

Early Detection of COVID-19 Hotspots Using Spatio-Temporal Data

Shixiang Zhu , Alexander Bukharin , Liyan Xie , *Member, IEEE*, Khurram Yamin , Shihao Yang, Pinar Keskinocak, and Yao Xie , *Member, IEEE*

Abstract—Recently, the Centers for Disease Control and Prevention (CDC) has worked with other federal agencies to identify counties with increasing coronavirus disease 2019 (COVID-19) incidence (hotspots) and offers support to local health departments to limit the spread of the disease. Understanding the spatio-temporal dynamics of hotspot events is of great importance to support policy decisions and prevent large-scale outbreaks. This paper presents a spatio-temporal Bayesian framework for early detection of COVID-19 hotspots (at the county level) in the United States. We assume both the observed number of cases and hotspots depend on a class of latent random variables, which encode the underlying spatio-temporal dynamics of the transmission of COVID-19. Such latent variables follow a zero-mean Gaussian process, whose covariance is specified by a non-stationary kernel function. The most salient feature of our kernel function is that deep neural networks are introduced to enhance the model's representative power while still enjoying the interpretability of the kernel. We derive a sparse model and fit the model using a variational learning strategy to circumvent the computational intractability for large data sets. Our model demonstrates better interpretability and superior hotspot-detection performance compared to other baseline methods.

Index Terms—COVID-19 hotspots, Gaussian processes, non-stationary kernel, spatio-temporal model.

I. INTRODUCTION

THE ongoing global pandemic caused by the coronavirus disease (COVID-19) has spread rapidly over more than 200 countries in the world since its emergence in 2019. Even the largest economies' resources have been strained due to the spread of COVID-19. Predicting potential hotspots ahead of time can play a significant role in deploying targeted interventions, such as testing, tracing, and isolation, and slow down the disease spread [1].

Large-scale, population-based testing can indicate regional hotspots, but at the cost of a delay between testing and actionable results. Accurately identifying changes in the infection rate




Fig. 1. Examples of COVID-19 hotspots identified by CDC [1]. In general, these hotspots define the onset of a local outbreak of COVID-19.

requires sufficient testing coverage of a given population, which can be costly and requires substantial testing capacity. Regional variation in testing access can also hamper the ability of public health organizations to detect rapid changes in infection rates. Recent studies [2] aimed at estimating the spread of COVID-19 by forecasting the number of confirmed cases or the number of deaths. However, these methods failed to provide a satisfactory case prediction accuracy. Therefore, there is a high unmet need for tools and methods that can facilitate the timely and accurate identification of infection hotspots and enable policymakers to act effectively with minimal delay [3].

The Centers for Disease Control and Prevention (CDC) with other federal agencies have identify counties with a significant increase in COVID-19 incidence (hotspots) [1], which offers a unique opportunity to investigate the spatio-temporal dynamics between the identified hotspots. Fig. 1 gives some real examples of the hotspots at four different counties in the state of Georgia. The identified hotspots indicate the relative temporal increases in confirmed cases and mark the onset of local outbreaks.

In this paper, we propose an effective COVID-19 hotspot detection framework that utilizes the hotspot data and multiple other data sources, including community mobility, to enhance hotspot detection accuracy. We assume the hotspot and number of cases in the same location depending on common priori factors, represented by a latent spatio-temporal random variable. This latent variable is modeled by a Gaussian process, whose covariance is characterized by an interpretable non-stationary kernel. We note that the non-stationarity of our kernel plays a pivotal role in the success of our model because the spread of the

Manuscript received May 30, 2021; revised October 31, 2021; accepted November 2, 2021. Date of publication February 28, 2022; date of current version April 18, 2022. The work of Shixiang Zhu, Liyan Xie, and Yao Xie was supported in part by NSF under Grant DMS-2134037, Grant CCF-1650913, Grant CMMI-2015787, Grant DMS-1938106, and Grant DMS-1830210. The guest editor coordinating the review of this manuscript and approving it for publication was Tuomas Virtanen. (Corresponding author: Yao Xie.)

The authors are with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: shixiang.zhu@gatech.edu; abukharin3@gatech.edu; lxie49@gatech.edu; khurram.yamin@gatech.edu; shihao.yang@isye.gatech.edu; pinar@isye.gatech.edu; yao.xie@isye.gatech.edu).

Digital Object Identifier 10.1109/JSTSP.2022.3154972

virus shows heterogeneous spatial correlation across different regions. For example, the virus is likely to spread more slowly in a sparsely populated area such as rural Nebraska compared to a densely populated area such as New York City. We formulate our kernel function using carefully crafted feature functions incorporating neural networks, which provide greater flexibility in capturing the complex dynamics of the spread of COVID-19 while still being highly interpretable. To tackle the computational challenge of the Gaussian process with a large-scale data set, we also derive a sparse model and fit the model efficiently via a variational learning strategy. The remainder of the paper is organized as follows. We first discuss the literature relevant to COVID-19 hotspot detection and other related work. We describe the data sets in Section II. We introduce the proposed hotspot detection framework in Section III. We present an efficient computation strategy and the learning algorithm for our detection framework in Section IV. Finally, we present the interpretation of our model and the numerical results on COVID-19 data in Section V.

Related work: Hotspot detection is closely related to the prediction of confirmed cases. Therefore, we first review some prediction methods for completeness. Compartmental models are mathematical modeling of infectious diseases and have been widely used in epidemiology. In simple SIR models, [4], the population is assigned to compartments with labels S (susceptible), I (infectious), and R (recovered), respectively. The transition rates between compartments are typically modeled using differential equations. Extensions and variants of SIR models include the SIRD model [5], [6] which considers deceased individuals, and the SEIR model [7]–[9] which considers exposed periods, to name a few. Compartmental models work well when applied to large regions/populations, such as a state or a country because they assume a fixed/closed population. However, populations between geographic areas, such as counties, may interact with each other in the desired high-resolution modeling. Therefore, we use a spatio-temporal model that is more flexible and can capture the spread between different counties.

