

# A Decision-theoretic Model of Disease Surveillance and Control and a Prototype Implementation for the Disease Influenza

Michael Wagner,<sup>1</sup> Gregory Cooper,<sup>1</sup> Fuchiang Tsui,<sup>1</sup> Jeremy U. Espino,<sup>1</sup> Hendrik Harkema,<sup>1</sup> John Levander,<sup>1</sup> Ricardo Villamarin,<sup>1</sup> Nicholas Millett,<sup>1</sup> Shawn T. Brown,<sup>2,3</sup> Anthony Gallagher<sup>1</sup>

<sup>1</sup>Center for Advanced Study of Informatics in Public Health, Department of Biomedical Informatics, <sup>2</sup>Graduate School of Public Health, <sup>3</sup>Pittsburgh Supercomputing Center, University of Pittsburgh, USA

<sup>1</sup>mmw1@pitt.edu

**Abstract**— This paper first describes a decision-theoretic model of disease surveillance and control. It then describes a prototype system for influenza monitoring based on the model. The decision-theoretic model connects disparate work in epidemiological modelling and disease control under a uniform mathematical formulation. We expect that this model will stimulate new avenues of research in both fields.

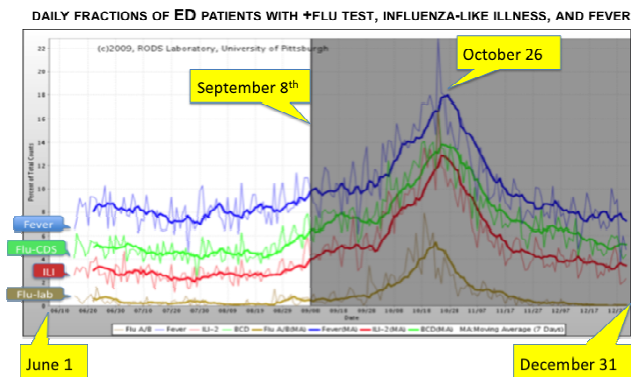
## I. INTRODUCTION

### A. The Problem of Selecting an Optimal Control Strategy for the 2009 H1N1 Influenza Outbreak

In early September 2009, U.S. public health officials were preparing for a second wave of the H1N1 influenza epidemic.

A key problem for the officials was their uncertainty about when the second wave would begin, when it would peak, and how many individuals would be infected. They were also uncertain about when the new H1N1 vaccine would become available and the quantity that would be produced over time.

A part of their uncertainty resulted from limitations in existing disease surveillance methods. For example, Fig. 1 shows the output of our Real-time Outbreak and Disease Surveillance (RODS) monitoring system [1, 2] for the period surrounding the H1N1 outbreak in Allegheny County, Pennsylvania (henceforth abbreviated as AC). The unshaded area in Fig. 1 shows the surveillance data available on September 8.



**Fig. 1.** Influenza surveillance report for Allegheny County, PA during the 2009 H1N1 epidemic. Shown are daily fractions (%) of ED patients. The thick lines are 7-day moving averages. The unshaded area shows the information available on Sept 8, 2009. (See also text)

Fig. 1 shows rising trends in all types of influenza surveillance data on September 8, 2009. The fractions of emergency department (ED) patients with definitive laboratory tests for influenza (Flu-lab), febrile illnesses (Fever), and two measures of influenza-like illnesses (ILI and Flu-CDS) are all increasing.

The shaded area in Fig.1 shows the subsequent time course of the epidemic in AC, which peaked in late October. The epidemic infected approximately 21% of the population [3]. Existing methods could not predict the subsequent time course from the data in September.

In addition to uncertainty about the epidemic and the vaccine supply, officials were also uncertain about the efficacy of non pharmaceutical control measures, such as school closing.

Despite the uncertainty about the epidemic, the availability of vaccine, and the efficacy of school closure, officials had to make decisions: Should schools be closed [4, 5]? Should vaccination be prioritized to certain groups? [3, 6]? Should adjuvants be used to increase the vaccine supply [7]?

