

Optimizing Anthrax Outbreak Detection Using Reinforcement Learning

Masoumeh T. Izadi and David L. Buckeridge

McGill University, 1140 Pine Ave West
Montreal, Quebec Canada H3A 1A3

mtabae@cs.mcgill.ca, david.buckeridge@mcgill.ca

Abstract

The potentially catastrophic impact of a bioterrorist attack makes developing effective detection methods essential for public health. In the case of anthrax attack, a delay of hours in making a right decision can lead to hundreds of lives lost. Current detection methods trade off reliability of alarms for early detection of outbreaks. The performance of these methods can be improved by modern disease-specific modeling techniques which take into account the potential costs and effects of an attack to provide optimal warnings. We study this optimization problem in the reinforcement learning framework. The key contribution of this paper is to apply Partially Observable Markov Decision Processes (POMDPs) on outbreak detection mechanism for improving alarm function in anthrax outbreak detection. Our approach relies on estimating the future benefit of true alarms and the costs of false alarms and using these quantities to identify an optimal decision. We present empirical evidence illustrating that the performance of detection methods with respect to sensitivity and timeliness is improved significantly by utilizing POMDPs.

Introduction

High-risk decision problems under uncertainty have considerable impact in many areas in public health. However, there has been relatively little work in this area. The very real threat of bioterrorism has accelerated the critical need for timely detection of outbreaks. As a result, the need for precise modeling and analysis of decisions faced by surveillance systems for providing the optimal warnings is becoming more acute. In the particular case of anthrax attack, delays of hours in making a decision to intervene can lead to hundreds of lives lost (Kaufmann, Meltzer, & Schmid. 1997; Wagner *et al.* 2001) and millions of dollars of additional expenses. Current studies of surveillance systems have demonstrated that a good detection algorithm can discover a disease outbreak before individual cases are diagnosed clinically. However, making a decision as to whether the partial information from a surveillance system reflects a real outbreak, is a challenge. Detection methods often can be adjusted to increase or decrease the sensitivity of their detection. These methods generally use a threshold which can be

tuned to increase sensitivity. However, improvement in sensitivity usually occurs at the cost of lower specificity. Intuitively, the optimal level of sensitivity relative to specificity depends on a specific disease. Although surveillance systems with low specificity generate many false alarms, which may ultimately be ignored by public health personnel. At the specificity of 0.9, (Buckeridge *et al.* 2006) show syndromic surveillance detected anthrax outbreaks on average one day before clinical case finding confirmed the outbreaks. This specificity is considered relatively low since it corresponds to 1 false alarm every 10 days.

The uncertainty as to whether an unusual pattern in the surveillance data is caused by anthrax makes it difficult to decide on whether to signal an alarm. The low specificity of detection methods makes it difficult to decide how credible are the resulting alarm signals. The sequential nature of the problem, the time-criticality of response decision making under these uncertain conditions, and the high risk of delays suggest strongly the need for a formal decision model to guide public health responses to the results of detection methods. Examples of decisions raised in response to anomalies in surveillance data include: whether or not we should wait to collect more data; whether we should investigate more information resources; or we should confirm an outbreak when we receive an alarm.

In this paper, we address precise modeling and analysis of decisions faced by surveillance systems for providing optimal warning of an epidemic. As a specific example we consider anthrax outbreaks. Our approach to tackle this problem is motivated by a principal observation that quantifying the potential costs and effects of an attack and the cost and effectiveness of interventions can be used as important criteria for optimizing the alarm function. We formulate the decision making problem for anthrax outbreak detection in POMDPs (Kaelbling, Littman, & Cassandra. 1998). In decision theoretic planning, POMDPs are well known as the most realistic model for decision making under uncertainty in dynamical systems. They have been widely proposed in modeling decision making in various domains such as medical decision making, mining engineering, robotics, and many more (Cassandra 1998b). Our POMDP model of a surveillance system accounts for the normal situation and different states of the attack after anthrax release until the time that an attack can be detected through clinical diagnosis of affected individ-

uals. The result of a detection method is used to provide observations to the POMDP model as noisy sensation about these true but unobservable states. The POMDP model performs further analysis on these results and optimizes the appropriate strategy to take in response to the output from the detection method. The enhanced detection approach described here can be coupled with any traditional outbreak detection method to optimize the way that the surveillance systems process alarm function.