Much work has been done on predicting the number of COVID-19 cases and deaths at the national level or state level, without considering the spatial correlation across smaller regions [10]–[16]. Machine learning-based approaches have also been considered in [17]. Some work [18] attempts to use neural networks to model the accumulative number of confirmed cases. Recurrent neural network-based methods [19], [20] have been applied to model the temporal dynamics of the COVID-19 outbreak. Moreover, online COVID-19 forecasting tools include the COVID-19 simulator [21] and the COVID-19 Policy Alliance developed by a group in MIT. In this paper, hotspot detection is a binary classification problem, which differs in nature from the regression investigated by these studies. Besides accurate prediction and detection, understanding the spatial spread underlying the COVID-19 outbreak is also of great importance. An interpretable spatial model can help the government develop efficient public health policies to slow the spread during the early stages of COVID-19. Compared with literature that focuses on prediction, studies evaluating the spatial spread of the COVID-19 pandemic are still limited [22]. Previously, the spatial spread

has been studied for the outbreak of severe acute respiratory syndrome (SARS) in Beijing, and mainland China [23], [23]–[27] using only limited or localized data. In [28], the multivariate Hawkes process has been applied to model the conditional intensity of new confirmed COVID-19 cases and deaths in the U.S. at the county-level, without considering the influence from the big cities (main transportation hubs) and other important demographic factors. In [29], two types of county-level predictive models are developed based on the exponential and linear model, respectively. It focuses on modeling the dynamics of cumulative death counts. In [30], graph neural networks are adopted to capture the spatio-temporal dynamics between various features; however, a common disadvantage of the neural network-based methods is the lack of interpretability, which hinders from further understanding the mechanism underlying the COVID-19 spread.

Few studies have so far been conducted to investigate COVID-19 hotspots and their early detection. Similar to the CDC's definition, a recent study [31] considers a sudden increase in the number of cases in a specific geographical region. Unlike hotspot detection, they focus on estimating disease prevalence using logistic regression based on both symptoms and swab test results. In [32], hotspots are defined as spots with the highest incidence rate. This paper adopts statistical and spatial analysis to determine the spatial distribution and spatial clustering patterns of the COVID-19 incidence rate. To identify the COVID-19 hotspots, the Getis-Ord spatial statistic [33] was then applied. In another work [34], topological data analysis was applied to identify the hotspots of COVID-19 infections, which is defined as regions with higher case counts than their surrounding areas. However, no quantitative results, such as the estimation of confirmed cases or the prediction of future hotspots, were provided in [34].

Many studies used the Gaussian process for COVID-19 case prediction. In [35], a Gaussian process regression model is applied to mortality rate prediction in India. Unlike our model, [35] does not consider any spatial factors in the spread of COVID-19 and instead predicts cases on a national level. Some recent work [36], [37] have used Gaussian process models with a squared exponential kernel to forecast cases in the United States. However, these works provide state and national-level forecasts less granular than the city- or county-level forecasts produced by our model. In addition, [36], [37] use a stationary kernel, meaning that the kernel cannot adapt to different spatial patterns in different locations like the non-stationary spatial kernel discussed in this paper.

II. COVID-19 DATA DESCRIPTION

The data sets we used in our study include the number of cases and deaths, COVID-19 hotspots identified by the Centers for Disease Control and Prevention (CDC), and community mobility provided by Google. The study period is from March 15, 2020, to January 17, 2021, consisting of 50 weeks and 3,144 US counties. We excluded the data after February 2021, when a large-scale COVID-19 vaccine rollout had been launched across




Fig. 2. Snapshots of hotspots identified by CDC. The black circles indicate the counties that have been identified as hotspots in that week.

the United States, which significantly shifted the dynamics of the COVID-19 spread.

Confirmed cases and deaths: We used the data set from The New York Times (NYT) [38]¹ which includes two parts: (i) *confirmed cases* are counts of individuals whose coronavirus infections were confirmed by a laboratory test and reported by a federal, state, territorial, or local government agency; (ii) *confirmed deaths* are individuals who have died and meet the definition for a confirmed COVID-19 case. In practice, we have observed periodic weekly oscillations in daily reported cases and deaths, which could have been caused by testing bias (higher testing rates on certain days of the week). To reduce such bias, we aggregate the number of cases and deaths of each county *by week*.

Hotspots: On May 7, 2020, the CDC and other federal agencies began identifying counties with increasing COVID-19 hotspots to better understand transmission dynamics and offer targeted support to health departments in affected communities. The CDC identified hotspots daily starting on January 22, 2020, among counties in U.S. states and the District of Columbia by applying standardized criteria developed through a collaborative process involving multiple federal agencies [1], [39]. In general, hotspots were defined based on relative temporal increases in the number of cases. To match the temporal resolution with the number of cases and deaths, we expand the definition of a hotspot from daily-level to weekly-level. A week is identified as a hotspot if it contains at least one hotspot day identified by CDC. The weekly number of counties meeting hotspot criteria peaked in early April, decreased and stabilized during mid-April–early June, then increased again during late June–early July. The percentage of counties in the South and West Census regions meeting hotspot criteria increased from 10% and 13%, respectively, during March–April to 28% and 22%, respectively, during June–July. Fig. 2 gives snapshots of the identified hotspots at two particular weeks.

Community mobility: The COVID-19 Community Mobility Reports [40] record people’s movement by county daily, across various categories such as retail and recreation, groceries and pharmacies, parks, transit stations, workplaces, and residential. The data shows how visitors to (or time spent in) categorized places change compared to the baseline days (in percentage). The negative percentage means that the level of mobility is lower




Fig. 3. Overview of Google mobility data in two selected categories: workplace and transit on two different days. Counties in red and blue indicate their mobility is lower and higher than the normal level, respectively. The mobility level varies over time and space due to local government policy changes in response to COVID-19.

than the baseline, and the positive percentage represents the opposite. The mobility on a baseline day represents a normal value for that day of the week. This mobility report sets the baseline as the median value from the five weeks from January 3rd to February 6th, 2020. Similar to the two data sets mentioned above, we aggregate each county’s mobility data by week. Examples of two categories, transit stations, and workplaces, are shown in Fig. 3.

III. COVID-19 HOTSPOT DETECTION FRAMEWORK

This section presents our hotspot detection framework, consisting of two spatio-temporal models: confirmed cases and hotspots. Consider weeks $\mathcal{T} = \{t = 1, \dots, T\}$ starting from March 15, 2020 to January 17, 2021 and locations (counties) $\mathcal{S} = \{i = 1, \dots, I\}$, with latitude and longitude $s_i \in \mathcal{S} \subset \mathbb{R}^2$, $i \in \mathcal{S}$, where $\mathcal{S}$ represents the space of geographic coordinate system (GCS). The two models, respectively, focus on weekly confirmed cases $y_{it} \in \mathbb{Z}_+$ and identified hotspots $h_{it} \in \{0, 1\}$ of COVID-19 at location $i \in \mathcal{S}$ and time $t \in \mathcal{T}$, where $h_{it} = 1$ if there is a hotspot at location i and time t , and 0, otherwise. CDC [1] defined the hotspots based on relative temporal increases in the number of cases, i.e., the occurrence of the hotspots depends on the spatio-temporal correlation across different locations and over time, and *not* on the mean number of cases (see the observation in Fig. 1). Hence, we capture the correlation between y_{it} and h_{it} by connecting these two models in the spatio-temporal space (t, s_i) through a latent spatio-temporal random variable $f(t, s)$, characterized by a Gaussian process (GP) with zero mean and covariance specified by a kernel function k .

