

Tracking disease outbreaks from sparse data with Bayesian inference

Bryan Wilder,¹ Michael Mina², Milind Tambe¹

¹ John A. Paulson School of Engineering and Applied Sciences, Harvard University

² T.H. Chan School of Public Health, Harvard University

bwilder@g.harvard.edu, mmina@hsph.harvard.edu, milind_tambe@harvard.edu

Abstract

The COVID-19 pandemic provides new motivation for a classic problem in epidemiology: estimating the empirical rate of transmission during an outbreak (formally, the time-varying reproduction number) from case counts. While standard methods exist, they work best at coarse-grained national or state scales with abundant data, and struggle to accommodate the partial observability and sparse data common at finer scales (e.g., individual schools or towns). For example, case counts may be sparse when only a small fraction of infections are caught by a testing program. Or, whether an infected individual tests positive may depend on the kind of test and the point in time when they are tested. We propose a Bayesian framework which accommodates partial observability in a principled manner. Our model places a Gaussian process prior over the unknown reproduction number at each time step and models observations sampled from the distribution of a specific testing program. For example, our framework can accommodate a variety of kinds of tests (viral RNA, antibody, antigen, etc.) and sampling schemes (e.g., longitudinal or cross-sectional screening). Inference in this framework is complicated by the presence of tens or hundreds of thousands of discrete latent variables. To address this challenge, we propose an efficient stochastic variational inference method which relies on a novel gradient estimator for the variational objective. Experimental results for an example motivated by COVID-19 show that our method produces an accurate and well-calibrated posterior, while standard methods for estimating the reproduction number can fail badly.

Introduction

A key goal for public health is effective surveillance and tracking of infectious disease outbreaks. This work is motivated in particular by the COVID-19 pandemic but the methods we describe are applicable to other diseases as well. A central question is to estimate the empirical rate of transmission over time, often formalized via the reproduction number R_t , $t = 1 \dots T$. R_t describes the expected number of secondary infections caused by someone infected at time t . Accurate estimates of R_t are critical to detect emerging outbreaks, forecast future cases, and measure the impact of interventions imposed to limit spread.

R_t is typically estimated using daily case counts, i.e., the number of new infections detected via testing each day. Standard methods, including prominent dashboards developed for COVID-19, provide accurate estimates under idealized conditions for the observation of cases and have been successfully used at a national or state level where many observations are available and sampling variation averages out (Abbott et al. 2020; Flaxman et al. 2020; Systrom, Vladek, and Krieger 2020). However, successful reopening will require programs which track spread at the level of particular colleges, workplaces, or towns, where *partial observability* poses several challenges. *First*, only a small number of infected people may be tested. It is estimated that only about 10% of SARS-COV-2 infections in the US result in a confirmed test (Havers et al. 2020) and we could expect even fewer in populations with a high prevalence of asymptomatic or mild infections (e.g., college students). *Second*, the biological properties of the test play an important role. For example, a PCR test which detects viral RNA will show positive results at different times than an antibody or antigen test. Further, there can be substantial heterogeneity across individuals. *Third*, testing programs may collect samples in a particular way which impacts the observations. For example, one suggestion for schools and workplaces to reopen is to institute regular surveillance testing of a fraction of the population in order to detect outbreaks and catch asymptomatic carriers (Larremore et al. 2020). The observations will depend on the fraction of the population enrolled in testing (potentially small due to budget constraints) along with the sampling design (e.g., cross-sectional vs longitudinal).

This paper presents *GPRT*, a novel Bayesian approach to estimating R_t which accounts for partial observability in a flexible and principled manner (illustrated in Figure 1). This method yields well-calibrated probabilistic estimates (the posterior distribution). Our model places a Gaussian process (GP) prior over R_t , allowing it to be an arbitrary smooth function. Then, we explicitly model the sampling process which generates the observations from the true trajectory of infections. While this substantially improves accuracy (as we show experimentally) it creates a much more difficult inference problem than has been previously considered. Specifically, our model contains tens or hundreds of thousands of discrete latent variables, preventing the application of out-of-the-box methods. Moreover, the values of

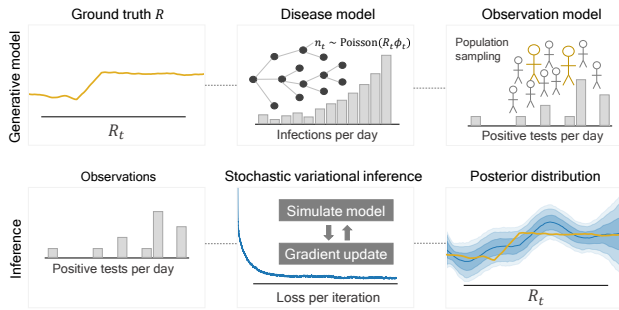


Figure 1: Illustration of our GPRt method. Top row: the generative model GPRt posits for the observed data. Bottom row: the inference process to recover a posterior over R_t .

many variables are tightly correlated in the posterior distribution, further complicating inference. To make inference computationally tractable, we propose a novel stochastic variational inference method, enabled by a custom stochastic gradient estimator for the variational objective. Extensive experiments show that our method recovers an accurate and well-calibrated posterior distribution in challenging situations where previous methods fail.

Related work

There is a substantial body of work which attempts to infer unknowns in a disease outbreak. A frequent target for inference is the *basic* reproduction number R_0 (Majumder and Mandl 2020; Riou and Althaus 2020; Wilder et al. 2020); by contrast, we attempt the more challenging task of estimating a reproduction number which can vary arbitrarily over time. There is a literature of both classic methods for estimating R_t (Wallinga and Teunis 2004; Cori et al. 2013; Campbell et al. 2019) and newer methods developed for the COVID-19 pandemic and implemented in popular dashboards (Abbott et al. 2020; Systrom, Vladdek, and Krieger 2020). None of these methods incorporate partial observability and we empirically demonstrate that GPRt improves substantially over methods in each category. Another strand of literature develops maximum likelihood or particle filter estimates of the parameters of an epidemiological model (King et al. 2008; Dureau, Kalogeropoulos, and Baguelin 2013; Cazelles, Champagne, and Dureau 2018). However, their work focuses on accommodating a complex model of the underlying disease dynamics; by contrast, we develop methods for probabilistically well-grounded inferences under a complex observation model. There is also a great deal of computational work more broadly concerned with disease control. Examples include optimization problems related to vaccination or quarantine decisions (Saha et al. 2015; Zhang and Prakash 2014; Zhang et al. 2015), machine learning methods for forecasting (without recovering a probabilistic view of R_t) (Chakraborty et al. 2014; Rekatsinas et al. 2015), and agent-based simulations of disease dynamics (Barrett, Eubank, and Marathe 2008). Our work complements this literature by allowing inference of a distribution over R_t from noisy data, which can serve as the input that parameterizes an optimization problem or simulation.