Partially Observable Markov Decision Processes

In this section we review the POMDP framework and illustrate solving sequential decision problems in POMDPs. Formally, a POMDP is defined by the following components: a finite set of hidden states S , a finite set of actions A , a finite set of observations Z , a transition function $T : S \times A \times S \rightarrow [0, 1]$, such that $T(s, a, s')$ is the probability that the agent will end up in state s' after taking action a while in state s , an observation function $O : A \times S \times Z \rightarrow [0, 1]$, such that $O(a, s', z)$ gives the probability that the agent receives observation z after taking action a and reaching state s' , an initial belief state b_0 , which is a probability distribution over the set of hidden states S and a reward function $R : S \times A \times S \rightarrow \mathfrak{R}$, such that $R(s, a, s')$ is the immediate reward received when the agent takes action a in hidden state s and ends up in state s' . Additionally, there can be a discount factor, $\gamma \in (0, 1)$, which is used to give less weight to rewards received further in the future.

Solving POMDPs

The goal of a POMDP agent is to find a long term plan or *policy* for acting in such a way as to maximize the total expected reward received. The best such plan is called an *optimal policy* or an *optimal solution* for the POMDP. The agent in a POMDP does not have knowledge of the hidden states, it only perceives the world through noisy observations as defined by the observation function O . Hence, the agent must keep a complete history of its actions and observations, or a sufficient statistic of this history, in order to act optimally. The sufficient statistic in a POMDP is the *belief state* b , which is a vector of length $|S|$ specifying a probability distribution over hidden states. The elements of this vector, $b(i)$, specify the conditional probability of the decision making agent being in state s_i , given the initial belief b_0 and the history (sequence of actions and observations) experienced so far. After taking action a and receiving observation z , the POMDP agent updates its belief state b' using Bayes' Rule:

$$\begin{aligned} b'(s') &= P(s'|b, a, z) \\ &= \frac{O(a, s', z) \sum_{s \in S} b(s) T(s, a, s')}{P(z|a, b)} \end{aligned}$$

The denominator is a normalizing constant and is given by the sum of the numerator over all values of $s' \in S$.

A policy is a mapping $\pi : B \rightarrow A$. The amount of total expected reward that a decision maker can accumulate over its lifetime as given by the horizon h and following a policy π is called the *value function* of π . Most of the POMDP

algorithms are based on estimating a value function. A value function V^π of the policy π defines the value for each belief state under policy π .

$$V^\pi : \mathcal{B} \rightarrow \mathfrak{R} \quad (1)$$

The value function assigns to each belief state b the expected value of the total reward the agent can get in the future, given that its starting point is b . The optimal policy π^* in particular is the one that maximizes the total expected future reward:

$$\pi^*(b) = \arg \max_{\pi} E \left[\sum_{t=0}^{h-1} \gamma^t r_{t+1} | b \right] \quad (2)$$

Finding optimal policies for POMDPs is generally difficult. We can transform a POMDP into a belief state MDP. Under this transformation, the belief state b becomes the (continuous) state of the MDP (Markov Decision Process). The problem here is that there are infinite number of belief states b , so solving the above equation in exact form is very difficult. For further details on POMDPs solution methods we refer interested readers to (Lovejoy 1991; Cassandra. 1998a).

Considerable research effort has been devoted to the design of approximation algorithms for POMDPs to make these models practical. Recently, algorithms have been proposed which take advantage of the fact that, for most POMDP problems, a large part of the belief space is not experienced by the agent and the actual belief states have a sparse probability distribution. Such approaches, which are known as point-based methods (Pineau, Gordon, & Thrun 2003), consider only a finite set of belief points and plan for those points only. Point-based algorithms rely on the fact that performing many fast approximate updates often results in a more useful value function than performing a few exact updates. This algorithm and its variations gained much success in solving POMDP problems which are orders of magnitude large or more difficult than the ones solvable by exact solution methods.

