Bayesian Epidemic Detection in Multiple Populations

Katherine Shatskikh*, Michael Ludkovski

Department of Statistics and Applied Probability, University of California Santa Barbara, Santa Barbara, CA 93106-3110, U.S.A.

Abstract

Traditional epidemic detection algorithms make decisions using only local information. We propose a novel approach that explicitly models spatial information fusion from several metapopulations. Our method also takes into account cost-benefit considerations regarding the announcement of epidemic. We utilize a compartmental stochastic model within a Bayesian detection framework which leads to a dynamic optimization problem. The resulting adaptive, nonparametric detection strategy optimally balances detection delay vis-a-vis probability of false alarms. Taking advantage of the underlying state-space structure, we represent the stopping rule in terms of a detection map which visualizes the relationship between the multivariate system state and policy making. It also allows us to obtain an efficient simulation-based solution algorithm that is based on the Sequential Regression Monte Carlo (SRMC) approach of Gramacy and Ludkovski (SIFIN, 2015). We illustrate our results on synthetic examples and also quantify the advantages of our adaptive detection relative to conventional threshold-based strategies.

Keywords: Biosurveillance; quickest detection; regression Monte Carlo; stochastic compartmental models;

1. Introduction

Infectious disease epidemics intrinsically unfold across both space and time. As a result, biosurveillance algorithms need to integrate spatio-temporal data. This is especially so in the context of statistical inference, whereby syndromic surveillance at neighboring locales carries additional information that can be fused for improved decision making in terms of initiating and organizing epidemic counter-measures. A crucial first step for response strategies is to identify, or detect, in real-time the epidemic outset. In this article, we propose a methodology that allows for such *optimal* decision-making with spatial information fusion. Specifically, we investigate a model that combines quickest

---

*Corresponding author

Email addresses: `shatskikh@pstat.ucsb.edu` (Katherine Shatskikh), `ludkovski@pstat.ucsb.edu` (Michael Ludkovski)

detection with a spatial metapopulation setup, integrating information received from multiple geographic domains. To reflect the inherent uncertainty in epidemic evolution (which is amplified under partial information), we develop a stochastic compartmental (or state-space) epidemic model, which allows us to generate adaptive, nonparametric detection rules. Extant approaches largely propose heuristic detection strategies, concentrating primarily on the inferential aspect of the statistical model [1, 2, 3, 4]. For instance, a typical approach is to announce an epidemic as soon as the estimated number of infecteds in the local population is above a fixed $\bar{I}$. In contrast, we dynamically optimize the detection strategy, to come up with a "best" detection rule within our mechanistic outbreak model.

Traditional compartmental epidemic models deal with a single population; the spatial aspect is treated by building a series of such single-population models that are estimated/forecasted independently. This is also a common surveillance approach, especially for recurring infectious epidemics, such as influenza-like illness (ILI), dengue fever, or measles. For example, in the US the existing biosurveillance systems for flu operate primarily at the state level and are siloed across states. This limitation of existing practice was brought into sharp relief during the 2014 Ebola outbreak in West Africa. The epidemic has been accompanied by a dearth of reliable information, leading to extreme spread in forecasts regarding the future course of the outbreak. In addition, numerous statistical methods [5, 6, 7] were put forth attempting to infer in "real-time" the actual size and parameters of the outbreak in different locales. However, nearly all these methods were single-population, so that when trying for example to infer the number of Ebola infecteds in Liberia, only Liberian data was utilized, completely ignoring similar and highly relevant data from neighboring Guinea and Sierra Leone. Similarly, at the more granular provincial level, data from neighboring provinces was generally not used during estimation procedures.

For a less dramatic and perhaps more statistically convenient example, we discuss the yearly influenza outbreaks in United States. Figure 1 illustrates the spatial dynamics of ILI during the 2012-13 flu season. As can be observed, the peak of the outbreak varied significantly (up to 6-8 weeks difference) across different parts of the country. Nevertheless, there is a clear propagation, making spatial information fusion desirable. Figure 1 indicates that the current, single-population based detection protocols are not sufficient; for instance the fact that there are increased ILI levels in Arizona is ought to be taken into account when trying to detect or forecast the epidemic start in California. A further important remark is that the illustrated spatial spread is year-specific, and in other years rather different patterns may be observed.

1.1. Contributions

In this paper we formulate and analyze an epidemic detection problem within a multi-population paradigm. To do so, we develop a reduced compartmental model that extends the classical Susceptible-Infected-Recovered (SIR) setup to two population pools. Pools are interpreted as distinct geographic regions, e.g. states or counties. To fix ideas, we consider the situation where the epidemic

Week 49, 2012	Week 1, 2013	Week 4, 2013

Figure 1: Spread of Influenza during the 2012-13 Flu season according to FluView CDC data. The colors represent weekly ILI activity levels in terms of percentage of doctor visits attributed to ILI relative to low-season baseline. Green indicates at/below mean, while shades of red indicate outbreak activity (with darkest color corresponding to eight or more standard deviations above the mean). Weeks are numbered from January 1, and are 12/3-9/2012 (Week 49), 12/31/2012-1/6/2013 (Week 1) and 1/21-27/2013 (Week 4), respectively. Data source: http://www.cdc.gov/flu/weekly/pastreports.htm.

begins in Pool 1 and subsequently may be transmitted to Pool 2 via infecteds that travel between the two pools. The aim of the policy-maker is to detect, as soon as possible and in online fashion, the onset of epidemic in Pool 2.

To capture the inferential aspect, we assume that full information is available about the outbreak in Pool 1, but only partial information about Pool 2. As a result, one has to make imperfect decisions and in particular address the canonical trade-off between making announcements too early (so called "false alarms") and making decisions too late ("detection delay"). Indeed, if the detection is too late, then a certain number of infections would be missed and it would be harder to stop the epidemic from spreading. If the detection is premature, human, financial and reputational resources would be wasted. Therefore, a careful trade-off between those costs should be done to balance costs due to epidemic morbidity and costs arising from policy actions. We then use the above cost analysis to quantify decision-making quality and to define optimality of detection strategies.

Mathematically, we cast the online detection problem as a dynamic optimization problem, connecting to the classical dynamic programming formulation [8] in control theory. A major challenge with dynamic programming (which is perhaps the prime reason for the lack in its uptake in the biosurveillance community) is computational bottlenecks due to the curse of dimensionality. Indeed, the above optimization problem is nontrivial from several directions. First, because the underlying system is stochastic, the optimal solution is *adaptive*, i.e. a function of the current system state. Consequently, there is no simple description to the resulting detection strategy which is instead summarized through a detection *map* that translates system states into optimal detection decisions. Second, the nonlinear dynamics of the SIR model preclude analytic solutions. Crucially, there are no analytic expressions for the future distribution of the system state, which necessitates the use of numerical approximations to solve the optimization problem. Third, because the system state is multivariate and too large to enumerate, the corresponding integrals are computationally demanding.

However, taking advantage of the detection strategy structure, which requires simply announcing at each stage whether the epidemic has reached Pool 2 or not, we implement an efficient numerical algorithm. Specifically, we rely on the recent Sequential Regression Monte Carlo (SRMC) method of [9], which blends modern statistical tools, including nonparametric regression and sequential design, with approximate dynamic programming, to drastically mitigate issues of computational efficiency.

The main contributions of this work are then threefold. First, we propose and analyze a multi-population extension of the classical SIR model, as well as a reduced version suitable for the Bayesian detection framework. Second, we develop and adapt an extension of the sequential regression Monte Carlo (SRMC) approach to efficiently solve the dynamic optimization problem. Third, we present a detailed investigation into the performance of the designed strategy, in particular in comparison to conventional threshold-based strategies.

The organization of the paper is as follows. Section 2 formalizes the mathematical aspects of our model, including the detection setup. The stochastic dynamics of the outbreak are rigorized in Section 3. Section 4 presents numerical illustrations of our method as well as comparison with other methods. Section 5 then describes the Sequential Regression Monte Carlo algorithm that we developed for our setup. Section 6 provides the conclusion and the discussion on future extensions of our framework.

1.2. Spatial Stochastic Epidemic Models

Mathematical models of infectious disease epidemics have become an important tool in the arsenal of public health policy. In an idealized world, detection reduces to the mathematical problem of clustering, tracking the health status of the surveyed individuals and identifying unusual aberrations in either the temporal or spatial dimensions. In reality, there is the additional aspect of missing information which necessitates the application of statistical inference algorithms, as well as a mathematical model for the epidemic. In the context of online inference, a simple mechanistic approach that allows for maximum tractability continues to be the most popular, and is also adopted here. Specifically, we rely on the formalism of an SIR model [10] that implies proportional homogenous mixing between infecteds and susceptibles within a population pool. Spatial heterogeneity is captured by incorporating meta-populations, also known as patch models [11, 12, 13]. The multi-patch approach partitions the global population into distinct discrete regions or pools, allowing for local spread of the epidemic within each pool, as well as global transmission that is specified via a mobility matrix. As in [11, 13] we assume that susceptibles are stationary, while infecteds can move or travel between the pools, creating cross-infections.

Alternative frameworks for epidemic spread include point process models [14], and network models [15] that provide more nuanced interaction between individuals to mimic existing social structures, such as households, schools, and workplaces. At even more detail, agent-based models [16] generate microsimulations that provide a detailed synthetic view for each individual and their social interactions. Such models can also incorporate precise travel patterns

4

[17]. However, the latter paradigms are geared towards realistic *forecasting* of epidemic progress and are less suited for online detection due to intractable inference in terms of observed data and the computational expenses in generating micro-scenarios.

A variety of approaches exist for constructing outbreak detection rules, see for example the recent survey by Shmueli and Burkom [18], and the monograph by Lawson [19]. Quality control methods [20] introduced in the 1950s form the simplest class of rules and continue to be common. Other heuristics include moving-average tools [21], various scan statistics [22, 14], and branching-process approximations [23]. More explicit cost-benefit analysis for the trade-off between false alarms and detection delay can be applied using the Cumulative Sum (CUSUM) framework [24]. CUSUM also underlies the early aberration response system (EARS) employed by the Centers for Disease Control [25]. Alternatively, Bayesian methods allow to further assess the uncertainty involved in decision-making based on partial information. Two main types are hidden Markov models [26, 27] and Bayesian hierarchical models [1, 28]. The Bayesian paradigm translates epidemic data into the posterior probability of an outbreak. To convert the latter into a detection rule, one typically employs a simple threshold strategy. For example, in [1], the authors recommend "an alert for action if the posterior probability is larger than 70%". We further refine this approach by deriving *optimal*, non-parametric detection strategies based on the inputted cost-benefit parameters.

Detection can be seen as a basic form of epidemic response, and indeed our computational methodology can be extended to this more general problem. In that sense, this paper extends the first author's previous work on stochastic control methods for controlling epidemics [29, 30]. Similar to [29], we design a Bayesian dynamic optimization algorithm for biosurveillance decision policy. Other mathematically oriented studies that consider optimal control of epidemics include [31, 32].

In the context of detection with limited information, a spatial epidemic model requires information fusion. Fusion of information channels for the purpose of biosurveillance has been an area of intense research in the past decade. On the one hand, novel information sources, such as social media [33] or internet data [34] have created new opportunities for syndromic surveillance. On the other hand, developments in statistical fusion techniques [18, 35, 36] have led to new ways of integrating multivariate information streams. In particular, there has been a lot of interest in online Bayesian approaches [2, 3, 23, 34, 37] that allow for predictive modeling and forecasting of epidemics. The above models all focus on a single homogenous population with the different surveillance channels complementing each other. In contrast, we consider multiple underlying population pools each with a distinct, but co-dependent information channel. In terms of explicitly accounting for spatial propagation, our work is closest to [38] who considered a spatial "wave" model for an epidemic. In the present article, we connect this framework to the SIR context, modeling epidemic spread across geographically-based population pools. The resulting decision strategy provides insights into integrating data from multiple spatial locales for the purposes of

5

detection, cf. Section 6 below.