Analysts addressed these decision problems with model-based analyses. In particular, they employed mathematical models of H1N1 epidemics to project the future course of the epidemic, measured as number of people sick, hospitalized, and other outcomes of importance such as number of ICU beds occupied. They ran simulations under different assumptions about both the epidemic and the control measures. The assumptions about the epidemic ranged over its timing, infectiousness, existing immunity, and severity of disease. The assumptions about control measures ranged over efficacy, supply constraints, administration capacity, cost of illness, and indirect economic impacts. The analysts ran thousands of simulations to understand the sensitivity of their results to the assumptions about the epidemic and the control strategies [8, 9].

### B. Epidemic Models

An *epidemic model* is a dynamic model of the transmission of a disease in a population. It is a type of population model.

The simplest epidemic model—the compartment model—is a state transition network that models the disease state of a population over time. The states (compartments) represent subpopulations such as those *Susceptible*, *Exposed*, *Infectious*, and *Recovered*. State transitions are defined by a set of differ-

ence or differential equations that model the dynamics of disease transmission [10].

Compartment models are a family of models that differ based on the disease states they can model, the transmission parameters, and the initialization of the compartments. For example, SEIR models, are a class of compartment models that share the characteristic of having compartments for the disease states *Susceptible*, *Exposed*, *Infectious*, and *Recovered*. Within the SEIR class are an infinite number of models.

An instance of a SEIR model can be fully specified by a small set of numbers. In particular, a SEIR model can be specified by (1) three of the four states, (2) the total population (the fourth state is the difference between the sum of the three states and the population), and (3) the set of parameters that define the transition probabilities. In our work, we specify the total population; the number of individuals in the S, I, and R compartments at time zero; and the parameters  $R_0$ , infectious period, and latent period. If a SEIR model is to be synchronized with a real epidemic, which is necessary for our purposes, the start date of the real epidemic is also necessary.

Compartment models can be extended to model the effects of some disease control strategies. A vaccination control measure, for example, can be modelled in the state transition functions by a transition that moves people directly from *susceptible* to the *recovered* state (immune) at a rate equal to the rate of effective vaccination of the population.

Compartment models can be extended to model sociodemographic differences in a population by stratifying the compartments (e.g., replacing compartments such as *susceptible* with compartments for *susceptible children* and *susceptible adults*). The most stratified epidemic simulations are agent-based, in which compartments are replaced by agents. Each agent can be in one of a set of mutually exclusive and exhaustive states such as *Susceptible*, *Exposed*, *Infectious*, or *Recovered*; the transitions between the states result from interactions with other agents, the environment, and the passage of time.

### C. The General Problem of Control Strategy Optimization

The decision problem illustrated by the example of 2009 H1N1 influenza is not limited to influenza. There are hundreds of biological agents that can cause epidemics [11]. As a second example, a suspected aerosol release of the organism *B. anthracis* would raise questions and decisions that are hard to address optimally at present. For example: What is the probability that a release occurred? What are the spatial distribution and exposure levels in the affected areas? What regions should be cordoned off to prevent further exposures? Are there sufficient antibiotics, hospital beds and respirators available or should a request be made to the President to mobilize the national strategic stockpile?

The ability for officials to control many epidemic diseases could potentially be improved by a more formal decision-theoretic approach to disease surveillance and control. Such an approach would be capable of representing the knowledge and uncertainties about the epidemics, control strategies, and the cost functions that decision makers are trying to optimize.

## II. A GENERAL DECISION-THEORETIC FORMULATION

Fig. 2 is a general decision-theoretic formulation of disease surveillance and control. It states that the optimal control strategy is the control strategy that maximizes the expected utility computed by model averaging over the outcomes predicted by all possible models of the epidemic under that control strategy.

This general model can be applied to epidemic models that are compartmental or agent-based. The only requirement is that there is a utility function over the outputs of interest, such as number of individuals who became sick, of the model; and a method for computing the probability of the epidemic model, given the available surveillance data.

