking some dependencies, and the Akaike information criterion [28] to build initial models for each time series, and then selected only significant predictors. Using the BIC criterion [28], produced the same results. This selection procedure returned the four models in Figure 4, one for each of the four time series. These models are described by the following set of equations:

$$\begin{aligned}\mu_{CH_t} &= 64.8846 + 0.5827CH_{t-1} + 24.7580x_{1t} \\ \mu_{AH_t} &= 17.6199 + 0.1755CH_{t-1} + 0.5528AH_{t-1} \\ \mu_{\log(ILI_t)} &= 0.6919\log(ILI_{t-1}) + 0.0056CH_{t-2} + 0.1502x_{1t} - 0.1267x_{2t} \\ \log(\mu_{P\&I_t}) &= 3.5542 + 0.0045P\&I_{t-1} + 0.0031AH_{t-1} + 0.0006ILI_{t-1} - 0.0014CH_{t-3} + 0.1153x_{1t} + 0.1370x_{2t}\end{aligned}$$

where the variables x_1 and x_2 are periodical functions defined as $x_1 = \sin(2\pi t/52)$ and $x_2 = \cos(2\pi t/52)$ and the probability distributions of the variables CH, AH, ILI and P&I at time t , given their parent nodes, are specified as

$$\begin{aligned}CH_t | CH_{t-1} &\sim N(\mu_{CH_t}, \sigma_{CH}^2) \\ AH_t | AH_{t-1}, CH_{t-1} &\sim N(\mu_{AH_t}, \sigma_{AH}^2) \\ \log(ILI_t) | \log(ILI_{t-1}), CH_{t-2} &\sim N(\mu_{\log(ILI_t)}, \sigma_{ILI}^2) \\ P\&I_t | P\&I_{t-1}, AH_{t-1}, ILI_{t-1}, CH_{t-3} &\sim P(\log(\mu_{P\&I_t}))\end{aligned}$$

where $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 , and $P(\mu)$ denotes the Poisson distribution with mean μ . Table I gives further details of the parameter estimates.

The overall dependency model was built joining the four regression models by their common predictors using standard path analysis [29], and is reported in Figure 5. Each node in the directed graphs represents one of the four variables at the time point defined by the subscript and the arrows define the temporal dependencies that are quantified by the equation above.

We validated this model by examining the posterior probability of each selected temporal model versus the others. We first approximated the logarithm of the marginal likelihood of each regression model m , say l_m by the Akaike information criterion [30]. Assuming uniform prior probabilities on the set of possible models, we then computed the posterior probability of each model m was $p_m = \exp(l_m)/k$, with the normalizing constant $k = \sum_m p_m$. The dependency structure of the selected model is very strong as shown by the plot in Figure 6 that depicts the posterior probability of 400 models describing the dependency of the number of P&I deaths on CH and AH data, for different time lags. The surface peaks when the time lag is 3 weeks for CH data and 1 week for AH data and, compared to the other temporal dependencies, there is very strong evidence that this model gives the best fit (posterior probability almost 1).

The temporal dependency structure in Figure 5 confirms some of the results suggested by the cross-correlation analysis. CH data are predictive of AH data with a 1 week lag, they are predictive of ILI data with a 2 week lag, and of the number of P&I deaths with a 3 week lag. This result confirms that children presenting to the ED with respiratory syndromes provide earlier signals of the spread of influenza epidemics compared to ILI data currently monitored through federal surveillance programs. It is worth noting that neither AH data, nor ILI data or P&I deaths are predictive of CH data, thus confirming that pediatric patients presenting to the ED are the first group to show symptoms of influenza. Both AH and ILI data are predictive of P&I deaths with a 1 week lag. Note that AH data are not predictive of ILI data, once we condition on CH data observed 2 weeks earlier, thus confirming that CH data contain earlier and stronger signals of influenza morbidity than AH data. Note also the negative effect of CH_{t-3} in the regression model fitted to P&I data, that may reflect the fact that children's data contributes to the morbidity of the disease but less to mortality.