2. Quickest Detection

2.1. Mathematical Model

We work with a state-space model, denoting by $\mathfrak{X}_t$ the epidemic state at times $t = 0, 1, 2, \ldots$. A typical length of one time period in biosurveillance is a week. The precise components of $\mathfrak{X}$ will be specified later; abstractly $\mathfrak{X}$ is taken to be a stochastic Markov process taking values in a state space $\mathcal{X} \subset \mathbb{R}^d$, and summarizes information about both Pool 1 and Pool 2. In particular, $\mathfrak{X}$ contains information about the number of infecteds $I_t^{(k)}$ in Pool $k = 1, 2$ at time t. The transition kernel of $\mathfrak{X}$ is assumed to be time-stationary and is denoted by $p_s(\mathbf{x}|\mathbf{y}) \equiv P(\mathfrak{X}_{t+s} = \mathbf{x}|\mathfrak{X}_t = \mathbf{y})$, $\mathbf{x}, \mathbf{y} \in \mathcal{X}$.

The aim of the policy maker is to detect the onset of epidemic in Pool 2. A detection strategy is probabilistically represented as a dynamic "alarm" which announces an outbreak in Pool 2, based on information gathered so far. Only a single announcement is allowed; once announced, the detection problem is assumed to be over. The set of such detection strategies is expressed through the set $\mathcal{S}$ of $\mathcal{F}$-stopping times, where $\mathcal{F}_t = \sigma(\mathfrak{X}_{0:t})$ is the information filtration generated by $\mathfrak{X}$ by time t. A strategy $\tau \in \mathcal{S}$ is a random variable taking values in $\tau \in \{0, 1, 2, \ldots\}$, such that $\{\tau = t\} \in \mathcal{F}_t$ (this requirement captures the fact that τ must be "online" in terms of the information available so far). Thanks to the Markov property of $\mathfrak{X}$, the structure of τ can be summarized via a detection map. Indeed, at each time-step there is the binary decision to either "announce" an outbreak (subset $\mathfrak{S}$), or wait for another period (subset $\mathfrak{C}$). Since the evolution of $\mathfrak{X}$ is stationary in time, the corresponding partition of the state space is also independent of t. Dynamically, this implies that τ announces the epidemic the first time that the state $\mathfrak{X}$ enters the region $\mathfrak{S} \subset \mathcal{X}$,

$$\tau = \inf\{t : \mathfrak{X}_t \in \mathfrak{S}\}. \tag{1}$$

Equation (1) gives a one-to-one correspondence between detection strategies τ and detection maps $\mathfrak{S}$. In other words, the detection strategies we consider are of online feedback type, based on the trajectory of $\mathfrak{X}$.

As mentioned, the dynamic optimization objective consists in optimally trading off the concern of premature announcements against any potential delays. These conflicting costs are measured through the immediate stopping cost $d(\mathbf{x})$ and the cost of waiting. The immediate costs are linked to the penalty for false alarms, specified by a given constant C_{FA}. We assume that C_{FA} is paid if and only if the epidemic has not yet reached Pool 2, so that

$$d(\mathbf{x}_0) := C_{\text{FA}} \cdot \mathbf{1}_{\{I_0^{(2)}=0\}}. \tag{2}$$

Waiting costs are assumed to be proportional to detection delay, i.e. the time between the outbreak reaching Pool 2 and outbreak announcement. Define θ to

be the time when the second population gets infected from the first population, i.e.

$$\theta := \inf\{t : I_{t-1}^{(2)} = 0 \text{ and } I_t^{(2)} > 0\}.$$

Then the detection delay is $\max(\tau - \theta, 0)$ and carries cost $C_{\text{Delay}} \max(\tau - \theta, 0)$. This is equivalent to charging waiting costs of $C_{\text{Delay}} \mathbf{1}_{\{I_t^{(2)} > 0\}}$ at each step until surveillance is terminated at the random instant τ, so that total waiting costs on $[0, \tau]$ are

$$c(\mathfrak{X}_{0:\tau}) := \sum_{s=0}^{\tau-1} C_{\text{Delay}} \mathbf{1}_{\{I_s^{(2)} > 0\}} + C_{\text{FA}} \mathbf{1}_{\{I_\tau^{(2)} = 0\}}. \tag{3}$$

We will refer to the costs $d(\cdot)$ and $c(\cdot)$ as the immediate cost and the future cost, respectively.

Remark 1. *Note that detection costs are intrinsically defined in terms of the count of infecteds in Pool 2, $I^{(2)}$, which is assumed to be unavailable to the policy-maker. Below we will operationalize (2) and (3) by taking conditional expectation with respect to information that is available, see (18)-(17).*

The aim of outbreak detection is to pinpoint θ, i.e. ideally one takes $\tau = \theta$. However, this is not possible if only partial information is available about $\mathfrak{X}$, specifically about $I_t^{(2)}$. When τ and θ are different, C_{Delay} penalizes the event $\{\tau > \theta\}$, and C_{FA} penalizes $\{\tau < \theta\}$. The cost structure in (3) is then a dynamic counterpart of the usual Type-I and Type-II errors in hypothesis testing.

2.2. Detection Problem

Our detection problem is formalized as minimizing the expected future cost over all possible stopping times τ [39], i.e. an optimal stopping problem. Namely, we define the value function V as

$$V(\mathbf{x}_0) := \inf_{\tau \in S} E\left[c(\mathfrak{X}_{0:\tau}) | \mathfrak{X}_0 = \mathbf{x}_0\right], \tag{4}$$

where $\mathbf{x}_0$ is the initial state. Assuming the infimum in (4) is achieved, the dynamic programming principle [39] implies

$$V(\mathbf{x}_0) = \min\left(d(\mathbf{x}_0), E\left[V(\mathfrak{X}_1) | \mathfrak{X}_0 = \mathbf{x}_0\right]\right), \tag{5}$$

where the conditional expectation operator is

$$E[V(\mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}_0] = \int V(\mathbf{x}) p_1(\mathbf{x}|\mathbf{x}_0) d\mathbf{x}.$$

The minimum operator in (5) corresponds to the idea that it is optimal to declare an outbreak if the immediate cost is smaller than the future cost, i.e. the likelihood of false alarms is dominated by the cost of waiting. The former case

is equivalent to the expectation of the value function at time 1 being greater than the immediate cost, and therefore we may classify the stopping region via

$$\mathfrak{S} := \{\mathbf{x} : E\left[V(\mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right] - d(\mathbf{x}) > 0\}. \tag{6}$$

Hence, in terms of the above detection map, our goal is to optimally partition $\mathcal{X} = \mathfrak{S} \cup \mathfrak{C}$ into two regions, such that $\mathfrak{S}$ consists of all initial states $\mathbf{x}_0$ where it is optimal to declare the epidemic, and $\mathfrak{C}$ is its complement, where it is optimal to wait.

2.3. Reduction to a Model Predictive Control Problem

The characterization in (5) is implicit, since it features $V(\cdot)$ on both sides of the expression. Specifically, the value function V corresponds to a fixed point [8] of the functional operator $\mathcal{L}$, defined by $(\mathcal{L}v)(\mathbf{x}) := \min\left(d(\mathbf{x}), E\left[v(\mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right]\right)$. To solve for $V(\mathbf{x})$, a basic strategy is then to apply Picard-type fixed-point iterations. In other words, given some initial guess $V^{(0)}(\mathbf{x})$, we build a sequence of approximations via $V^{(k)} := \mathcal{L}V^{(k-1)}$, or explicitly,

$$V^{(k)}(\mathbf{x}_0) = \min\left(d(\mathbf{x}_0), E\left[V^{(k-1)}(\mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}_0\right]\right). \tag{7}$$

However, to guarantee the convergence of $V^{(k)}$ does not appear tractable, and the practical performance of (7) is very sensitive to the initial guess $V^{(0)}$. To circumvent this challenge, we rely on the concept of model predictive control. To wit, we introduce an auxiliary parameter t which can be intuitively thought of as *forward time*. The value functions $V(t, \cdot)$ and detection maps $\mathfrak{S}_t$ are now also indexed by t. We start with the trivial initial condition $V(0, \mathbf{x}) := d(\mathbf{x})$, which corresponds to $\mathfrak{S}_0 \equiv \mathcal{X}$. Next, mimicking the classical dynamic programming on finite horizon, we define

$$V(t, \mathbf{x}_0) := \min\left(d(\mathbf{x}_0), E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}_0\right]\right), \quad t = 1, 2, \ldots. \tag{8}$$

Define the Q-value, also known as costs-to-go by

$$q(t, \mathbf{x}) := E[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}]. \tag{9}$$

Then the stopping set at iteration t is

$$\mathfrak{S}_t := \{\mathbf{x}_0 \in \mathcal{X} : q(t, \mathbf{x}_0) - d(\mathbf{x}_0) > 0\}, \quad t = 1, 2, \ldots. \tag{10}$$

We may "unroll" the expectation encoded in $V(t-1, \mathfrak{X}_1)$ to write

$$q(t, \mathbf{x}_0) = E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}_0\right] = E\left[c(\mathfrak{X}_{0:\tau^{(t)}})|\mathfrak{X}_0 = \mathbf{x}_0\right], \tag{11}$$

where $\tau^{(t)} = \min\{s \geq 1 : \mathfrak{X}_s \in \mathfrak{S}_{t-s}\}$. This justifies the interpretation of $q(t, \cdot)$ as costs-to-go, since $c(\mathfrak{X}_{0:\tau^{(t)}})$ are indeed the future costs associated with *not* stopping immediately.

Figure 2 illustrates the first step of the recursion (8) at $t = 1$. In the plot we compare

$$E\left[V(0, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right] = E[d(\mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}]$$

8
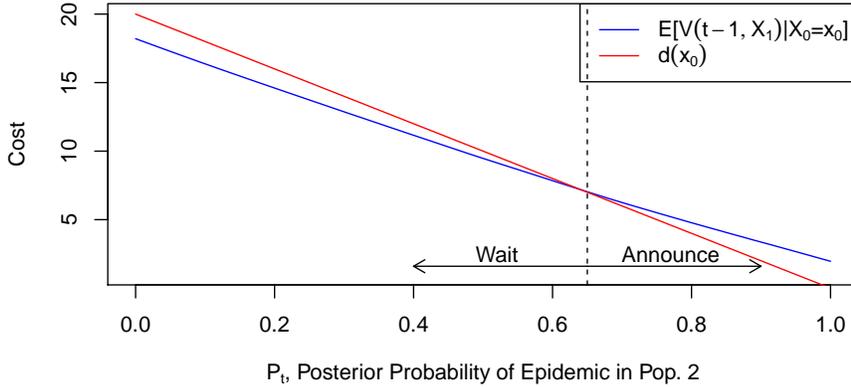
Figure 2: Detection strategy at iteration $t = 1$. The example is based on the model of Section 4, with parameters in Table 1. In the plot, the state of Pool 1 is held fixed at $S_0^{(1)} = 1990$, $I_0^{(1)} = 10$.

against $d(\mathbf{x})$. As discussed, the epidemic is announced when $E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right] > d(\mathbf{x})$ (the right side of the plot). In the opposite case (the left side of the plot), the optimal decision is to wait. As shown by the Figure, the structure of the decision map is driven by the regions where these two quantities are equal to each other, which corresponds to the detection *boundary*,

$$\partial\mathfrak{S}_t := \{\mathbf{x} : E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right] = d(\mathbf{x})\}. \tag{12}$$

The stopping region $\mathfrak{S}_t$ is our detection rule at iteration step t. It can be characterized as the optimal detection rule among all strategies in $\mathcal{S}^{(t)} = \{\tau \in \mathcal{F} : \tau \leq t\}$ that are upper-bounded by t (By construction, $\tau^{(t)} \leq t$). As $t \to \infty$, we have that the set of admissible rules expands $\mathcal{S}^{(t)} \nearrow \mathcal{S}$, and hence we expect that $\mathfrak{S}_t \to \mathfrak{S}$ and $V(t, \mathbf{x}) \to V(\mathbf{x})$. Intuitively, for large t, the recursively defined (8) converges to a stationary case that ought to be the fixed point defining $V(x)$ in (4). The above convergence can be improved via model predictive control (also known as receding horizon control) [40] which applies the fixed detection map $\hat{\mathfrak{S}}^{(k)}$, rather than the time-dependent $\hat{\mathfrak{S}}_t$ at each step, cf. Section 5.1.