The goal is to find the optimal pair of these two models that best predict the hotspots and the cases for one week ahead. We refer to the proposed framework as the spatio-temporal Gaussian process (STGP).

¹One reason we use the NYT data rather than Johns Hopkins (JHU) data <https://coronavirus.jhu.edu/map.html> is that JHU data have retrospective data revision (when state update the COVID-19 definition, or have data error, etc), while NYT data never revise its history.

A. Spatio-Temporal Gaussian Process (STGP) Models

For the notational simplicity, we first denote the spatio-temporal coordinate (t, s_i) by $\mathbf{x}_{it} \in \mathcal{X}$, where $\mathcal{X} := \mathcal{T} \times \mathcal{S} \subset \mathbb{R}^3$ represents the spatio-temporal space. For any subset $\mathbf{X} \subseteq \mathcal{X}$ with N spatio-temporal coordinates, the set of function variables $\mathbf{f} := \{f(\mathbf{x})\}_{\mathbf{x} \in \mathbf{X}}$ has joint zero-mean Gaussian distribution

$$p(\mathbf{f}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{XX}), \quad (1)$$

where $\mathbf{K}_{XX}$ is a $N \times N$ matrix and its entries are pairwise evaluations of $k(\mathbf{x}, \mathbf{x}')$, $\forall \mathbf{x}, \mathbf{x}' \in \mathbf{X}$.

Case model: We define a spatio-temporal model for the confirmed cases in the following form:

$$y_{it} = \mu_{it} + f(\mathbf{x}_{it}) + \epsilon_{it}, \quad i \in \mathcal{I}, t \in \mathcal{T}, \quad (2)$$

where $\epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is assumed to be i.i.d. normally distributed; μ_{it} is the mean of number of confirmed cases at time t in location i . For a set of N observed spatio-temporal coordinates $\mathbf{X}$, we denote the number of confirmed cases and their means as $\mathbf{y} := \{y_{it}\}_{\mathbf{x}_{it} \in \mathbf{X}}$ and $\boldsymbol{\mu} := \{\mu_{it}\}_{\mathbf{x}_{it} \in \mathbf{X}}$, respectively. We assume the mean of the number of confirmed cases at a particular location relates to the covariates in its nearby counties according to an underlying undirected graph $\mathcal{G} = (\mathcal{I}, \mathcal{E})$, where $\mathcal{I}$ is the set of vertices representing all the locations, and $\mathcal{E} \subseteq \{(i, j) \in \mathcal{I}^2\}$ is a set of undirected edges representing the connections between locations. There is an edge between two vertices whenever the corresponding locations are geographically adjacent. Let $\boldsymbol{\eta}_{it} := [\eta_{it1}, \dots, \eta_{itl}, \dots, \eta_{itL}]^\top \in \mathbb{R}^L$ denote the data of these covariates at location $i \in \mathcal{I}$ and time $t \in \mathcal{T}$, and let $\boldsymbol{\omega}_{it} := [\omega_{it1}, \dots, \omega_{itl}, \dots, \omega_{itL}]^\top \in \mathbb{R}^L$ denote the parameters of the corresponding covariates; L denotes the number of features. In practice, we use the number of confirmed cases, the number of deaths, and six community mobilities variables in the past two weeks as the input covariates with $L = 16$. Formally, we define μ_{it} as

$$\mu_{it} = \sum_{\tau \in \mathcal{H}_t} \sum_{j: (i, j) \in \mathcal{E}} \boldsymbol{\eta}_{j\tau}^\top \boldsymbol{\omega}_{j\tau}, \quad \forall i \in \mathcal{I}, t \in \mathcal{T}, \quad (3)$$

where $\mathcal{H}_t = \{\tau : t - d \leq \tau < t\}$ represent the recent history with memory depth $d < T$.

Hotspot model: We express the conditional probability of the hotspots $\mathbf{h} := \{h_{it}\}_{\mathbf{x}_{it} \in \mathbf{X}}$ for a set of spatio-temporal coordinates $\mathbf{X}$ as:

$$p(\mathbf{h}|\mathbf{f}) = \prod_{\mathbf{x}_{it} \in \mathbf{X}} \mathcal{B}(h_{it}|\phi(f(\mathbf{x}_{it}))), \quad (4)$$

where $\mathcal{B}(h_{it}|\phi(f(\mathbf{x}_{it}))) = \phi(f(\mathbf{x}_{it}))^{h_{it}}(1 - \phi(f(\mathbf{x}_{it})))^{1-h_{it}}$ is the likelihood for the Bernoulli distribution and ϕ is a sigmoid function.

Learning objective: We aim to detect the hotspot while taking advantage of the information that has been recorded in the number of confirmed cases. To this end, we learn the model by optimizing the following combined objective:

$$\max_{\boldsymbol{\theta} \in \Theta} \ell(\boldsymbol{\theta}) := \ell_h(\boldsymbol{\theta}) + \delta \ell_y(\boldsymbol{\theta}), \quad (5)$$

where $\delta > 0$ controls the ratio between two objectives and $\boldsymbol{\theta} \in \Theta$ is the set of parameters defined in the kernel k . The $\ell_y(\boldsymbol{\theta}) := \log p(\mathbf{y})$ denotes the log marginal likelihood of observed confirmed cases and $\ell_h(\boldsymbol{\theta}) := \log p(\mathbf{h})$ denotes the log marginal likelihood of observed hotspots. We note that log marginal likelihood of cases in the second term plays a key role in “regularizing” the model by leveraging the information in the case records as shown in Appendix A². We also present the 5-fold cross-validation that quantitatively measures the F_1 score of the hotspot detection and the mean square error of the case prediction with different δ in Appendix A. The result confirms that the appropriate choice of δ can significantly improve the performance of hotspot detection.

B. Spatio-Temporal Deep Neural Kernel

We discuss the choice of the kernel function k in this subsection. Standard GP models use a stationary covariance, in which the covariance between any two points is a function of their Euclidean distance. However, stationary GPs fail to adapt to variable smoothness in the function of interest. This is of particular importance in geophysical and other spatial data sets, in which domain knowledge suggests that the function may vary more quickly in some parts of the input space than in others. For example, COVID-19 is likely to be spreading slower than in sparsely versus densely populated regions. Here, we consider the following non-stationary spatio-temporal kernel:

$$k(t, t', s, s') = \nu(t, t') \cdot \left(\sum_{r=1}^R w_{s'}^{(r)} v^{(r)}(s, s') \right), \quad (6)$$

where $\nu(t, t')$ is a stationary kernel that captures temporal correlation between time t and t' ; $v^{(r)}(s, s')$ is a component of the non-stationary spatial kernel which evolves over the space and $w_{s'}^{(r)}$ is the corresponding weight satisfying $\sum_{r=1}^R w_{s'}^{(r)} = 1$. R is the number of components considered. By likening the relationship between the spatial kernel component to that of the Gaussian component in the Gaussian mixture, we seek to enhance the representative power of our kernel by adding more independent components to the spatial kernel.