Model

We now introduce a model for a disease process with a time-varying reproduction number. Subsequently, we introduce example models for how the observations are generated from the disease process which our framework can support.

Disease model

We use a standard stochastic model of disease transmission similar to other R_t estimation methods (Wallinga and Teunis 2004; Cori et al. 2013; Campbell et al. 2019); our contribution is a more powerful inference methodology which can accommodate complex observation models alongside a GP prior. Let $\mathbf{R} = [R_1 \dots R_T]$ be the vector with R_t for each time. From \mathbf{R} , the disease model defines a distribution over a vector $\mathbf{n} = [n_1 \dots n_T]$ with the number of people newly infected each day. The main idea is that, over the course of a given individual’s infection, they cause a Poisson-distributed number of new infections with mean determined by \mathbf{R} . Specifically, if on day t individual i has been infected for h^i days the expected number of new infections caused by i that day is

$$\lambda_t^i = w_{h^i} R_t.$$

Here w_{h^i} gives the level of infectiousness of an individual h^i days post-infection. w is normalized so that $\sum_h w_h = 1$. For later convenience, we will define $\phi_t = \sum_{i=1}^N w_{h^i}$ to be the total infectiousness in the population before scaling by R_t (N is the total population size).

Each day, each infected individual i draws $n_t^i \sim \text{Poisson}(\lambda_t^i)$ other individuals to infect. We also incorporate infections from outside the population, with a mean of γ such infections per day. We assume the rate of external infection is constant with respect to our time but our model could be extended to a time-varying γ . We treat γ mostly as a nuisance parameter: our true objective is to infer \mathbf{R} , but doing so requires accounting for the potential that some detected cases are due to infections from outside the population. We define $n_t = \sum_{i=1}^N n_t^i + \text{Poisson}(\gamma)$ to be the total number of new infections. Since ϕ is fixed given the time series \mathbf{n} , we denote it as a function $\phi(\mathbf{n})$. We denote the probabilistic disease model induced by a specific choice of \mathbf{R} and γ as $M(\mathbf{R}, \gamma)$ and the draw of a time series of infections from the model as $\mathbf{n} \sim M(\mathbf{R}, \gamma)$.

Observation model

We now depart from the standard disease model used in previous work and describe a wide-ranging set of examples for how our framework can accommodate models of the process which generates the observed data from the latent (unknown) true infections. Previous work assumes either perfect observability or else the simplest of the three observation models we describe below (uniform undersampling). Our focus is where observations are generated by a medical test which confirms the presence of the pathogen of interest. Individuals have some probability of being tested at different times (depending on the testing policy adopted) and then test positive with a probability which depends on the biological characteristics of the disease and the test in question.

Modeling tests The kind of test employed determines when an individual is likely to test positive during the course of infection. For example, for COVID-19, PCR tests are commonly used to detect SARS-COV-2 RNA. They are highly sensitive and can detect early infections. Most infected individuals become PCR-negative within the week or two following infection as viral RNA is cleared (Kucirka et al. 2020). By contrast, serological tests detect the antibodies produced after infection. An individual is not likely to test serologically positive until a week or more post-infection, but may then continue to test serologically positive for months afterwards (Iyer et al. 2020). The observable data generated by a serological testing program is likely quite different than a PCR testing program since the time-frame and variance of when individuals test positive differs strongly. A range of other examples are possible (e.g., antigen tests) and can be easily incorporated into our framework.

Our model adopts a generic representation of a particular test as a distribution D over t_{convert} , the number of days post-infection when an individual begins to test positive and t_{revert} , the number of days post-infection when they cease to test positive. For an infected individual i , we write $t_{\text{convert}}^i, t_{\text{revert}}^i \sim D$. Our method only assumes the ability to sample from D , meaning that we can directly plug in the results of lab studies assessing the properties of a test. $t_{\text{convert}}^i, t_{\text{revert}}^i$ are unobserved: we only get to see if an individual tests positive at a given point in time, not the full range of times that they *would* have tested positive.

Next, we describe a series of example models for how and when individuals are tested, which reveal observations depending on the status $(t_{\text{convert}}^i, t_{\text{revert}}^i)$ for each person who is tested. For convenience, we let $t_{\text{convert}}^i = \infty$ for an individual who is never infected. Note however that we can model false negatives by having D sometimes set $t_{\text{convert}}^i = \infty$ for an infected individual, or false positives by returning finite t_{convert}^i for an uninfected individual. We denote the number of observed positive tests on day t as x_t . An observation model is a distribution over \mathbf{x} given \mathbf{n} , denoted $\mathbf{x} \sim \text{Obs}(\mathbf{n})$. Each setting below describes one such distribution.

Uniform undersampling In this setting, each individual who is infected is tested independently with some probability p_{test} (e.g., if they individually decide whether to seek a test). To model this process, we introduce two new sets of latent random variables. First, a binary variable $z^i \sim \text{Bernoulli}(p_{\text{test}})$, indicating whether individual i is tested. Second, a delay c^i , giving the number of days between t_{convert}^i and when individual i is tested. We can integrate out the z^i and obtain the following conditional distribution for the observed number of positive tests x_t :

$$x_t | c, t_{\text{convert}} \sim \text{Binomial} \left(\sum_{i=1}^N 1[t = t_{\text{convert}}^i + c^i], p_{\text{test}} \right)$$

where $1[\cdot]$ denotes the indicator function of an event. However, we cannot analytically integrate out t_{convert} and c .

Cross-sectional testing Here, a uniformly random sample of s_t individuals are tested on each day t . This models a random screening program (e.g., testing random employees

each day as they enter a workplace). In this case, we have

$$x_t | t_{\text{convert}}, t_{\text{revert}} \sim \text{Binomial} \left(s_t, \frac{1}{N} \sum_{i=1}^N 1[t_{\text{convert}}^i \leq t < t_{\text{revert}}^i] \right)$$

This expression provides the likelihood of \mathbf{x} after conditioning on the latent variables $t_{\text{convert}}^i, t_{\text{revert}}^i$, though there is no closed-form expression conditioning only on \mathbf{n} .