Outbreak Detection in Surveillance Systems

The objective of detection algorithms in public health surveillance systems is to recognize from sequential input data (e.g. medical visits, absenteeism from work, drug consumption) the occurrence of an event such as an epidemic. A detection method may be as simple as comparing the amplitude of the signal with a threshold. If this value is above the pre-specified threshold of the algorithm then the algorithm indicates an alarm for a detected outbreak.

The accuracy of the detection method is reported using various ratios such as sensitivity, specificity. Sensitivity is the probability of an alarm given an outbreak, $P(A|O) = \frac{n(A, O)}{n(O)}$. Specificity is the probability of no alarm given that there is no outbreak, $P(A^-|O^-) = \frac{n(A^-, O^-)}{n(O^-)}$. Timeliness is also treated as a property of a detection method. Timeliness can be measured by: *detection time - actual time of the event*. The timeliness of a method is usually improved by adjustment of its threshold and at the expense of its other parameters. The outputs of an ideal surveillance system should

be those needed for response. However, these methods have to usually trade-off between timeliness and sensitivity on the one hand, versus specificity or false alarms on the other hand. The primary challenge to interpreting the output of the surveillance systems is the signal noise, or unpredictability, that prevents accurate modeling of the data and leads to errors in the model's predictions. These errors appear as noise that may cause false positives and false negatives. False positives occur when noise spikes in the model's predictions are detected as possible outbreaks, lowering the system's overall specificity. False negatives occur when noise in the model's predictions masks the effects of actual outbreaks, lowering overall sensitivity (Reis, Pagano, & Mandl 2003).

Popular methods for outbreak detection include simple and exponential weighted moving averages, applied either directly to the data or to residuals obtained by comparing observed data to expected data (Thacker & Berkelman 1998; Box & Jenkins. 1976; Reis & Mandl 2003). A fundamental challenge of detection systems is that if we increase the sensitivity of the system and improve the timeliness of detection, then the number of false alarms will increase. Unfortunately, these systems have low sensitivity during the first few days after a release of anthrax.

Formulating the Anthrax Outbreak Problem in POMDP Framework

The epidemic curve for anthrax by days after exposure is assumed to be < 1 day, 0% of cases; 1 day, 5%; 2 days, 20%; 3 days, 35%; 4 days, 20%; 5 days, 10%; 6 days, 5%; and 7 or more days, 5% (Messelson *et al.* 1994; Benenson AS 1995). The mean time for clinical detection of anthrax is estimated to be between 3-4 days following a release of 0.1 kg of anthrax spore in an urban area (Adamou *et al.* 2006) and (Buckeridge *et al.* 2006). Therefore, if a surveillance system takes longer than 3-4 days to detect an outbreak, then the system may not be very helpful. There is always a small probability of starting an outbreak. In a normal situation (no outbreak) we assume a probability of an attack $P = 0.01$. We use this prior knowledge in the transition function of the model which is discussed later in this section. Here we build the model and its parameters based on experts opinions, and the results from simulation studies in the literature for anthrax attacks. Figure 1 depicts the POMDP model we designed for this problem.

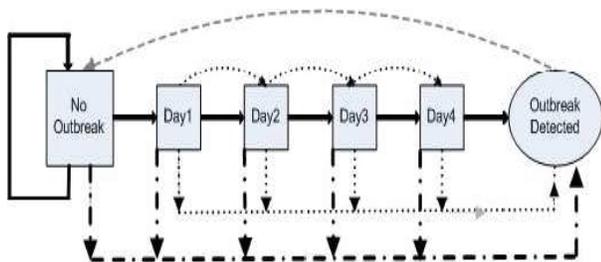


Figure 1: The POMDP model for anthrax outbreak detection.