Stationary temporal kernel: We define the kernel function that characterizes the temporal correlation between $t, t' \in \mathcal{T}$ as an stationary Gaussian function:

$$\nu(t, t') = \exp \left\{ -\frac{1}{2\sigma_\nu^2} \|t - t'\|^2 \right\},$$

where $\sigma_\nu \in \mathbb{R}_+$ is the bandwidth parameter. This kernel function hypothesizes that the virus' transmission is highly related to its recent history and their correlation will decay exponentially over time.

Non-stationary spatial kernel: To account for non-stationarity, we now allow the smoothing kernel to depend on spatial location s . For ease of discussion and simplicity of notation, we omit the superscript r in $v^{(r)}(s, s')$ and $w_{s'}^{(r)}$, and present the structure of a single non-stationary spatial kernel

²Online appendix is available at <https://arxiv.org/abs/2106.00072>




Fig. 4. An example of the randomly generated focus points ψ and their corresponding covariance Σ . The horizontal coordinate represents different focus points realizations; the vertical coordinate represents the value of these realizations. The ellipses portray the shape of the corresponding covariance for the kernel $\kappa_s(\cdot)$ associated with that location s .

component. We use $\kappa_s(\cdot)$ to denote a kernel which is centered at the point s and whose shape is a function of location s . Once $\kappa_s(\cdot)$ is specified for all $s \in \mathcal{S} \subseteq \mathbb{R}^2$, the correlation between two points s and s' is then

$$v(s, s') \propto \int_{\mathbb{R}^2} \kappa_s(u) \kappa_{s'}(u) du. \quad (7)$$

Because of the constructive formulation under the moving average specification, the resulting correlation function $v(s, s')$ is certain to be positive definite. We favor working with the kernels $\kappa_s(\cdot)$ rather than directly with the correlation function $v(s, s')$ since this makes it difficult to ensure positive symmetry definiteness for all s and s' . Following the idea of [41], [42], we define each $\kappa_s(\cdot)$ to be a normal kernel centered at s with spatially varying covariance matrix Σ_s . In this case given the parameterized Σ_s and $\Sigma_{s'}$, the correlation function is given by an easy to compute formula

$$v(s, s') \propto \frac{|\Sigma_s + \Sigma_{s'}|^{-\frac{1}{2}}}{2\pi} \exp \left\{ -\frac{1}{2} (s' - s)^\top (\Sigma_s + \Sigma_{s'})^{-1} (s' - s) \right\}.$$

The derivation of this formula can be found in Appendix B.

To assure that the kernel $\{\kappa_s(\cdot)\}$ vary smoothly over space $\mathcal{S}$, we parameterize Σ_s and then allow the parameters to evolve with location. For this paper we will focus on a geometrically based specification which readily extends beyond the use of the Gaussian kernel considered here.

There is a one-to-one mapping from a bivariate normal distribution to its one standard deviation ellipse, so we define a spatially varying family of ellipses which, in turn, defines the spatial distribution for Σ_s . Let the two focus points in $\Psi \subset \mathbb{R}^2$ denoted by $\psi_s := (\psi_x(s), \psi_y(s)) \in \Psi$ and $-\psi_s := (-\psi_x(s), -\psi_y(s)) \in \Psi$ define an ellipse centered at s with fixed area A . This then corresponds to the Gaussian kernel with covariance matrix Σ_s defined by

$$\Sigma_s = \lambda^2 \begin{pmatrix} Q + \frac{\|\psi_s\|^2}{2} \cos 2\alpha & \frac{\|\psi_s\|^2}{2} \sin 2\alpha \\ \frac{\|\psi_s\|^2}{2} \sin 2\alpha & Q - \frac{\|\psi_s\|^2}{2} \cos 2\alpha \end{pmatrix}, \quad (8)$$

where $\alpha = \tan^{-1}(\psi_y(s)/\psi_x(s))$, $Q = \sqrt{4A^2 + \|\psi_s\|^4 \pi^2 / 2\pi}$, and λ is a scaling parameter that controls the overall intensity of the covariance. Fig. 4 shows a series of randomly generated focus points and their resulting ellipses. This demonstrates how the spatially distributed




Fig. 5. An illustration of the deep neural network that maps an arbitrary spatial location s to its covariance Σ_s and the corresponding weight w_s .




Fig. 6. An examples of the spatial kernel with two components $\sum_r w_s^{(r)} v^{(r)}(\cdot, s)$ evaluated at the same location s . This instance is constructed using two different kernel κ , which are parameterized by two randomly generated φ_1 and φ_2 .

pairs $\psi_s := (\psi_x(s), \psi_y(s))$ give rise to a spatially distributed covariance matrix Σ_s . The derivation of (8) can be found in Appendix C.

Neural network representation for focus points: Here we represent the mapping $\varphi : \mathcal{S} \rightarrow \Psi \times [0, 1]$ from the location space $\mathcal{S}$ to the joint space of focus point Ψ and the weight $[0, 1]$ using a deep neural network. To be specific, the input of the network is the location s , and the output of the networks is the concatenation of the corresponding focus points ψ_s of that location and the weight w_s defined in (6). The architecture of the neural network has been described in Fig. 5. In Fig. 6, we also demonstrate two specific instances of the resulting spatial kernel v given two different κ . This implies that the neural network φ encodes the non-homogeneous geographical information across the region in spreading the virus.

IV. EFFICIENT COMPUTATION FOR LARGE-SCALE DATA SET

There are two major challenges in learning the model and calculating the objective defined in (5). First, the GP approach is notoriously intractable for large data sets since the computations require the inversion of a matrix of size $N \times N$, which scales as $O(N^3)$ [43]. In this study, the data set includes 3,144 counties and more than 50 weeks extending from March 2020 to January 2021 ($N = 3,144 \times 50$). Second, the inference of the posterior distribution of the hotspot $p(\mathbf{f}|\mathbf{h})$ requires the calculation of integral $\int p(\mathbf{h}|\mathbf{f})p(\mathbf{f})d\mathbf{f}$, which is an intractable integration.