Longitudinal testing In this setting, a single sample from the population is chosen up front and every individual in the sample is tested every d days. We again denote the total number of individual tested on day t as s_t , but note that now the group of individuals who are tested repeats every d days. Longitudinal testing offers different (and potentially more revealing) information than cross-sectional testing since when an individual first tests positive, we know that they did *not* test positive d days ago. However, it complicates inference by introducing correlations between the test results at different time steps. Let A_t denote the set of individuals who are tested at time t . We assume that the complete sample $\bigcup_{t=1}^d A_t$ is chosen uniformly at random from the population, with the chosen individuals then randomly partitioned between the d days. We have

$$x_t = \sum_{i \in A_t} 1[t_{\text{convert}}^i > t - d \text{ and } t_{\text{convert}}^i \leq t \text{ and } t_{\text{revert}}^i > t]$$

where $t_{\text{convert}}^i > t - d$ captures that i was not positive on their previous test. This introduces correlations between x_t and x_{t-d} , so there is not a simple closed-form expression for the distribution of the time series \mathbf{x} even after conditioning on t_{convert}^i and t_{revert}^i . (as there is in the cross-sectional case). We will instead build a flexible framework for inference which can just as well use a kind of sample of the log-likelihood.

Inference problem

We will place a Gaussian process (GP) prior over \mathbf{R} , resulting in the following generative model:

$$\begin{aligned} \mathbf{R} &\sim \mathcal{GP}(1, \mathcal{K}), \gamma \sim \text{Exp}(\bar{\gamma}) \\ \mathbf{n} &\sim M(\mathbf{R}, \gamma) \\ \mathbf{x} &\sim \text{Obs}(\mathbf{n}). \end{aligned}$$

where $\mathcal{GP}(1, \mathcal{K})$ denotes a Gaussian process with constant mean 1 and kernel \mathcal{K} and $\text{Exp}(\bar{\gamma})$ is an exponential prior on γ with mean $\bar{\gamma}$. Given the observation \mathbf{x} , our goal is to compute the resulting posterior distribution over \mathbf{R} and γ . However, is complicated by the fact that \mathbf{x} is determined by a large number of discrete latent variables, primarily \mathbf{n} (the time series of infections) and $\{t_{\text{convert}}^i, t_{\text{revert}}^i\}_{i=1}^N$, the times when each individual tests positive. A typical strategy for inference in complex Bayesian models is Markov Chain Monte Carlo (MCMC). However, MCMC is difficult to apply because of tight correlations between the values of variables over time: due to the GP prior, we expect values of \mathbf{R} to be closely correlated between timesteps, and successive values of \mathbf{n} are highly correlated via the model M . Formulating good proposal distributions for high-dimensional, tightly correlated random variables is notoriously difficult

and has presented problems for GP inference via MCMC in other domains (Titsias, Lawrence, and Rattray 2008).

The other main approach to Bayesian inference is *variational inference*, where we attempt to find the best approximation to the posterior distribution within some restricted family. Modern variational inference methods, typically intended for deep models such as variational autoencoders (Kingma and Welling 2013), use a combination of autodifferentiation frameworks and the reparameterization trick to differentiate through the variational objective (Kucukelbir et al. 2017). This process is highly effective for models with only continuous latent variables. However, our model has many thousands of discrete latent variables which cannot be reparameterized in a differentiable manner. Typical solutions to this problem would be to either integrate out the discrete variables or to replace them with a continuous relaxation (Jang, Gu, and Poole 2017; Vahdat et al. 2018). Neither solution is attractive in our case – the structure of the model does not allow us to integrate out the discrete variables analytically, while a continuous relaxation is infeasible because our latent variables have a strict integer interpretation (every infection requires in a particular individual becoming test-positive at particular points in time).

The last resort to differentiate through discrete probabilistic models is the score function estimator (Paisley, Blei, and Jordan 2012), which is often difficult to apply due to high variance. GPRt uses a combination of techniques which exploit the structure of infectious disease models to develop an estimator with controlled variance. *First*, we develop a more tractable variational lower bound which is amenable to stochastic optimization. *Second*, we hybridize the reparameterization and score function estimators across different parts of the generative model to take advantage of the properties of each component. *Third*, we develop techniques to sample low-variance estimates of the log-likelihood for each of the observation models introduced earlier. These techniques are introduced in the next section.

GPRt: variational inference algorithm

We now derive *GPRt*, a novel variational inference method for R_t estimation. GPRt approximates the true (uncomputable) posterior over (\mathbf{R}, γ) via a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix Σ . μ_t is the posterior mean for R_t while $\Sigma_{t,t'}$ gives the posterior covariance between R_t and $R_{t'}$. μ_γ is the mean for γ and $\Sigma_{\gamma,\cdot}$ gives its covariance with R . The diagonal $\Sigma_{t,t}$ gives the variance of the posterior over R at each time, capturing the overall level of uncertainty. The aim is to find a $\boldsymbol{\mu}$ and Σ which closely approximate the true posterior. Let $q(\mathbf{R}, \gamma | \boldsymbol{\mu}, \Sigma)$ denote the variational distribution. Let p be the true generative distribution, where $p(\mathbf{R}, \gamma, \mathbf{x})$ is the joint distribution over \mathbf{x} and (\mathbf{R}, γ) , $p(\mathbf{R}, \gamma)$ is the prior over (\mathbf{R}, γ) , and $p(\mathbf{R}, \gamma | \mathbf{x})$ is the posterior over (\mathbf{R}, γ) after conditioning on \mathbf{x} .

The aim of variational inference is to maximize a lower bound on the total log-probability of the evidence \mathbf{x} :

$$\begin{aligned} \log p(\mathbf{x}) \geq & \mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log p(\mathbf{R}, \gamma)] + \mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log p(\mathbf{x} | \mathbf{R}, \gamma)] \\ & - \mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log q(\mathbf{R}, \gamma | \boldsymbol{\mu}, \Sigma)] \end{aligned}$$

where the right-hand side is referred to as the *Evidence Lower Bound (ELBO)*. Our goal is to maximize the ELBO via gradient ascent on the parameters $\boldsymbol{\mu}$ and Σ . This requires us to develop an estimator for the gradient of each term in the ELBO. The first term is the negative cross-entropy between q and the prior $p(\mathbf{R}, \gamma)$. Because both q and p have simple parametric forms, this can be easily computed and differentiated. The last term is the entropy, which is similarly tractable. The middle term is the expected log-likelihood. Developing an estimator for the gradient of this term is substantially more complicated and will be our focus. In fact, for computational tractability we will actually develop an estimator for a lower bound on the expected log-likelihood; substituting this lower bound into the ELBO still gives a valid lower bound on $\log p(\mathbf{x})$ and so is a sensible objective.