3. Epidemic Model

3.1. Multiple Population SIR model

A susceptible-infected-recovered (SIR) model provides an aggregate "gravity" view of the epidemic by focusing on three basic types of individuals in the

9
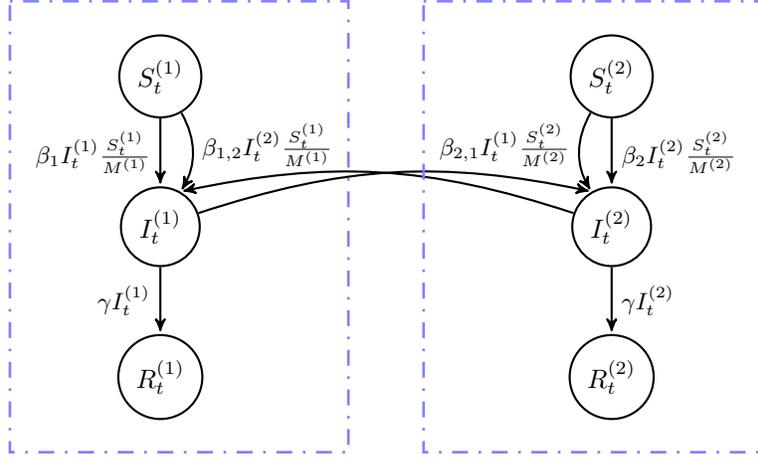
Figure 3: Stochastic SIR epidemic model in two populations.

population: susceptible, infected and recovered. Susceptible individuals are the ones who haven't experienced the disease yet. Interaction between an infected and susceptible individuals can lead to an infection. Thus, contacts stochastically generate new infecteds who in turn can further infect other susceptible individuals. After some time an infected individual recovers and becomes immune (i.e. becomes a Recovered): he/she can no longer infect others or get infected.

While the detection problem is specified at the discrete instances $t = 1, 2, \ldots$, for describing outbreak dynamics it is more convenient to work with continuous-time dynamical systems. As in [10, Ch. 6], we thus first recall the multi-type stochastic SIR model in continuous time. The overall epidemic state at epoch $t \in \mathbb{R}_+$ is denoted by $\{\mathbf{S}_t, \mathbf{I}_t, \mathbf{R}_t\}$, where $\mathbf{S}_t = \{S_t^{(1)}, \ldots, S_t^{(K)}\}$, $\mathbf{I}_t = \{I_t^{(1)}, \ldots, I_t^{(K)}\}$ and $\mathbf{R}_t = \{R_t^{(1)}, \ldots, R_t^{(K)}\}$ are vectors denoting the count of susceptible, infected and recovered individuals in each of $1 \leq k \leq K$ meta-populations. We assume that the pool size of each meta-population is fixed at $M^{(k)} = S_t^{(k)} + I_t^{(k)} + R_t^{(k)}$. As a result, we omit further mention of $R_t^{(k)}$ since it can be found from $R_t^{(k)} = M^{(k)} - S_t^{(k)} - I_t^{(k)}$.

The continuous evolution of the state process $\{\mathbf{S}_t, \mathbf{I}_t\} \in \{(\mathbf{s}, \mathbf{i}) : s_k + i_k \leq M^{(k)} \ \forall k \leq K\}$ is described through Markov chain or stochastic kinetic system language. Namely, the epidemic state is piecewise constant in time. Next, there are $2K + K(K-1)$ possible transitions, described by the reaction channels [41]:

$$
\left\{
\begin{array}{lll}
\text{Infection} & S^{(k)} + I^{(k)} \to 2I^{(k)} & \text{w/rate} \quad \beta_k I^{(k)} \frac{S^{(k)}}{M^{(k)}} \\
\text{Transmission} & S^{(k)} + I^{(k')} \to I^{(k)} + I^{(k')} & \text{w/rate} \quad \beta_{k,k'} I^{(k')} \frac{S^{(k)}}{M^{(k)}} \\
\text{Recovery} & I^{(k)} \to \emptyset & \text{w/rate} \quad \gamma I^{(k)}
\end{array}
\right\}
$$
(13)

where each reaction is further indexed by $1 \leq k \leq K$ and $k' \leq K$, $k' \neq k$.

The first transition represents an infection of a susceptible individual by an infected individual from the same pool k. This transition happens at rate $\beta_k I_t^{(k)} \frac{S_t^{(k)}}{M^{(k)}}$, where $1 \le k \le K$ and β_k is a contact rate of infected and susceptible individuals within the k-th meta-population.

The second transition is a transmission: an infection of a susceptible individual from pool k by an infected individual from a different pool k'. The frequency of such infections is $\beta_{k,k'} I_t^{(k')} \frac{S_t^{(k)}}{M^{(k)}}$, where $1 \le k, k' \le K$ and $\beta_{k,k'}$ is a contact rate of infected and susceptible individuals from different populations. Since contacts between individuals from different populations are less frequent, $\beta_{k,k'} \ll \beta_{k'}$. To reduce the number of parameters, we thus assume that cross-population interactions occur at rate $\beta_{k,k'} \equiv \alpha \beta_{k'}$, where α is the proportion of "travelers" in each pool. Thus, cross-contacts happen at the fraction α of a contact rate within one population; a typical range is $\alpha \in [0.01, 0.2]$.

The last transition in (13) is a recovery and subsequent immunity of an infected individual in a population k. The rate of transition is $\gamma I_t^{(k)}$, where γ is a recovery rate, independent of the pool index k. This can be interpreted as individuals staying infected for an Exponentially distributed time with mean $1/\gamma$.

In this paper we focus on two-population models, positing that the outbreak begins in Pool 1 and may subsequently spread to Pool 2, where it is to be detected. Accordingly, we will be fusing information from Pool 1 and Pool 2 to identify the onset of epidemic in Pool 2. We assume that the two pools have similar characteristics, so that all parameters are homogenous in $k = 1, 2$. Thus, the two-population SIR model, shown in Figure 3 has 3 parameters α – the mixing parameter between two populations, β – the within-pool contact rate of infected and susceptible individuals, and γ – the recovery rate. These parameters are assumed to be known and are a function of the modeled disease family (e.g. influenza or dengue fever), the demographics and public health characteristics of the populations, and the travel patterns across pools.

3.2. Partial Observation

Under full observations the detection problem (4) would be trivial, since one can directly track $I_t^{(2)}$ and declare an outbreak as soon as there any infecteds in the second pool. However, realistically $I^{(2)}$ is not observed. Some of the reasons include mis-diagnoses among infecteds, patients not seeking care, false positives, mis-reporting or lack of reporting of epidemiological data, etc. Consequently, we assume that the true size of the S/I/R compartments in Pool 2 is not known. To simplify the presentation, we assume that $I_t^{(1)}$ *is observed* in Pool 1, perhaps due to better epidemiological surveillance in that pool.

In our detection problem, the main event of interest is the presence of any infecteds in Pool 2, $\{I_t^{(2)} > 0\}$. Accordingly, we consider $\tilde{P}_t = P(I_t^{(2)} > 0 | \mathcal{G}_t)$, the posterior probability that the epidemic started in the second population given the limited knowledge about it available by time t, here summarized by some information set $\mathcal{G}_t$. Depending on assumptions about the observations

structure, $\tilde{P}_t$ may be available in closed form (e.g. through Bayesian conjugate updating [30, 2]) or may have to be only approximately computed through e.g. particle filtering methods [3, 4]. The latter method, which computes the whole posterior distribution $\pi_t \sim I_t^{(2)}|\mathcal{G}_t$, is computationally expensive, while conjugate updating requires carrying several sufficient statistics about the posterior of $I_t^{(2)}$. In either case, $\tilde{P}_t$ on its own is not Markovian, and hence does not possess simple dynamics. Therefore we propose a model that works with a simplified, Markovian version of $\tilde{P}_t$, which we denote as P_t.
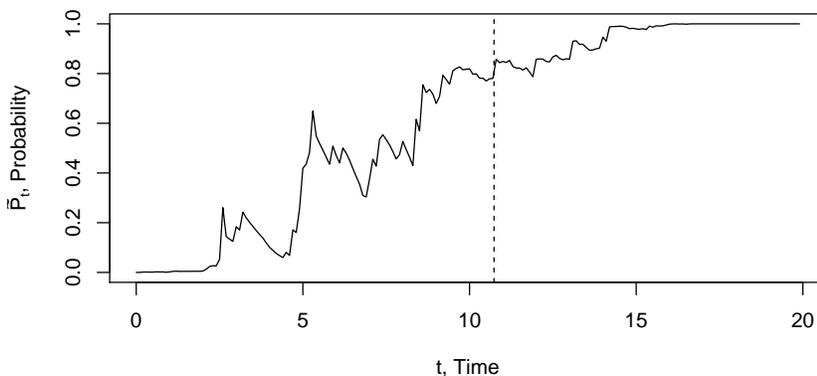
Figure 4: Posterior probability that the epidemic started in the second population $\tilde{P}_t$ vs. time t, where the dashed line represents the actual start time θ of outbreak in Pool 2. The plot was constructed using particle filtering using the following model parameter values: $\beta = 0.75$, $\alpha = 0.01$, $\gamma = 0.5$, $M^{(1)} = M^{(2)} = 2000$.

Figure 4 shows a sample scenario of the evolution of $\tilde{P}_t$ in a partially observed framework. The plot was generated using particle filtering and used the two-pool model (13) with noisy Poisson-type observations in each pool [42]. We observe that $\tilde{P}_t$ tends to drift up (i.e. posterior probability of outbreak increases over time) and eventually hits 1.

3.3. Reduced Model

Our reduced model consists of the state of epidemic in the first population $\left\{ S_t^{(1)}, I_t^{(1)} \right\}$ and a process P_t that is interpreted as the probability that the epidemic reached Pool 2 conditional on the information $\mathcal{G}_t = \sigma(S_{0:t}^{(1)}, I_{0:t}^{(1)})$ from Pool 1. The first two components $S_t^{(1)}$ and $I_t^{(1)}$ come from a one-population SIR model (see the definition of SIR model for $K = 1$ in Section 3.1). To prescribe the dynamics of the pseudo-posterior P_t, we decompose the event $\{I_t^{(2)} > 0\} \equiv \{\theta \leq t\}$ into two cases: the event that the epidemic already started at time $t - 1$ (i.e. $\theta \leq t - 1$), and the event that it starts at $t = \theta$. We

also add some stochastic noise to denote exogenous fluctuations in our posterior estimates regarding the second pool. In total, we thus assume that

$$P_t = P_{t-1} + P(I_{t-1}^{(2)} = 0 \text{ and } I_t^{(2)} > 0 | \mathcal{G}_{t-1}) + \delta_t, \tag{14}$$

where δ_t are i.i.d. noise terms. Intuitively, the probability of outbreak has a positive drift over time, and the drift is precisely the posterior probability of the outbreak beginning during the current period, $\{\theta \in [t-1, t]\}$.

