

Semiparametric Bayesian inference for the transmission dynamics of COVID-19 with a state-space model

Tianjian Zhou^{a,*}, Yuan Ji^b

^a Department of Statistics, Colorado State University, 851 Oval Dr, Fort Collins, CO 80523-1877, USA

^b Department of Public Health Sciences, University of Chicago, Chicago, IL, USA

ARTICLE INFO

Keywords:

Effective reproduction number
Forecasting
Gaussian process
Infectious disease
Parallel tempering
SIR model

ABSTRACT

The outbreak of Coronavirus Disease 2019 (COVID-19) is an ongoing pandemic affecting over 200 countries and regions. Inference about the transmission dynamics of COVID-19 can provide important insights into the speed of disease spread and the effects of mitigation policies. We develop a novel Bayesian approach to such inference based on a probabilistic compartmental model using data of daily confirmed COVID-19 cases. In particular, we consider a probabilistic extension of the classical susceptible-infectious-recovered model, which takes into account undocumented infections and allows the epidemiological parameters to vary over time. We estimate the disease transmission rate via a Gaussian process prior, which captures nonlinear changes over time without the need of specific parametric assumptions. We utilize a parallel-tempering Markov chain Monte Carlo algorithm to efficiently sample from the highly correlated posterior space. Predictions for future observations are done by sampling from their posterior predictive distributions. Performance of the proposed approach is assessed using simulated datasets. Finally, our approach is applied to COVID-19 data from six states of the United States: Washington, New York, California, Florida, Texas, and Illinois. An R package BaySIR is made available at <https://github.com/tianjianzhou/BaySIR> for the public to conduct independent analysis or reproduce the results in this paper.

1. Introduction

The outbreak of Coronavirus Disease 2019 (COVID-19), caused by Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2), was declared a pandemic on March 11, 2020 by the World Health Organization. As of September 3, 2020, the number of confirmed COVID-19 cases worldwide has exceeded 26 million, and the death toll has surpassed 864,000. In order to control the spread of the virus, countries around the world have implemented unprecedented non-pharmaceutical interventions, such as case isolation, closure of schools, stay-at-home orders, banning of mass gatherings, and local and national lockdowns. At the same time, social distancing and mask wearing by the public also contribute to the containment of COVID-19.

Researchers have made substantial efforts to study the transmission dynamics of COVID-19, evaluate the effects of government interventions, and forecast infection and death counts. The modeling approaches taken by existing works can be broadly categorized into three groups: (i) curve fitting, (ii) compartmental modeling, and (iii) agent-based modeling. Curve fitting approaches, by definition, fit a curve to the observed number of confirmed cases or deaths. For example, an

early model (IHME COVID-19 health service utilization forecasting team and)[22] uses a Gaussian error function to model the cumulative death rate at a specific location, and Woody et al. [50] use negative binomial regression to model the daily deaths. Compartmental modeling approaches consider a partition of the population into compartments corresponding to different stages of the disease, and characterize the transmission dynamics of the disease by the flow of individuals through compartments. Examples of compartmental modeling approaches include Aguilar et al. [1], Chen and Qiu [8], Flaxman et al. [12], Giordano et al. [16], Gu et al. [18], Li et al. [32], Sun et al. [42], Wang et al. [45], Wang et al. [46], Wu et al. [51], Zhang et al. [53], among many others. Finally, agent-based modeling approaches (e.g., [17,21]) use computer simulations to study the dynamic interactions among the agents (e.g., people in epidemiology) and between an agent and the environment.

In this paper, we develop a novel semiparametric Bayesian approach to modeling the transmission dynamics of COVID-19, which is critical for characterizing disease spread. We aim to address a few issues related to the COVID-19 pandemic. First, we provide estimation of key epidemiological parameters, such as the effective reproduction number

* Corresponding author.

E-mail address: tianjian.zhou@colostate.edu (T. Zhou).

<https://doi.org/10.1016/j.cct.2020.106146>

Received 3 July 2020; Received in revised form 3 September 2020; Accepted 7 September 2020

Available online 15 September 2020

1551-7144/ © 2020 Elsevier Inc. All rights reserved.

of COVID-19. The Bayesian framework allows us to elicit informative priors for parameters that are difficult to estimate due to lack of data based on clinical characteristics of COVID-19, and also offers coherent uncertainty quantification for the parameter estimates. Our second goal is to make predictions about the future trends of the spread of COVID-19 (e.g., future case counts), which will be done by calculating the posterior predictive distributions for the future observations. Although such predictions are technically straightforward, we avoid over-interpretation of the predictions because they rely on extrapolation of highly unpredictable human behaviors and the number of diagnostic tests that will be deployed. Nevertheless, such predictions may be useful for the public and decision makers to understand the trends and future impacts of COVID-19 based on current rates of transmission. We shall see this in our case studies later. Our analysis will be based on a probabilistic compartmental model motivated by the classical susceptible-infectious-recovered model [27]. Therefore, our approach belongs to the compartmental modeling group. We will use data of daily confirmed COVID-19 cases reported by the Center for Systems Science and Engineering at Johns Hopkins University (JHU CSSE) [9]. We provide an R package BaySIR, available at <https://github.com/tianjianzhou/BaySIR>, that can be used to conduct independent analysis of COVID-19 data or reproduce the results in this paper.

The proposed Bayesian approach attempts to improve COVID-19 modeling in at least four aspects. First, we explicitly model the number of *undocumented infections*, which is only considered by some, but not all, existing works. Due to the potentially limited testing capacity and the existence of pre-symptomatic and asymptomatic COVID-19 cases [19,41], many infected individuals may not have been detected as having the disease. Therefore, modeling of undocumented infections is essential for accurate inference. Second, we estimate the disease transmission rate via Gaussian process regression (GPR), a semiparametric regression method. The GPR approach is highly flexible and captures nonlinear and non-monotonic relationships without the need of specific parametric assumptions. Third, we develop a parallel-tempering Markov chain Monte Carlo (PTMCMC) algorithm to efficiently sample from the posterior distribution of the epidemiological parameters, which leads to improvements in convergence and mixing compared to a standard MCMC procedure. We find that standard MCMC cannot produce reliable inference due to poor mixing. Lastly, we rigorously assess our approach through simulation studies, sensitivity analyses, cross-validation and goodness-of-fit tests. Such validations provide insights into the modeling of COVID-19 data, not only for our approach, but also for others based on similar assumptions such as the popular compartmental modeling approaches.

The remainder of the paper is organized as follows. In Section 2, we provide a brief review of the susceptible-infectious-recovered (SIR) compartmental model. In Section 3, we develop a probabilistic state-space model for COVID-19 motivated by the classical SIR model. In Section 4, we present strategies for posterior inference. In Section 5, we carry out simulation studies to assess the performance of our method in estimating the epidemiological parameters. In Section 6, we apply our method to COVID-19 data from six states of the United States (U.S.): Washington, New York, California, Florida, Texas, and Illinois. We conclude with a discussion in Section 7.

2. Review of the susceptible-infectious-recovered model

We start with a review of the susceptible-infectious-recovered (SIR) model [27,48], a simple type of compartmental model. The purpose of this review is to introduce the reader to the basics of epidemic modeling and motivate our proposed approach.

Consider a closed population of size N . Here, “closed” means that N does not vary over time. It is a good approximation for a fast-spreading and less fatal pandemic like COVID-19. The SIR model divides the population into the following three compartments:

- (S) Susceptible individuals: those who do not have the disease but

may be infected;

- (I) Infectious individuals: those who have the disease and are able to infect the susceptible individuals;

- (R) Recovered/removed individuals: those who had the disease but are then removed from the possibility of being infected again or spreading the disease. Here, the removal can be due to several possible reasons, including death, recovery with immunity against reinfection, and quarantine and isolation from the rest of the population.

At time t ($t \geq 0$), denote by S_t , I_t and R_t the numbers of individuals in the S, I and R compartments, respectively, and write $\mathbf{V}_t = (S_t, I_t, R_t)$. We have $S_t + I_t + R_t = N$.

2.1. Deterministic SIR models

The classical SIR model [27] describes the flow of people from S to I to R via the following system of differential equations:

$$\frac{dS_t}{dt} = -\frac{\beta}{N} S_t I_t, \quad \frac{dI_t}{dt} = \frac{\beta}{N} S_t I_t - \alpha I_t, \quad \frac{dR_t}{dt} = \alpha I_t \quad (1)$$

here, β is the *disease transmission rate*, and α is the *removal rate*. The rationale behind the first equation in Eq. (1) is as follows: suppose each infectious individual makes *effective contacts* (sufficient for disease transmission) with β others per unit time; therefore, $\beta S_t/N$ of these contacts are with susceptible individuals per unit time, and as a result, I infectious individuals lead to a rate of new infections $(\beta S_t/N) \cdot I$. The third equation in Eq. (1) describes that the infectious individuals leave the infective class at a rate of αI . The second equation in Eq. (1) follows immediately from the first and third equations. The parameters β and α are determined according to the natural history of the disease. The quantities $\mathcal{R}_0 = \beta/\alpha$ and $\mathcal{R}_e = (\beta S_0)/(\alpha N)$ are referred to as the *basic reproduction number* and *effective reproduction number*, respectively, where S_0 is the initial number of susceptibles at time $t = 0$.