Gradient estimator

The essential problem is that computing the log-likelihood of \mathbf{R} requires integrating out the discrete latent variable \mathbf{n} induced by the disease spread model, which is computationally intractable. The aim of this section is to develop the following stochastic estimator:

Theorem 1. *Let L be the Cholesky factor of Σ , $\xi \sim N(0, I)$, and $\begin{bmatrix} \mathbf{R} \\ \gamma \end{bmatrix} = \boldsymbol{\mu} + L\xi$. Let $\mathbf{n} \sim M(\mathbf{R}, \gamma)$. Finally, define*

$$\hat{\nabla} = \nabla_{\boldsymbol{\mu}, L} \log M(\mathbf{n} | \mathbf{R}, \gamma) \log p(\mathbf{x} | \mathbf{n}).$$

There exists a function $g(\boldsymbol{\mu}, \Sigma)$ with $\mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log p(\mathbf{x} | \mathbf{R}, \gamma)] \geq g(\boldsymbol{\mu}, \Sigma) \forall \boldsymbol{\mu}, \Sigma$ and $\mathbb{E}[\hat{\nabla}] = \nabla g(\boldsymbol{\mu}, \Sigma)$.

Essentially, Theorem 1 states that $\hat{\nabla}$ is an unbiased estimator for a lower bound on the expected log-likelihood, exactly what we need to optimize a lower bound on $\log p(\mathbf{x})$ by stochastic gradient methods. Moreover, as we will highlight below, we can efficiently compute the terms of $\hat{\nabla}$ via a combination of leveraging the structure of the disease model to apply autodifferentiation tools and novel sampling methods for the observation model. We now derive $\hat{\nabla}$.

Proof. Expanding the dependence of \mathbf{x} on \mathbf{n} we can rewrite the log-likelihood as

$$\mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log p(\mathbf{x} | \mathbf{R}, \gamma)] = \mathbb{E}_{\mathbf{R}, \gamma \sim q(\boldsymbol{\mu}, \Sigma)} \left[\log \left(\mathbb{E}_{\mathbf{n} \sim M(\mathbf{R}, \gamma)} [p(\mathbf{x} | \mathbf{n})] \right) \right]$$

It is not clear how to develop a well-behaved gradient estimator for this expression because we wish to differentiate with respect to the parameters governing two nested expectations, one within the log. However, via Jensen’s inequality, we can derive the lower bound

$$\mathbb{E}_{\mathbf{R}, \gamma \sim q} [\log p(\mathbf{x} | \mathbf{R}, \gamma)] \geq \mathbb{E}_{\mathbf{R}, \gamma \sim q(\boldsymbol{\mu}, \Sigma)} \left[\mathbb{E}_{\mathbf{n} \sim M(\mathbf{R}, \gamma)} [\log p(\mathbf{x} | \mathbf{n})] \right],$$

pushing the log inside the expectation. We will substitute this bound into the ELBO, obtaining a valid lower bound to maximize. The key advantage is that our new lower bound admits an efficient stochastic gradient estimator. We start

with the inner expectation and attempt to compute a gradient with respect to \mathbf{R} (which controls the distribution of the simulation results \mathbf{n}). Using *score function estimator* gives

$$\begin{aligned} & \nabla_{\mathbf{R}, \gamma} \mathbb{E}_{\mathbf{n} \sim M(\mathbf{R}, \gamma)} [\log p(\mathbf{x}|\mathbf{n})] \\ &= \mathbb{E}_{\mathbf{n} \sim M(\mathbf{R}, \gamma)} [\nabla_{\mathbf{R}, \gamma} \log M(\mathbf{n}|\mathbf{R}, \gamma) \log p(\mathbf{x}|\mathbf{n})] \end{aligned}$$

which expresses the gradient with respect to (\mathbf{R}, γ) in terms of the gradient of the probability density of the disease model M with respect to (\mathbf{R}, γ) . It turns out that $\log M(\mathbf{n}, \phi|\mathbf{R}, \gamma)$ can be easily computed. Recall that $n_t = \sum_{i=1}^N n_t^i + \text{Poisson}(\gamma)$, where $n_t^i \sim \text{Poisson}(R_t \phi_t^i)$. Using the Poisson superposition theorem, we have that $n_t \sim \text{Poisson}(\sum_{i=1}^N R_t \phi_t^i + \gamma)$ (while ϕ_t is a deterministic function of $n_1 \dots n_{t-1}$). Accordingly, we have that

$$\begin{aligned} \log M(\mathbf{n}|\mathbf{R}, \gamma) &= \sum_{t=1}^T \log \Pr[n_t | n_1 \dots n_T] \\ &= \sum_{t=1}^T n_t \log(R_t \phi_t(\mathbf{n}) + \gamma) - e(R_t \phi_t(\mathbf{n}) + \gamma) - \log n_t! \end{aligned}$$

where the second line substitutes the Poisson log-likelihood. This expression can be easily differentiated with respect to \mathbf{R} and γ in closed form. Accordingly, we obtain an unbiased estimate of the gradient of our lower bound by sampling $\mathbf{n} \sim M(\mathbf{R}, \gamma)$ and computing

$$\nabla_{\mathbf{R}, \gamma} \log M(\mathbf{n}|\mathbf{R}, \gamma) \log p(\mathbf{x}|\mathbf{n}).$$

This suffices to estimate the gradient with respect to (\mathbf{R}, γ) . However, our goal is to differentiate with respect to $\boldsymbol{\mu}$ and Σ , which control the distribution over (\mathbf{R}, γ) . Fortunately, \mathbf{R} and γ are continuous. So, we can exploit the reparameterization trick by writing (\mathbf{R}, γ) as a function of a random variable whose distribution is fixed. Specifically, since Σ is positive semi-definite, it has a Cholesky decomposition $\Sigma = LL^T$ (for convenience, we actually optimize over L instead of Σ). Sampling a standard normal variable $\xi \sim N(0, I)$ and letting $\begin{bmatrix} \mathbf{R} \\ \gamma \end{bmatrix} = \boldsymbol{\mu} + L\xi$ is equivalent to sampling $\mathbf{R}, \gamma \sim N(\boldsymbol{\mu}, \Sigma)$. We rewrite the lower bound as

$$g(\boldsymbol{\mu}, \Sigma) = E_{\xi \sim N(0, I)} \left[\mathbb{E}_{\mathbf{n} \sim M(\mathbf{R}(\xi), \gamma(\xi))} [\log p(\mathbf{x}|\mathbf{n})] \right]$$

where $\boldsymbol{\mu}$ and L appear only as parameters of the deterministic function expressing \mathbf{R} and γ in terms of ξ , instead of in the distribution of a random variable. Taking a sample from each of the expectations and substituting the score function estimator now gives the desired expression for $\hat{\nabla}$. \square