The economical impact of an attack used in our POMDP

model is based on the analysis reported in (Kaufmann, Meltzer, & Schmid. 1997). Figure 2 shows the cumulative economic impact of a large release of aerosolized B.anthraxis created from this analysis. The authors consider the impact of an attack on a suburb of a major city, with 100,000 people exposed in the target area. This is based on cost estimation of deaths, the costs of hospitalizations, the costs of outpatient visits. The rewards and costs are related to:

- Cost of a single false positive and false negative;
- Cost of intervention (depends on the population size, cost per person, and implementation time);
- Cost of a single day delay;
- Detection benefit = the number of death \times future earnings + number of hospital days \times cost of 1 hospital day + number of outpatient visits \times cost of 1 outpatient visits intervention costs;
- Intervention costs = cost per person \times number of cases seeking care

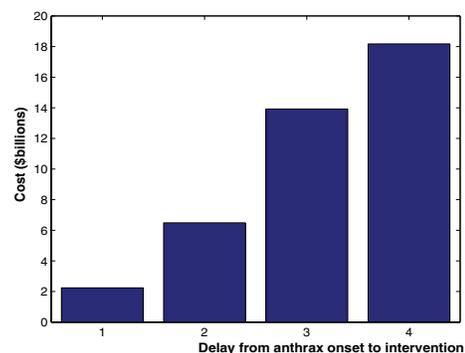


Figure 2: The estimate of accumulative preventable loss for detection of anthrax per day after a release.

Model Parameters

- **The state space:** we propose to consider a state space consisting of six states: *NoOutbreak*, *Day1*, *Day2*, *Day3*, *Day4* and *OutbreakDetected* to reflect the assumption that at an outbreak will be detected clinically within 5 days.
- **The action space:** we consider 4 possible strategies available at each time. These are *declare an outbreak* as a safe but very expensive option; a somewhat cheaper option of *more systematic studies* to gather extra information from external sources (for example more patient files in emergency departments); *more investigation* can take up to a few hours of a human expert to review carefully data already in hand ; and the last option is not to do anything and *wait*. The space of actions is based on standard strategies in epidemiology (Gregg. 1996) and discussions with epidemiologists.

- **The observation space:** at each instant of time (a day), we perceive observations from a detection method which reflect an alarm condition. These observations which are dependent on the underlying state of the model are some informative statistics output from detection algorithm. The output of a detection system includes a sequentially updated probabilistic assessment of the threats being monitored. The distribution of p-value, one-step-ahead daily forecast of respiratory syndrome counts, and cumulative sum for detection of positive deviation in the forecast residuals are commonly available assessments in detection algorithms. In this model, we have considered two observations, *suspicious* and *non-suspicious* to reflect the binary output received from a detection algorithm.
- **Transition functions:** There is a small chance of moving to the first day of an attack from the NoOutbreak state under any action. We consider the probability of $P = 0.01$ for this case. The transition through consequent days of an outbreak by choosing to wait is performed naturally. A systematic study may take 1 day to give some probability of an attack and the investigation option takes only a few hours of a human expert. Human decision makers are subject to biases that lead to suboptimal decisions, especially when they are dealing with rare events, uncertainty, and high cost options (Kahneman, Slovic, & Tversky. 1982). Therefore, the systematic study and the investigation options can reduce the uncertainty about the state of the outbreak by only a small amount. This amount increases as the outbreak progresses. After comprehensive discussions with domain experts we decide to consider an extra 10 percent sensitivity for the investigation action and an extra 30 percent sensitivity for the systematic study. This means that the probability of the attack being detected after an investigation on day1 will be 0.1 and after systematic study this would be 0.3.
- **Observation functions:** as the outbreak progresses the detection method provides more reliable information on whether or not there is an outbreak. At any outbreak state, for defined observation *suspicious*, the observation function is defined by the sensitivity of the detection method used. For the second observation the noise is defined as $1 - \text{specificity}$ at normal states.
- **Reward-Cost functions:** There is a reward/cost associated with each action at each underlying state of the system. In the NoOutbreak state, if we choose not to do anything, we do not incur any cost. For other situations we use previous studies modeling anthrax (Braithwaite, Fridsma, & Roberts 2006; Kaufmann, Meltzer, & Schmid. 1997) to incorporate more realistic information to the model. We have used the difference between preventable losses in each day and in the consecutive day presented in Figure 2 as the reward for declaring an outbreak on each day of the outbreak. For the investigation option we have assumed that it takes almost half a day to perform this action. Therefore, in this case we assigned half of the reward for transiting to outbreak detected states from each day of the outbreak. We also assumed that a systematic