We note that analysts who were working during the 2009 H1N1 outbreak were also solving an optimization problem; however, they did not model average over the probability of the epidemic models under consideration, which we denote as  $P(M_{Pop})$ , where the subscript *Pop* indicates that the model is of a population of interest. Their analytic method was to compute the expected number of sick under an expert-determined base-case (“most likely”) scenario for the epidemic and control measures, and conduct sensitivity analyses around the base-case. Thus, the optimization was local and the search method heuristic.

$$\arg \max_{CS} \sum_{M_{Pop}} P(M_{Pop} | E) \times U(M_{Pop, CS})$$

Available control strategies    Epidemic models    Disease surveillance    Economic model    Outcomes

Fig. 2 Decision-theoretic formulation of the problem of disease surveillance and control

## III. PROTOTYPE IMPLEMENTATION

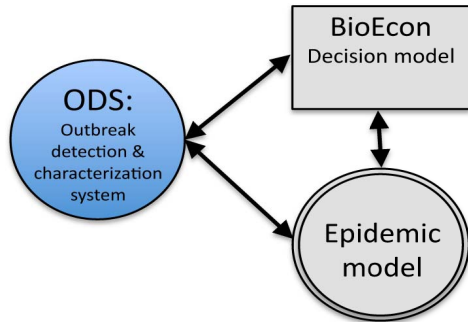
We developed a prototype implementation of the general decision-theoretic formulation discussed in the previous section for the influenza monitoring and control. We have deployed this system in AC. In the prototype, we use a SEIR influenza model that is capable of modelling a vaccine intervention, although our approach is not limited to SEIR models or vaccine interventions.

The system comprises three major components, as depicted in Fig. 3. The Outbreak Detection and Characterization System (ODS) identifies likely SEIR models, given surveillance data obtained from electronic medical records. It computes and outputs the distribution  $P(M_{Pop} | E)$  in Fig. 2, where  $E$  is the disease surveillance data available about individuals in the population.

The BioEcon decision-analytic component constructs and solves decision models (trees) representing alternative control strategies. It solves the expression in Fig. 2.

Both BioEcon and ODS use the third component, the epidemic model, but in different ways. BioEcon uses the epi-

demographic model to project the effects of a control measure into the future to compute the outcomes required by its utility function (labelled ‘Economic models’ in Fig. 2). In contrast, ODS uses the epidemic model as the starting point for constructing a probabilistic model of the relationship between population incidence of disease and observable disease surveillance data.



**Fig. 3** High-level schematic of the implemented influenza monitoring system. The processing begins with ODS obtaining prior probabilities of model parameters (e.g.,  $R_0$ ) from BioEcon, which may obtain these distributions from an epidemiologist end-user. ODS samples from these distributions many times to obtain possible SEIR models to evaluate and calls the epidemic model simulator with the sampled parameters to obtain an incidence curve for the epidemic in AC, which is its starting point for computing  $P(E | M_{Pop})$ . When the ODS algorithm completes these computations, it passes the (large) set of probable epidemic models with their probabilities to BioEcon. BioEcon then constructs a decision model for one or more user-selected control strategies, as in Fig. 5. To evaluate the decision model, BioEcon calls the epidemic model simulator for each model provided to it by ODS for each control strategy. BioEcon computes the expected utility of each control strategy by model averaging over the utilities it computes for the set of probable epidemic models provided by ODS when run under each control strategy.

#### A. ODS

The function of ODS is to compute the distribution  $P(M_{Pop}|E)$ , where  $M_{Pop}$  is a variable that takes as its values a large number of SEIR models. Each SEIR model is represented as a vector of seven SEIR model parameters—a set of parameters that uniquely identify a SEIR model.

ODS takes three inputs: (1) the likelihoods of influenza for every patient that is seen in monitored EDs. These likelihoods are provided by a Bayesian case detection system described in [1]; (2) The prior probability of influenza; and (3) prior probability distributions over the seven input parameters for the SEIR epidemic model, which we set to uniform over a range defined by the literature or expert knowledge.