To circumvent these two issues, we derive sparse models for both cases and hotspots similar to [44]–[46], where their log




Fig. 7. A diagram of our graphical model. The gray and white nodes represent the observed and latent variables, respectively; the black dots represent the input variables; the black boxes represent the prior distribution. We use the dashed box to highlight the joint distribution of $\mathbf{f}$ and $\mathbf{u}$ defined in (9).

marginal likelihood is computationally tractable for large data sets and they do not require an analytical expression for inferring the non-Gaussian posterior distribution. First, we define a small set of inducing variables that aim to best approximate the training data. Then we adopt a variational learning strategy for such sparse approximation and jointly infer the inducing inputs and other model parameters by maximizing a lower bound of the true log marginal likelihood [44], [46]. Since the learning strategy can be applied to the above two models, we use $\mathbf{y}$ to represent both cases and hotspots for notational brevity in the following discussion. Lastly, the objective is jointly learned by performing stochastic gradient descent.

A. Variational Inference for Sparse Gaussian Process

Unlike the exact GP approaches approximating the true covariance by the Nyström approximation [43], we desire a sparse method that directly approximates the posterior GP's mean and covariance function. Now we introduce a small set of M auxiliary inducing variables $\mathbf{u}$ evaluated at the pseudo-inputs $\mathbf{Z} := \{\mathbf{z} \in \mathcal{X}\}$; $\mathbf{Z}$ can be a subset of the training inputs or auxiliary pseudo-points [47]. $\mathbf{u}$ are function points drawn from the same GP prior as the training functions $\mathbf{f}$ in (1), so the joint distribution can be written as

$$p([\mathbf{f}, \mathbf{u}]^\top) = \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{XX} & \mathbf{K}_{XZ} \\ \mathbf{K}_{XZ}^\top & \mathbf{K}_{ZZ} \end{bmatrix}\right), \quad (9)$$

where $\mathbf{K}_{ZZ}$ is formed by evaluating the kernel function pairwise at all pairs of inducing points in $\mathbf{Z}$, and $\mathbf{K}_{XZ}$ is formed by evaluating the kernel function across the data points $\mathbf{X}$ and inducing points $\mathbf{Z}$ similarly. Fig. 7 presents the diagram of our graphical model, consisting of observed variables $\mathbf{y}, \mathbf{h}$, latent variable $\mathbf{f}$, and the introduced auxiliary variable $\mathbf{z}$.

To obtain computationally efficient inference, we approximate the posterior distribution $p(\mathbf{f}, \mathbf{u}|\mathbf{y})$ over random variable vector $\mathbf{f}$ and $\mathbf{u}$ by a variational distribution $q(\mathbf{f}, \mathbf{u})$. We assume this variational distribution $q(\mathbf{f}, \mathbf{u})$ can be factorized as $q(\mathbf{f}, \mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$. To jointly determine the variational parameters and model parameters, the variational evidence lower bound (ELBO) substitutes for the marginal likelihood $\ell_y(\boldsymbol{\theta})$ and $\ell_h(\boldsymbol{\theta})$ defined in (5):

$$\log p(\mathbf{y}) \geq \mathbb{E}_{q(\mathbf{f})} [\log p(\mathbf{y}|\mathbf{f})] - \text{KL}[q(\mathbf{u})||p(\mathbf{u})], \quad (10)$$

where $\text{KL}[q||p]$ denotes the Kullback–Leibler (KL) divergence between two distributions q and p [48]. We have defined: $q(\mathbf{f}) := \int p(\mathbf{f}|\mathbf{u})q(\mathbf{u})d\mathbf{u}$ and assume $q(\mathbf{u}) := \mathcal{N}(\mathbf{m}, \mathbf{S})$, which is the most common way to parameterize the prior distribution of inducing variables in terms of a mean vector $\mathbf{m}$ and a covariance matrix $\mathbf{S}$. To ensure that the covariance matrix remains positive definite, we represent it using a lower triangular form $\mathbf{S} = \mathbf{L}\mathbf{L}^\top$. This leads to the following analytical form for $q(\mathbf{f})$:

$$q(\mathbf{f}) = \mathcal{N}(\mathbf{A}\mathbf{m}, \mathbf{K}_{XX} + \mathbf{A}(\mathbf{S} - \mathbf{K}_{ZZ})\mathbf{A}^\top),$$

where $\mathbf{A} = \mathbf{K}_{XZ}\mathbf{K}_{ZZ}^{-1}$. In classification or regression, we also factorize the likelihood as $p(\mathbf{y}|\mathbf{f}) = \prod_{n=1}^N p(y_n|f_n)$ for the ease of computation in (10). Therefore, the ELBO objective can be rewritten as

$$\ell_{\text{ELBO}}(\boldsymbol{\theta}, \mathbf{Z}, \mathbf{m}, \mathbf{S}) := \sum_{n=1}^N \mathbb{E}_{q(f_n)} [\log p(y_n|f_n)] - \text{KL}[q(\mathbf{u})||p(\mathbf{u})]. \quad (11)$$

In practice, the one dimensional integrals of the log-likelihood in (11) can be computed by Gauss-Hermite quadrature [49]. In contrast to directly maximizing the marginal log likelihood defined in (5), computing this objective and its derivatives can be done in $O(NM^2)$ time. The derivation of the ELBO can be found in Appendix D.

B. Prediction With Variational Posterior

To make one-week ahead predictions for the hotspots and the number of confirmed cases, we first need to derive the posterior distribution of prediction $p(\mathbf{f}|\mathbf{y}, \mathbf{h})$ given the past observation. Suppose we have the spatio-temporal coordinates $\mathbf{X}_t := \{\mathbf{x}_{j\tau}\}_{j \in \mathcal{J}, \tau \leq t}$ and their observations $\mathbf{y}_t := \{y_{j\tau}\}_{j \in \mathcal{J}, \tau \leq t}$, $\mathbf{h}_t := \{h_{j\tau}\}_{j \in \mathcal{J}, \tau \leq t}$ until time t and the optimal inducing points $\mathbf{Z}$. We assume that the unobserved future data comes from the same generation process. Therefore, for all the locations at time $t+1$, i.e., $\mathbf{X}_* := \{\mathbf{x}_{j,t+1}\}_{j \in \mathcal{J}}$, we first estimate their means according to (3) denoted by $\boldsymbol{\mu}_* := \{\mu_{j,t+1}\}_{j \in \mathcal{J}}$, then the distribution of one-week-ahead prediction $\mathbf{f}_* := \{\hat{f}_{j,t+1}\}_{j \in \mathcal{J}}$ is given by

$$p(\mathbf{f}_*|\mathbf{y}_t, \mathbf{h}_t) = \mathcal{N}(\mathbf{A}_*\mathbf{m}, \mathbf{A}_*\mathbf{S}\mathbf{A}_*^\top + \mathbf{B}_*), \quad (12)$$

where $\mathbf{A}_* = \mathbf{K}_{*Z}\mathbf{K}_{ZZ}^{-1}$ and $\mathbf{B}_* = \mathbf{K}_{**} - \mathbf{K}_{*Z}\mathbf{K}_{ZZ}^{-1}\mathbf{K}_{*Z}^\top$. The $\mathbf{K}_{*Z}$ denotes a $I \times M$ matrix and its entries are pairwise evaluations of $k(\mathbf{x}_*, \mathbf{z})$ where $\mathbf{x}_* \in \mathbf{X}_*$ and $\mathbf{z} \in \mathbf{Z}$. The derivation of the predictive posterior can be found in Appendix E. The prediction for the number of cases and the probability of hotspots therefore can be made by plugging (12) into (2) and (4), respectively. We consider that a detector would raise an alarm if the hotspot probability h_{it} for location i at time t is above the pre-set threshold ζ_i . This threshold is chosen for each location by a grid search in $[0, 1]$. For each location, the threshold with the largest in-sample F_1 score is chosen.