study takes one day to perform. Therefore, we consider the reward of the next day for transiting to outbreak detected state up on taking this action. The penalty of not detecting the outbreak and transition to the next state of the outbreak is determined by the corresponding loss at that state. There is a penalty equal to treatment expenses for maximum number of people seeking care, for choosing to confirm an outbreak from a clear states. In computing the number of people seeking care and number of deaths. we considered a population of 100,000 exposed, as in (Kaufmann, Meltzer, & Schmid. 1997). The authors modeled the costs and benefits based on clinical and experimental findings with respect to the disease progression and treatment options.

Discussion

We argued that the alarm function in outbreak detection methods can be formulated as a decision problem. However, there are other aspects of disease outbreak which can also be formulated as decision problems and can be solved by optimization techniques. This include data gathering form a verity of sources at different expenses: as which data are more useful for early detection? The response strategies and treatment options after the outbreak confirmed are decision making problems which requires precise modeling and optimization. Public health researchers are motivated to apply mathematical frameworks in signal detection theory and decision theory to provide better analysis and measurement in this area (Wagner *et al.* 2001; Wagner, Moore, & Aryel 2006). Probabilistic graphical models have been suggested in related work in detecting epidemics. (LeStrat & Carrat. 1999) proposed detecting epidemic and non-epidemic phases of influenza by HMMs using a mixture of Gaussian distributions. (Rath, Carreras, & Sebastiani 2003) also proposed using a 2-state HMM, where non-epidemic rates are modeled with an exponential distribution, and epidemic rates with a Gaussian distribution. We note that (Das, Moore, & Schneider. 2004) previously suggested the use of POMDP belief states for a decision to signal an alarm in surveillance systems. However, the authors did not explain the design of the model components including the reward function which motivates the progression of POMDP solution methods. A simple 2-3 state POMDP introduced by the authors can not explain the status of a surveillance system under different diseases. Clearly, the reward function and the state space definitions are domain dependent factors and have to be estimated carefully for the disease monitored by the surveillance systems. To date, there have been few economic analyses of biosurveillance problems (Kaufmann, Meltzer, & Schmid. 1997; Meltzer, Cox, & Fukuda 2005; Wagner, Wallstrom, & Onisko 2005). Achieving this level of analysis is considered as a long-term goal for biosurveillance research (Wagner, Moore, & Aryel 2006).

Empirical evaluations

To work with a detection method and guide response strategies, in our experimental setup we used a moving average method, applied to residuals from a time-series model, to

provide the observation for the POMDP model. The policy obtained from solving the POMDP model is the one with minimum amount of cost and maximum amount of possible reward. The exact solution methods were unable to solve this POMDP. We have used a point-base approximation method introduced in (Izadi, Precup, & Azar. 2006) to solve this POMDP in a few seconds.

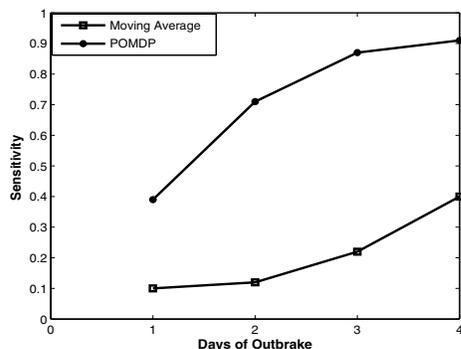


Figure 3: The timeliness of anthrax outbreak detection method with and without using POMDPs: sensitivities during different days of the outbreak for attacks that resulted in 10 additional visits.