In some applications, it may be convenient to consider a discrete-time approximation of the differential equations in Eq. (1), which can be expressed as follows:

$$\begin{aligned} S_t &= S_{t-1} - \beta S_{t-1} I_{t-1}/N \\ I_t &= (1 - \alpha) I_{t-1} + \beta S_{t-1} I_{t-1}/N \\ R_t &= R_{t-1} + \alpha I_{t-1} \end{aligned} \quad (2)$$

for $t = 1, 2, \dots$. This discretization replaces the derivatives in Eq. (1) by the differences per unit time.

The SIR models given by Eqs. (1) and (2) are both *deterministic* models, meaning that their behaviors are completely determined by their initial conditions and parameter values.

2.2. Stochastic SIR models

The deterministic SIR models are appealing due to their simplicity. However, the spread of disease is naturally stochastic. The disease transmission between two individuals is random rather than deterministic. Therefore, a stochastic formulation of the SIR model may be preferred for epidemic modeling, because it allows one to more readily capture the randomness of the epidemic process.

In a stochastic SIR model, $\{\mathbf{V}_t : t \geq 0\}$ is treated as a stochastic process. A commonly used formulation is as follows ([15; 38; 4]). Suppose that an infectious individual makes effective contacts with any given individual in the population at times given by a Poisson process of rate β/N , and assume all these Poisson processes are independent of each other. Therefore, the expected number of effective contacts made by each infectious individual is β per unit time. Furthermore, suppose each infectious individual remains so (before being removed) for a period of time, known as the *infectious period*. Lastly, assume that the length of the infectious period for each individual is independent and follows an exponential distribution with mean α^{-1} . It can be shown that $\{\mathbf{V}_t : t \geq 0\}$ is a Markov process with transition probabilities:

$$\begin{aligned}
\text{Infection: } & \Pr[\mathbf{V}_{t+\delta} = (s-1, i+1, r) \mid \mathbf{V}_t = (s, i, r)] = \beta s i \delta / N + o(\delta) \\
\text{Removal: } & \Pr[\mathbf{V}_{t+\delta} = (s, i-1, r+1) \mid \mathbf{V}_t = (s, i, r)] = \alpha i \delta + o(\delta)
\end{aligned} \tag{3}$$

here, δ is a small increment in time.

2.3. State-space SIR models

There are, of course, other ways to model the uncertainty of the epidemic process. Probabilistic state-space modeling approaches that build on deterministic models have recently been popular in the statistics literature [11,36,37]. A state-space SIR model typically consists of two components: An evolution model for the epidemic process, and an observation model for the data. As an example, the model in Osthus et al. [37] has the form

$$\text{Evolution: } \mathbf{V}_t \sim p[\mathbf{V}_t \mid f(\mathbf{V}_{t-1}, \beta, \alpha), \kappa]$$

$$\text{Observation: } \tilde{I}_t \sim p(\tilde{I}_t \mid I_t, \lambda)$$

for $t = 1, 2, \dots$. In the evolution model, $f(\mathbf{V}_{t-1}, \beta, \alpha)$ is the solution to Eq. (1) at time t with a initial value of $\mathbf{V}_{t-1}$ at time $(t-1)$ and parameters β and α , and $\mathbf{V}_t$ is assumed to be centered at $f(\cdot)$ with its variance characterized by κ . In other words, κ measures the derivation of $\mathbf{V}_t$ from the solution given by the deterministic model. In the observation model, $\tilde{I}_t$ is the number of patients seen with the disease reported by healthcare providers, which can be thought of as a proxy to the true number of infectious individuals I_t . The observation $\tilde{I}_t$ is assumed to be centered at I_t with variance characterized by λ . State-space epidemic models are quite flexible and are in general more computationally manageable compared to stochastic epidemic models as in Eq. (3).

The SIR model can be extended in many different ways, such as by considering vital dynamics (births and deaths) and demographics, adding more compartments to the model, and allowing more possible transitions across compartments. For example, the susceptible-exposed-infectious-recovered (SEIR) model includes an additional compartment for exposed individuals who are exposed to the disease but are not yet infectious, and the susceptible-infectious-recovered-infectious (SIRS) model allows recovered individuals to return to a susceptible state. These extensions may better capture the characteristics of the disease under consideration. For a comprehensive review of deterministic epidemic models, see, for example, Anderson and May [3], Hethcote [20] or Brauer [6]. For a comprehensive review of stochastic epidemic models, see, for example, Becker and Britton [5], Andersson and Britton [4] or Allen [2].

3. Proposed model for COVID-19

We now turn to our proposed model for the COVID-19 data, which belongs to the state-space model category (Section 2.3). Our approach integrates the discrete-time deterministic SIR model (Eq. (2)) and semiparametric Bayesian inference. To capture some unique features of COVID-19, we consider the following extensions of the classical SIR model. First, we split the infectious individuals into two subgroups: undocumented infectious individuals and documented infectious individuals. The reason is that many people infected with SARS-CoV-2 have not been tested for the virus thus are not detected or reported as having the infection [32]. Second, we allow some epidemiological parameters (such as the disease transmission rate β) to be time-varying to reflect the impact of mitigation policies such as stay-at-home orders and the change of public awareness of the disease over time. For the COVID-19 application, many existing works have considered time-varying epidemiological parameters, such as Flaxman et al. [12], Gu et al. [18], Sun et al. [42], and Wang et al. [46]. We discuss details next.

3.1. Model for the epidemic process

Consider the transmission dynamics of COVID-19 in a specific

country or region (e.g., a state, province or county). For simplicity, we consider a closed population (with no immigration and emigration) and also ignore nature births and deaths. Let N denote the population size. At any time point, we assume that each individual in the population precisely belongs to one of the following four compartments:

(S) Susceptible individuals who do not have the disease but are susceptible to it;

(UI) Undocumented infectious individuals who have the disease and may infect the susceptible individuals. However, they have not been detected as having the disease for several possible reasons. For example, they may have limited symptoms and are thus not tested for the disease;

(DI) Documented infectious individuals who have been confirmed as having the disease and are capable of infecting the susceptible individuals;

(R) Removed individuals who had the disease but are then removed from the possibility of being infected again or spreading the disease.

We further assume that the infectious individuals (including both the UI- and DI-individuals) infect the S-individuals with a transmission rate of β . After being infected, a S-individual first becomes an UI-individual before being detected as a DI-individual. All the infectious (UI- and DI-) individuals recover or die with a removal rate of α . Those UI-individuals who have not been removed are diagnosed with the disease with a diagnosis rate of γ . In total, there are four possible transitions across compartments: S to UI, UI to R, UI to DI, and DI to R. See Fig. 1. Note that it is possible to assume different transmission rates for the UI- and DI-individuals, or to further split the UI and DI compartments into smaller subgroups (e.g., quarantined, hospitalized, etc.) with each subgroup having its distinct transmission rate. It is also possible to consider an extra compartment for the exposed (but not yet infectious) individuals as in the SEIR model. Here, we use a more parsimonious model without the exposed compartment for simplicity and characterize the average transmission rate for all infectious individuals with a single parameter β (β depends on time, which will be clear later). Finally, we assume recovery from COVID-19 confers immunity to reinfection, although there is only limited evidence for this assumption [28,34].

We define day $t = 0$ as the date when the 100th case is confirmed in the country/region under consideration, and index subsequent dates by $t = 1, 2, \dots, T$, where T is the current date. The reason for choosing day 0 in this way is because we believe the transmission dynamics of the disease is more trackable after a sufficient number of infectious individuals are reported in the country/region, although the choice of “the 100th case” is arbitrary and can be modified. Denote by S_t, I_t^U, I_t^D and R_t the numbers of individuals belonging to compartments S, UI, DI and R on day t , respectively. We have $S_t + I_t^U + I_t^D + R_t \equiv N$. The transmission rate and diagnosis rate are allowed to vary over time and are hereafter denoted by β_t and γ_t , respectively. The number of individuals diagnosed with the disease between day $(t-1)$ and day t is observed and is denoted by B_{t-1} . This is our data. We propose modeling the transmission dynamics of COVID-19 over time by the following equations:




Fig. 1. Compartmental model for COVID-19. We consider four compartments and four possible transitions across compartments. The number under each arrow indicates the transition rate between two compartments.

$$\begin{aligned}
S_t &= S_{t-1} - \beta_{t-1} S_{t-1} (I_{t-1}^U + I_{t-1}^D) / N \\
I_t^U &= (1 - \alpha) I_{t-1}^U + \beta_{t-1} S_{t-1} (I_{t-1}^U + I_{t-1}^D) / N - B_{t-1} \\
I_t^D &= (1 - \alpha) I_{t-1}^D + B_{t-1} \\
R_t &= R_{t-1} + \alpha (I_{t-1}^U + I_{t-1}^D)
\end{aligned} \quad (4)$$

for $t = 1, \dots, T$. Denote by $\mathbf{V}_t = (S_t, I_t^U, I_t^D, R_t)$. The epidemic process, $\{\mathbf{V}_t, t = 0, 1, \dots, T\}$, is determined by its initial value $\mathbf{V}_0$, the parameters $\{\beta_t, \alpha\}$, and the observations $\{B_t\}$. Rigorously speaking, $\mathbf{V}_t$ should be a vector of non-negative integers, but for computational convenience, we relax this restriction and only require it to be a vector of non-negative real numbers. Model (4) is a simple extension of (2) by adding a component of I^U , the undocumented infections, and by incorporating the observed daily new cases B_{t-1} into the equations. Later, we introduce a model for the observation B_{t-1} to complete the state-space model.