The above formula reveals that predictive posterior distribution only depends on the inducing variables $\mathbf{u}$ at learned spatio-temporal coordinates $\mathbf{Z}$ and does not depend on the $\mathbf{f}$ at training coordinates. This shows that all the information from

Algorithm 1: Learning Algorithm for the COVID-19 Hotspot Detection Framework.

Initialization: Randomly initialize $\theta, \mathbf{Z}, \mathbf{m}, \mathbf{S}$;
Input: Data set $\mathbf{X}, \mathbf{y}, \mathbf{h}$; Number of iterations B ; Batch size n ;
for $b = \{1, \dots, B\}$ **do**
 Sample a subset $\mathbf{X}_b, \mathbf{y}_b, \mathbf{h}_b$ with n points from $\mathbf{X}, \mathbf{y}, \mathbf{h}$, respectively;
 Calculate ELBO of ℓ_h and ℓ_y based on (11) given data $\mathbf{X}_b, \mathbf{y}_b, \mathbf{h}_b$;
 Calculate the gradient of (5) w.r.t. θ ;
 Calculate the natural gradient of (5) w.r.t. $\mathbf{Z}, \mathbf{m}, \mathbf{S}$;
 Ascend the gradient of $\theta, \mathbf{Z}, \mathbf{m}, \mathbf{S}$;
end

the training data has been summarized by the proposed posterior distribution $q(\mathbf{u})$ defined in Section IV-A and the prediction for future weeks can be carried out efficiently.

C. Stochastic Gradient Descent Based Optimization

Now we describe our learning algorithm. The optimal parameters of the proposed model can be found by maximizing the combined objective (5) using gradient-based optimization. However, the full gradient evaluation can still be expensive to carry out. With a sparse prior (inducing variables), even though we can tackle the computational challenge in inverting a big matrix, evaluating the gradient of the first term in (11) still requires the full data set, which is memory-intensive if the size of the data set N is too large. To alleviate the problem of expensive gradient evaluation, we adopt a stochastic gradient-based method [45] and only compute the gradient of the objective function with respect to the GP's parameters denoted by $\theta := \{\{\varphi^{(r)}\}_{r=1, \dots, R}, \sigma_\nu\}$, evaluated on a random subset of the data at each iteration.

Additionally, the conventional stochastic gradient descent algorithm assumes that the parameters' loss geometry is Euclidean. This is a non-ideal assumption for the parameters of many distributions, e.g., Gaussian. Here we follow the idea of [45], [50] and apply adapted stochastic gradient descent to the variational parameters $(\mathbf{Z}, \mathbf{m}, \mathbf{S})$ in our GP model by taking steps in the direction of the approximate natural gradient. These gradients are computed by the usual gradient re-scaled by the inverse Fisher information. The Kullback–Leibler divergence is used to measure the “closeness” in the variational distribution space. Our learning algorithm is summarized in Algorithm 1.

V. RESULTS

This section reports the numerical results of our study. In the following examples, we consider the spatial kernel with $R = 4$ components and fit the model with $M = 500$ inducing variables using the COVID-19 data set described in Section II. For each spatial kernel component r , we choose a three-layer neural network with 64 nodes per layer to represent its mapping $\varphi^{(r)}$ through cross-validation (Appendix G). We first evaluate the explanatory power of the proposed framework by

investigating the learned spatial kernel function using COVID-19 data. We demonstrate the interpretable components of our model and visualize the spatio-temporal correlation across regions discovered by our fitted model. Then we examine the result of the hotspot detection and the case prediction by visualizing the one-week-ahead predictions and their distribution. We emphasize that our model not only generates accurate predictions, but also quantifies the uncertainty about the predictions. Lastly, we compare our method with six other commonly-used binary classification approaches by evaluating their out-of-sample predictive performance. The inputs to the hotspot prediction model are past county-level case and death records, identified hotspots, and community mobility information. For ease of presentation, we only focus on the counties in the contiguous United States.

A. Model Interpretation

The proposed framework offers a unique opportunity for understanding the dynamics of the spread of COVID-19 utilizing the carefully crafted kernel design. In this experiment, we fit the model using the entire COVID-19 data and then visually examine the learned spatial kernel.

We first visualize the learned kernel induced feature $\kappa_s^{(r)}$ of each spatial kernel component in Fig. 8, which portrays the spatial pattern of virus' propagation. Recall that, for any arbitrary s , $\kappa_s^{(r)}$ is a normal kernel centered at s with spatially varying covariance matrix $\Sigma_s^{(r)}$, which can be uniquely represented by its focus points. Here, we connect two focus points with a red line at each location and plot them on the map. Length and rotation angle of the red line at s represents the strength and direction of the influence of location s , respectively, jointly determining the shape of its covariance matrix $\Sigma_s^{(r)}$. The color depth indicates the weight $w_s^{(r)}$, representing the “significance” of location s in the kernel component r .

To intuitively interpret the learned spatial kernel, we also visualize the kernel evaluation given one of its inputs, i.e., $\sum_{r=1}^R w_s^{(r)} \nu^{(r)}(s, \cdot)$. Such kernel evaluation represents the spatial correlation (or sphere of influence) of a particular location spreading the virus. The area in darker red signals that the corresponding counties are more likely to become the next hotspot if an outbreak is detected in the region of interest. In Fig. 9, we observe that these major metropolitan areas have a substantially different spatial correlation with their neighboring regions due to the non-stationarity of the spatial kernel. For example, as one of the nation's major economic and transportation hubs, New York has a significant impact on the entire Eastern United States, while Atlanta only has a regional influence in the Southeastern United States. Chicago and Los Angeles, the second and third most populous cities in the United States, can extend their influences to the entire north and south of the country, respectively. Such interpretability of our spatial kernel function could be of particular importance for the policymakers or the local authorities, suggesting that counties with more extensive influence on other regions should receive more attention in epidemic prevention. We also note that increasing the number of spatial kernel components could further enhance the flexibility and the interpretability of the model (Appendix F); however, due to the need for additional parameters in the neural networks, the




Fig. 8. Visualizations of the learned kernel induced feature $\kappa_s^{(r)}$ using COVID-19 data set. Each panel shows one of four kernel components, where the line segment is the edge that connects two focus points of s , indicating the shape and the rotation of the kernel at that location; the shaded area shows the intensity of the corresponding weight $w_s^{(r)}$ at location s ; the darker the region, the larger the weight.