Figure 1 and Figure 2 summarize our experimental results utilizing the POMDP model. The results for the moving average method were extracted from the results in (Reis, Pagano, & Mandl 2003). We considered a fixed specificity of 0.97 for all cases. The results reported for the case of POMDP are based on averages of 10 independent runs of the POMDP generative model over 5000-day time period. Over this period, we examined 257 outbreak on average.

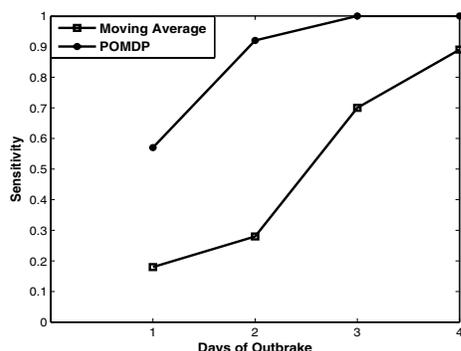


Figure 4: The timeliness of anthrax outbreak detection method with and without using POMDPs: sensitivities during different days of the outbreak for attacks that resulted in 20 additional visits.

The timeliness of different approaches with respect to the actual days of the outbreak is shown in the figures. All outbreaks were detected by our approach prior to the fourth day of the attacks when the detection method provides a better sensitivity. The other model was based on the moving average method which provides low sensitivity, can not de-

tect all the outbreaks prior to day 4. Of course not all outbreaks are detected by the moving average methods at the specificity equal to 0.97 in general. The POMDP utilizing approach yielded much higher sensitivity for both outbreak sizes, while maintaining the same specificity, resulting in much better overall performance.

Conclusion and future work

The events surrounding an outbreak due to bioterrorism will unfold rapidly. The response decision making must be formalized in advance of an attack into a decision policy that can be applied without bias or delay during a crisis. In this paper, we discussed the development of an optimal surveillance alarm function for anthrax outbreak. The empirical evaluation of our approach shows dramatic improvements over traditional outbreak detection methods. Our promising results suggest further directions for research, including consideration of outbreaks due to other diseases. Infectious threats such as SARS and human H5N1 influenza infections have prompted the development of detection systems that respond in a timely way to emerging epidemics, allowing authorities to respond at the earliest possible stage. Worldwide developments concerning biological weapons and terrorism are driving force for improving public health surveillance and outbreak response. It is worth mentioning that our approach applies not only to surveillance for outbreaks caused by terrorists, but also to naturally occurring outbreaks both in the community and in hospitals. In future work, we intend to apply our proposed model to routinely encountered infectious diseases such as influenza. Working with more frequent threats such as water contamination or influenza makes this application potentially useful for routinely encountered public health problems.

To apply our approach to any surveillance system, the set of actions appropriate to the disease in question must be determined and the accuracy of effects of actions must be quantified. It is important to have economic estimates of the effect of different actions as a function of time for those threats. The appropriate state space which corresponds to the time window for the specific disease must also be identified. Timely detection is not a key requirement of every public health threat, but it is for many. We found a rich literature on quantitative and qualitative analysis of timeliness requirements for the case of inhalational anthrax. This type of analysis is needed to do more research and development on reliable and early warning systems for other types of threats.

Acknowledgments

This research was supported by Canada Research Chair Grant in Public Health Informatics, and CDC grant from Centers for Disease Control and Preventions R09LM008705-01.