ODS then samples from distributions representing the input parameters to the SEIR epidemic model. In particular, ODS samples from distributions for *infectious period*, *latent period*,  $R_0$ , *initial number infected*, and *start date*. ODS computes the posterior probability of each sampled model to obtain the distribution  $P(M_{Pop}|E)$ . ODS computes these probabilities by Bayesian inversion from a model that predicts the observed surveillance data in EDs from the *population* incidence in AC as predicted by the sampled  $M_{Pop}$ . A special case SEIR model is  $M_0$ , a *no-outbreak SEIR model*, which ODS can use to compute the probability that an influenza outbreak exists as  $1 - P(M_0|E)$ .

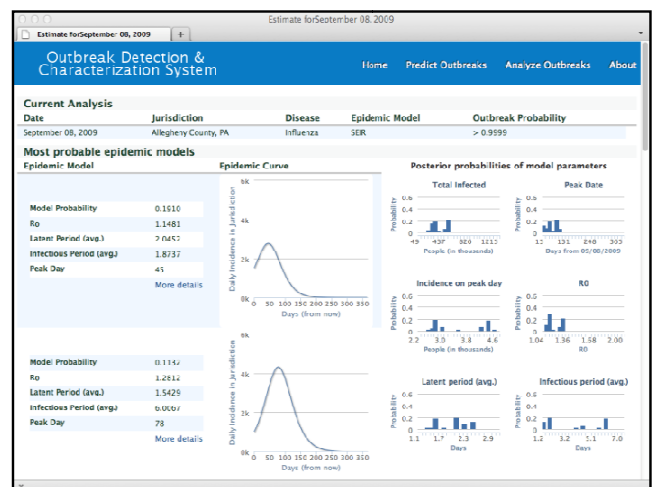
More formally, let  $M_{Pop}$  denote an epidemic model of the entire population in a region that is being monitored for an outbreak of disease. In our current application,  $M_{Pop}$  is a SEIR model. We would like to infer a distribution over such models given evidence about patients who seek care at emergency departments in the region. Let  $M_{ED}$  denote the disease states of all the ED patients during the monitoring period. Let  $E$  designate the clinical evidence that is available about the ED patients during the monitoring period.

At a high level, ODS is based on the following equation, which is an instance of Bayes’ theorem:

$$P(M_{Pop} | E) = \frac{\sum_{M_{ED}} P(E | M_{ED}) \cdot P(M_{ED} | M_{Pop}) \cdot P(M_{Pop})}{\int_{M_{Pop}} \sum_{M_{ED}} P(E | M_{ED}) \cdot P(M_{ED} | M_{Pop}) \cdot P(M_{Pop}) dM_{Pop}}$$

The sum is taken over all possible disease states of all the ED patients being monitored. The number of terms in the sum is therefore very large; however, we are able to take advantage of some basic mathematical techniques, such as the application of the binomial distribution, to compute the sum efficiently. ODS approximates the integral in the equation by sampling  $M_{Pop}$ , which leads to the integral becoming a sum. The terms  $P(M_{ED} | M_{Pop})$  and  $P(E | M_{ED})$  represent key modelling components of ODS. The term  $P(M_{Pop})$  represents a prior probability distribution over the parameters in a SEIR model of the population. The independence assumption in the equation is that  $P(E | M_{ED}, M_{Pop}) = P(E | M_{ED})$ , which expresses that in predicting ED patient data, knowledge of the disease status of the population at large is irrelevant, once we have knowledge of the disease status of the ED patients.

Fig. 4 (left) shows the two most probable models computed by ODS using the ED surveillance data available up through September 7, 2009.



**Fig. 4 ODS.** This screen shows the SEIR influenza models (only two most probable out of 100000 sampled), given the likelihoods of influenza from CDS, for all Allegheny County patients seen in monitored EDs from May through September 7, 2009. On the epidemic curves, day zero is September 8, 2009.