Fig. 9. Examples of the learned spatial kernel $\sum_{r=1}^R w_s^{(r)} v^{(r)}(\cdot, s)$ with four components evaluated at four major metropolitan areas in the U.S.. These maps show the spatial influence of these area to other region of the U.S.. The color depth indicates the intensity of the kernel value; the darker the color the higher the kernel's value.




Fig. 10. Temporal view of one-week-ahead and county-wise hotspot probability $p(h_x)$ suggested by our fitted model ($\delta = 10^{-5}$) using COVID-19 data. The first 6 panels represent major metropolitan areas and the last two panels represent less populated counties in the United States.

computational time dramatically increases when $R \geq 3$, with minimal performance improvement.

B. Hotspot Detection

We evaluate the one-week-ahead in-sample prediction accuracy of our proposed hotspot detection framework at the county level. We first fit the model using the entire data set from March 15, 2020, to January 17, 2021, which contains 3,144 counties and 50 weeks in total. The in-sample prediction for time t is obtained by feeding the data before t into the fitted model and predicting the one-week-ahead hotspots. We test our model with different δ values and perform cross-validation to identify the optimal δ . In Fig. 10, we report the in-sample prediction results for eight representative locations, which include six major metropolitan cities and two sparsely populated counties. The shaded gray area indicates the number of cases reported in that location, and the black star indicates the time of the identified hotspot. The

solid red line represents the corresponding estimated hotspot probability. The hotspot probability resulting from our model is considerably high whenever a genuine hotspot occurs and considerably low otherwise, which confirms the effectiveness of our framework. In Fig. 11, we visualize the prediction results on the map to intuitively examine the predictive performance from the spatial perspective. We select four particular weeks representing different stages of the COVID-19 pandemic in 2020. The black dot indicates the genuinely identified hotspot, and the color depth indicates the hotspot probability suggested by our fitted model. As we can see, our method can capture the spatial occurrence of these hotspots nicely, in which regions with sparsely distributed hotspots usually have a lower probability. In comparison, other regions with densely distributed hotspots have a higher probability. We emphasize that our hotspot detection framework can provide a realistic prediction that varies smoothly over time and space due to our GP assumption. This can be extremely useful when we try to make a continuous prediction




Fig. 11. Spatial view of one-week-ahead and county-wise hotspot probability $p(\mathbf{h}_*)$ suggested by our fitted model ($\delta = 10^{-5}$) using real COVID-19 data. This figure presents examples at four particular weeks, where the color depth indicates the probability of predicted hotspots and the black circles represent the hotspots given in the data.




Fig. 12. Temporal view of one-week-ahead and county-wise case prediction $\mathbf{y}_*$ suggested by our fitted model ($\delta = 10^{-5}$). This figure presents eight examples for major metropolitan area (first six panels) and less populated counties (last two panels) in the United States.




Fig. 13. Spatial view of the confidence interval of one-week-ahead and county-wise case prediction $\mathbf{y}_*$ at four particular weeks. The color depth indicates the length of 95% confidence interval of the prediction; the darker the region, the more uncertain the prediction becomes.

or estimate the likelihood of a hotspot to happen at an arbitrary spatio-temporal coordinate.

C. Case Prediction and Uncertainty Quantification

Our proposed framework also provides case prediction and uncertainty quantification besides hotspot detection. In Fig. 12, we present the predicted case number $\mathbf{y}_*$ over time as well as its confidence interval for the eight same locations that appeared previously. The black dash line represents the real reported cases, and the solid blue line represents the prediction $\mathbf{y}_*$ suggested by our case model. The one and two- σ regions are highlighted by the dark and light blue shaded areas. As we can see, the prediction result captures the general trend of the case records, which confirms that the case model can successfully extract useful information from the cases that will be used to regularize the hotspot model by optimizing (5). We note that the estimated confidence interval reflects the uncertainty level of our prediction for both cases and hotspots since the confidence interval only depends on the latent spatio-temporal variable $\mathbf{f}$. In Fig. 13, we show the confidence interval over the map for

four different weeks. The color depth indicates the uncertainty level (the length of the confidence interval) for that location. This intuitively tells us which area we are confident in making predictions and how this confidence changes over space.

D. Comparison With Baselines

We adopt standard performance metrics, including precision, recall, and F_1 score. This choice is because hotspot detection can be viewed as a binary classification problem. We aim to identify a hotspot for a particular location at a particular week in the data. Define the set of all identified hotspots as U , the set of detected hotspots using our method as V . Then precision P and recall R are defined as $P = |U \cap V|/|V|$, $R = |U \cap V|/|U|$, where $|\cdot|$ is the number of elements in the set. The F_1 score combines the *precision* and *recall*: $F_1 = 2PR/(P + R)$ and the higher F_1 score the better. Since numbers of hotspots in real data are highly sparse (comparing to the total number of spatio-temporal coordinates), we do not use the ROC curve (true positive rate versus false-positive rate) in our setting. The

TABLE I
F₁ SCORE OF OUT-OF-SAMPLE HOTSPOT DETECTIONS

	Precision	Recall	F ₁ score
Perceptron	0.424	0.242	0.308
Logistic	0.564	0.178	0.270
Linear SVM	0.622	0.064	0.117
k-NN	0.517	0.398	0.450
Kernel SVM	0.599	0.360	0.450
Decision Tree	0.537	0.293	0.340
STGP ($\delta = 10^{-5}$)	0.457	0.968	0.621

evaluation procedure is described as follows. Given the observed hotspot and other covariates (cases, deaths, and mobility) until t , we perform detection for all the locations at time $t + 1$. If the detected hotspot were indeed identified as a genuine hotspot by CDC, then it is a success. Otherwise, it is a misdetection. In our data, there are $50 \times 3144 = 157,200$ spatio-temporal coordinates in total, and 12,000 of them were identified as genuine hotspots.

We compare the hotspot detection results of our proposed method and several standard methods in binary classification, including perceptron, logistic regression, linear support vector machine (SVM), k -nearest neighbor (k -NN), kernel SVM with Gaussian kernel, and decision tree; see [51] for a detailed review of those machine learning algorithms and see Appendix H for hyperparameter choices. Table I shows the F_1 score for the out-of-sample prediction at county-level using our method. The result confirms that our model significantly outperforms other baseline methods.