With time-varying disease transmission rates, the basic reproduction number and effective reproduction number are also functions of time. That is, $\mathcal{R}_0(t) = \beta_t / \alpha$ and

$$\mathcal{R}_e(t) = (\beta_t S_t) / (\alpha N)$$

here, $\mathcal{R}_e(t)$ is interpreted as the rate of secondary infections generated by each infectious case at time t , scaled by the length of the infectious period (α^{-1}). If $\mathcal{R}_e(t) < 1$ for $t \geq t^*$, then the number of infectious individuals ($I_t^U + I_t^D$) will monotonically decrease after time t^* , because each infectious individual will only be able to infect less than 1 other during the course of his/her infectious period. In other words, an $\mathcal{R}_e(t) < 1$ indicates containment of the disease. Due to the important role of $\mathcal{R}_e(t)$ in characterizing disease spread, we consider the estimation of $\mathcal{R}_e(t)$ as our main interest.

3.2. Model for the observed data

Our observations only consist of the daily new confirmed COVID-19 cases, B_t . Assume that on day t , the UI-individuals who have not been removed are diagnosed with the disease with a diagnosis rate of γ_t . Mathematically, this means $B_t = \gamma_t (1 - \alpha) I_t^U$, where γ_t is between 0 and 1. We consider the logit transformation of γ_t , $\tilde{\gamma}_t = \text{logit}(\gamma_t) \triangleq \log[\gamma_t / (1 - \gamma_t)]$. Other transformations, such as the probit and complementary log-log transformations, can also be specified in the BaySIR package. Empirically we find the proposed model to be robust to different specifications of the link function (see appendix C). We assume a prior transformation

$$\tilde{\gamma}_t \sim N(\mathbf{y}_t^\top \boldsymbol{\eta}, \sigma_\gamma^2) \quad (5)$$

where $\mathbf{y}_t$ is a vector of covariates that are thought to be related to the diagnosis rate. In other words, the sampling model for B_t can be written as

$$\text{logit} \left[\frac{B_t}{(1 - \alpha) I_t^U} \right] | I_t^U, \alpha \sim N(\mathbf{y}_t^\top \boldsymbol{\eta}, \sigma_\gamma^2) \quad (6)$$

In the simulation studies and real data analyses, we use a simple choice of $\mathbf{y}_t = 1$, assuming the mean diagnosis rate is a constant. It is possible to include other covariates in $\mathbf{y}_t$, such as the number of tests (available at the COVID Tracking Project, <https://covidtracking.com/>), but empirically we find it hard to detect the effects of these covariates. In the BaySIR package, the user has the option to include any covariates. The parameters $\boldsymbol{\eta}$ and σ_γ^2 are the regression coefficients and variance term, respectively, where σ_γ^2 captures random fluctuations of confirmed case counts and report errors.

For some countries and regions, the numbers of recoveries and deaths are also available, and one may think of using them as the observed number of removed individuals. We choose not to use these data for two reasons. First, many infected individuals, even with confirmed disease, are not hospitalized, and their recoveries are not recorded. In other words, the reported number of recoveries and deaths is a

significant underestimate of the size of the removed population. Second, according to Wölfel et al. [49] and He et al. [19], the ability of a COVID-19 patient to infect others becomes negligible several days before the patient recovers or dies, suggesting that “removal” in our application is not equivalent to “recovery or death”.

3.3. Prior specification

In what follows, we discuss prior specification for the initial condition and parameters. Due to the limited amount of observable information, many latent variables and parameters in the proposed model are unidentifiable. See Appendix A for a detailed discussion with an example showing that two epidemic processes with distinct parameters lead to exactly the same observed data. We note that this problem is pervasive in most existing methods, and a typical solution to the problem is to prespecify some parameter values based on prior knowledge. Here, we elicit informative priors for some parameters based on the clinical characteristics of COVID-19, which favor more clinically plausible estimates.

3.3.1. Initial condition

The initial condition of the epidemic process refers to the vector $\mathbf{V}_0 = (S_0, I_0^U, I_0^D, R_0)$. We assume that there are no removed individuals on day 0, i.e., $R_0 = 0$. As a result, the number of DI-individuals on day 0, I_0^D , equals to the cumulative number of confirmed cases on that day and is observed. We further assume

$$I_0^U / I_0^D \sim \text{Ga}(\nu_1, \nu_2)$$

where $\text{Ga}(\nu_1, \nu_2)$ refers to a gamma distribution with shape and rate parameters ν_1 and ν_2 , respectively. We set $\nu_1 = 5$ and $\nu_2 = 1$, such that $E(I_0^U / I_0^D) = 5$. This choice is based on the findings in Li et al. [32] that 86% of all infections were undocumented at the beginning of the epidemic in China. Lastly, note that $S_0 = N - I_0^U - I_0^D - R_0$.

3.3.2. Transmission rate

The disease transmission rate β_t must be non-negative. We consider $\tilde{\beta}(t) = \log(\beta_t)$ and assume

$$\tilde{\beta}(t) \sim \text{GP}[m(t), C(t, t')]$$

where $\text{GP}[m(t), C(t, t')]$ refers to a Gaussian process (GP) with mean function $m(t)$ and covariance function $C(t, t')$. The GP [39] is a very flexible prior model for a stochastic process. It enables one to capture potential non-linear relationships between t and $\tilde{\beta}(t)$ without the need to impose any parametric assumptions. Specifically, for any $t_1, \dots, t_n \geq 0$, the vector $(\tilde{\beta}(t_1), \dots, \tilde{\beta}(t_n))^\top$ follows a multivariate Gaussian distribution with mean $(m(t_1), \dots, m(t_n))^\top$ and covariance matrix $\mathbf{C}$ with the (i, j) -th entry being $C(t_i, t_j)$. For applications of GP to epidemic modeling, see, for example, Xu et al. [52] and Kyraios and O'Neill [29].

We specify $m(t)$ and $C(t, t')$ as below:

$$m(t) = \mathbf{x}_t^\top \boldsymbol{\mu}, \quad C(t, t') = \sigma_\beta^2 \rho^{|t-t'|} \quad (7)$$

here, $\mathbf{x}_t$ is a vector of covariates that are thought to be related to the transmission rate, and $\boldsymbol{\mu}$ is a vector of regression coefficients. In the simulation studies and real data analyses, we use $\mathbf{x}_t = (1, t)^\top$, which contains an intercept term and the time. Other covariates, such as indicators for mitigation policies at time t , may also be included in $\mathbf{x}_t$. Nevertheless, in practice, we find our GP model with a time trend is sufficient to capture the change of $\tilde{\beta}(t)$ over time and the potential effects of mitigation policies and public awareness. Users of our software may include other covariates using the R package BaySIR. The variance parameter σ_β^2 characterizes the amplitude of the difference between $\tilde{\beta}(t)$ and $m(t)$, and the correlation parameter ρ characterizes the correlation between $\tilde{\beta}(t)$ and $\tilde{\beta}(t')$ for any t and t' . We note that based on our specification of the covariance function, our GP model is

equivalent to a first-order autoregressive model. Indeed, autoregressive models of any orders are discrete-time equivalents of GP models with Matérn covariance functions [40].

We place the following priors on μ , σ_β and ρ :

$$\mu \sim N(\mu^*, \Sigma_\mu), \quad \sigma_\beta^2 \sim \text{Inv} - \text{Ga}(11, 1), \quad \rho \sim \text{Beta}(4, 1)$$

such that $E(\sigma_\beta^2) = 0.1$ and $E(\rho) = 0.8$. Here, $\text{Inv} - \text{Ga}(\cdot, \cdot)$ refers to an inverse gamma distribution, and $\text{Beta}(\cdot, \cdot)$ refers to a beta distribution. The prior choices for σ_β^2 and ρ shrink $\tilde{\beta}(t)$ toward its mean function (i.e., a linear regression model) and impose a strong prior correlation between the transmission rates for two consecutive days. For the prior of μ , we use $\mu^* = (-1.31, 0)^\top$ and $\Sigma_\mu = \text{diag}(0.3^2, 1^2)$, where $\text{diag}(\cdot)$ represents a diagonal matrix. In this way, the prior median of the basic reproduction number on day 0 is 2.5 (with 95% credible interval 1.4 to 4.5), assuming the infectious period is 9.3 days. This is based on the findings in Li et al. [31] and Wu et al. [51]. The prior also induces a mild shrinkage (toward 0) for the regression coefficient of the time trend.

3.3.3. Removal rate

The removal rate is between 0 and 1. The inverse of the removal rate, α^{-1} , corresponds to the average time to removal after infection. We assume

$$\alpha^{-1} \sim \text{Ga}(\nu_1^\alpha \nu_2^\alpha) \cdot \mathbf{1}(\alpha^{-1} \geq 1)$$

We take $\nu_1^\alpha = 325.5$ and $\nu_2^\alpha = 35$, such that $E(\alpha^{-1}) = 9.3$ with prior 95% credible interval between 8.3 and 10.3 days. The mean infectious period of 9.3 days is chosen based on the findings in He et al. [19], who estimated that the infectiousness of COVID-19 starts from around 2.3 days before symptom onset and declines quickly within 7 days after symptom onset.