Using Theorem 1, our final gradient estimator will sample b values for ξ , run the model M once for each of the resulting values of \mathbf{R} to sample \mathbf{n} , and then compute

$$\frac{1}{b} \sum_{k=1}^b \nabla_{\boldsymbol{\mu}, L} \log M(\mathbf{n}(k)|\mathbf{R}(\xi(k)), \gamma(\xi(k))) \log p(\mathbf{x}|\mathbf{n}(k)),$$

easily accomplished with standard autograd tools given the closed-form expressions for $\mathbf{R}(\xi)$, $\gamma(\xi)$, and $\log M(\mathbf{n}|\mathbf{R})$. In practice, we also use the mean $\log p(\mathbf{x}|\mathbf{n}(k))$ as a simple control variate to reduce variance (Sutton and Barto 1998).

Computing the likelihood

We now turn to the task of computing the log-likelihood function $\log p(\mathbf{x}|\mathbf{n})$, which measures the log-likelihood of observing the sequence of positive test results \mathbf{x} given \mathbf{n} new infections per day. Unfortunately, the log-likelihood is not available in closed form for any of the settings that we consider because it depends on additional latent variables (e.g., t_{convert} , t_{revert} , C , or A). We will show that it suffices to develop an estimator which lower-bounds the log-likelihood and that such estimators can be efficiently implemented for each of the observation models we consider. Specifically, denote the collection of latent variables used in a particular observation model as α . Then, we have

$$\log p(\mathbf{x}|\mathbf{n}) = \log \left(\mathbb{E}_{\alpha} [p(\mathbf{x}|\mathbf{n}, \alpha)] \right),$$

which presents a similar difficulty as in developing our earlier lower bound: sampling α to approximate the inner expectation does not result in an unbiased estimator due to the outer log. Using Jensen's inequality in the same way gives

$$\log p(\mathbf{x}|\mathbf{n}) \geq \mathbb{E}_{\alpha} [\log p(\mathbf{x}|\mathbf{n}, \alpha)],$$

and so substituting the right-hand side into our variational objective preserves validity of the lower bound. The RHS has the crucial advantage that we can now develop an unbiased estimator by drawing a single sample of α , which can then be substituted into the stochastic gradient estimator of Theorem 1. That is, for each of the simulation results $n(1) \dots n(b)$ we sample a corresponding value for the latent variables, $\alpha(1) \dots \alpha(b)$ and use the gradient estimator

$$\frac{1}{b} \sum_{k=1}^b \nabla_{\boldsymbol{\mu}, L} \log M(\mathbf{n}(k)|\mathbf{R}(\xi(k)), \gamma(\xi(k))) \log p(\mathbf{x}|\mathbf{n}(k), \alpha(k))$$

This works without issue for the uniform undersampling and cross-sectional models where we can obtain a closed form for the log likelihood after conditioning on the appropriate latent variables. However, the longitudinal testing model presents additional complications. In particular, after sampling the latent variables t_{convert} , t_{revert} , and A_t , the number of positive tests becomes deterministic quantity. Denote this simulated trajectory of positive tests \tilde{x} . If $\tilde{x} = x$, then $p(\mathbf{x} | t_{\text{convert}}, t_{\text{revert}}, A) = 1$ and otherwise $p(\mathbf{x} | t_{\text{convert}}, t_{\text{revert}}, A) = 0$. This renders the above gradient estimator useless because $\log p(\mathbf{x}|\mathbf{n}(k), \alpha(k)) = -\infty$ unless the simulated trajectory *exactly* matches the observed data (a very low-probability event). While $-\infty$ is technically a valid lower bound for the variational objective, it is not very useful for optimization. Essentially, we need to develop a lower-variance estimator where the lower bound is more useful.

We now present one such improved estimator. The intuition is that we can marginalize out a great deal of the randomness in the naive estimator by only revealing the results of random draws determining A_t a single individual at a time. We start by sampling t_{convert} and t_{revert} . Note that we can expand $\log p(\mathbf{x}|\mathbf{n}, t_{\text{convert}}, t_{\text{revert}}) = \sum_{t=1}^T \log p(x_t | x_1 \dots x_{t-1}, \mathbf{n}, t_{\text{convert}}, t_{\text{revert}})$ and consider the likelihood at each day t after conditioning on the results observed on previous days. To compute an estimate for this