Fig. 4 (right) shows posterior distributions for several quantities of interest to decision makers, including the total number of individuals expected to be infected during the course of the epidemic, the peak date, the incidence of influenza (new cases) on the peak date, the reproductive rate of the epidemic ( $R_0$ ), the mean latent period of the disease, and its mean infectious period. These six distributions are computed from the 269 most likely epidemic models found by ODS—given the disease surveillance data—whose cumulative posterior probability summed to an arbitrary threshold probability  $> 0.99995$ .

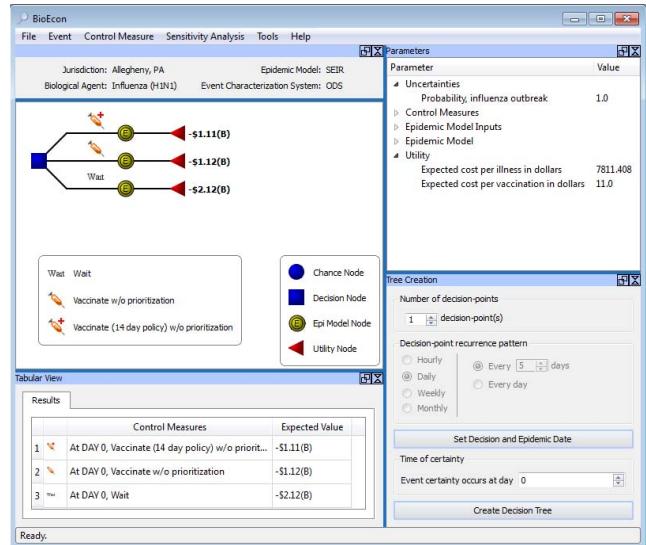
We described the ODS computation for influenza for expository purposes. However, the above approach could be applied to other diseases and it can be generalized to use other types of epidemic models, including segmented compartment models and agent-based models. Additionally, ODS used uniform prior probability distributions for SEIR model parameters in this analysis. For example, the prior distribution for  $R_0$  was uniform over the range 1.1 to 1.9. In practical application, we would expect that experts would use a more informative prior distribution over these parameters.

## B. BioEcon

BioEcon constructs and solves decision models (trees) representing alternative epidemic control strategies (Fig 5). It solves the expression shown in Fig. 2, obtaining the distribution  $P(M_{pop}|E)$  from ODS, and the result  $M_{pop, CS}$  from an epidemic model simulator configured with the SEIR parameters of each epidemic model sampled by ODS and the control measure parameters.

BioEcon constructs the structure of a decision model semi-automatically. A user specifies the epidemic model (SEIR or agent-based), the set of control measures to study, the points in time that a decision can be revisited (if desired), and a utility function. BioEcon then constructs a decision model for every combination of control measures, which we refer to as control strategies, subject to logical constraints such as two school closure policies cannot coexist. The logical constraints are represented declaratively as properties of control measures.

BioEcon contains representations of vaccine control measures—both vaccine supply and administration capacity—and school closure policies. BioEcon represents this knowledge using an object-oriented representation. For example, the class representing a vaccination control measure has the following attributes: jurisdiction (e.g., Allegheny County), supply schedule, vaccine administration capacity, efficacy, and lists of other control measures that it can run concurrent with, follow, or precede. BioEcon can acquire and store this information for multiple jurisdictions, each of which can have different capacities. Note that we believe that information about control measures should be acquired and stored in a response management system. Our project focuses, therefore, on defining representations, not on developing a large knowledge base of control measure information.



**Fig. 5 BioEcon.** An automatically generated decision model for H1N1 influenza in AC, Sept 8, 2009 (retrospective analysis). The upper panel, left, shows the generated decision tree. Beneath the tree, is a tabular display of the expected utilities of the decision alternatives. The panels on the right show all parameters (top) and allow the creation of sequential decision models and manipulation of decision date and epidemic  $t$  when the epidemic timing is not being set by ODS (bottom).