Diagnosis rate. We place the following standard weakly informative priors on η and σ_η^2 , the regression coefficients and variance term in the diagnosis rate model (Eq. (5)):

$$\eta \sim N(\eta^*, \Sigma_\eta), \quad \sigma_\eta^2 \sim \text{Inv} - \text{Ga}(1, 1)$$

when y_t only has an intercept term, we use $\eta \sim N(0, 1^2)$.

4. Inference

4.1. Posterior sampling

Let $\theta = \{I_0^U, \beta, \alpha, \mu, \sigma_\beta, \rho, \eta, \sigma_\eta^2\}$ denote all model parameters and hyperparameters, where $\beta = (\beta_0, \beta_1, \dots, \beta_T)$, and let $\mathbf{B} = (B_0, B_1, \dots, B_T)$ be the vector of daily increments in confirmed cases. The joint posterior distribution of θ is given by

$$\pi(\theta | \mathbf{B}, I_0^D) \propto \left[\prod_{t=0}^T \phi(\tilde{y}_t | \mathbf{y}_t^\top \eta, \sigma_\eta^2) \right] \cdot \pi^*(\theta)$$

where $\phi(\cdot | \mu, \sigma^2)$ denotes the density function of a normal distribution with mean μ and standard deviation σ^2 , and $\pi^*(\theta)$ represents the prior density of θ . Recall that $\tilde{y}_t = \text{logit}\{B_t / [(1 - \alpha)I_t^U]\}$.

We use a Markov chain Monte Carlo (MCMC) algorithm (see, e.g., [33]), in particular the Gibbs sampler, to simulate from the posterior distribution and implement posterior inference. Metropolis-Hastings steps are used when the conditional posterior distribution of a parameter is not available in closed form. The regular Gibbs sampler is not very efficient in our application because of the strong correlations among the model parameters. This issue was also noted by Osthus et al. [37]. We therefore use parallel tempering (PT) to improve the convergence and mixing of the Markov chains [14]. Consider J parallel Markov chains with a target distribution of

$$\pi_j(\theta_j | \mathbf{B}, I_0^D) \propto \left[\prod_{t=0}^T \phi(\tilde{y}_{t,j} | \mathbf{y}_t^\top \eta_j, \sigma_{\eta_j}^2) \right]^{1/\Delta_j} \cdot \pi^*(\theta_j)$$

for the j -th chain, where Δ_j is the *temperature*. The temperatures $\{\Delta_1, \Delta_2, \dots, \Delta_J\}$ are decreasing with $\Delta_J = 1$. Thus the target distribution of the J -th chain is the original posterior $\pi(\theta | \mathbf{B}, I_0^D)$. At each MCMC iteration, we first independently update all J chains based on Gibbs transition probabilities. Then, for $j = 1, 2, \dots, J - 1$, we propose a swap between θ_j and θ_{j+1} and accept the proposal with probability

$$p_{\text{swap}}(\theta_j, \theta_{j+1}) = 1 \wedge \left[\prod_{t=0}^T \frac{\phi(\tilde{y}_{t,j+1} | \mathbf{y}_t^\top \eta_{j+1}, \sigma_{\eta_{j+1}}^2)}{\phi(\tilde{y}_{t,j} | \mathbf{y}_t^\top \eta_j, \sigma_{\eta_j}^2)} \right]^{\frac{1}{\Delta_j} - \frac{1}{\Delta_{j+1}}}$$

The draws from the J -th chain are kept. A chain with a higher temperature can more freely explore the posterior space, and the swap proposal allows interchange of states between adjacent chains. Therefore, the PT scheme helps the Markov chain avoid getting stuck at local optima. In Appendix B, we demonstrate the advantage of the PT scheme with an example.

In the simulation studies and real data analyses, we run $J = 10$ parallel Markov chains with a temperature of $\Delta_j = 1.5^{10-j}$ for the j -th chain. We run MCMC simulation for 50,000 iterations, discard the first 20,000 draws as initial burn-in, and keep one sample every 30 iterations. This leaves us a total of 1000 posterior samples.

4.2. Predictive inference

In addition to the estimation of epidemiological parameters, one may be interested in the prediction of a future observation, which can be achieved by sampling from its posterior predictive distribution. As an example, let $\mathbf{B}^* = (B_{T+1}, \dots, B_{T+T^*})$ denote the vector of daily confirmed cases for future days $t = T + 1, \dots, T + T^*$. The posterior predictive distribution of $\mathbf{B}^*$ is given by

$$\pi(\mathbf{B}^* | \mathbf{B}, I_0^D) = \int \pi(\mathbf{B}^* | \theta, \mathbf{B}, I_0^D) \cdot \pi(\theta | \mathbf{B}, I_0^D) d\theta \quad (8)$$

Sampling from Eq. (8) involves computing $\pi(\tilde{\beta}^* | \mathbf{B}, I_0^D) = \int \pi(\tilde{\beta}^* | \tilde{\beta}) \cdot \pi(\tilde{\beta} | \mathbf{B}, I_0^D) d\tilde{\beta}$ for $\tilde{\beta}^* = (\tilde{\beta}_{T+1}, \dots, \tilde{\beta}_{T+T^*})$. We have

$$\tilde{\beta}^* | \tilde{\beta} \sim N[\mathbf{X}^* \mu + \mathbf{C}^* \mathbf{C}^{-1}(\tilde{\beta} - \mathbf{X} \mu), \mathbf{C}^{**} - \mathbf{C}^* \mathbf{C}^{-1} \mathbf{C}^*]$$

where $\mathbf{X} = (\mathbf{x}_0, \dots, \mathbf{x}_T)^\top$, $\mathbf{X}^* = (\mathbf{x}_{T+1}, \dots, \mathbf{x}_{T+T^*})^\top$, $\mathbf{C}^*$ is a $T^* \times (T + 1)$ matrix with the (i, j) -th entry being $C(T + i, j - 1)$, and $\mathbf{C}^{**}$ is a $T^* \times T^*$ matrix with the (i, j) -th entry being $C(T + i, T + j)$. This is based on a GP prediction rule [39].

5. Simulation studies

We assess the performance of the proposed method in estimating the epidemiological parameters by applying it to simulated epidemic time series. Consider a closed population of size $N = 20,000,000$. We assume the initial condition on day 0 is $I_0^D = 100$, $I_0^U = 800$, $R_0 = 0$, and $S_0 = N - I_0^D - I_0^U$. We set the removal rate $\alpha = 9.3^{-1}$. For the transmission rate, we consider the following three scenarios:

(Scn. 1) $\beta_t = b \cdot \alpha / [(t + 1)^c - a]$, where a, b and c are chosen such that $\mathcal{R}_0(0) = 3$, $\mathcal{R}_0(14) = 2$ and $\mathcal{R}_0(49) = 1$;

(Scn. 2) $\beta_t = \alpha \cdot \exp[a \cdot \sin(0.2t) - bt + c]$, where a, b and c are chosen such that $\mathcal{R}_0(0) = 2.5$, $\mathcal{R}_0(14) = 2.2$ and $\mathcal{R}_0(49) = 1$;

(Scn. 3) $\beta_t = \alpha \cdot \exp[\log(2.5) - 0.4 \cdot [(t/20)]]$, where $[a]$ represents the largest integer that is smaller than a .

Recall that $\mathcal{R}_0(t) = \beta_t / \alpha$. In all the scenarios, $\mathcal{R}_0(t) \rightarrow 0^+$ as $t \rightarrow \infty$. For scenario 2, $\mathcal{R}_0(t)$ is non-monotonic, and for scenario 3, $\mathcal{R}_0(t)$ is discontinuous. Next, we generate $\tilde{y}_t \sim N[\text{logit}(0.2), 0.25^2]$ and $\eta_t = 1 - \exp(-\exp(\tilde{y}_t))$. Finally, for each scenario, we generate a hypothetical epidemic process for 80 days according to Eq. (4) with $B_t = \gamma_t(1 - \alpha)I_t^U$. We keep $\mathbf{B} = (B_0, \dots, B_T)$ and I_0^D as our observations ($T = 79$). The simulated datasets, shown in Fig. 2 (upper panel), are similar to a real COVID-19 dataset (e.g., Fig. 3).

We fit the proposed model to the simulated datasets using the




Fig. 2. The upper panel shows the simulated daily confirmed cases for the three scenarios. The lower panel shows the estimated time-varying effective reproduction numbers (solid black line), 95% credible intervals (grey band), and simulation truth (dashed red line) for the three scenarios. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

PTMCMC algorithm. Fig. 2 (lower panel) shows a comparison of the estimated time-varying effective reproduction numbers (posterior medians) with the simulation truth. The simulation truth is nicely recovered, and the 95% credible intervals of $\mathcal{R}_e(t)$'s always cover the true values. To further quantify the accuracy of our $\mathcal{R}_e$ estimates, we define

$$\text{Bias} = \frac{1}{T+1} \sum_{t=0}^T |\hat{\mathcal{R}}_e(t) - \mathcal{R}_e(t)|$$

$$\text{RMSE} = \frac{1}{T+1} \sum_{t=0}^T \sqrt{\frac{1}{L} \sum_{\ell=1}^L [\mathcal{R}_e^{(\ell)}(t) - \mathcal{R}_e(t)]^2}$$

$$\text{Coverage} = \frac{1}{T+1} \sum_{t=0}^T \mathbf{1}[\mathcal{R}_e^{(0.025)}(t) \leq \mathcal{R}_e(t) \leq \mathcal{R}_e^{(0.975)}(t)]$$

here, RMSE stands for the (average) root mean square error, $\hat{\mathcal{R}}_e(t)$ denotes the estimated $\mathcal{R}_e$ (posterior median) at time t , $\{\mathcal{R}_e^{(\ell)}(t), \ell = 1, \dots, L\}$ represents the set of posterior samples of $\mathcal{R}_e(t)$, and $\mathcal{R}_e^{(0.025)}(t)$

Table 1

Summary of simulation results. For each scenario, columns 2–4 report the bias, root mean square error (RMSE) and coverage of the $\mathcal{R}_e$ estimates on one simulated dataset, and columns 5–7 show averages of these metrics over 100 repeat simulations with their standard deviations in subscripts.