The model in Figure 5 suggests that the spread of respiratory infections, and particularly influenza, in a community begins in children who report to the ED with respiratory syndromes. This observation is confirmed by the absence of predictive signal of ILI and P&I mortality data on CH data. Further support for this conjecture is provided by using the dynamic model for prediction of influenza morbidity and mortality a few weeks in advance and by quantifying the predictive effect of pediatric patients. For this assessment, we created competitive models by iteratively dropping each of the predictors from the selected model. We trained the competitive models on the data from June 1998 to the end of October 1999 and then computed their 2-week ahead predictions of ILI and P&I data from November 1999 to October 2000. We chose this test set as the virus circulating during the influenza season 1999–2000 had significant morbidity. After each prediction, we iteratively updated the model parameters using all data seen. Because we use non-Gaussian distributions, we had to resort to stochastic computations for the predictive inference. We used the implementation of Gibbs sampling in Winbugs 1.4 to compute the forecasts [31].

Figure 7 reports the BUGS code for this '2-week-ahead' forecast with the model in Figure 5. The model describes the dependency structure between the four data streams assuming that we have initial values at week $t - 2$ and we wish to forecast P&I mortality data, and ILI morbidity data at week t . To initialize the dynamic network, we need data for the parent nodes of the four variables at time $t - 2$. Therefore, we provide CH data at week $t - 4$, and

AH data, ILI data and P&I data at week $t - 2$ to initialize the conditional distribution of P&I at week $t - 1$. Similarly, we use CH data at week $t - 2$ to initialize the conditional distribution of CH at week $t - 1$; CH data at week $t - 3$ together with AH data at week $t - 2$ will initialize the conditional distribution of AH at week $t - 1$, while CH data at week $t - 3$, and ILI data at week $t - 2$ will initialize the conditional distribution of ILI at week $t - 1$. The parameters that specify the transition distributions, as well as the time point t and the values of the covariates x_1 and x_2 are provided via the data file, and are iteratively updated for each weekly forecast.

Plate 2 shows the number of patients with ILI symptoms recorded in MA between November 1999 and October 2000 by the sentinel physicians in the CDC influenza surveillance program (black line). The red line depicts the 2-week ahead prediction computed by the dynamic model, while the green line shows the 2-week ahead prediction when CH data are removed from the model. The pale blue line is the 2-week ahead prediction when ILI data are removed from the model. Besides the very close match between true and predicted values provided by our dynamic model, it is apparent that CH data alone are sufficient to forecast the dynamics of ILI data, while the model in which CH data are removed suffers of a delay. Statistical analysis of the forecasting error of our model showed no significant difference between observed and predicted values (t -test = 0.79, p -value = 0.21), while removing ILI data increases the average forecasting error that is still not significant (t -test = 1.2394, p -value = 0.1104), and removing CH data makes the average forecasting error significantly different from 0 (t -test = 3.6795, p -value = 0.0002818). Similar tests were conducted to assess the predictive effect of both CH and AH data on influenza deaths and confirmed the early signature of syndromic data. For example, using only ILI data to predict PI increases the average prediction error by about 12 per cent. Using AH data only increases the average prediction error by 9 per cent and using CH data increases the average prediction error by 4 per cent.

5. DISCUSSION

The increasing emphasis on biosurveillance, motivated by recent epidemic and bioterrorism events, has stimulated the electronic collection of different sources of data that are supposed to contain early signs of disease outbreaks. One of the major challenges that researchers in biosurveillance face at the moment is the integration of these data streams into models that can indeed alert public health officials of impending disease outbreaks. To achieve these objectives, we need models that can integrate information from different data streams and that can be used for accurate forecasting.

In agreement with the work of other investigators [32], our analysis shows that dynamic Bayesian networks provide a simple but effective modelling tool for the integration of several time series data into a dynamic model. These results are in agreement with the work of other investigators. The modularity of dynamic Bayesian networks makes it possible to learn them from data comprising several variables using standard multiple regression techniques. Furthermore, the availability of software for stochastic computations with directed graphical models makes probabilistic inference with dynamic Bayesian networks very efficient, and removes the need to impose restrictive assumptions on the dependency

structure or probability distribution of the network variables such as those imposed on Gaussian networks [27]. Although in this paper we have applied dynamic Bayesian networks to model four data streams, the approach can be extended to include other types of data that can be informative for influenza surveillance, such as school/work absenteeism, or syndromic data collected at other EDs [33].

An issue created by the analysis of several time series is the fact that different data streams may be measured in different time intervals. In our application, syndromic data are measured daily, while ILI and P&I data are measured weekly. For convenience, we analysed the variations in the weekly frequency of syndromic data, with a potential loss of information and timeliness. The extension of this methodology to the analysis of time series with different time frequency is an open issue that deserves further investigation.