From the SIR dynamics, the probability that $\{\theta \in [t-1, t]\}$ conditional on Pool-1 observations up to previous stage $t-1$, is equal to the probability that an infected from Pool 1 interacts with a susceptible from Pool 2, times the conditional probability that $\{\theta > t-1\}$. The former happens with rate $\alpha \beta I_s^{(1)} \frac{S_s^{(2)}}{M^{(2)}}$, $s \in [t-1, t]$, while the latter event is the complement of $\{\theta \le t-1\}$ and hence has probability $1 - P_{t-1}$. Using the fact that conditional on $\{\theta \ge t\}$, $M^{(2)} = S_{t-1}^{(2)}$, and making the transition rate constant on $[t-1, t]$ we obtain

$$P(\theta \in [t-1, t] | \mathcal{G}_{t-1}) \simeq \alpha \beta I_{t-1}^{(1)} (1 - P_{t-1}). \tag{15}$$

To guarantee $P_t \in [0, 1]$ is interpretable as probability we confine it to $[0, 1]$, yielding

$$P_t := \begin{cases} 0 \vee (P_{t-1} + \alpha \beta I_{t-1}^{(1)} (1 - P_{t-1}) + \delta_t) \wedge 1, & \text{if } P_{t-1} \ne 1 \\ 1, & \text{if } P_{t-1} = 1 \end{cases} \tag{16}$$

In our simulations we use centered Gaussian noise $\delta_t \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, \sigma_\delta^2)$ with variance σ_δ^2, however it can take any distribution. Note that $P_t = 1$ is an absorbing state, representing certainty that the outbreak reached Pool 2, while $P_t = 0$ is a boundary case, since even if it is certain that the outbreak is currently not in Pool 2, it can still get cross-infected in the future. Similar features hold for the true posterior probability $\tilde{P}_t$, cf. Figure 4. Alternative models for probability of outbreak P_t, are discussed in Section 6.

Remark 2. *Note that* (16) *is in discrete-time; to connect to the continuous-time dynamics of SIR one could take the limit as the time increment goes to zero, obtaining a diffusive model $dP_t = \alpha \beta I_t^{(1)} (1 - P_t) \, dt + \delta dW_t$ where (W_t) is a Brownian motion. However, since detection is assumed to take place only at instances $t = 1, 2, \ldots$, we prefer to work with* (16) *as is.*

3.4. Detection within the Reduced Model

To sum up, the developed reduced 2-pool model has a 3-dimensional state $\{\mathfrak{X}\}_t = \left(S_t^{(1)}, I_t^{(1)}, P_t \right)$ with state space

$$\mathcal{X} := \{(s, i, p) : s, i \in \mathbb{N}, s + i < M^{(1)}, p \in [0, 1]\}.$$

Figure 5 shows a few sample trajectories of $\mathfrak{X}$ to illustrate the resulting dynamics.
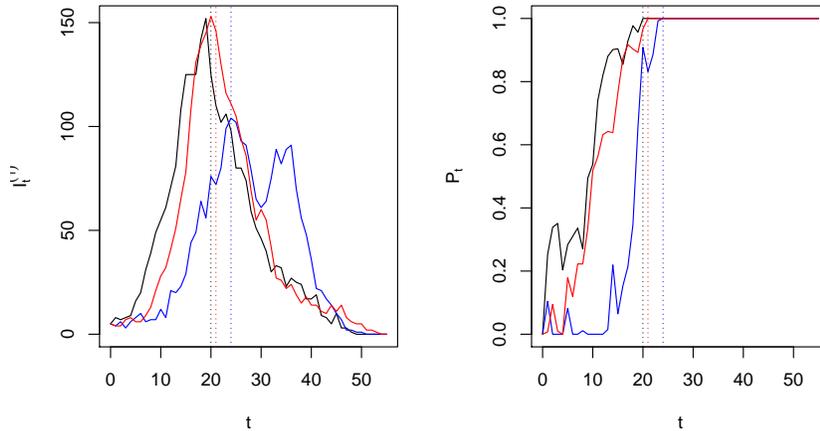
Figure 5: Three sample trajectories of $\mathfrak{X}$ with the initial condition $S_0^{(1)} = 1995$, $I_0^{(1)} = 5$, $P_0 = 0$ and outbreak parameters from Table 1. Left panel is the plot of $\{I_t^{(1)}\}$, the number of infecteds in the first population, and right panel is the plot of $\{P_t\}$, the posterior probability that the epidemic started in the second population. The vertical dotted lines represent times when P_t hits 1 and outbreak becomes certain.

Our detection problem (10) relies on the computation of the immediate and future expected costs $E\left[c(\mathfrak{X}_{0:\tau})|\mathfrak{X}_0\right]$ and $d(\mathfrak{X}_0)$. Re-writing the definitions of immediate and future costs (2) and (3) in terms of the event $\{I_t^{(2)} > 0\}$, and taking conditional expectation we obtain:

$$d(\mathfrak{X}_0) := C_{\text{FA}}(1 - P_0), \tag{17}$$

$$c(\mathfrak{X}_{0:\tau}) := \sum_{s=0}^{\tau-1} C_{\text{Delay}} P_s + C_{\text{FA}}(1 - P_\tau), \tag{18}$$

where $\tau \in \mathcal{S}$. Rather than in terms of the unobserved $I^{(2)}$, the above expressions are now given in terms of the component P_t, allowing to measure detection costs within the $\mathfrak{X}$-model. Notice that $d(\mathfrak{X}_0)$ is a function of P_0 and $c(\mathfrak{X}_{0:\tau})$ is a function of the future trajectory $\{P_s, s = 0, \ldots, \tau\}$.

Our goal is to find the detection maps $\hat{\mathfrak{S}}_t$ for $t = 1, 2, \ldots$, defined recursively in (10). To do so, at each step we need to evaluate $E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right]$ and $d(\mathbf{x})$. The immediate cost $d(\mathbf{x})$ can be computed exactly via (17). However, the expectation $E\left[V(t-1, \mathfrak{X}_1)|\mathfrak{X}_0 = \mathbf{x}\right]$ can not be computed analytically since there are no closed-form expressions for the distribution of $\mathfrak{X}_{0:\tau}$. In Section 5.3 we present the sequential Regression Monte Carlo approach which offers an efficient way to empirically estimate $\hat{\mathfrak{S}}_t$ based on synthetically generated epidemic scenarios. We then use Model Predictive Control to estimate the stationary detection map $\mathfrak{S}$.

14

4. Case Study

To illustrate the dynamic detection strategy within our 2-pool model, in this section we present a detailed case study. Table 1 summarizes the parameters used. Epidemic parameters are taken to be $\beta = 0.75$ and $\gamma = 0.5$. Thus, the initial reproduction ratio is $\mathcal{R}_0 = \beta/\gamma = 1.5$, which is a moderately infectious epidemic. We assume that the pool mixing parameter is $\alpha = 0.01$, which is reasonable for pools representing well-separated cities or counties. The inference noise in (16) is taken to be Gaussian with variance $\delta_t \sim \mathcal{N}(0, \sigma_\delta^2 = 0.01^2)$. For the detection costs in (17)-(18), we take without loss of generality $C_{\text{Delay}} = 1$ and fix $C_{\text{FA}} = 20$. As we will see, this corresponds to a moderate penalty for false alarms.

Epidemic:	$M^{(1)} = 2000$	$S_0^{(1)} = M^{(1)} - I_0^{(1)}$	$\sigma_\delta = 1/100$
	$\beta = 0.75$	$\alpha = 0.01$	$\gamma = 0.5$
Costs/Penalties:	$C_{\text{FA}} = 20$	$C_{\text{Delay}} = 1$	

Table 1: Outbreak and costs parameters for the case study of Section 4. σ_δ refers to the noise in P, cf. (16).
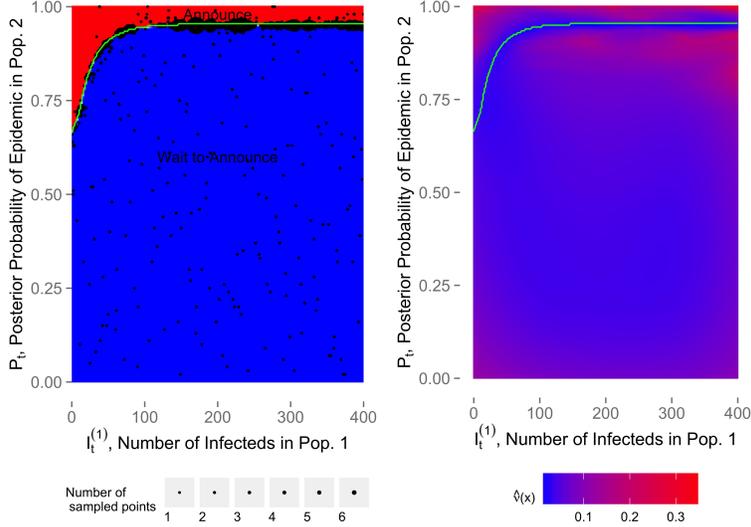
Figure 6: Left panel: detection rule $\mathfrak{S}_{20}^{LP}$ in terms of $I^{(1)}$ and P. The detection boundary $\partial\mathfrak{S}_{20}^{LP}$ is shown with the solid curve. We also show the experimental design $\mathcal{Z}$ that was used, illustrated with the scatterplot. Size of pixels corresponds to the number of times that neighborhood was sampled. Right panel: standard errors $\hat{v}(\mathbf{x})$ from (27). Observe lower standard errors in regions where the design $\mathcal{Z}$ is more dense.

So far the case study features a three-dimensional state $\{S^{(1)}, I^{(1)}, P\}$, so that the resulting detection maps are in 3-D. To aid visualization, we consider

a variant with a reduced dimension. Namely, we drop the component $S^{(1)}$ measuring the number of infecteds in Pool 1. Indeed, at the early stages of the outbreak the ratio $S_t^{(1)}/M^{(1)}$ is approximately one. As a result, one may assume that the rate of infections in Pool 1 is simply $\beta I_t^{(1)}$, which corresponds to the classical branching process epidemic model [10]. It is known [43] that this approximation remains valid up to $t = O(\log(M^{(1)}))$ by which time, $I_t^{(1)} = O(\sqrt{M^{(1)}})$.; therefore it works especially well in large populations, and hence is termed a large-population (LP) approximation. The LP model only has two dimensions, $\mathfrak{X}' := \{I^{(1)}, P\}$ allowing to plot the corresponding 2-D stopping set $\mathfrak{S}^{LP}$.

Figure 6 shows $\mathfrak{S}^{LP}$ generated under the conditions of Table 1 and the above large population assumption. As expected, epidemic detection is triggered once the posterior probability P_t of $\{I_t^{(2)} > 0\}$, is high enough. However, we observe that detection is also highly sensitive to values of $I_t^{(1)}$; for instance detection is progressively delayed as $I_t^{(1)}$ gets bigger. This dependence between the two pools in terms of decision making illustrates the underlying cross-pool information fusion. Intuitively, detection should take place once P_t is high enough. However, conditional on a fixed P_t, larger number of Pool 1 infecteds makes an impending outbreak in Pool 2 more likely, lowering waiting costs. Hence, the detection boundary curves in $I^{(1)}$. Mathematically, recall that in (16), the growth rate of P increases in $I^{(1)}$. As a result, for large values of $I_t^{(1)}$, one may expect that the next-stage P_{t+1} will also be large, i.e. move into the "Announce" region quicker. This again lowers the waiting costs and therefore delays announcement.