Outbreak setting									
Longitudinal	PCR				Serological				
	0.5%	1%	2%	5%	0.5%	1%	2%	5%	
WT	0.481 ± 0.147	0.405 ± 0.11	0.395 ± 0.0947	0.401 ± 0.113	0.445 ± 0.153	0.415 ± 0.117	0.423 ± 0.103	0.438 ± 0.147	
Cori	1.74 ± 0.774	1.18 ± 0.573	0.806 ± 0.373	0.546 ± 0.212	1.55 ± 0.757	0.97 ± 0.577	0.621 ± 0.362	0.455 ± 0.191	
EpiNow	0.329 ± 0.211	0.265 ± 0.145	0.25 ± 0.135	0.267 ± 0.168	0.308 ± 0.226	0.25 ± 0.171	0.232 ± 0.136	0.225 ± 0.134	
GPRt	0.228 ± 0.0713	0.2 ± 0.055	0.183 ± 0.0579	0.186 ± 0.0692	0.237 ± 0.0805	0.232 ± 0.0667	0.218 ± 0.0669	0.216 ± 0.0712	
Cross-sectional	0.05%	0.1%	0.2%	0.5%	0.05%	0.1%	0.2%	0.5%	
WT	0.474 ± 0.149	0.396 ± 0.132	0.369 ± 0.102	0.358 ± 0.101	0.472 ± 0.136	0.484 ± 0.135	0.501 ± 0.132	0.509 ± 0.123	
Cori	1.3 ± 0.676	0.859 ± 0.442	0.554 ± 0.184	0.502 ± 0.197	1.12 ± 0.51	0.825 ± 0.363	0.664 ± 0.246	0.584 ± 0.167	
EpiNow	0.306 ± 0.199	0.277 ± 0.174	0.294 ± 0.184	0.302 ± 0.205	0.25 ± 0.146	0.29 ± 0.181	0.269 ± 0.153	0.295 ± 0.174	
GPRt	0.215 ± 0.063	0.178 ± 0.0509	0.177 ± 0.049	0.172 ± 0.0471	0.262 ± 0.09	0.265 ± 0.076	0.249 ± 0.0791	0.238 ± 0.0732	
Uniform underreporting	1%	2%	5%	10%	1%	2%	5%	10%	
WT	0.395 ± 0.105	0.389 ± 0.106	0.377 ± 0.111	0.382 ± 0.104	0.407 ± 0.0937	0.425 ± 0.114	0.429 ± 0.133	0.408 ± 0.139	
Cori	0.892 ± 0.552	0.614 ± 0.355	0.412 ± 0.162	0.38 ± 0.108	0.98 ± 0.553	0.587 ± 0.245	0.431 ± 0.149	0.38 ± 0.139	
EpiNow	0.311 ± 0.193	0.31 ± 0.186	0.359 ± 0.231	0.394 ± 0.245	0.254 ± 0.136	0.212 ± 0.128	0.251 ± 0.139	0.267 ± 0.16	
GPRt	0.204 ± 0.0806	0.22 ± 0.0878	0.181 ± 0.0677	0.181 ± 0.0467	0.26 ± 0.0948	0.259 ± 0.0839	0.222 ± 0.0865	0.233 ± 0.0807	
Random trend setting									
Longitudinal	PCR				Serological				
	0.5%	1%	2%	5%	0.5%	1%	2%	5%	
WT	0.427 ± 0.149	0.345 ± 0.101	0.321 ± 0.101	0.292 ± 0.104	0.398 ± 0.118	0.322 ± 0.0851	0.288 ± 0.0937	0.26 ± 0.0646	
Cori	1.28 ± 0.678	0.872 ± 0.512	0.622 ± 0.326	0.392 ± 0.159	1.04 ± 0.666	0.57 ± 0.325	0.358 ± 0.151	0.28 ± 0.0768	
EpiNow	0.332 ± 0.233	0.321 ± 0.195	0.337 ± 0.232	0.349 ± 0.244	0.364 ± 0.225	0.291 ± 0.207	0.304 ± 0.185	0.296 ± 0.209	
GPRt	0.199 ± 0.0745	0.187 ± 0.0652	0.181 ± 0.0551	0.157 ± 0.0476	0.232 ± 0.0733	0.216 ± 0.0754	0.213 ± 0.0699	0.194 ± 0.0694	
Cross-sectional	0.05%	0.1%	0.2%	0.5%	0.05%	0.1%	0.2%	0.5%	
WT	0.392 ± 0.123	0.335 ± 0.107	0.319 ± 0.101	0.284 ± 0.096	0.381 ± 0.128	0.404 ± 0.118	0.406 ± 0.13	0.396 ± 0.115	
Cori	0.94 ± 0.535	0.581 ± 0.217	0.478 ± 0.159	0.411 ± 0.117	0.68 ± 0.242	0.61 ± 0.219	0.513 ± 0.137	0.485 ± 0.117	
EpiNow	0.359 ± 0.202	0.356 ± 0.191	0.421 ± 0.225	0.383 ± 0.215	0.362 ± 0.242	0.438 ± 0.216	0.436 ± 0.265	0.456 ± 0.264	
GPRt	0.192 ± 0.068	0.182 ± 0.0641	0.168 ± 0.0512	0.149 ± 0.0467	0.242 ± 0.0889	0.246 ± 0.0877	0.233 ± 0.0925	0.221 ± 0.0788	
Uniform underreporting	1%	2%	5%	10%	1%	2%	5%	10%	
WT	0.285 ± 0.095	0.275 ± 0.0884	0.267 ± 0.107	0.259 ± 0.109	0.316 ± 0.0943	0.3 ± 0.0917	0.287 ± 0.0892	0.282 ± 0.0908	
Cori	0.558 ± 0.368	0.396 ± 0.187	0.326 ± 0.128	0.281 ± 0.109	0.527 ± 0.298	0.382 ± 0.205	0.303 ± 0.108	0.276 ± 0.105	
EpiNow	0.348 ± 0.244	0.315 ± 0.179	0.383 ± 0.238	0.336 ± 0.201	0.308 ± 0.173	0.356 ± 0.236	0.318 ± 0.22	0.349 ± 0.237	
GPRt	0.172 ± 0.0694	0.17 ± 0.0632	0.163 ± 0.071	0.181 ± 0.0682	0.2 ± 0.0793	0.213 ± 0.0831	0.206 ± 0.0848	0.211 ± 0.0742	

Table 1: Mean absolute error of each method averaged over instances and time points for each setting, along with standard deviation of the absolute error. ‘‘PCR’’ and ‘‘Serological’’ denote settings where the observations are generated by the respective testing method. Individual column headings give the percentage of the population enrolled in testing.

sum, we introduce a new object, the series of matrices C^t . At each time t , $C^t[t_1, t_2]$ denotes the number of individuals who have $t_{\text{convert}} = t_1$, $t_{\text{revert}} = t_2$, and have not yet actually tested positive by time t . Since A_t is selected uniformly at random from the population, independent of the infection process, the x_t individuals who test positive on day t are drawn uniformly at random from the set of all individuals who converted between days $t - d$ and t , and who have not yet reverted. Let n_{draws} denote the number of individuals in A_t who have not yet tested positive by time t and $n_{\text{conv}} = \sum_{t_1=t-d}^t \sum_{t_2=t+1}^T C^t[t_1, t_2]$ denote the number of individuals who are ‘‘eligible’’ to test positive at time t . Now $x_t | x_1 \dots x_{t-1}, \mathbf{n}, C^t$ follows a binomial distribution with n_{draws} draws and success probability $\frac{n_{\text{conv}}}{N - \sum_{i=1}^{t-1} x_i}$. Accordingly, the log-likelihood $\log p(x_t | x_1 \dots x_{t-1}, \mathbf{n}, C^t)$ can be computed in closed form. After this, we can sample $C^t | C^{t+1}$ by selecting a uniformly random individual to remove from C^{t+1} . We can view this as iteratively revealing the test-positive members of A_t after conditioning on the se-

quence of previous test results, instead of sampling the entire set up front as in the naive method.