BioEcon contains a simple utility function for influenza and it is not our intention to develop utility functions for all possible epidemic diseases affecting human, plants, and animals. Rather, our focus in BioEcon is on representing the attributes (arguments) of multi-attribute utility functions, not on representing the functions. In particular, we have not developed extensive representations of cost information within BioEcon; instead, we use Excel spreadsheets and other tools to develop detailed economic models. We store rolled up costs in BioEcon as components of its utility functions.

At present, BioEcon passes two attributes to the influenza utility function—number of people vaccinated and total number infected. The utility function is:

$$U(v,s) = v(-\$11) + s(-\$7811.41),$$

where  $v$  is number of people vaccinated,  $-\$11$  is the cost of vaccination,  $s$  is number sick, and  $-\$7811.41$  is the average cost per sick person, which equals the cost of illness and the loss of productivity.

BioEcon handles sensitivity analysis over the uncertainty about the epidemic inherently, by model averaging over all epidemic models received from ODS. It allows an end-user to perform sensitivity analyses over the uncertainty about control measures, including costs, date of availability, schedules, capacities for vaccine administration. It supports probabilistic, one-way and two-way sensitivity analyses.

Note that the expected utilities shown in Fig. 5 were computed using the most probable epidemic model from ODS (the first model in Figure 4); they are not the results of model averaging over the set of SEIR models produced by ODS.

### C. Epidemic Models

BioEcon currently has access to two influenza epidemic simulators: a simulator that can run the SEIR models described above with vaccination control strategies, and an agent-based simulator that can run influenza models with vaccination, school closure, and other social-distancing control strategies.




Fig. 6. **Apollo Web Service.** End-user applications like BioEcon submit configuration objects to the epidemic model and receive output objects containing the results of the model run (e.g., an epidemic curve).

BioEcon obtains access to epidemic model simulators via Web services. We built an intermediary called the Apollo Web Server to reduce the effort of epidemic modellers to offer such services to end-user applications such as BioEcon (Fig. 6). We are currently working to standardize the vocabulary and message syntax for epidemic model configuration and result reporting.

A simple end-user application that demonstrates the basic functionality of the Apollo Web Service is located at <http://research.rods.pitt.edu/apollo/>

### CONCLUSIONS

The decision-theoretic model of disease surveillance extends the decision framework that was used in the 2009 H1N1 epidemic to include an explicit representation of the uncertainty about an epidemic. Our particular implementation uses the epidemiological knowledge about influenza epidemics, and of other types of epidemic diseases that is represented in epidemic models when inferring the probability that an outbreak of that disease exists, and when inferring its key characteristics.

Prior work using Bayesian algorithms for disease surveillance has had an emphasis on detection of epidemics rather than their characterization. Examples of temporal methods include [12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23]. Bayesian spatial approaches include that by Neill, Moore, and Cooper [32], who extended Kulldorff's spatial scan statistic to produce posterior probabilities of influenza in geographical sub-regions. A multivariate generalization was developed in [24]. Spatio-temporal approaches include the WSARE 3.0 algorithm [25], the PANDA algorithm for detecting anthrax outbreaks [26], the PCTS algorithm for detecting outbreaks of

all CDC Category A diseases that are of special concern for biosurveillance [27], and a Bayesian hierarchical model to detect anomalously high levels of influenza [28]. In previous research, we developed Bayesian algorithms [26, 27] that employed a data likelihood approach, similar to the method we describe here. However, they were based only on chief complaints as evidence.

Our approach to outbreak (a synonym for epidemic) detection and characterization (OD&C) has important features not present in previous work. First, instead of analyzing counts of data to estimate an epidemic curve [29, 30], we use a flexible and more general approach that models probabilistically the available evidence, such as the rich set of patient findings in ED reports. The approach reflects the intrinsic synergy between individual patient diagnosis and population OD&C. In particular, OD&C is derived based on past probabilistic patient diagnoses. In turn, the diagnosis of a newly arriving patient is based on prior probabilities that are derived from probabilistic inference over current OD&C models. To our knowledge, no prior research (either Bayesian or non-Bayesian) has taken such an integrated approach to patient diagnosis and population OD&C.