Another open issue is the effect of the time-invariance assumption. The predictive accuracy in our evaluation provides evidence of the good performance of our model and, therefore, of the fact that the time-invariance assumption is reasonable. In general application, the effect of this assumption should be assessed by using for example residual plots to investigate temporal patterns of the residuals.

From the epidemiology point of view, our analysis identifies in the number of pediatric patients who present to the ED with respiratory syndromes an early quantitative indicator of influenza epidemics. Compared to hospitalization rates [34] or school absenteeism [35] used for monitoring or detecting ongoing influenza outbreaks, pediatric patients with respiratory syndromes are able to predict upcoming epidemics 2–3 weeks in advance. The pediatric sentinel signal is early enough to allow ample time to step up vaccination [36, 37] and implement other effective control measures in the community to reduce the burden of illness and mortality.

Acknowledgments

Contract/grant sponsor: Alfred P. Sloan Foundation; contract/grant number: 2002-12-1

Contract/grant sponsor: NIAID; contract/grant number: U19 AI62627

This work was supported by the Alfred P. Sloan Foundation (2002-12-1) and by a pilot grant originating from the Blood Center of Wisconsin, via NIAID grant U19 AI62627.

References

1. Arias E, Smith BL. Deaths: preliminary data for 2001. *National Vital Statistics Reports*. 2003; 51:1–45.
2. Simenson L, Fukuda K, Schonberg LB, Cox NJ. The impact of influenza epidemics on hospitalizations. *Journal of Infectious Diseases*. 2000; 181:831–837. [PubMed: 10720501]
3. Bridges CB, Harper SA, Fukuda K, Uyeki TM, Cox NJ, Singleton JA. Advisory committee on immunization practices. Prevention and control of influenza. Recommendations of the advisory committee on immunization practices (ACIP). *MMWR Recommendation and Reports*. 2003; 52:1–34.
4. Schoenbaum SC. Economic impact of influenza. The individual's perspective. *American Journal of Medicine*. 1987; 82:26–30. [PubMed: 3109239]

5. Meltzer MI, Cox NJ, Fukuda K. The economic impact of pandemic influenza in the united states: priorities for intervention. *Emerging Infectious Diseases*. 1999; 5:659–671. [PubMed: 10511522]
6. Patriarca PA, Cox NJ. Influenza pandemic preparedness plan for the united states. *Journal of Infectious Diseases*. 1997; 176(Suppl 1):S4–S7. [PubMed: 9240686]
7. Liand KS, Guan Y, Wang J, Smith GJ, Xu KM, Duan L, Rahardjo AP, Puthavathana P, Buranathai C, Nguyen TD, Estoepangestie AT, Chaisingh A, Auewarakul P, Long HT, Hanh NT, Webby RJ, Poon LL, Chen H, Shortridge KF, Yuen KY, Webster RG, Peiris JS. Genesis of a highly pathogenic and potentially pandemic h5n1 influenza virus in eastern asia. *Nature*. 2004; 430:209–213. [PubMed: 15241415]
8. Rath, T.; Carreras, M.; Sebastiani, P. Automated detection of influenza epidemics with Hidden Markov Models. In: Berthold, MR.; Lenz, HJ.; Bradley, E.; Kruse, R.; Borgelt, C., editors. *Advances in Intelligent Data Analysis V; 5th International Symposium on Intelligent Data Analysis, IDA 2003; Berlin, Germany. August 28–30, 2003; New York: Springer; 2003. p. 521-531. Proceedings*
9. Serfling RE. Methods for current statistical analysis of excess pneumonia-influenza deaths. *Public Health Reports*. 1963; 78:494–506. [PubMed: 19316455]
10. Mandl K, Overhage JM, Wagner MM, Lober WL, Sebastiani P, Mostashari F, Pavlin JA, Gesteland PH, Treadwell T, Koski E, Hutwagner L, Buckeridge DL, Aller RD, Grannis S. Implementing syndromic surveillance: a practical guide informed by the early experience. *Journal of the American Medical Informatics Association*. 2004
11. Goldenberg A, Shmueli G, Caruana RA, Fienberg SE. Early statistical detection of anthrax outbreaks by tracking over-the-counter medication sales. *Proceedings of the National Academy of Sciences of the United States of America*. 2002; 99:5237–5240. [PubMed: 11959973]
12. Gesteland, PH.; Wagner, MM.; Chapman, WW.; Espino, JU.; Tsui, FC.; Gardner, RM.; Rolfs, RT.; Dato, V.; James, BC.; Haug, PJ. Rapid deployment of an electronic disease surveillance system in the state of Utah for the 2002 olympic winter games. *Proceedings of the Annual AMIA Fall Symposium; 2002; p. 285-289.*
13. Reis BY, Pagano M, Mandl KD. Using temporal context to improve biosurveillance. *Proceedings of the National Academy of Sciences of the United States of America*. 2003; 100:1961–1965. [PubMed: 12574522]
14. Lazarus R, Kleinman K, Dashevsky I, Adams C, Kludt P, DeMaria A, Platt R. Use of automated ambulatory-care encounter records for detection of acute illness clusters, including potential bioterrorism events. *Emerging Infectious Diseases*. 2002; 8:753–760. [PubMed: 12141958]
15. Lazarus R, Vercelli D, Palmer LJ, Klimecki WJ, Silverman EK, Richter B, Riva A, Ramoni MF, Martinez FD, Weiss ST, Kwiatkowski DJ. SNPs in innate immunity genes: abundant variation and potential role in complex human disease. *Immunological Reviews*. 2003; 190:9–25. [PubMed: 12493003]
16. Tsui FC, Espino JU, Dato VM, Gesteland PH, Hutman J, Wagner MM. Technical description of RODS: a real-time public health surveillance system. *Journal of the American Medical Informatics Association*. 2003; 10:399–408. [PubMed: 12807803]
17. Wang L, Ramoni MF, Mandl KD, Sebastiani P. Factors affecting the performance of a syndromic surveillance system. *Artificial Intelligence in Medicine*. 2005; 34:269–278.
18. Reis, BY.; Mandl, K. Integrating syndromic surveillance data across multiple locations: effects on outbreak detection performance. *Proceedings of the Annual AMIA Fall Symposium; 2003; p. 549-553.*
19. Russell, S.; Norvig, P. *Artificial Intelligence: A Modern Approach*. 2. Prentice-Hall; Englewood Cliffs, NJ: 2003.
20. Lauritzen, SL. *Graphical Models*. Oxford University Press; Oxford, U.K: 1996.
21. Castillo, E.; Gutierrez, JM.; Hadi, AS. *Expert Systems and Probabilistic Network Models*. Springer; New York, NY: 1997.
22. Lauritzen SL, Spiegelhalter DJ. Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society, Series B*. 1988; 50:157–224.