4.1. Evaluating Detection Rules

Figure 7 shows dynamic decision-making in the LP model through a collection of generated trajectories of $\mathfrak{X}' = \{I_t^{(1)}, P_t\}$ and their corresponding detection times τ^{LP}, the first time the state process $\mathfrak{X}'$ enters the stopping set $\mathfrak{S}^{LP}$. We observe that the trajectories generally move north-east, as both P and $I^{(1)}$ tend to increase. However, the rate at which they grow and the precise direction are uncertain and vary across scenarios. Consequently, at detection, both $P_{\tau^{LP}}$ and $I_{\tau^{LP}}^{(1)}$ have a nontrivial distribution.

To better understand the detection map $\mathfrak{S}^{LP}$, we analyze the resulting detection strategy given by τ^{LP} and compare it to alternatives. Two classes of simpler detection rules are Threshold-P and Threshold-t. The Threshold-P strategy announces an outbreak as soon as $P_t \geq \bar{P}$ for a given threshold $\bar{P}$. Hence, it acts solely based on local (posterior) information about Pool 2. This mimics the CDC policy [25] of announcing an epidemic when the number of infecteds in Population 2 crosses some pre-specified level. In contrast to the fused detection strategy with a curved detection boundary which jointly takes into account both P_t and $I_t^{(1)}$, Threshold-P rule only uses P_t for detection decisions, yielding a flat, horizontal detection boundary in Figure 7. The threshold-t strategy is a simple non-adaptive strategy that announces at the fixed stage $\bar{t}$.

16
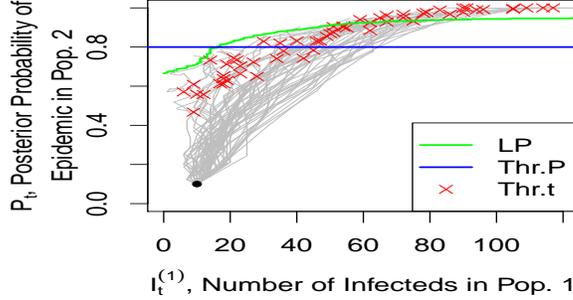
Figure 7: Fifty sampled epidemic trajectories $\{I_t^{(1)}, P_t\}, t = 1, \ldots, \tau$ emanating from the initial state $I_0^{(1)} = 10$ and $P_0 = 0.1$. We show the LP detection boundary (namely $\partial \mathfrak{S}_{20}^{LP}$), as well as a threshold strategy that announces epidemic as soon as $P_t \geq \bar{P} = 0.8$. Lastly, the red crosses denote the locations of the trajectories at $t = 8$, which is the basis of the alternate Threshold-t strategy.

It is illustrated in Figure 7 where we record the joint distribution of $I_{\bar{t}}^{(1)}, P_{\bar{t}}$ at $\bar{t} = 8$.

	Detection time τ		Realized Cost Q		PFA $E[1 - P_\tau]$
	Mean	StDev.	Mean	StDev.	
Optimal	8.86	2.59	6.53	1.70	8.2%
LP	9.32	2.95	6.57	1.81	6.4%
Threshold-P	7.88	2.85	7.03	1.58	15.3 %
Threshold-t	8.00	N/A	7.18	2.21	14.4 %

Table 2: Comparison of Optimal, Large Population(LP), Threshold-P with $\bar{P} = 0.8$ and Threshold-t with $\bar{t} = 8$ strategies. Statistics are based on 1000 synthetic trajectories of $\{I^{(1)}, S^{(1)}, P\}$, where $Q = c(\mathfrak{X}_{0:\tau^{(t)}})$.

Returning to the full 3-D model with state $\mathfrak{X}$ we evaluate the resulting optimal detection strategy τ^* and proceed to compare its performance against the other potential detection rules discussed above. Specifically, the first two alternatives are a Threshold-P rule with $\bar{P} = 0.8$ (declare an epidemic if its probability is above 80%) and a Threshold-t strategy with $\bar{t} = 8$. The latter was found to be the best strategy among those that declare outbreak at a fixed stage. The last alternative is the LP strategy τ^{LP} from last section. Recall that τ^{LP} makes decisions while ignoring $S^{(1)}$. In that sense, when applied to the full 3-D model, it gives a simplified, but still adaptive, detection rule. To recap, Threshold-t strategy is completely non-adaptive; Threshold-P only relies on P_t; LP relies on $\{I_t^{(1)}, P_t\}$, and Optimal strategy uses all of $\{S_t^{(1)}, I_t^{(1)}, P_t\}$.

To compare the performance of the above competing strategies, we fixed the initial condition at $S_0^{(1)} = 1990, I_0^{(1)} = 10$ and $P_0 = 0.1$, so that there

17
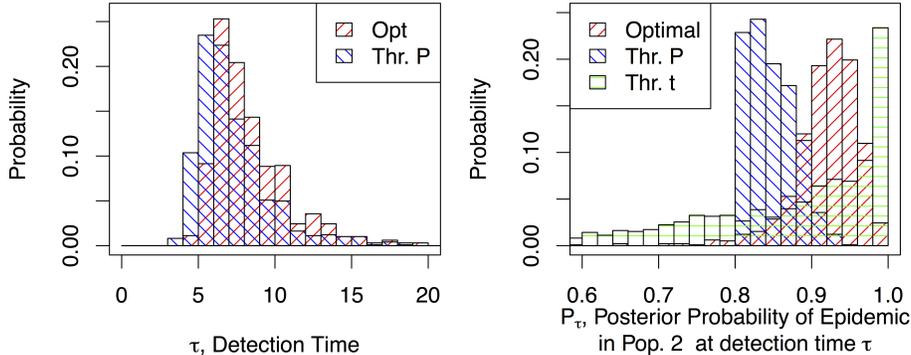
Figure 8: Summary statistics of different detection strategies constructed from 1000 sample epidemic trajectories. The LP detection strategy is from Figure 6. Right: Distribution of detection times τ; Left: Distribution of posterior probability of outbreak in Pool 2 at detection time, P_τ.

are 10 infecteds in Pool 1 and 10% prior probability of epidemic already in Pool 2. Then we simulated 1000 epidemic trajectories $\{\mathbf{x}_{0:\tau}^n\}$, $n = 1, \ldots, 1000$, emanating from this fixed initial condition up to the detection time τ (which depends in turn on the strategy used). Table 2 then presents the resulting summary statistics based on these frozen 1000 trajectories (note that there are no analytic formulas to obtain these metrics, so we have to resort to simulation).

The comparison is done in terms of several different metrics, including realized detection costs $c(\mathfrak{X}_{0:\tau^{(t)}})$, distribution of detection times τ, and frequency of false alarms, represented by $d(\mathfrak{X}_\tau) = 1 - P_\tau$ in our setup. As expected, the Optimal strategy with detection time τ^* that directly optimizes the cost-benefit in the full model performs best. The corresponding expected costs are $V(\mathbf{x}_0) \simeq 6.53$, with average detection time $E[\tau^*] \simeq 8.86$. It outperforms the Threshold-P strategy by about 7% in terms of reducing detection costs, and the Threshold-t strategy by about 9%. These are nontrivial cost savings which highlight the benefit of information fusion. Table 2 also shows that the 2-D LP approximation performs well in this example, generating very similar expected costs. At least for this case study, detection happens early enough that the branching process approximation of the outbreak works fine.

Recall that our model is stochastic and generates adaptive detection strategy. Hence the detection time τ^* is a random variable. As shown in Table 2, the corresponding standard deviation $StDev(\tau^*) \simeq 2.6$ is substantial. This illustrates the sub-optimality of the Threshold-t strategy that stops at a fixed $\bar{t}$ with $StDev(\bar{t}) = 0$ trivially. Not surprisingly, the ability to delay or speed up outbreak announcements based on latest data are crucial for optimizing policy making. We also note that compared to the Threshold-P strategy, the Optimal

18

strategy tends to announce later, $E[\tau^*] \simeq 8.86 > 7.88 \simeq E[\tau^{Thr-P}]$, this is also confirmed by the respective histograms of τ^* and τ^{Thr-P} in Figure 8. However, we emphasize that the detection rules do not have a clear ordering. In other words, the random variables τ^*, τ^{Thr-P}, etc., cannot be directly compared.

A complementary metric of detection quality is provided by the probability of false alarms, $PFA := E[1 - P_\tau]$. For the optimal strategy we find that $PFA^* = 8.2\%$. In contrast, for Threshold-P strategy, we have $PFA^{Thr-P} = 15.3\%$. Note that because we use a discrete-time model, at time of detection P_τ will strictly exceed the threshold $\bar{P} = 0.8$, hence $PFA^{Thr-P} < 1 - \bar{P}$. The histograms of P_τ are shown in Figure 8 and confirm the qualitative difference among the detection strategies. The Threshold-P strategy only stops once $P_t > \bar{P}$, so that P_τ has support on roughly $[0.8, 0.9]$. In contrast, the adaptive Optimal (and LP) strategies, have a much wider range for P_τ. In particular, sometimes epidemics are announced even before P_t hits the level 0.8.
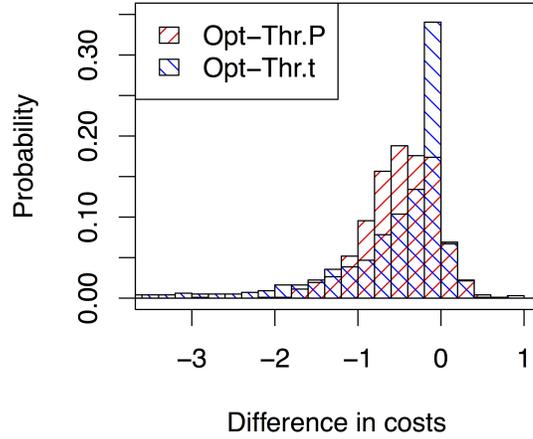
Figure 9: Relative realized detection costs across different strategies. The histogram shows the distribution of the difference in costs along the 1000 simulated trajectories, namely $c(\mathbf{x}_{0:\tau^*}) - c(\mathbf{x}_{0:\tau^{Thr-P}})$, and $c(\mathbf{x}_{0:\tau^*}) - c(\mathbf{x}_{0:\tau^{Thr-t}})$.

To further quantify the improvement provided by the Optimal detection rule, Figure 9 gives a scenario-by-scenario comparison of relative realized detection costs. Note that in hindsight, τ^* may sometimes perform worse that τ^{Thr-P} or even τ^{Thr-t}. Figure 8 plots the histogram of the difference in costs for each trajectory $\mathbf{x}_{0:t}^n$, $n = 1, \ldots, 1000$, namely $c(\mathbf{x}_{0:\tau^*})$, $c(\mathbf{x}_{0:\tau^{Thr-P}})$, and $c(\mathbf{x}_{0:\tau^{Thr-t}})$. We find that the costs computed with Optimal/LP strategies are smaller than costs computed with Threshold strategies for more than 80% of the trajectories.

To sum up, we observe material improvement from using Optimal detection rule in this case study. Moreover, the obtained detection rule is substantially different from the thresholding protocol. On the one hand, the adaptive detection

time τ^* exhibits a wide spread and is highly non-constant across trajectories. On the other hand, the posterior probability of false alarms P_{τ^*} is also strongly variable. As a result, the average frequency of false alarms is drastically lowered relative to Threshold-P strategy, reducing overall expected costs.