Second, our approach represents a general Bayesian framework for modeling OD&C. It can be applied with many different types of disease outbreak models including SEIR (Susceptible, Exposed, Infectious, and Recovered) model [10], agent-based, and outdoor-substance-release (OSR) models [31].

We expect that the decision-theoretic model will influence the fields of epidemic modelling and disease surveillance. The requirements of the decision-theoretic model will serve to increase the clarity about the parameters to which real decisions are sensitive and thereby inform research agendas in both fields. The new requirement for epidemic models, disease surveillance systems and decision models to 'talk' to each other will drive standardization of the interfaces between these components, especially the standardization of terminology and syntax required for interoperability.

### ACKNOWLEDGMENT

This research was funded by grants 5R01LM009132-02 from the National Library of Medicine for "Decision Making in Biosurveillance" and grant 1P01HK000086 from the Centers for Disease Control and Prevention to the University of Pittsburgh for support of a CDC Center of Excellence in Public Health Informatics.

### REFERENCES

1. Tsui F, Wagner MM, Cooper G, Que J, Harkema H, Dowling J, et al. Probabilistic case detection for disease surveillance using data in electronic medical records. *Online Journal of Public Health*. 2011;to appear Dec. 2012.
2. Tsui F-C, Espino JU, Dato VM, Gesteland PH, Hutman J, Wagner MM. Technical description of RODS: A real-time public health surveillance system. *J Am Med Inform Assoc*. 2003;10(5):399-408.
3. Lee BY, Brown ST, Cooley P, Grefenstette JJ, Zimmerman RK, Zimmer SM, et al. Vaccination deep into a pandemic wave potential mechanisms for a "third wave" and the impact of vaccination. *Am J Prev Med*. 2010 Nov;39(5):e21-9.