Experimental results

We test the performance GPRt vs standard baselines on a wide variety of settings. We choose three baselines which have been recommended by leading epidemiologists as methods of choice for COVID-19 (Gostic et al. 2020). First is the *Wallinga-Teunis (WT)* method (Wallinga and Teunis 2004), which uses the distribution of the time between an infected person and their secondary infections to simulate possible who-infected-who scenarios, each of which induces a particular R_t . WT assumes that cases are observed exactly and that there is no delay in observation. Second is the method of *Cori et al. (Cori)* (Thompson et al. 2019; Cori et al. 2013) which computes a Bayesian posterior distribution in a similar Poisson branching process infection model. The difference is that their method does not place a GP prior over R_t (instead the posterior factorizes over time) and does

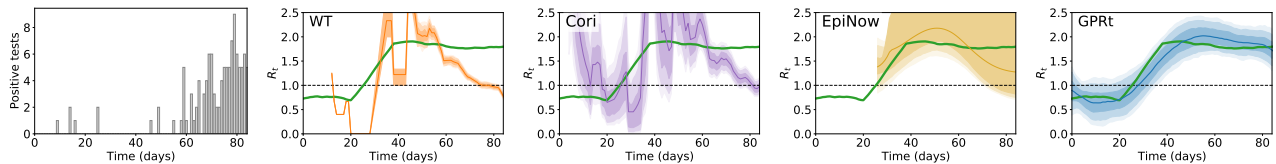


Figure 2: Observed x_t and the distribution over R_t returned by each method on an example in the outbreak setting with longitudinal sampling, $d = 14$, and a 1% sample. The green line gives the ground truth R_t .

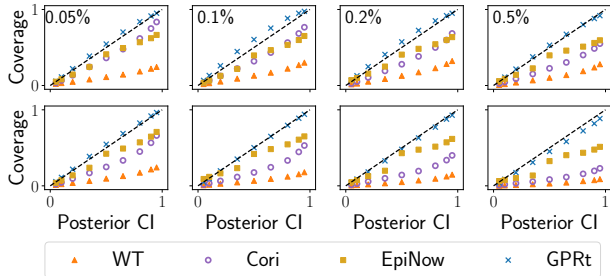


Figure 3: Calibration of each method for cross-sectional testing. Top row: PCR. Bottom row: serological. The number in the upper-left hand corner of each column gives % tested per day. Each individual plot shows the calibration of each method for that setting. Each (x, y) point gives the fraction of the ground truth data points (y) which are covered by the method’s posterior interval at level x . So, e.g., a point placed at $(0.6, 0.7)$ would indicate that 60% of the ground truth data points fell into the method’s 70% credible interval. The dashed diagonal line shows perfect calibration and points lying closer to this line indicate better calibration.

not model the sampling method or delays. For both WT and Cori, we apply a common heuristic to correct for time-to-reporting delays, which is shift the method’s predictions by the mean delay. Third is *EpiNow* (Abbott et al. 2020), a MCMC method recently developed for COVID-19 which places a GP prior over R and accounts for the delay distribution, but does not model partial observability.

We test each method in an array of settings, with different distributions for both the true value of R_t and the observations. We include two different settings for the ground truth R_t . First, the *outbreak* setting where R starts below 1 and rises above 1 at a random time. Second, the *random trend* setting where R follows a linear trend which changes randomly at multiple points in time. Details of the settings and other experimental parameters are in the appendix.

We also include different observation models characterized by the test used, the sampling method, and the sample size. We include both *PCR* and *serological* tests, using previously estimated distributions for D (Kucirka et al. 2020; Iyer et al. 2020). We also include three sampling models introduced earlier: *uniform underreporting*, *cross-sectional*, and *longitudinal*. Finally, for each of the four combinations of tests and sampling method, we include four different sample sizes. Many sizes model a challenging setting with sparse observations, representing highly limited testing capacity.

Note that the sample sizes evaluated are different for each method because they have different interpretations, e.g. 1% in the cross-sectional case means sampling 1% of the population each day while in the longitudinal case it would mean 1% every d days. For each setting, Table 1 shows the mean absolute error between the posterior mean R produced by each method and the ground truth. Each entry averages over 100 instances. For longitudinal testing we use $d = 14$; results for other values are very similar (see appendix).

Across almost all settings, our method has lower MAE than any baseline, often by a substantial margin (reducing error by a factor of 2-10x). Notably, GPRt performs well even with extremely limited data (e.g., when testing 0.05% of the population per day or when 1% of infections are observed). Performance improves with more data, but the gains are limited (e.g., 0.02-0.04 MAE), indicating that our method is able to make effective use of even very sparse data.

Figure 2 shows a representative example. The observed data is quite sparse, with 0-8 positive tests observed per day. Our method recovers a posterior which closely tracks the ground truth. WT produces an estimate which is correlated with the ground truth but has many fluctuations and overly tight confident intervals. Cori is not appropriate for data this sparse and produces a widely fluctuating posterior. EpiNow does not return an estimate for much of the time series, only estimating the part with denser observations (we gave the baselines an advantage by only evaluating their MAE where they returned an estimate). Moreover, even in the higher-observation portion, it is less accurate than GPRt.

Finally, Figure 3 shows *calibration*, a metric which evaluates the entire posterior (not just the mean). Intuitively, calibration reflects that, e.g., 90% of the data should fall into the 90%-credible interval of the posterior. Calibration is critical for the posterior distribution to be interpretable as a valid probabilistic inference, and for it to be useful in downstream decision making. Figure 3 shows the fraction of the data which is covered by the credible intervals of each method. This figure shows cross-sectional testing in the outbreak setting, but results for other settings are very similar (see appendix). *GPRt is close to perfectly calibrated (the dotted diagonal line) while the baseline methods are not well calibrated.* The baselines suffer from two problems. First, as to be expected from their higher MAE, they are biased and so their credible intervals often exclude the truth. Second, they are over-confident: paradoxically, their calibration worsens with increased data since the larger sample size makes them more confident in their erroneous prediction. We conclude that GPRt offers uniquely well-calibrated inferences.

Acknowledgments

This work was supported by the Army Research Office (MURI W911NF1810208) and NIH directors DP5 award 1DP5OD028145-01.