4.2. Effect of Detection Cost Parameters
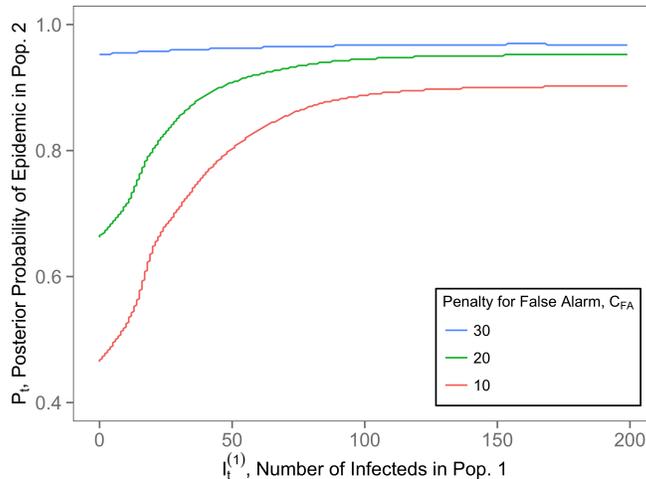
Figure 10: Boundaries of detection maps $\partial \mathfrak{S}_{20}^{LP}$ constructed based on different penalties for false alarm, C_{FA}.

C_{FA}	τ^*		Cost		$PFA = E[1 - P_{\tau^*}]$
	Mean	StDev.	Mean	StDev.	
10	6.84	1.62	5.32	0.99	21.4 %
20	8.87	2.60	6.54	1.71	8.3%
30	9.61	2.79	7.21	2.22	5.3%

Table 3: Summary statistics of the Optimal detection strategy τ^* for different false alarm penalties C_{FA}.

The main parameter in our quickest detection setup is the ratio of the cost of false alarms and the cost of detection delay, $C_{\mathrm{FA}}/C_{\mathrm{Delay}}$. A high ratio penalizes premature announcements and requires more care in the assessment of the potential outbreak in Pool 2. A low ratio invites more aggressive actions. To better understand the role of this ratio, in Figure 10 we show several detection boundaries $\partial \mathfrak{S}^{LP}$ corresponding to varying C_{FA}, while $C_{\mathrm{Delay}} = 1$ is kept fixed. As expected, a lower C_{FA} enlarges the Announce set $\mathfrak{S}$. In particular, the boundary $\partial \mathfrak{S}$ shifts down and to the right. As a result, starting from a fixed location $(I_0^{(1)}, P_0)$, the stopping set $\mathfrak{S}$ will be reached sooner, so that τ decreases

(in the sense of stochastic dominance for the corresponding random variables). This is confirmed in Table 3 that reports statistics for τ^* and various C_{FA}. We find that $E[\tau^*] = 8.86$ when $C_{\text{FA}} = 20$, but is only $E[\tau^*] = 6.84$ for $C_{\text{FA}} = 10$. Simultaneously, the frequency of premature announcements PFA will increase. The precise relationship is however nonlinear. Lowering C_{FA} from 20 to 10, the PFA rises dramatically to about 21% from 8%. Conversely, raising C_{FA} to 30 only reduces PFA to 5.3%. A common approach in the decision literature is to select a priori a desired level of PFA (say $PFA = 10\%$) and then numerically solve the inverse problem to obtain the corresponding C_{FA} and hence the corresponding detection rule $\mathfrak{S}$.

5. Numerical Implementation

To find the detection maps $\hat{\mathfrak{S}}_t$ for $t = 1, 2, \ldots$, defined recursively in (10) we use approximate dynamic programming techniques. In particular, we rely on the Regression Monte Carlo approach [44, 9] to approximate the conditional expectation map over $\mathbf{x} \in \mathcal{X}$.

5.1. Regression Monte Carlo

For the remainder of this section the auxiliary "time" variable t is fixed and the goal is to approximate the conditional expectation $q(t, \mathbf{x}) := E[c(\mathfrak{X}_{0:\tau^{(t)}} | \mathfrak{X}_0 = \mathbf{x}]$ in (11). Recall that at step t, detection rules are restricted to satisfy $\tau^{(t)} \leq t$. The Regression Monte Carlo technique approximates $q(t, \cdot)$ by a predicted surrogate value $\hat{q}(t, \cdot)$ which is based on a statistical regression framework.

The surrogate prediction is built using data simulated from the specified model. To do so, a design $\mathcal{Z} := \{\mathbf{x}_0^n, \ n = 1, \ldots, N\}$ of N locations is first generated. Next, we generate the corresponding scenarios $\{\mathfrak{X}_{0:t}^n\}$ with the initial value $\mathfrak{X}_0^n = \mathbf{x}_0^n$, one scenario for each initial location. Define

$$\tau_t^n := \min\{s \geq 1 : \mathfrak{X}_s^n \in \mathfrak{S}_{t-s}\}, \tag{19}$$

which leads to path-wise waiting costs $q^n := c(\mathfrak{X}_{0:\tau_t^n}^n)$ using formula (3) on the n-th scenario. The aggregate dataset is

$$Z = \{(\mathbf{x}_0^n, q^n), n = 1, \ldots, N\}. \tag{20}$$

The construction of $\hat{q}(t, \cdot)$ then involves response surface modeling, i.e. determining the relationship between the initial condition $\mathbf{x}$ and the mean of the sampled $Q | \mathbf{x} \equiv c(\mathfrak{X}_{0:\tau_t})$. Statistically, we start with

$$Q | \mathbf{x} = q(t, \mathbf{x}) + \epsilon, \tag{21}$$

where $q(t, \cdot)$ is the true response surface, $Q = c(\mathfrak{X}_{0:\tau^{(t)}})$ are random scenario-based costs, and ϵ are mean-zero residuals with variance σ^2 arising from Monte Carlo simulations. Empirically, (21) translates into regressing $\{q^n\}$ on $\{\mathbf{x}_0^n\}$,

$n = 1, \ldots, N$; this step is discussed in section 5.2. After determining $\hat{q}$, and using (10) the estimated detection rule $\hat{\tilde{\mathfrak{S}}}_t$ is

$$\hat{\tilde{\mathfrak{S}}}_t := \{\mathbf{x} : \hat{q}(t, \mathbf{x}) - d(\mathbf{x}) > 0\}. \qquad (22)$$

The above provides a recipe to obtain an (approximate) $\hat{\tilde{\mathfrak{S}}}_t$ using the collection of detection rules $\hat{\tilde{\mathfrak{S}}}_{1:t-1}$. Iterating over t, yields the sequence of detection maps $\hat{\tilde{\mathfrak{S}}}_t$ for $t = 1, 2, \ldots$. We recall that as $t \to \infty$, we expect $\hat{\tilde{\mathfrak{S}}}_t$ to stabilize and tend to a time-invariant detection map. Such convergence is illustrated in Figure 11, where we trace the boundaries $\partial \tilde{\mathfrak{S}}_t$ for $t = 1, \ldots, 20$. Convergence takes hold after about 15 iterations and suggests that $\hat{\tilde{\mathfrak{S}}}_{20} \simeq \mathfrak{S}$; this is what we used for Figures 6-10 where the boundary of $\hat{\tilde{\mathfrak{S}}}_{20}$ was taken as the final output of the Algorithm.

The detection rule (19) is time-dependent since it utilizes a new $\hat{\tilde{\mathfrak{S}}}_{t-s}$ at each stage s. Model predictive control simplifies this feature with a time-invariant rule that simply utilizes $\hat{\tilde{\mathfrak{S}}}_{t-1}$ (that we relabel as $\hat{\tilde{\mathfrak{S}}}^{(t-1)}$ for typographical distinction). Indeed, as Figure 11 shows, the early maps $\hat{\tilde{\mathfrak{S}}}_1, \hat{\tilde{\mathfrak{S}}}_2, \ldots$, are not as accurate as $\hat{\tilde{\mathfrak{S}}}_{t-1}$ for t large, so it makes sense to completely "forget" them and rely just on the last iteration step. Accordingly, we implement a blend of (7) and (8) by first using (19) over $t = 1, 2, \ldots, t^*$ and then switching to a receding-horizon rule

$$\tau_t^{MPC} := \min\{s \geq 1 : \mathfrak{X}_s \in \mathfrak{S}_{t-1}\}, \qquad t = t^*, t^* + 1, \ldots. \qquad (23)$$

The above MPC iterations are terminated once $\hat{q}(t, \mathbf{x})$ and $\hat{q}(t+1, \mathbf{x})$ do not change much, namely $\|\hat{q}(t, \cdot) - \hat{q}(t+1, \cdot)\|_{L^\infty} < Tol$ for a specified tolerance level Tol.

5.2. Regression Model

Because we have limited a priori knowledge about the structure of the detection rule, it is preferable to work with a nonparametric regression architecture for $q(t, \mathbf{x})$. (For example a linear regression model for q would imply that $\mathfrak{S}$ in (22) is defined through linear constraints, i.e. forms a simplex in $\mathcal{X}$.) In addition, nonparametric regression is typically more robust for dealing with the non-Gaussian residuals ϵ that arise in our model.

There are numerous nonparametric regression frameworks that can be used, including splines, Gaussian processes, or generalized additive models; see e.g. the classic monograph [45]. Note that even though $\mathbf{x} \mapsto V(\mathbf{x})$ is continuous, some discontinuous response surfaces might also be helpful, such as random forests or dynamic trees [9]. In the present article we take up a variant of local linear regression, known as Loess. Loess fits weighted linear regression models to localized subsets of data, determined using a kernel function, specifically a k-nearest-neighbor algorithm [46]. Compared to classical linear models, Loess better handles outliers and heteroscedasticity, and also does not make assumptions about the global shape of the response surface.
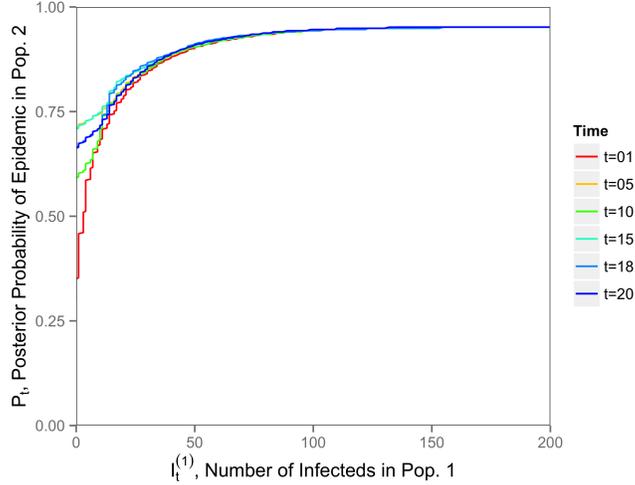
Figure 11: Convergence of the detection boundaries $\partial\hat{\mathfrak{S}}_t$ over $t = 1$ to $t = 20$ for the 2-D LP detection rule from Section 4.