4. Brown ST, Tai JH, Bailey RR, Cooley PC, Wheaton WD, Potter MA, et al. Would school closure for the 2009 H1N1 influenza epidemic have been worth the cost?: a computational simulation of Pennsylvania. *BMC Public Health*. 2011;11:353.
5. Lee BY, Brown ST, Cooley P, Potter MA, Wheaton WD, Voorhees RE, et al. Simulating school closure strategies to mitigate an influenza epidemic. *J Public Health Manag Pract*. 2010 May-Jun;16(3):252-61.
6. Lee BY, Brown ST, Korch GW, Cooley PC, Zimmerman RK, Wheaton WD, et al. A computer simulation of vaccine prioritization, allocation, and rationing during the 2009 H1N1 influenza pandemic. *Vaccine*. 2010 Jul 12;28(31):4875-9.
7. Lee BY, Ercius AK, Smith KJ. A predictive model of the economic effects of an influenza vaccine adjuvant for the older adult (age 65 and over) population. *Vaccine*. 2009 Apr 6;27(16):2251-7.
8. Lee BY, Wiringa AE. The 2009 H1N1 influenza pandemic: a case study of how modeling can assist all stages of vaccine decision-making. *Hum Vaccin*. 2011 Jan 1;7(1):115-9.
9. Chao DL, Matrajt L, Basta NE, Sugimoto JD, Dean B, Bagwell DA, et al. Planning for the control of pandemic influenza A (H1N1) in Los Angeles County and the United States. *Am J Epidemiol*. 2011 May 15;173(10):1121-30.
10. Vynnycky E, White R. *An introduction to infectious disease modelling*: Oxford University Press; 2010.
11. Wagner MM, Dato V, Dowling JN, Allswede M. Representative threats for research in public health surveillance. *J Biomed Inform*. 2003 Jun;36(3):177-88.
12. Burkom HS, Ramac-Thomas L, Babin S, Holtry R, Mnatsakanyan Z, Yund C. An integrated approach for fusion of environmental and human health data for disease surveillance. *Statistics in Medicine*. [10.1002/sim.3976]. 2011;30(5):470-9.
13. Harvey AC. The Kalman filter and its applications in econometrics and time series analysis. *Methods of Operations Research*. 1982;44(1):3-18.
14. Le Strat Y, Carrat F. Monitoring epidemiologic surveillance data using hidden Markov models. *Statistics in medicine*. 1999;18(24):3463-78.
15. Mnatsakanyan ZR, Burkom HS, Coberly JS, Lombardo JS. Bayesian information fusion networks for biosurveillance applications. *Journal of the American Medical Informatics Association*. 2009;16(6):855-.
16. Nobre FF, Monteiro ABS, Telles PR, Williamson GD. Dynamic linear model and SARIMA: a comparison of their forecasting performance in epidemiology. *Statistics in medicine*. 2001;20(20):3051-69.
17. Ong JBS, Chen MIC, Cook AR, Lee HC, Lee VJ, Lin RTP, et al. Real-Time Epidemic Monitoring and Forecasting of H1N1-2009 Using Influenza-Like Illness from General Practice and Family Doctor Clinics in Singapore. *PLoS ONE*. 2010;5(4):e10036.
18. Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*. 1989;77(2):257-86.
19. Rath T, Carreras M, Sebastiani P. Automated detection of influenza epidemics with hidden Markov models. *Advances in Intelligent data analysis V*. 2003:521-32.
20. Sebastiani P, Mandl KD, Szolovits P, Kohane IS, Ramoni MF. A Bayesian dynamic model for influenza surveillance. *Statistics in Medicine*. 2006;25(11):1803-16.
21. Shiryaev AN. *Optimal stopping rules*: Springer; 1978.
22. Stroup DF, Thacker SB. A Bayesian approach to the detection of aberrations in public health surveillance data. *Epidemiology*. 1993:435-43.
23. Watkins R, Eagleson S, Veenendaal B, Wright G, Plant A. Disease surveillance using a hidden Markov model. *BMC medical informatics and decision making*. 2009;9(1):39-.
24. Neill DB, Cooper GF. A multivariate Bayesian scan statistic for early event detection and characterization. *Machine Learning*. 2010;79(3):261-82.
25. Wong W-K, Moore A, Cooper G, Wagner M, editors. *Bayesian Network Anomaly Pattern Detection for Disease Outbreaks*. Proceedings of the Twentieth International Conference on Machine Learning; 2003; Menlo Park, California: AAAI Press.
26. Cooper GF, Dash DH, Levander JD, Wong W-K, Hogan WR, Wagner MM, editors. *Bayesian biosurveillance of disease outbreaks*. Proceedings of the Twentieth Annual Conference on Uncertainty in Artificial Intelligence; 2004; Banff, Canada. 1036855: AUAI Press.
27. Jiang X, Cooper GF. A Bayesian spatio-temporal method for disease outbreak detection. *Journal of the American Medical Informatics Association*. 2010;17(4):462-.
28. Chan T-C, King C-C, Yen M-Y, Chiang P-H, Huang C-S, Hsiao CK. Probabilistic Daily ILI Syndromic Surveillance with a Spatio-Temporal Bayesian Hierarchical Model. *PLoS ONE*. 2010;5(7):e11626.
29. Ong JB, Chen MI, Cook AR, Lee HC, Lee VJ, Lin RT, et al. Real-time epidemic monitoring and forecasting of H1N1-2009 using influenza-like illness from general practice and family doctor clinics in Singapore. *PLoS One*. 2010;5(4):e10036.
30. Skvortsov A, Ristic B, Woodruff C, editors. *Predicting an epidemic based on syndromic surveillance*. Proceedings of the Conference on Information Fusion; 2010.
31. Hogan WR, Cooper GF, Wallstrom GL, Wagner MM, Depinay JM. The Bayesian aerosol release detector: An algorithm for detecting and characterizing outbreaks caused by an atmospheric release of *Bacillus anthracis*. *Statistics in Medicine*. 2007;26(29):5225-52.
32. Neill DB, Moore AW, Cooper GF. A Bayesian Spatial Scan Statistic. *Advances in Neural Information Processing Systems*. 2006;18:1003-10.