Scn.	A single dataset			Average over 100 datasets		
	Bias	RMSE	Coverage	Bias	RMSE	Coverage
1	0.029	0.192	1.000	0.050 0.012	0.210 0.019	1.000 0.000
2	0.077	0.267	1.000	0.093 0.017	0.280 0.013	1.000 0.000
3	0.102	0.308	1.000	0.104 0.015	0.287 0.020	0.999 0.003

and $\mathcal{R}_e^{(0.975)}(t)$ represent the 2.5% and 97.5% posterior quantiles of $\mathcal{R}_e(t)$, respectively. Table 1 (columns 2–4) reports these accuracy measurements for the three simulated datasets. The small estimation errors and RMSEs indicate good recovery of the truth.

We also carry out sensitivity analyses to explore how the choice of




Fig. 3. Observed number of daily confirmed cases (solid red line) for six U.S. states: Washington, New York, California, Florida, Texas, and Illinois. The dashed vertical lines correspond to the start dates of statewide stay-at-home orders and state reopening plans. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the link function (Eq. (6)) and priors can affect the performance of the proposed method. Details of the sensitivity analyses are reported in Appendix C. In general, our method is robust to different specifications of the link function. The choice of the priors, on the other hand, may have an impact on the parameter estimates, because of parameter unidentifiability issues (see Appendix A).

Next, to evaluate the frequentist properties and reproducibility of the proposed method, under each simulation scenario, we repeatedly generate 100 datasets with the assumed sampling model using different random seeds. We apply our method to each simulated dataset and calculate the bias, RMSE and coverage of the $\mathcal{R}_e$ estimates. The results over repeated simulations are summarized in Table 1 (columns 5–7), which do not give rise to concerns regarding the reproducibility of the proposed method.

Lastly, we explore the computation time of our MCMC implementation for epidemic processes of different lengths. As expected, running time increases with the length of the epidemic processes. See Appendix C and Fig. C.2 for more details. For an epidemic process of 80 days, the MCMC implementation takes around 540 s (based on 50,000 MCMC iterations using an Intel Xeon E5-2650 v4 2.20 GHz processor).

6. Case studies

To illustrate the practical application of the proposed method, we carry out data analysis based on daily counts of confirmed COVID-19 cases reported by JHU CSSE [9]. This is the B_t in our model. We limit our analysis to six U.S. states (Washington, New York, California, Florida, Texas, and Illinois) to keep the paper in reasonable length. The reader can carry out independent analysis for other states, countries or regions using the R package BaySIR. The populations of these states are obtained from U.S. Census Bureau [43].

6.1. Estimation of the effective reproduction number

Fig. 3 shows the observed number of daily confirmed cases for the six states, and Fig. 4 shows the estimated $\mathcal{R}_e(t)$. The start dates of statewide stay-at-home orders and state reopening plans are also displayed in Fig. 4 for reference (data source: [35] and [47]). The estimated initial $\mathcal{R}_e$ ranges from 2.7 to 4.4. Specifically, $\mathcal{R}_e(0) = 2.8, 4.4, 2.7, 3.4, 2.9$ and 3.0 for Washington, New York, California, Florida, Texas and Illinois, respectively. During the early stage of the outbreak, the $\mathcal{R}_e$ generally has a decreasing trend. We suspect that the decline in

$\mathcal{R}_e$ may be associated with the implementation of mitigation policies (e.g., statewide stay-at-home orders, shown in Fig. 4) and the increase of public awareness. Starting from April, the $\mathcal{R}_e$ for these states is maintained around or below 1, indicating (partial) containment of the disease. However, with the gradual lift of stay-at-home orders and re-opening of businesses, we can clearly observe rebounds of $\mathcal{R}_e$ for some states (e.g., Florida) since May. For all the states, we can observe local fluctuations of $\mathcal{R}_e$ over time, which may potentially be attributed to some unobserved factors such as social distancing fatigue. Our analysis is preliminary and does not lead to definitive conclusions about whether a specific intervention is effective in controlling disease spread. Due to the issue of (potentially unmeasured) confounding, it is very challenging to draw causal inference about the effectiveness of an intervention. Nevertheless, our analysis can shed light on the transmission dynamics of COVID-19 and may be used as a reference for decision-makers.

6.2. Test of fit

We carry out the Bayesian χ^2 test [25] to assess the goodness-of-fit of our model using Illinois data as an example. First, we choose quantiles $0 \equiv a_0 < a_1 < \dots < a_{G-1} < a_G \equiv 1$, with $p_g = a_g - a_{g-1}$, $g = 1, \dots, G$. As suggested by Johnson [25], we use $(a_0, \dots, a_5) = (0, 0.2, 0.4, 0.6, 0.8, 1)$, so $p_g = 0.2$ and $G = 5$. Next, let $\theta^{(t)}$ be a posterior sample of the model parameters θ , and let $m_g(\theta^{(t)})$ denote the number of observations (i.e., B_t 's) such that $\text{logit}\{B_t / [(1 - \alpha^{(t)})I_t^{U(t)}]\}$ falls between the a_{g-1} and a_g quantiles of the distribution $N(y_t^\top \eta^{(t)}, \sigma_y^{(t)})$. Let

$$\omega(\theta^{(t)}) = \sum_{g=1}^G \left[\frac{m_g(\theta^{(t)}) - (T+1)p_g}{\sqrt{(T+1)p_g}} \right]^2$$

Then, under the null hypothesis of a good model fit, the statistic ω should follow a χ^2 -distribution with $G - 1 = 4$ degrees of freedom. A quantile-quantile plot of the posterior samples of ω against the expected order statistics from a χ^2_4 distribution (Appendix Fig. D.1) shows that ω plausibly comes from a χ^2_4 distribution. In addition, we find the proportion of posterior samples of ω exceeding the 95% quantile of a χ^2_4 distribution to be 0.043. There is no evidence of a lack of fit.




Fig. 4. Estimated time-varying effective reproduction numbers (solid black line) for six U.S. states: Washington, New York, California, Florida, Texas, and Illinois. The start date in each graph is the date when the 100th case is confirmed in the state. The grey band represents the 95% posterior credible interval. The dashed vertical lines correspond to the start dates of statewide stay-at-home orders and state reopening plans. The dashed horizontal line represents an $\mathcal{R}_e$ of 1.




Fig. 5. Within-sample forecasts for Illinois using 20-day, 40-day or 60-day training data. The upper panel shows the observed daily confirmed cases (solid red line), and posterior medians (dashed line) and 95% credible intervals (grey band) for $(\mathbf{B}^{\text{te}} | \mathbf{B}^{\text{tr}}, I_0^D)$. The upper bounds of the credible intervals are truncated for better display. The lower panel shows the posterior medians (solid red line) and 95% credible intervals (red band) for $[\mathcal{R}_e(0), \dots, \mathcal{R}_e(t^*) | \mathbf{B}^{\text{tr}}, I_0^D]$, and posterior medians (dashed line), posterior draws (thin grey lines) and 95% credible intervals (blue band) for $[\mathcal{R}_e(t^* + 1), \dots, \mathcal{R}_e(T) | \mathbf{B}^{\text{tr}}, I_0^D]$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Out-of-sample forecasts for Illinois in the next 30 days. (a) Observed daily confirmed cases (solid red line), and posterior medians (dashed line) and 95% credible intervals (grey band) for $(\mathbf{B}^* | \mathbf{B}, I_0^D)$. (b) Posterior medians (solid red line) and 95% credible intervals (red band) for $[\mathcal{R}_e(0), \dots, \mathcal{R}_e(T) | \mathbf{B}, I_0^D]$, and posterior medians (dashed line), posterior draws (thin grey lines) and 95% credible intervals (blue band) for $[\mathcal{R}_e(T + 1), \dots, \mathcal{R}_e(T + T^*) | \mathbf{B}^{\text{tr}}, I_0^D]$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

6.3. Forecasts

6.3.1. Retrospective forecasts

As described in Section 4.2, the proposed method can be used to predict a future observation based on its posterior predictive distribution. To evaluate the forecasting performance of the proposed model, we conduct within-sample forecasts using Illinois as an example. Specifically, we split the observations $\mathbf{B}$ into a training set $\mathbf{B}^{\text{tr}}$ and a testing set $\mathbf{B}^{\text{te}}$, where $\mathbf{B}^{\text{tr}} = (B_0, B_1, \dots, B_{t^*})$ and $\mathbf{B}^{\text{te}} = (B_{t^*+1}, B_{t^*+2}, \dots, B_T)$. We consider three different scenarios, $t^* \in \{19, 39, 59\}$, so that the training set consists of observations for 20, 40 and 60 days, respectively. We first sample from the posterior distribution of the parameters evaluated on the training set, $\pi(\theta | \mathbf{B}^{\text{tr}}, I_0^D)$, and then sample from the posterior predictive distribution of the testing observations, $\pi(\mathbf{B}^{\text{te}} | \mathbf{B}^{\text{tr}}, I_0^D)$.