The Loess response model is of the form

$$\hat{q}(t, \mathbf{x}) = \sum_{i=1}^{r} \hat{\beta}_i(\mathbf{x}) B_i(\mathbf{x}) \tag{24}$$

where $B_i(\cdot)$ is the set of r pre-specified basis functions and $\hat{\beta}_i$ are estimated regression coefficients at $\mathbf{x}$. Given input matrix $\vec{X}$ and matching response vector Q, $\hat{\beta}$ is fitted using local least-squares minimization

$$\hat{\beta}(\mathbf{x}) := \arg\min_{\vec{\beta} \in \mathbb{R}^r} K_\lambda(\mathbf{x}, \vec{X})(Q - B(\vec{X})^T \vec{\beta})^2, \tag{25}$$

where $K_\lambda(\mathbf{x}, \vec{X})$ is the weighting kernel. The idea behind the kernel is to base the predicted $\hat{q}(t, \mathbf{x})$ on the samples in the neighborhood of $\mathbf{x}$, weighted by their distance from $\mathbf{x}$ [45, Sec. 2.8.2]. The size of the neighborhood is controlled by the smoothing parameter λ. If $\lambda < 1$, only a proportion λ of the samples will be used in fitting. The smaller λ, the more "wiggly" the fit $\hat{q}(t, \cdot)$ is going to be since fewer samples are used in computing $\hat{\beta}(\mathbf{x})$. Loess can be viewed as a special kernel regression method, with the prediction being a weighted average of the responses q^n: $\hat{q}(t, \mathbf{x}) = \sum_n l_n(\mathbf{x})q^n$ for the equivalent kernel $l(\cdot)$. In our numerical examples, we use the implementation of Loess provided in the R by the built-in package stats [47], which uses a tri-cubic kernel and linear, first-order basis functions; the smoothing parameter was $\lambda = 0.4$.

5.3. Experimental Design

The aim of the response surface is to maximize the accuracy of $\hat{\mathfrak{S}}_t$. This is equivalent to maximizing model fidelity along the boundary of the detection

23

map. Statistically, for a localized response surface, accuracy is primarily driven by the local density of the input data that is specified by the experimental design $\mathcal{Z}$. Hence, to maximize our confidence regarding the boundary of $\mathfrak{S}_t$ in (22), we generate appropriate, adaptively chosen experimental designs $\mathcal{Z}$. This is achieved using the Sequential RMC framework introduced by [9]. SRMC uses tools from active learning/Bayesian optimization to gradually *grow* the design $\mathcal{Z}$ so as to zoom-in to the boundary of $\hat{\mathfrak{S}}_t$. This is done by first quantifying the accuracy of the existing response surface, and then adding new design sites so as to maximize information gain. See [9, 48] for details. The SRMC approach is illustrated in Figure 6 where the adaptively generated experimental design $\mathcal{Z}$ (of size 2000 in the figure) is highly concentrated around the detection boundary $\partial\mathfrak{S}$. This targeted sampling of outbreak scenarios allows for more efficient estimation, in particular lowering the local standard errors $\hat{v}(\mathbf{x})$ along $\partial\hat{\mathfrak{S}}_t$, cf. the right panel of Figure 6.

In (22) the boundary of $\hat{\mathfrak{S}}_t$ corresponds to the regions of $\mathcal{X}$ where the cost difference between immediate detection and waiting is zero. Hence, we aim to have more design points in regions where $\{\hat{q}(t, \mathbf{x}) - d(\mathbf{x}) \simeq 0\}$. To this end, we define the "posterior" measure of response surface accuracy via

$$p(\mathbf{x}) := \Phi\left(\frac{-|\hat{q}(t, \mathbf{x}) - d(\mathbf{x})|}{\sqrt{\hat{v}(\mathbf{x})}}\right), \qquad (26)$$

where Φ is the standard normal cdf and the predictive variance $\hat{v}$ measures the standard error of the surrogate prediction,

$$\hat{v}(\mathbf{x}) = \hat{\sigma}^2(\mathbf{x}) \|l(\mathbf{x})\|^2, \qquad (27)$$

with $\hat{\sigma}^2(\mathbf{x})$ the estimated variance of ϵ around $\mathbf{x}$ in (21), see [45, Sec 6.1.2].

The motivation for (26) is that $p(\mathbf{x})$ mimics the Bayesian posterior probability of estimating the wrong *sign* (conditional on the samples in $\mathcal{Z}$) of $q(t, \mathbf{x}) - d(\mathbf{x})$, assuming that the posterior distribution is Gaussian with the empirical mean $\hat{q}(t, \mathbf{x})$ and variance $\hat{v}(\mathbf{x})$.

The defined metric $p(\cdot)$ serves as a guide to augment new design locations. Namely, it defines an acquisition function $w(\mathbf{x})$ for greedily growing $\mathcal{Z}$, similar to active learning methods [49]. The acquisition function is highest in the regions where $p(\mathbf{x})$ is close to 0.5 which correspond to $\partial\hat{\mathfrak{S}}_t$. Our main choice is

$$w^{\min}(\mathbf{x}) = \min\left[p(\mathbf{x}), 1 - p(\mathbf{x})\right]. \qquad (28)$$

Alternatives include the Gini weights $w^{\text{gini}}(\mathbf{x}) = p(\mathbf{x})(1 - p(\mathbf{x}))$ and Entropic weights $w^{\text{Ent}}(\mathbf{x}) = -p(\mathbf{x})\log p(\mathbf{x}) - (1 - p(\mathbf{x}))\log(1 - p(\mathbf{x}))$.

To speed up the response surface modeling, which requires refitting of $\hat{q}(t, \cdot)$ multiple times, we used batch steps, incrementally working with designs $\mathcal{Z}^{(N)}$ of size $N = N_0, N_0 + N', \ldots, N^{end}$. At each sequential design iteration, an additional N' design points $\{\mathbf{x}_0^n\}_{n=N+1}^{N+N'}$ are added to existing $\mathcal{Z}^{(N)}$. Those are sampled multinomially in proportion to the acquisition function $w(\cdot)$ from a candidate set X_{finite}. Both the initial design $\mathcal{Z}^{(N_0)}$ and the candidate sets

X_{finite} are generated using Latin hypercube sampling (LHS) of size D from $\mathcal{X}$. The overall procedure, summarized in Algorithm A.2, finally refits at each iteration the Loess model for $\hat{q}$ (and hence $\mathfrak{S}_t$), grows the experimental design $\mathcal{Z}^{(N+N')} = \mathcal{Z}^{(N)} \cup \{\mathbf{x}_0^n\}_{n=N+1}^{N+N'}$ and recomputes the acquisition function (28). As the design size gets larger, we expect that the implied empirical estimate $\partial\hat{\mathfrak{S}}_t^{(N)}$ gets closer to the true $\partial\mathfrak{S}_t$.

Remark 3. *One can apply standard, non-sequential RMC by skipping the inner while loop (steps 7-15) in Algorithm A.2. This reduces to building a response model on a pre-specified (possibly randomized) design $\mathcal{Z} := \{\mathbf{x}_0^n\}_{n=1}^{N_0}$, keeping all other steps as is.*

For the detection map in Figure 6 in Section 4 we used an initial design of $N_0 = 200$, which was grown over 10 iterations with $N' = 200$ to a final design of $N^{end} = 2000$. The acquisition function was $w^{\min}$ and the candidate sets X_{finite} of size $D = 2500$ were generated with LHS. Since detection happens while $I^{(1)}$ is still relatively small, we restricted the response surface regression domain to $I^{(1)} \in \{0, 1, \ldots, 400\}$, $S^{(1)} \in \{1000, \ldots, 2000\}$. Lastly we note that the method is still computationally intensive, with the bulk of the effort spent on generating $T \cdot N^{end}$ scenarios of $\mathfrak{X}$; running times (on a 8-core 2.27GHz machine with 12GB of RAM) were about 20 minutes.

6. Discussion

We have presented a framework for optimal detection of epidemics in a coupled meta-population model. Our approach explicitly takes into account cost-benefit considerations regarding announcement of an epidemic, as well as spatial dependence across susceptible pools. Given the information about two populations and characteristics of the infection, our algorithm produces the full detection map which can then be used repeatedly. We demonstrate that information about the epidemic in one pool can be used to lower the detection costs in another pool, realizing savings compared to traditional threshold-type detection methods.

Since our dynamic optimization approach is entirely simulation-based it is unusually flexible. Indeed, the precise underlying epidemic model of $\mathfrak{X}$ is not crucial, since Algorithm A.2 only requires the ability to generate its trajectories. In fact, the computational complexity of our algorithms is tied not to the dynamics of $\mathfrak{X}$ but to its dimensionality. In Section 4, we had $\dim(\mathfrak{X}) = 3$; based on our experience with RMC in [30, 9], the present approach can straightforwardly handle up to 6-8 dimensions. In high dimensions, extra care must be applied for generating the experimental designs $\mathcal{Z}$ since the concept of neighborhoods underlying nonparametric local regression breaks down. For example, Loess regression performs poorly if the dimension of the data is higher than 4-5.

The presented SIR framework gives a basic mechanistic description of disease progression that is obviously not very realistic. More sophisticated versions might allow for further compartments (such as Exposed or Diseased individuals),

age stratification, and heterogeneity among the meta-populations. One could also include further transitions beyond (13), such as immigration $\emptyset \to I^{(k)}$, immunity lapses $R^{(k)} \to S^{(k)}$, or vaccination $S^{(k)} \to R^{(k)}$. Introducing immigration would allow for endogenous epidemic in Pool 2, removing the assumption that outbreaks always start in Pool 1 and then spread to Pool 2. The constant transition rates used can be replaced with seasonal patterns, stochastic fluctuations [50], or hierarchical Markov structures [2].

Alternatively, one can also imagine more sophisticated models for the outbreak pseudo-posterior P_t – recall that the proposed one was largely for convenience than any realism. For example, the Gaussian noise δ_t in the dynamics of P_t that was used in the case study may be better modeled via a Beta distribution (which arises naturally as conjugate to the Poisson increments of the fully-observed stochastic SIR model [2]). Overall, the key requirement is the Markov structure which makes it possible to use regression against $\mathfrak{X}$ to describe the detection rule. The Markovianity requirement can be partly relaxed if one is willing to accept approximately-optimal solutions. Indeed, one can always project the optimal detection rule into the smaller space of rules that only depend on some subset $\mathfrak{X}'$; in other words restricting the detection map to only take into account some of the state-space dimensions. This idea was already discussed in Section 4 where we described the sub-optimal LP strategy.

Second, one may modify the cost structures (2)-(3) to better capture the desired detection goals. The presented costs were motivated by their classical analogues in sequential change-point detection, but might not be the most appropriate for public health contexts. For example, there is some leeway in what constitutes an outbreak. In (3), the threshold was zero, i.e. even a single infected individual in Pool 2 was reportable. One can use thresholds $\bar{I}$ other than zero, so that an outbreak is reportable only when $I_t^{(2)} > \bar{I}$, otherwise an announcement is treated as premature. Similarly, the waiting cost in (3) was constant; it may be more realistic to make it proportional to $I_t^{(2)}$, which would correspond to fixed costs per infected.

For such more general formulations, the costs $d(\mathfrak{X})$ and $c(\mathfrak{X})$ would no longer be functions of P_t, and one would need to work with the full posterior distribution π_t of $I_t^{(2)}|\mathcal{G}_t$. The RMC framework could still be usable, namely we may use particle filtering [51] to obtain π_t along a simulated trajectory of the underlying epidemic model. Certainly, particle filtering can become computationally expensive, making efficient inference essential. We refer to [3, 2] for some recent implementation strategies in this direction that specifically target epidemic models. The integrated sequential inference + optimization model would then allow to treat a partially observed version of a K-pool SIR model of (13), and ultimately a larger-scale setup such as influenza surveillance across all 50 states, cf. Figure 1.

Acknowledgement

[1] T.-C. Chan, C.-C. King, M.-Y. Yen, P.-H. Chiang, C.-S. Huang, C. K. Hsiao, Probabilistic daily ILI syndromic surveillance with a spatio-temporal Bayesian hierarchical model, PLoS One 5 (7) (2010) e11626.

[2] J. Lin, M. Ludkovski, Sequential Bayesian inference in Hidden Markov stochastic kinetic models with application to detection and response to seasonal epidemics, Statistics and Computing 24 (6) (2014) 1047–1062.