References

Adamou, C.; Cooper, G. F.; Wong, W.-K.; Dowling, J. N.; and Hogan, W. R. 2006. Modeling clinician detection time

- of a disease outbreak due to inhalational anthrax. *Advances in Disease Surveillance*. 1(3).
- Benenson AS, e. 1995. Control of communicable diseases manual. In *16th ed. Washington (DC): American Public Health Association*.
- Box, G. E., and Jenkins., G. M. 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Fransisco.
- Braithwaite, S.; Fridsma, D.; and Roberts, M. 2006. The cost-effectiveness of strategies to reduce mortality from an intentional release of aerosolized anthrax spores. *Medical Decision Making*. 26(2):182–93.
- Buckeridge, D. L.; Owens, D. K.; Switzer, P.; Frank, J.; and Musen., M. A. 2006. Evaluating detection of an inhalational anthrax outbreak. *Emerging Infectious Diseases* 12(12):1942–1949.
- Cassandra., A. 1998a. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. Ph.D. Dissertation, Brown University.
- Cassandra, A. 1998b. A survey of pomdp applications. *Working Notes of AAAI, Fall Symposium on Planning with Partially Observable Markov Decision Processes* 17–24.
- Das, K.; Moore, A.; and Schneider., J. 2004. Belief state approach to signaling alarms in surveillance systems. *Knowledge Discovery and Data Mining*.
- Gregg., M. B. 1996. *Field Epidemiology*. Oxford university Press.
- Izadi, M. T.; Precup, D.; and Azar., D. 2006. Belief selection in point-based planning algorithms for pomdps. *International Conference on Artificial Intelligence* 383–394.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra., A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101.
- Kahneman, D.; Slovic, P.; and Tversky., A. 1982. *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Kaufmann, A.; Meltzer, M.; and Schmid., G. 1997. The economic impact of a bioterrorist attack: are prevention and postattack intervention programs justifiable? *Emerging Infectious Diseases* 3(2):83–94.
- LeStrat, Y., and Carrat., F. 1999. Monitoring epidemiologic surveillance data using hidden markov models. *Statistics in Medicine* 24(18):3463–3478.
- Lovejoy, W. S. 1991. A survey of algorithmic methods for partially observed markov decision processes. *Annals. Operations Research* 28(1-4):47–66.
- Meltzer, M. I.; Cox, N. J.; and Fukuda, K. 2005. The economic impact of pandemic influenza in the united states: Priorities for intervention. *Emerging Infectious Diseases* 5(5):659–671.
- Messelson, M.; Guillemin, J.; Hugh-Jones, M.; Langmuir, A.; Popova, I.; and Shelokov, A. 1994. The sverdlovsk anthrax outbreak of 1979. *Science* 266:1202–1208.
- Pineau, J.; Gordon, G.; and Thrun, S. 2003. Point-based value iteration: An anytime algorithms for POMDPs. In *International Joint Conference on Artificial Intelligence*, 1025–1032.
- Rath, T. M.; Carreras, M.; and Sebastiani, P. 2003. Automated detection of influenza epidemics with hidden markov models. *IDA* 521–532.
- Reis, B., and Mandl, K. 2003. Time series modeling for syndromic surveillance. *BMC Medical Information Decision Making* 3(2).
- Reis, B. Y.; Pagano, M.; and Mandl, K. D. 2003. Using temporal context to improve biosurveillance. In *Proceedings of National Academy of Science USA*, volume 100, 1961–1965.
- Thacker, S. B., and Berkelman, R. L. 1998. National hiv behavior surveillance among injecting drug users in los angeles county. *Epidemiology Review* 10:164–190.
- Wagner, M. M.; Tsui, F.-C.; Espino, J. U.; Dato, V. M.; Sittig, D. F.; Caruana, R. A.; McGinnis, L. F.; Deerfield, D. W.; Druzdzal, M. J.; ; and Fridsma, D. B. 2001. The emerging science of very early detection of disease outbreaks. *Journal of Public Health Manag. Pract.* 7(6):51–59.
- Wagner, M. M.; Moore, A. W.; and Aryel, R. M., eds. 2006. *Handbook of Biosurveillance*. Elsevier.
- Wagner, M.; Wallstrom, G.; and Onisko, A. 2005. Issue a boil-water advisory or wait for definitive information? a decision analysis. In *Annual Fall Symposium of the American Medical Informatics Association*, 774–778.