Fig. 5 shows the forecasting results for the three scenarios. The projection produces a wide range of uncertainty; the upper bound of the 95% credible intervals (i.e., The 97.5% percentile of $\pi(\mathbf{B}^{\text{te}} | \mathbf{B}^{\text{tr}}, I_0^D)$) can reach 350,000 which is much higher than the actual number of daily confirmed cases. Therefore, the upper bound is truncated in the figure for better display. To better understand the forecasting behavior of the proposed model, the predictions of future $\mathcal{R}_e(t)$'s are also displayed. Using 20-day training data, the median of $\pi(\mathbf{B}^{\text{te}} | \mathbf{B}^{\text{tr}}, I_0^D)$ underestimates the actual observations, although the 95% credible interval covers the observed values for a long period of time. In general, prediction of an

epidemic process is challenging, especially when the epidemiological parameters vary over time [23]. To see this, notice that there is a rebound of $\mathcal{R}_e(t)$ around April 21, which cannot be captured by the GP prediction rule with 20-day training data. Since the stay-at-home order is still in effect on April 21, this rebound cannot be captured by policy-related covariates either. To summarize, future predictions are made based on extrapolation of the current trend, and if the trend changes unexpectedly, the predictions will be inaccurate.

With more training data, the prediction accuracy improves, as seen in Fig. 5(b, c). Using 60-day training data, the median of $\pi(\mathbf{B}^{\text{te}} | \mathbf{B}^{\text{tr}}, I_0^D)$ matches well with the actual observations in the subsequent month, although the prediction still fails to capture the rebound starting from June 15. Lastly, the short-term predictions (within, say, the next 10 days) are reasonably accurate in all the scenarios.

6.3.2. Prospective forecasts

To make future predictions, we first sample from $\pi(\theta | \mathbf{B}, I_0^D)$ and then sample from $\pi(\mathbf{B}^* | \mathbf{B}, I_0^D)$; recall that $\mathbf{B}^* = (B_{T+1}, \dots, B_{T+T^*})$. Fig. 6 shows the projected daily confirmed cases and $\mathcal{R}_e(t)$'s for Illinois in the next 30 days (i.e., $T^* = 30$). The projections are based on the assumption that the decreasing trend of $\mathcal{R}_e(t)$ continues. With the lift of the stay-at-home order and the reopening of businesses, it is possible that $\mathcal{R}_e(t)$ will rebound, thus caution is needed in interpreting the forecasting results.

7. Discussion

We developed a Bayesian approach to statistical inference about the transmission dynamics of COVID-19. We proposed to estimate the disease transmission rate using GPR, which captures nonlinear and non-monotonic trends without the need of specific parametric assumptions. A PTMCMC algorithm was used to efficiently sample from the posterior distribution of the epidemiological parameters. Case studies based on the proposed method revealed the overall decreasing trend of $\mathcal{R}_e$ in six U.S. states (Washington, New York, California, Florida, Texas, and Illinois), which may be associated with the implementation of mitigation policies and the increasing public awareness of the disease. Projections for future case counts can be made based on extrapolation, although caution is needed in interpreting the forecasting results.

Extensions of the proposed compartmental model can be made in a number of ways. First, while recovery from COVID-19 is assumed to confer immunity to reinfection, some recent evidence may suggest that such reinfection is possible [26]. Therefore, it may be desirable to allow recovered individuals to become susceptible again with a small probability. Second, as described in Section 3.1, it is possible to further split the UI and DI compartments and to incorporate an exposed compartment. Third, we may also split the removed compartment into recovered and deceased compartments. See, for example, Giordano et al. [16], Zhang et al. [53] and Aguilar et al. [1]. Considering more compartments and transitions will make the model more realistic. However, by adding complexity to the current parsimonious model, sampling, estimation and model unidentifiability problems are likely exacerbated [7,37]. A possible way out could be to utilize more observable information, such as numbers of recoveries and hospitalizations. Nevertheless, not every country/region has these data (or accurate measurements of these data) available, and we chose to model only daily confirmed cases to keep our method general enough and applicable to most countries/regions.

The proposed model as in Eq. (4) is a state-space model motivated by a deterministic SIR model. A future direction is to consider a stochastic epidemic model. For example, a model similar to Lekone and Finkenstädt [30] may be used,

$$\begin{aligned} S_t &= S_{t-1} - A_{t-1}, & I_t^U &= I_{t-1}^U + A_{t-1} - B_{t-1} - C_{t-1} \\ I_t^D &= I_{t-1}^D + B_{t-1} - D_{t-1}, & R_t &= R_{t-1} + C_{t-1} + D_{t-1} \end{aligned}$$

where

$$\begin{aligned} A_t &\sim \text{Bin}[S_t, 1 - e^{-\beta_t(I_t^U + I_t^D)/N}], & C_t &\sim \text{Bin}(I_t^U, 1 - e^{-\alpha}) \\ B_t | C_t &\sim \text{Bin}(I_t^U - C_t, 1 - e^{-\eta}), & D_t &\sim \text{Bin}(I_t^D, 1 - e^{-\alpha}) \end{aligned}$$

Compared to Eq. (4), this model may better reflect the stochastic nature of the epidemic process. The cost is increased computational complexity.

In our models for the diagnosis rate (Eq. (5)) and transmission rate (Eq. (7)), we allow incorporation of covariates. Currently, we only considered an intercept term and a time trend, because empirically we found it hard to identify the effects of other covariates. While the number of tests is likely to be related to the diagnosis rate, the US COVID-19 testing data are quite noisy. Within the same US state, it is common to see large swings of reported test numbers and test positive rates within adjacent or few days. Take New Jersey as an example: for adjacent dates, the test number swings from $< 3,000$ to $> 50,000$, with test positive rate varying from > 0.95 to < 0.05 , respectively. Furthermore, the number of UI individuals and testing policies (e.g., who is

eligible for testing) play an important role in the relationship between the number of tests and diagnosis rate, but these factors are either unobservable or very hard to quantify. In addition, the implementation and lift of government interventions are relevant time-varying covariates for the transmission rate. However, the wide variety of interventions (e.g., travel restriction, school closure, mask wearing, etc.) and the spatio-temporal heterogeneity of their implementations (e.g., within the same state, some cities are reopened earlier than the others, and different businesses are reopened at different stages) make it hard to quantify their effects. Moreover, the degree of public compliance with government interventions is arguably a more important covariate, but it is unmeasurable. As a result, due to Ockham's razor [24], we resort to the simpler model and did not include many potential covariates in our model. More efficient ways to incorporate covariates, potentially based on model selection or variable selection techniques, are worth further investigation.

The incorporation of other covariates will also have an impact on the predictive inference. On one hand, additional covariates will make case predictions more challenging. For example, prediction of future case numbers may be modeled by first predicting future covariates (e.g., future number of tests and future government interventions), and then by using a semiparametric regression incorporating the predicted covariate values. This will lead to complex inference and challenging computation, which must be addressed in the methodology development. On the other hand, additional covariates may improve the accuracy of future predictions. For example, with only an intercept term and a time trend, the GP prediction rule is not able to capture possible case rebounds. However, if government interventions are included as covariates, and if it can be shown that the lift of interventions is associated with the increase in $\mathcal{R}_e$, then by assuming future relaxation of government interventions, the model is possible to predict case rebounds.

Our data analysis was carried out separately for each country/region. A nature extension is to model multiple countries/regions jointly using a hierarchical model to achieve borrowing of information, which usually leads to improvements in parameter estimations. We assumed that the population in each country/region is closed, ignoring immigration and emigration. Arguably, a more realistic model should take into account spatial spread of the disease, as seen in Li et al. [32]. Again, the main drawbacks to these extensions would be increased computation time.

As discussed in Appendix A, the parameters in model (4) are unidentifiable with only daily confirmed cases (B_t) observed, thus parameter estimates are sensitive to prior choices and modeling assumptions. Many existing models for COVID-19 share the same situation, which could partially explain why different studies may lead to substantially different estimates. For example, some consider the infectious period as the time from infection to recovery or death, which is around 20–30 days [44]. Under this definition, the estimated effective reproduction numbers would be higher (e.g., [1]). Therefore, when interpreting the results, it is important to recognize their reliance on underlying assumptions.

Lastly, since the proposed model (4) is a state-space model, it is of interest to further explore online and sequential algorithms for posterior sampling, such as sequential Monte Carlo [10,11]. In that way, when data at more time points become available, one can update the posterior in an efficient way rather than re-fitting the model to the complete data.