References

- Abbott, S.; Hellewell, J.; Thompson, R. N.; Sherratt, K.; Gibbs, H. P.; Bosse, N. I.; Munday, J. D.; Meakin, S.; Doughty, E. L.; Chun, J. Y.; et al. 2020. Estimating the time-varying reproduction number of SARS-CoV-2 using national and subnational case counts. *Wellcome Open Research* 5(112): 112.
- Barrett, C. L.; Eubank, S. G.; and Marathe, M. V. 2008. An Interaction-Based Approach to Computational Epidemiology. In *AAAI*, 1590–1593.
- Campbell, F.; Cori, A.; Ferguson, N.; and Jombart, T. 2019. Bayesian inference of transmission chains using timing of symptoms, pathogen genomes and contact data. *PLoS Computational Biology* 15(3): e1006930.
- Cazelles, B.; Champagne, C.; and Dureau, J. 2018. Accounting for non-stationarity in epidemiology by embedding time-varying parameters in stochastic models. *PLoS Computational Biology* 14(8): e1006211.
- Chakraborty, P.; Khadivi, P.; Lewis, B.; Mahendiran, A.; Chen, J.; Butler, P.; Nsoesie, E. O.; Mekaru, S. R.; Brownstein, J. S.; Marathe, M. V.; et al. 2014. Forecasting a moving target: Ensemble models for ILI case count predictions. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, 262–270. SIAM.
- Cori, A.; Ferguson, N. M.; Fraser, C.; and Cauchemez, S. 2013. A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology* 178(9): 1505–1512.
- Dureau, J.; Kalogeropoulos, K.; and Baguelin, M. 2013. Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics* 14(3): 541–555.
- Flaxman, S.; Mishra, S.; Gandy, A.; Unwin, H. J. T.; Mellan, T. A.; Coupland, H.; Whittaker, C.; Zhu, H.; Berah, T.; Eaton, J. W.; et al. 2020. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* 584(7820): 257–261.
- Gostic, K. M.; McGough, L.; Baskerville, E.; Abbott, S.; Joshi, K.; Tedijanto, C.; Kahn, R.; Niehus, R.; Hay, J. A.; De Salazar, P. M.; et al. 2020. Practical considerations for measuring the effective reproductive number, Rt. *medRxiv*.
- Havers, F. P.; Reed, C.; Lim, T.; Montgomery, J. M.; Klena, J. D.; Hall, A. J.; Fry, A. M.; Cannon, D. L.; Chiang, C.-F.; Gibbons, A.; et al. 2020. Seroprevalence of antibodies to SARS-CoV-2 in 10 sites in the United States, March 23-May 12, 2020. *JAMA Internal Medicine*.
- Iyer, A. S.; Jones, F. K.; Nodoushani, A.; Kelly, M.; Becker, M.; Slater, D.; Mills, R.; Teng, E.; Kamruzzaman, M.; Garcia-Beltran, W. F.; et al. 2020. Dynamics and significance of the antibody response to SARS-CoV-2 infection. *MedRxiv*.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*.
- King, A. A.; Ionides, E. L.; Pascual, M.; and Bouma, M. J. 2008. Inapparent infections and cholera dynamics. *Nature* 454(7206): 877–880.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. In *International Conference on Learning Representations*.
- Kucirka, L. M.; Lauer, S. A.; Laeyendecker, O.; Boon, D.; and Lessler, J. 2020. Variation in false-negative rate of reverse transcriptase polymerase chain reaction–based SARS-CoV-2 tests by time since exposure. *Annals of Internal Medicine*.
- Kucukelbir, A.; Tran, D.; Ranganath, R.; Gelman, A.; and Blei, D. M. 2017. Automatic differentiation variational inference. *The Journal of Machine Learning Research* 18(1): 430–474.
- Larremore, D. B.; Wilder, B.; Lester, E.; Shehata, S.; Burke, J. M.; Hay, J. A.; Tambe, M.; Mina, M. J.; and Parker, R. 2020. Test sensitivity is secondary to frequency and turnaround time for COVID-19 surveillance. *MedRxiv*.
- Majumder, M. S.; and Mandl, K. D. 2020. Early in the epidemic: impact of preprints on global discourse about COVID-19 transmissibility. *The Lancet Global Health* 8(5): e627–e630.
- Paisley, J.; Blei, D.; and Jordan, M. 2012. Variational Bayesian inference with stochastic search. *arXiv preprint arXiv:1206.6430*.
- Rekatsinas, T.; Ghosh, S.; Mekaru, S. R.; Nsoesie, E. O.; Brownstein, J. S.; Getoor, L.; and Ramakrishnan, N. 2015. Sourceeer: Forecasting rare disease outbreaks using multiple data sources. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, 379–387. SIAM.
- Riou, J.; and Althaus, C. L. 2020. Pattern of early human-to-human transmission of Wuhan 2019 novel coronavirus (2019-nCoV), December 2019 to January 2020. *Euro-surveillance* 25(4): 2000058.
- Saha, S.; Adiga, A.; Prakash, B. A.; and Vullikanti, A. K. S. 2015. Approximation algorithms for reducing the spectral radius to control epidemic spread. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, 568–576. SIAM.
- Sutton, R. S.; and Barto, A. G. 1998. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge.
- System, K.; Vladek, T.; and Krieger, M. 2020. rt.live. <https://github.com/rtcovidlive/covid-model>.
- Thompson, R.; Stockwin, J.; van Gaalen, R. D.; Polonsky, J.; Kamvar, Z.; Demarsh, P.; Dahlquist, E.; Li, S.; Miguel, E.; Jombart, T.; et al. 2019. Improved inference of time-varying reproduction numbers during infectious disease outbreaks. *Epidemics* 29: 100356.

Titsias, M. K.; Lawrence, N.; and Rattray, M. 2008. Markov chain Monte Carlo algorithms for Gaussian processes. *Inference and Estimation in Probabilistic Time-Series Models* 9.

Vahdat, A.; Macready, W.; Bian, Z.; Khoshaman, A.; and Andriyash, E. 2018. DVAE++: Discrete Variational Autoencoders with Overlapping Transformations. In *International Conference on Machine Learning*, 5035–5044.

Wallinga, J.; and Teunis, P. 2004. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology* 160(6): 509–516.

Wilder, B.; Charpignon, M.; Killian, J. A.; Ou, H.-C.; Mate, A.; Jabbari, S.; Perrault, A.; Desai, A.; Tambe, M.; and Majumder, M. S. 2020. Modeling Between-Population Variation in COVID-19 Dynamics in Hubei, Lombardy, and New York City. *SSRN* .

Zhang, Y.; Adiga, A.; Vullikanti, A.; and Prakash, B. A. 2015. Controlling propagation at group scale on networks. In *2015 IEEE International Conference on Data Mining*, 619–628. IEEE.

Zhang, Y.; and Prakash, B. A. 2014. Dava: Distributing vaccines over networks under prior information. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, 46–54. SIAM.