23. Pearl, J. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann; San Francisco, CA: 1988.
24. Cheng J, Druzdzel M. AIS-BN: an adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research*. 2000; 13:155–188.
25. Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1984; 6:721–741. [PubMed: 22499653]
26. Thomas, A.; Spiegelhalter, DJ.; Gilks, WR. Bugs: a program to perform Bayesian inference using Gibbs sampling. In: Bernardo, J.; Berger, J.; Dawid, AP.; Smith, AFM., editors. *Bayesian Statistics*. Vol. 4. Oxford University Press; Oxford, U.K: 1992. p. 837-842.
27. Cowell, RG.; Dawid, AP.; Lauritzen, SL.; Spiegelhalter, DJ. *Probabilistic Networks and Expert Systems*. Springer; New York, NY: 1999.
28. Hastie, T.; Tibshirani, R.; Friedman, J. *The Elements of Statistical Learning*. Springer; New York: 2001.
29. Hand, DJ.; Mannila, H.; Smyth, P. *Principles of Data Mining*. MIT Press; Cambridge, MA: 2001.
30. Kass RE, Raftery A. Bayes factors. *Journal of the American Statistical Association*. 1995; 90:773–795.
31. Spiegelhalter, DJ.; Thomas, A.; Best, NG. Computation on Bayesian graphical models (with discussion). In: Bernardo, JM.; Berger, J.; Dawid, AP.; Smith, AFM., editors. *Bayesian Statistics*. Vol. 5. Oxford University Press; Oxford, U.K: 1996. p. 407-425.
32. Wong WK, Cooper G, Dash D, Levander J, Dowling J, Hogan W, Wagner M. Use of multiple data streams to conduct Bayesian biologic surveillance. *MMWR*. 2005; 54(Suppl):63–69. [PubMed: 16177695]
33. Fienberg SE, Shmueli G. Statistical issues and challenges associated with rapid detection of bio-terrorist attacks. *Statistics in Medicine*. 2005; 24:513–529. [PubMed: 15678405]
34. Glezen WP, Decker M, Joseph SW, Mercready RG. Acute respiratory disease associated with influenza epidemics in Houston, 1981–1983. *Journal of Infectious Diseases*. 1987; 155:1119–1126. [PubMed: 3572033]
35. Lenaway DD, Ambler A. Evaluation of a school-based influenza surveillance system. *Public Health Reports*. 1995; 110:333–337. [PubMed: 7610226]
36. Nichol KL, Lind A, Margolis KL, Murdoch M, McFadden R, Hauge M, Magnan S, Drake M. The effectiveness of vaccination against influenza in healthy, working adults. *New England Journal of Medicine*. 1995; 333:889–893. [PubMed: 7666874]
37. Rennels MB, Meissner HC. Committee on infectious diseases. Technical Report: Reduction of the influenza burden in children. *Pediatrics*. 2002; 110:e80. [PubMed: 12456947]