Appendix A. Parameter identifiability

With only daily confirmed cases observed, the parameters in model (4) are unidentifiable. To see this, consider the following two epidemic processes (indexed by $j = 1$ and 2),

$$\begin{aligned}
S_{j,t} &= S_{j,t-1} - \beta_{j,t-1} S_{j,t-1} (I_{j,t-1}^U + I_{j,t-1}^D) / N \\
I_{j,t}^U &= (1 - \alpha_j) I_{j,t-1}^U + \beta_{j,t-1} S_{j,t-1} (I_{j,t-1}^U + I_{j,t-1}^D) / N - \gamma_{j,t-1} (1 - \alpha_j) I_{j,t-1}^U \\
I_{j,t}^D &= (1 - \alpha_j) I_{j,t-1}^D + \gamma_{j,t-1} (1 - \alpha_j) I_{j,t-1}^U \\
R_{j,t} &= R_{j,t-1} + \alpha_j (I_{j,t-1}^U + I_{j,t-1}^D)
\end{aligned}$$

for $t = 1, \dots, T$. The observation is the daily increment in confirmed cases, $B_{j,t} = \gamma_{j,t} (1 - \alpha_j) I_{j,t}^U$. These two processes give rise to identical observations $B_{1,t}$ and $B_{2,t}$ for all t , if

$$\gamma_{1,0} (1 - \alpha_1) I_{1,0}^U = \gamma_{2,0} (1 - \alpha_2) I_{2,0}^U \quad (9)$$

and

$$\begin{aligned}
&\gamma_{1,t} (1 - \alpha_1) [(1 - \alpha_1) (1 - \gamma_{1,t-1}) I_{1,t-1}^U + \beta_{1,t-1} S_{1,t-1} (I_{1,t-1}^U + I_{1,t-1}^D) / N] = \\
&\gamma_{2,t} (1 - \alpha_2) [(1 - \alpha_2) (1 - \gamma_{2,t-1}) I_{2,t-1}^U + \beta_{2,t-1} S_{2,t-1} (I_{2,t-1}^U + I_{2,t-1}^D) / N]
\end{aligned} \quad (10)$$

for $t = 1, \dots, T$. In other words, different sets of parameters can lead to exactly the same observed data. Even if we restrict that $(S_{1,0}, I_{1,0}^U, I_{1,0}^D, R_{1,0}) = (S_{2,0}, I_{2,0}^U, I_{2,0}^D, R_{2,0})$ (same initial conditions), $\gamma_{1,t} = \gamma_{2,t}$, and $\gamma_{2,t} = \gamma_{2,0}$ (constant diagnosis rate), for any $\alpha_1 \neq \alpha_2$ we can still solve Eqs. (9) and (10) and get distinct $\{\gamma_{1,t}, \beta_{1,t}\}$ and $\{\gamma_{2,t}, \beta_{2,t}\}$ that lead to the same observed data.

A specific example is given below. Consider a population size of $N = 20,000,000$. Suppose there are two epidemic processes with the same initial conditions, $I_{1,0}^U = I_{2,0}^U = 800$, $I_{1,0}^D = I_{2,0}^D = 100$, $R_{1,0} = R_{2,0} = 0$, and $S_{1,0} = S_{2,0} = N - 900$. Suppose further $\alpha_1 = 0.3$, $\alpha_2 = 0.05$, $\gamma_{1,t} = \gamma_{2,t} = \gamma_{2,0} = 0.2$, and $\gamma_{2,t} = \gamma_{2,0} \cdot (1 - \alpha_1) / (1 - \alpha_2) = 0.147$. Then, the parameters $\beta_{1,t}$ and $\beta_{2,t}$ can be chosen (Fig. A.1(a)) such that $\{B_{1,t}\}$ and $\{B_{2,t}\}$ are identical (Fig. A.1(b)). The resulting effective reproduction numbers for the two epidemic processes, $\mathcal{R}_e^j(t) = (\beta_{j,t} S_{j,t}) / (\alpha_j N)$, are shown in Fig. A.1(c) and are quite different. This example highlights that the parameters in (4) are unidentifiable in the absence of strong prior information.




Fig. A.1. An example of two epidemic processes giving rise to identical observations. Panel (a) shows the distinct transmission rates for the two processes. Panel (b) shows the identical observations given by the two processes. Panel (c) shows the distinct effective reproduction numbers for the two processes.

Appendix B. Posterior sampling: parallel tempering

To demonstrate the advantage of the PT scheme, we show in Fig. B.1 the Markov chains for I_0^U and η generated using or not using PT based on a simulated dataset. We evaluate the convergence of the chains using Geweke's diagnostic [13]. Under the null hypothesis of chain convergence, Geweke's z-score should follow a standard normal distribution. The z-score indicates lack of convergence for the chains generated without PT.




Fig. B.1. Markov chains for I_0^U and η using (a, c) or not using (b, d) parallel tempering. The posterior correlation of I_0^U and η is -0.82 . The value z_G refers to Geweke's z-score for convergence diagnostic. All chains are based on 50,000 iterations (discarding first 20,000 iterations as burn-in and keeping 1 draw every 30

iterations).

Appendix C. Simulation studies: sensitivity analysis and computation time

We carry out sensitivity analyses to explore how the choice of the link function (Eq. (6)) and priors can affect the performance of the proposed method. We consider the following four settings:

- (Set. 1) Replacing the default logit link for γ_t by the probit link;
- (Set. 2) Replacing the default logit link for γ_t by the complementary log-log link;
- (Set. 3) Replacing the default prior on α^{-1} by $\alpha^{-1} \sim \text{Ga}(46.5, 5)1(\alpha^{-1} \geq 1)$. This leads to a larger prior variance for α^{-1} compared to the default. Recall that the default prior is $\alpha^{-1} \sim \text{Ga}(325.5, 35)1(\alpha^{-1} \geq 1)$;
- (Set. 4) Replacing the default prior on α^{-1} by $\alpha^{-1} \sim \text{Ga}(700, 35)1(\alpha^{-1} \geq 1)$. This leads to a different prior mean for α^{-1} compared to the default.

We fit our model to the Simulation Scenario 1 dataset. Fig. C.1 shows the estimated time-varying effective reproduction numbers under the four settings. The estimates are robust to the choice of the link function (Fig. C.1(a, b)). Also, increasing the prior variance for α^{-1} does not lead to much change in the estimates (Fig. C.1(c)). Lastly, altering the prior mean for α^{-1} can lead to substantially different estimates (Fig. C.1(d)). This is due to parameter unidentifiability issues (Appendix A). Multiple solutions may explain the observed data equally well, thus the solutions that are more consistent with the prior would be preferred. Under Setting 4, the prior for α^{-1} is centered around 20, while the true $\alpha^{-1} = 9.3$. As a result, the parameter estimates deviate from the simulation truth.




Fig. C.1. Simulation Scenario 1. Estimated time-varying effective reproduction numbers (solid black line) with different link functions and priors for α^{-1} . The grey band represents the 95% posterior credible interval, and the dashed red line shows the simulation truth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

We also explore the computation time of our MCMC implementation for epidemic processes of different lengths. Fig. C.2 shows the computation time for epidemic processes of lengths from 10 days to 100 days. The computation is based on 50,000 MCMC iterations using an Intel Xeon E5-2650 v4 2.20 GHz processor.




Fig. C.2. Computation time (in seconds) of the MCMC implementation for epidemic processes of different lengths.

Appendix D. Case studies: test of fit

We carry out the Bayesian χ^2 test [25] to assess the goodness-of-fit of our model using Illinois data as an example. Under the null hypothesis of a good model fit, the statistic ω should follow a χ_4^2 distribution. Fig. D.1 shows a quantile-quantile plot of posterior samples of ω against expected order statistics from a χ_4^2 distribution. There is no evidence that ω deviates from a χ_4^2 distribution.




Fig. D.1. Quantile-quantile plot of posterior samples of the test statistic ω against expected order statistics from a χ_4^2 distribution for the Bayesian χ^2 test.