[3] D. M. Sheinson, J. Niemi, W. Meiring, Comparison of the performance of particle filter algorithms applied to tracking of a disease epidemic, Mathematical Biosciences 255 (2014) 21–32.

[4] A. Skvortsov, B. Ristic, Monitoring and prediction of an epidemic outbreak using syndromic observations, Mathematical Biosciences 240 (1) (2012) 12 – 19.

[5] Ebola Virus Disease in West Africa: The first 9 months of the epidemic and forward projections, New England Journal of Medicine 371 (16) (2014) 1481–1495, pMID: 25244186.

[6] G. Chowell, H. Nishiura, Transmission dynamics and control of Ebola virus disease (EVD): a review, BMC Medicine 12 (1) (2014) 196.

[7] D. Fisman, E. Khoo, A. Tuite, Early epidemic dynamics of the West African 2014 Ebola outbreak: estimates derived with a simple two-parameter model, PLoS Currents Outbreaks 6 (2014) Sep 8.

[8] D. P. Bertsekas, Dynamic programming and optimal control. Vol. I, 3rd Edition, Athena Scientific, Belmont, MA, 2005.

[9] B. Gramacy, M. Ludkovski, Sequential Design for Optimal Stopping Problems, SIAM Journal on Financial Mathematics (2015) to Appear.

[10] H. Andersson, T. Britton, Stochastic Epidemic Models and Their Statistical Analysis, Lecture Notes in Statistics Series, Springer Verlag, 2000.

[11] F. Ball, D. Clancy, The final size and severity of a generalised stochastic multitype epidemic model, Advances in Applied Probability (1993) 721–736.

[12] L. J. Allen, Y. Lou, A. L. Nevai, Spatial patterns in a discrete-time SIS patch model, Journal of Mathematical Biology 58 (3) (2009) 339–375.

[13] P. Neal, The basic reproduction number and the probability of extinction for a dynamic epidemic model, Mathematical Biosciences 236 (1) (2012) 31–35.

[14] D. Neill, Fast Bayesian scan statistics for multivariate event detection and visualization, Statistics in Medicine 30 (5) (2011) 455–469.

[15] M. J. Keeling, K. T. Eames, Networks and epidemic models, Journal of the Royal Society Interface 2 (4) (2005) 295–307.

[16] M. Ajelli, B. Gonçalves, D. Balcan, V. Colizza, H. Hu, J. Ramasco, S. Merler, A. Vespignani, Comparing large-scale computational approaches to epidemic modeling: Agent-based versus structured metapopulation models, BMC Infectious Diseases 10 (1) (2010) 190.

[17] L. A. Rvachev, I. M. Longini, A mathematical model for the global spread of influenza, Mathematical Biosciences 75 (1) (1985) 3–22.

[18] G. Shmueli, H. Burkom, Statistical challenges facing early outbreak detection in biosurveillance, Technometrics 52 (1) (2010) 39–51.

[19] A. B. Lawson, Bayesian disease mapping: hierarchical modeling in spatial epidemiology, CRC Press, 2013.

[20] B. J. Cowling, I. O. L. Wong, L.-M. Ho, S. Riley, G. M. Leung, Methods for monitoring influenza surveillance data, International Journal of Epidemiology 35 (2006) 1314–1321.

[21] H. Zhou, A. B. Lawson, EWMA smoothing and Bayesian spatial modeling for health surveillance, Statistics in Medicine 27 (28) (2008) 5907–5928.

[22] M. Kulldorff, R. Heffernan, J. Hartman, R. Assuncao, F. Mostashari, A space–time permutation scan statistic for disease outbreak detection, PLoS Med 2 (3) (2005) e59.

[23] H. Nishiura, Real-time forecasting of an epidemic using a discrete time stochastic model: a case study of pandemic influenza (h1n1-2009), BioMedical Engineering OnLine 10 (1) (2011) 15.

[24] O. Hadjiliadis, M. Ludkovski, H. Yang, Quickest detection in the Wiener disorder problem with post-change uncertainty, preprint (2015).

[25] L. Hutwagner, W. Thompson, G. Seeman, T. Treadwell, A simulation model for assessing aberration detection methods used in public health surveillance for systems with limited baselines, Statistics in Medicine 24 (4) (2005) 543–550.

[26] Y. LeStrat, F. Carrat, Monitoring epidemiologic surveillance data using hidden Markov models, Statistics in Medicine 18 (1999) 3463–3478.

[27] M. Martínez-Beneito, D. Conesa, A. López-Quílez, A. López-Maside, Bayesian Markov switching models for the early detection of influenza epidemics, Statistics in Medicine 27 (22) (2008) 4455–4468.

[28] P. Sebastiani, K. D. Mandl, P. Szolovits, I. S. Kohane, M. F. Ramoni, A Bayesian dynamic model for influenza surveillance, Statistics in Medicine 25 (11) (2006) 1803–1816.

[29] M. Ludkovski, J. Niemi, Optimal disease outbreak decisions using stochastic simulation, in: Proceedings of the Winter Simulation Conference, WSC '11, Winter Simulation Conference, 2011, pp. 3849–3858.

[30] M. Ludkovski, J. Niemi, Optimal dynamic policies for influenza management, Statistical Communications in Infectious Diseases 2(1) (2010) article 5 (electronic).

[31] M. W. Tanner, L. Sattenspiel, L. Ntaimo, Finding optimal vaccination strategies under parameter uncertainty using stochastic programming, Mathematical Biosciences 215 (2) (2008) 144–151.

[32] H. J. Wearing, P. Rohani, M. J. Keeling, Appropriate models for the management of infectious diseases, PLoS Medicine 2 (7) (2005) e174.

[33] A. Culotta, Towards detecting influenza epidemics by analyzing Twitter messages, in: Proceedings of the first workshop on social media analytics, ACM, 2010, pp. 115–122.

[34] V. Dukic, H. Lopes, N. Polson, Tracking epidemics with Google Flu Trends data and a state-space SEIR model, Journal of the American Statistical Association 107 (500) (2012) 1410–1426.

[35] D. Banks, G. Datta, A. Karr, J. Lynch, J. Niemi, F. Vera, Bayesian CAR models for syndromic surveillance on multiple data streams: Theory and practice, Information Fusion 13 (2012) 105–116.

[36] A. Noufaily, D. G. Enki, P. Farrington, P. Garthwaite, N. Andrews, A. Charlett, An improved algorithm for outbreak detection in multiple surveillance systems, Statistics in Medicine 32 (7) (2013) 1206–1222.

[37] W. Yang, A. Karspeck, J. Shaman, Comparison of filtering methods for the modeling and retrospective forecasting of influenza epidemics, PLoS Computational Biology 10 (4) (2014) e1003583.

[38] M. Ludkovski, Monte Carlo methods for adaptive disorder problems, in: R. A. Carmona, P. Del Moral, P. Hu, N. Oudjane (Eds.), Numerical Methods in Finance, Vol. 12 of Springer Proceedings in Mathematics, Springer Berlin Heidelberg, 2012, pp. 83–112.

[39] H. V. Poor, O. Hadjiliadis, Quickest detection, Vol. 40, Cambridge University Press, 2009.

[40] V. Nevistic, J. A. Primbs, Model predictive control: Breaking through constraints, in: 35th IEEE Conference on Decision and Control, 1996, pp. 3932–3937.

[41] D. J. Wilkinson, Stochastic Modelling for Systems Biology, Chapman & Hall/CRC, London, 2006.

[42] K. Shatskikh, M. Ludkovski, Bayesian inference and detection for a multi-population SIR model of stochastic epidemics, in preparation (2015).

[43] F. Ball, P. Donnelly, Strong approximations for epidemic models, Stochastic Processes and their Applications 55 (1) (1995) 1–21.

[44] D. Egloff, Monte Carlo algorithms for optimal stopping and statistical learning, Annals of Applied Probability 15 (2) (2005) 1396–1432.

[45] T. Hastie, R. Tibshirani, J. Friedman, The elements of statistical learning: data mining, inference and prediction, Springer, 2009.

[46] W. S. Cleveland, S. J. Devlin, Locally weighted regression: An approach to regression analysis by local fitting, Journal of the American Statistical Association 83 (403) (1988) pp. 596–610.

[47] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2013).

[48] R. Hu, M. Ludkovski, Bayesian sequential design for ranking response surfaces, preprint (2015).

[49] D. MacKay, Information-based objective functions for active data selection, Neural computation 4 (4) (1992) 590–604.

[50] D. He, E. Ionides, A. King, Plug-and-play inference for disease dynamics: measles in large and small populations as a case study, Journal of the Royal Society Interface 7 (43) (2010) 271–283.

[51] M. Ludkovski, A simulation approach to optimal stopping under partial observations, Stochastic Processes and Applications 119 (12) (2009) 4061–4087.

Appendix A. Algorithms

Algorithm A.1 Path and Cost Generation

Require: $\{\mathbf{x}_0^n\}_{n=1}^N$, $\mathfrak{S}_{0:t-1}$

1: **for** $n = 1, \ldots, N$ **do**
2:     $s \leftarrow 1$
3:     **while** $s \leq t$ **do**
4:         Simulate the next state $\mathbf{x}_s^n \sim p_1(\cdot | \mathbf{x}_{s-1}^n)$
5:         **if** $\mathbf{x}_s^n \in \mathfrak{S}_{t-s}$ **then** Break
6:         **end if**
7:         $s \leftarrow s + 1$
8:     **end while**
9:     $\tau_t^n \leftarrow s$
10:     Compute $q^n \equiv c(\mathbf{x}_{0:\tau_t^n}^n)$ using formula (3)
11: **end for**
12: **return** $\{(\mathbf{x}_0^n, q^n)\}_{n=1}^N$

Algorithm A.2 Sequential Regression Monte Carlo

Require: $C_{\text{FA}}, C_{\text{Delay}}, N_0, N', N^{end}$, D

1: $\hat{\mathfrak{S}}_0 \leftarrow \mathcal{X}$
2: **for** $t = 1, 2, \ldots$ **do**
3:     Generate experimental design $\{\mathbf{x}_0^n, \ n = 1, \ldots, N_0\}$
4:     Compute scenario costs $q^n = c(\mathfrak{X}_{0:\tau_t^n}^n)$ for $n = 1, \ldots, N_0$ using Algorithm A.1 and $\hat{\mathfrak{S}}_{0:t-1}$
5:     $Z \leftarrow \{(\mathbf{x}_0^n, q^n)\}_{n=1}^{N_0}$
6:     Regress $\{q^n\}$ on $\{\mathbf{x}_0^n\}$, $n = 1, \ldots, N_0$ using Loess (24)
7:     Initialize $N \leftarrow N_0$
8:     **while** $N < N^{end}$ **do**
9:         Generate X_{finite} of size D using Latin Hypercube Sampling on $\mathcal{X}$
10:         Compute the acquisition weights $w(\mathbf{x}) \ \forall \mathbf{x} \in X_{\text{finite}}$ via (28) and (26)
11:         Sample $\{\mathbf{x}_0^n\}_{n=N+1}^{N+N'}$ from X_{finite} using weights $w(\mathbf{x})$
12:         Simulate the costs q^n, $n = N+1, \ldots, N+N'$ using Algorithm A.1
13:         $Z \leftarrow Z \cup \{(\mathbf{x}_0^n, q^n)\}_{n=N+1}^{N+N'}$
14:         Update the Loess regression model (24) using the latest Z
15:         $N \leftarrow N + N'$
16:     **end while**
17:     $\hat{\mathfrak{S}}_t \leftarrow \{\mathbf{x} \in \mathcal{X} : \hat{q}(t, \mathbf{x}) - d(\mathbf{x}) > 0\}$, cf. (22)
18: **end for**