Figure 1.

Plots of the weekly number of respiratory syndrome cases at CH ($\cdot \cdot \cdot$) and AH ($- \cdot -$), patients with ILI symptoms ($—$) and the number of deaths from P&I ($- -$). The x -axis reports time in weeks from June 1998 to the end of September 2002.

Figure 2.

Left: Cross-correlation (y-axis) between the weekly number of CH and ILI data, deaths from P&I, and the weekly number of AH data for different week lags (x-axis). For each ordered pair of time series, the cross-correlation for x weeks lag represents the correlation between the first time series shifted back by x weeks when $x > 0$, while the second time series is shifted back by x weeks when $x < 0$. Right: Cross-correlation between the weekly number of AH data, ILI cases and deaths from P&I. The dotted line ($\cdot \cdot \cdot$) displays the cross-correlation between the series of ILI cases and deaths from P&I.

Figure 3.

A directed acyclic graph representing the temporal dependency of three categorical variables with states – and +. The conditional probability table shows the transition probabilities that quantify the dependency of Y_3 on its parent nodes.

Figure 4.

Regression models for each of the four time series (CH, AH, P&I and ILI). All nodes in the graphs represent stochastic variables and arrows represent temporal dependencies quantified by probability distributions. An arrow from a node A to a node B means that A leads B by the number of weeks identified by the subscript. For example, the graph (a) shows that the number of pediatric patients seen at the ED with respiratory syndrome at week $t - 1$ (CH_{t-1}) leads the number of pediatric patients at week t (CH_t). This is a simple autoregressive model of order 1. The graph (b) shows that the model describing the dynamics of P&I at week t has an autoregressive component ($P\&I_{t-1}$) as well as the three predictors CH_{t-3} , AH_{t-1} , and ILI_{t-1} . Therefore, CH data lead P&I data by 3 weeks, while both AH and ILI data lead the number of pneumonia and influenza deaths by 1 week only. Similarly, the graphs (c) and (d) show that CH data lead AH data by 1 week, and the number of patients with ILI symptoms by 2 weeks.

Figure 5.
The overall dynamic Bayesian network describing the interplay between the four data streams.

Figure 6.

Posterior probability (z -axis) for the models describing the dependency of the number of deaths for pneumonia and influenza on the number of pediatric and adult patients with respiratory syndromes for varying lags in weeks (x and y axes). The probability is maximum when the number of pediatric patients with respiratory syndrome is measured with a three week lag and the number of adult patients with respiratory syndrome is measured with 1 week lag and, compared to the other models, there is very strong evidence that this model gives the best fit (posterior probability almost 1).

```

model
{
  ### initial values of the time series
  d.t.2 <- pi[t-2]      ### P&I at week t-2
  e.i.t.2 <- exp(ili[t-2]); ##### ILI at week t-2 ILI are in log-scale
  a.t.2 <- bid[t-2]      ##### AH data at week t-2
  c.t.4 <- cc[t-4]; c.t.3<-cc[t-3]; ##### CH data at week t-4 and t-3

  ##### P&I
  d.t.1 ~ dpois( d.mean.1)    ### P&I at week t-1
  log(d.mean.1)<- int.d+theta.d*d.t.2+ theta.a*a.t.2+ theta.c*c.t.4+
                theta.i*e.i.t.2+theta.x.1*x.1[t-1]+theta.x.2*x.2[t-1]
  d.t ~ dpois( d.mean)      ### P&I at week t
  log(d.mean) <- int.d+theta.d*d.t.1+ theta.a*a.t.1+theta.c*c.t.3+
                theta.i*e.i.t.1+theta.x.1*x.1[t]+theta.x.2*x.2[t]

  ##### ILI
  i.t.2<-ili[t-2]
  e.i.t.1 ~ dlnorm( i.mean.1,i.tau)  ### ILI at week t-1
  i.mean.1 <- eta.i*i.t.2+eta.c*c.t.3+eta.x.1*x.1[t-1]
  i.t.1<-log(e.i.t.1)
  e.i.t ~ dlnorm( i.mean,i.tau)  ### ILI at week t
  i.mean <- eta.i*i.t.1+eta.c*c.t.2+eta.x.1*x.1[t]

  ##### AH
  a.t.1~ dnorm( a.mean.1, a.tau) ### AH at week t-1
  a.mean.1 <- int.a + beta.a*a.t.2+ beta.c*c.t.2
  a.t ~ dnorm( a.mean, a.tau)  ### AH at week t
  a.mean <- int.a + beta.a*a.t.1+ beta.c*c.t.1

  ##### CH
  c.t.1 ~ dnorm( c.mean.1, c.tau)    ### CH at week t-1
  c.mean.1 <- int.c + gamma.c*c.t.2+gamma.x.1*x.1[t-1]
  c.t ~ dnorm( c.mean, c.tau)      ### CH at week t
  c.mean <- int.c + gamma.c*c.t.1+gamma.x.1*x.1[t]
}