References

- [1] J.B. Aguilar, J.S. Faust, L.M. Westafer, J.B. Gutierrez, Investigating the impact of asymptomatic carriers on COVID-19 transmission, medRxiv (2020).
- [2] L.J. Allen, An introduction to stochastic epidemic models, in: F. Brauer, P. van den Driessche, J. Wu (Eds.), *Mathematical Epidemiology*, Springer, Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 81–130.
- [3] R.M. Anderson, R.M. May, *Infectious Diseases of Humans: Dynamics and Control*, Oxford University Press, Oxford, 1991.
- [4] H. Andersson, T. Britton, *Stochastic Epidemic Models and their Statistical Analysis*, Volume 151 of Lecture Notes in Statistics, Springer Science & Business Media, New York, 2000.
- [5] N.G. Becker, T. Britton, Statistical studies of infectious disease incidence, *J. R. Stat. Soc. Ser. B Stat Methodol.* 61 (2) (1999) 287–307.
- [6] F. Brauer, *Compartmental Models in Epidemiology*, Springer, Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 19–79.
- [7] A. Capaldi, S. Behrend, B. Berman, J. Smith, J. Wright, A.L. Lloyd, Parameter estimation and uncertainty quantification for an epidemic model, *Math. Biosci. Eng.* 9 (3) (2012) 553–576.
- [8] X. Chen, Z. Qiu, Scenario analysis of non-pharmaceutical interventions on global COVID-19 transmissions, arXiv Preprint (2020) arXiv:2004.04529.
- [9] E. Dong, H. Du, L. Gardner, An interactive web-based dashboard to track COVID-19 in real time, *Lancet Infect. Dis.* 20 (5) (2020) 533–534.
- [10] A. Doucet, N. De Freitas, N. Gordon, An introduction to sequential Monte Carlo methods, in: A. Doucet, N. De Freitas, N. Gordon (Eds.), *Sequential Monte Carlo Methods in Practice*, Springer Science & Business Media, New York, 2001, pp. 3–14.
- [11] V. Dukic, H.F. Lopes, N.G. Polson, Tracking epidemics with Google flu trends data and a state-space SEIR model, *J. Am. Stat. Assoc.* 107 (500) (2012) 1410–1426.
- [12] S. Flaxman, S. Mishra, A. Gandy, H.J.T. Unwin, T.A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J.W. Eaton, et al., Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe, *Nature* 584 (2020) 257–261.
- [13] J. Geweke, *Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments*, 196 Federal Reserve Bank of Minneapolis, Research Department, Minneapolis, MN, USA, 1991.
- [14] C.J. Geyer, Markov chain Monte Carlo maximum likelihood, *Computing Science and Statistics, Proceedings of the 23rd Symposium on the Interface*, Interface Foundation of North America, Fairfax Station, VA, 1991, pp. 156–163.
- [15] G.J. Gibson, E. Renshaw, Estimating parameters in stochastic compartmental models using Markov chain methods, *Mathe. Med. Biol. A J. IMA* 15 (1) (1998) 19–40.
- [16] G. Giordano, F. Blanchini, R. Bruno, P. Colaneri, A. Di Filippo, A. Di Matteo, M. Colaneri, Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy, *Nat. Med.* 26 (2020) 855–860.
- [17] J. Gomez, J. Prieto, E. Leon, A. Rodriguez, INFEKTA: a general agent-based model for transmission of infectious diseases: studying the COVID-19 propagation in Bogotá-Colombia, medRxiv (2020).
- [18] J. Gu, H. Yan, Y. Huang, Y. Zhu, H. Sun, X. Zhang, Y. Wang, Y. Qiu, S. Chen, Better strategies for containing COVID-19 epidemics—a study of 25 countries via an extended varying coefficient SEIR model, medRxiv (2020).
- [19] X. He, E.H. Lau, P. Wu, X. Deng, J. Wang, X. Hao, Y.C. Lau, J.Y. Wong, Y. Guan, X. Tan, et al., Temporal dynamics in viral shedding and transmissibility of COVID-19, *Nat. Med.* 26 (5) (2020) 672–675.
- [20] H.W. Hethcote, The mathematics of infectious diseases, *SIAM Rev.* 42 (4) (2000) 599–653.
- [21] N. Hoertel, M. Blachier, C. Blanco, M. Olsson, M. Massetti, M.S. Rico, F. Limosin, H. Leleu, A stochastic agent-based model of the SARS-CoV-2 epidemic in France, *Nat. Med.* 26 (2020) (2020) 1417–1421 Forthcoming.
- [22] IHME COVID-19 Health Service Utilization Forecasting Team, C.J. Murray, Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US State in the next 4 months, medRxiv (2020).
- [23] J.P. Ioannidis, S. Cripps, M.A. Tanner, Forecasting for COVID-19 has failed, *Int. J. Forecast.* (2020) Forthcoming.
- [24] W.H. Jefferys, J.O. Berger, Ockham's razor and Bayesian analysis, *Am. Sci.* 80 (1) (1992) 64–72.
- [25] V.E. Johnson, A Bayesian χ^2 test for goodness-of-fit, *Ann. Stat.* 32 (6) (2004) 2361–2384.
- [26] A. Joseph, Scientists are reporting several cases of COVID-19 reinfection—but the implications are complicated, *Stat News* (2020) URL: <https://www.statnews.com/2020/08/28/covid-19-reinfection-implications/> (accessed August 30, 2020).
- [27] W.O. Kermack, A.G. McKendrick, A contribution to the mathematical theory of epidemics, *Proc. R. Soc. Lond. Ser. A* 115 (772) (1927) 700–721.
- [28] R.D. Kirkcaldy, B.A. King, J.T. Brooks, COVID-19 and postinfection immunity: limited evidence, many remaining questions, *JAMA* 323 (22) (2020) 2245–2246.
- [29] T. Kypraios, P.D. O'Neill, Bayesian nonparametrics for stochastic epidemic models, *Stat. Sci.* 33 (1) (2018) 44–56.
- [30] P.E. Lekone, B.F. Finkenstädt, Statistical inference in a stochastic epidemic SEIR model with control intervention: ebola as a case study, *Biometrics* 62 (4) (2006) 1170–1177.
- [31] Q. Li, X. Guan, P. Wu, X. Wang, L. Zhou, Y. Tong, R. Ren, K.S. Leung, E.H. Lau, J.Y. Wong, et al., Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia, *N. Engl. J. Med.* 382 (2020) 1199–1207.
- [32] R. Li, S. Pei, B. Chen, Y. Song, T. Zhang, W. Yang, J. Shaman, Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2), *Science* 368 (6490) (2020) 489–493.
- [33] J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer Science & Business Media, 2008.
- [34] Q.-X. Long, B.-Z. Liu, H.-J. Deng, G.-C. Wu, K. Deng, Y.-K. Chen, P. Liao, J.-F. Qiu, Y. Lin, X.-F. Cai, et al., Antibody responses to SARS-CoV-2 in patients with COVID-19, *Nat. Med.* 26 (2020) 845–848.
- [35] S. Mervosh, D. Lu, V. Swales, See which states and cities have told residents to stay at home, *New York Times* (2020) URL: <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html> (accessed June 30, 2020).
- [36] D. Osthus, J. Gattiker, R. Priedhorsky, S.Y. Del Valle, Dynamic Bayesian influenza forecasting in the United States with hierarchical discrepancy (with discussion), *Bayesian Anal.* 14 (1) (2019) 261–312.
- [37] D. Osthus, K.S. Hickmann, P.C. Caragea, D. Higdon, S.Y. Del Valle, Forecasting seasonal influenza with a state-space SIR model, *Ann. Appl. Stat.* 11 (1) (2017) 202–224.
- [38] P.D. O'Neill, G.O. Roberts, Bayesian inference for partially observed stochastic epidemics, *J. R. Stat. Soc. Ser. A Stat. Soc.* 162 (1) (1999) 121–129.
- [39] C.E. Rasmussen, C.K. Williams, *Gaussian Process for Machine Learning*, MIT Press, 2006.
- [40] S. Roberts, M. Osborne, M. Ebdon, S. Reece, N. Gibson, S. Aigrain, Gaussian

- processes for time-series modelling, *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* 371 (2013) 20110550.
- [41] C. Rothe, M. Schunk, P. Sothmann, G. Bretzel, G. Froeschl, C. Wallrauch, T. Zimmer, V. Thiel, C. Janke, W. Guggemos, et al., Transmission of 2019-nCoV infection from an asymptomatic contact in Germany, *N. Engl. J. Med.* 382 (10) (2020) 970–971.
 - [42] H. Sun, Y. Qiu, H. Yan, Y. Huang, Y. Zhu, J. Gu, S.X. Chen, Tracking reproductivity of COVID-19 epidemic in China with varying coefficient SIR model, *J. Data Sci.* 18 (3) (2020) 455–472.
 - [43] U.S. Census Bureau, Annual Estimates of the Resident Population for the United States, Regions, States, and Puerto Rico, April 1, 2010 to July 1, 2019. URL: <https://www2.census.gov/programs-surveys/popest/tables/2010-2019/state/totals/nst-est2019-01.xlsx>.
 - [44] R. Verity, L.C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G. Cuomo-Dannenburg, H. Thompson, P.G. Walker, H. Fu, et al., Estimates of the severity of coronavirus disease 2019: a model-based analysis, *Lancet Infect. Dis.* 20 (6) (2020) 669–677.
 - [45] L. Wang, G. Wang, L. Gao, X. Li, S. Yu, M. Kim, Y. Wang, Z. Gu, Spatiotemporal dynamics, nowcasting and forecasting of COVID-19 in the United States, *arXiv Preprint* (2020) arXiv:2004.14103.
 - [46] L. Wang, Y. Zhou, J. He, B. Zhu, F. Wang, L. Tang, M.C. Eisenberg, P.X. Song, An epidemiological forecast model and software assessing interventions on COVID-19 epidemic in China, *J. Data Sci.* 18 (3) (2020) 409–432.
 - [47] Washington Post Staff, Where states are reopening after the U.S. shutdown, *Wash. Post* (2020) URL: <https://www.washingtonpost.com/graphics/2020/national/states-reopening-coronavirus-map/> (accessed June 30, 2020).
 - [48] H. Weiss, The SIR model and the foundations of public health, *Mater. Matemàtica* 2013 (3) (2013) 1–17.
 - [49] R. Wölfel, V.M. Corman, W. Guggemos, M. Seilmaier, S. Zange, M.A. Müller, D. Niemeyer, T.C. Jones, P. Vollmar, C. Rothe, et al., Virological assessment of hospitalized patients with COVID-2019, *Nature* 581 (7809) (2020) 465–469.
 - [50] S. Woody, M.G. Tec, M. Dahan, K. Gaither, M. Lachmann, S. Fox, L.A. Meyers, J.G. Scott, Projections for first-wave COVID-19 deaths across the US using social-distancing measures derived from mobile phones, *medRxiv* (2020).
 - [51] J.T. Wu, K. Leung, G.M. Leung, Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study, *Lancet* 395 (10225) (2020) 689–697.
 - [52] X. Xu, T. Kypraios, P.D. O'Neill, Bayesian non-parametric inference for stochastic epidemic models using Gaussian processes, *Biostatistics* 17 (4) (2016) 619–633.
 - [53] Y. Zhang, C. You, Z. Cai, J. Sun, W. Hu, X.-H. Zhou, Prediction of the COVID-19 outbreak based on a realistic stochastic model, *medRxiv* (2020).