```

Figure 7.

BUGS code used for the 2-week-ahead forecast. The parameters are updated at each iteration.

Plate 1.

Baseline model and epidemic threshold of P&I and ILI data used for detection and monitoring of influenza epidemics. The left image reports the P&I data from January 2000 to the end of 2004. The blue lines are the season baseline and epidemic threshold computed using cyclic regression. The red line depicts the percentage of P&I deaths. The right image plots, in red, the percentage of visits for ILI reported by sentinel providers for the influenza season 2003–2004 (week 40 of 2003–week 20 of 2004). The green and maroon lines are the ILI data for the influenza seasons 1999–2000 and 2002–2003. The 1999–2000 influenza season was characterized by a severe mortality and morbidity, while no influenza epidemic was detected for the season 2002–2003. The dashed line is the national epidemic threshold.

Plate 2.

Two-step-ahead prediction of ILI cases in Massachusetts between October 1999 and October 2000: the year of a major influenza epidemic. The black continuous line depicts the observed number of patients with ILI symptom (in logarithmic scale) between November 1999 and the end of October 2000. The red dotted line depicts the values predicted by the model using only the data available until 2 weeks before. The green dashed line shows the values predicted by the model when the past ILI data are ignored, and the blue dotted–dashed line depicts the values predicted by the same model using only the volume of pediatric patients with respiratory syndromes. The purple continuous line depicts the prediction based on the sinusoidal components only. The synchrony between the lines in black and blue shows that the signal provided by the pediatric respiratory syndrome data is sufficient to reconstruct the dynamics of influenza and to identify peaks of the epidemic with 2 weeks anticipation. The x -axis reports the week beginning with November 1999 till the end of October 2000.

Table I

Estimates and standard errors of the regression parameters for the four models fitted to CH, AH, ILI and P&I data when all the data are used to induce the network.

	Estimate	Std. error
<i>Regression model for P&I</i>		
(Intercept)	3.5542	0.0656
$P\&I_{t-1}$	0.0045	0.0007
AH_{t-1}	0.0031	0.0005
ILI_{t-1}	0.0006	0.0003
CH_{t-3}	-0.0014	0.0004
x_1	0.1153	0.0234
x_2	0.1370	0.0189
<i>Regression model for AH</i>		
(Intercept)	17.6199	4.4835
AH_{t-1}	0.5528	0.0530
CH_{t-1}	0.1755	0.0294
<i>Regression model for log(ILI)</i>		
$\log(ILI_{t-1})$	0.6919	0.0443
CH_{t-2}	0.0056	0.0008
x_1	0.1502	0.0579
x_2	-0.1267	0.0542
<i>Regression model for CH</i>		
(Intercept)	64.8846	7.9730
CH_{t-1}	0.5827	0.0505
x_1	24.7580	3.5840

The model fitted to P&I data is a log-linear model, so that the regression coefficients parameterize the logarithm of the mean. The model fitted to ILI data is in logarithmic scale. The variables x_1 and x_2 are periodical functions defined as $x_1 = \sin(2\pi t/52)$ and $x_2 = \cos(2\pi t/52)$. Note the negative effect of CH_{t-3} in the regression model fitted to P&I data, that may reflect the fact that children's data contributes to the morbidity of the disease but less to mortality.