VI. CONCLUSION

This paper proposes a Bayesian framework that combines hotspot detection and case prediction cohesively through a latent spatio-temporal random variable. The latent variable is modeled by a Gaussian process, where a flexible non-stationary kernel function characterizes its covariance. The framework has shown immense promise in modeling and predicting the COVID-19 hotspots in the United States. Our numerical study has also shown that the proposed kernel enjoys great representative power while being highly interpretable.

REFERENCES

- [1] A. M. Oster *et al.*, "Trends in number and distribution of COVID-19 hotspot counties—United States, Mar. 8–July 15, 2020," *Morbidity Mortality Weekly Rep.*, vol. 69, no. 33, pp. 1127–1132, 2020.
- [2] Centers for Disease Control and Prevention, "COVID-19 forecasts: Deaths," 2021. [Online]. Available: <https://www.cdc.gov/coronavirus/2019-ncov/covid-data/forecasting-us.html#state-forecasts>
- [3] H. Rossman *et al.*, "A framework for identifying regional outbreak and spread of COVID-19 from one-minute population-wide surveys," *Nat. Med.*, vol. 26, no. 5, pp. 634–638, 2020.
- [4] T. Harko, F. S. Lobo, and M. Mak, "Exact analytical solutions of the susceptible-infected-recovered (SIR) epidemic model and of the SIR model with equal death and birth rates," *Appl. Math. Comput.*, vol. 236, pp. 184–194, Jun. 2014. [Online]. Available: <https://doi.org/10.1016/j.amc.2014.03.030>
- [5] J. Fernández-Villaverde and C. I. Jones, "Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities," *Nat. Bur. Econ. Res.*, Working Paper 27128, May 2020. [Online]. Available: <https://doi.org/10.3386/w27128>
- [6] D. Caccavo, "Chinese and Italian COVID-19 outbreaks can be correctly described by a modified SIRD model," *medRxiv*, 2020, doi: [10.1101/2020.03.19.20039388](https://doi.org/10.1101/2020.03.19.20039388).
- [7] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Rev.*, vol. 42, no. 4, pp. 553–768, Jan. 2000. [Online]. Available: <https://doi.org/10.1137/S0036144500371907>
- [8] Z. Yang *et al.*, "Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions," *J. Thoracic Dis.*, vol. 12, no. 3, pp. 165–174, Mar. 2020.
- [9] C. Hou *et al.*, "The effectiveness of quarantine of Wuhan city against the Corona Virus Disease 2019 (COVID-19): A well-mixed SEIR model analysis," *J. Med. virol.*, vol. 92, no. 7, pp. 841–848, Jul. 2020. [Online]. Available: <https://doi.org/10.1002/jmv.25827>
- [10] Los Alamos National Laboratory, "LANL COVID-19 cases and deaths forecasts," 2021. [Online]. Available: <https://covid-19.bsvgateway.org/>
- [11] The University of Texas COVID-19 Modeling Consortium, "COVID-19 mortality projections for US States," 2021. [Online]. Available: <https://covid-19.tacc.utexas.edu/dashboards/us/>
- [12] Northeastern University, Laboratory for the Modeling of Biological and Socio-technical Systems, "COVID-19 modeling," 2021. [Online]. Available: <https://covid19.gleamproject.org/>
- [13] Institute for Health Metrics and Evaluation, "COVID-19 projections for the United States," 2021. [Online]. Available: <https://covid19.healthdata.org/united-states-of-america>
- [14] S. Zhu, A. Bukharin, L. Xie, M. Santillana, S. Yang, and Y. Xie, "High-resolution spatio-temporal model for county-level COVID-19 activity in the U.S.," *ACM Trans. Manage. Inf. Syst.*, vol. 12, no. 4, Sep. 2021, doi: [10.1145/3468876](https://doi.org/10.1145/3468876).
- [15] A. L. Bertozzi, E. Franco, G. Mohler, M. B. Short, and D. Sledge, "The challenges of modeling and forecasting the spread of COVID-19," *Proc. Nat. Acad. Sci.*, vol. 117, no. 29, pp. 16732–16738, Jul. 2020.
- [16] P. Ghosh, R. Ghosh, and B. Chakraborty, "COVID-19 in India: Statewise analysis and prediction," *JMIR Public Health Surveill.*, vol. 6, no. 3, 2020, Art. no. e20341.
- [17] M. Alazab, A. Awajan, A. Mesleh, A. Abraham, V. Jatana, and S. Alhyari, "COVID-19 prediction and detection using deep learning," *Int. J. Comput. Inf. Syst. Ind. Manage. Appl.*, vol. 12, pp. 168–181, 2020.
- [18] S. Tamang, P. Singh, and B. Datta, "Forecasting of COVID-19 cases based on prediction using artificial neural network curve fitting technique," *Glob. J. Environ. Sci. Manage.*, vol. 6, no. Special Issue (Covid-19), pp. 53–64, 2020.
- [19] M. Hawas, "Generated time-series prediction data of COVID-19's daily infections in Brazil by using recurrent neural networks," *Data Brief*, vol. 32, Oct. 2020, Art. no. 106175.
- [20] Z. Zhao, K. Nehil-Puleo, and Y. Zhao, "How well can we forecast the COVID-19 pandemic with curve fitting and recurrent neural networks?," *medRxiv*, 2020, doi: [10.1101/2020.05.14.20102541](https://doi.org/10.1101/2020.05.14.20102541).
- [21] J. Chhatwal *et al.*, "COVID-19 simulator: An interactive tool to inform COVID-19 intervention policy decisions in the United States," in *Value in Health*, vol. 23. New York, NY, USA: Elsevier Science Inc, 2020, pp. S555–S555.
- [22] C. Poirier *et al.*, "Real-time forecasting of the COVID-19 outbreak in chinese provinces: Machine learning approach using novel digital data and estimates from mechanistic models," *J. Med. Internet Res.*, vol. 22, no. 8, 2020, Art. no. e20285.
- [23] B. Meng, J. Wang, J. Liu, J. Wu, and E. Zhong, "Understanding the spatial diffusion process of severe acute respiratory syndrome in Beijing," *Public Health*, vol. 119, no. 12, pp. 1080–1087, Dec. 2005. [Online]. Available: <https://doi.org/10.1016/j.puhe.2005.02.003>
- [24] L.-Q. Fang *et al.*, "Geographical spread of SARS in mainland China," *Trop. Med. Int. Health*, vol. 14, no. s1, pp. 14–20, Oct. 2009. [Online]. Available: <https://doi.org/10.1111/j.1365-3156.2008.02189.x>
- [25] D. Kang, H. Choi, J.-H. Kim, and J. Choi, "Spatial epidemic dynamics of the COVID-19 outbreak in China," *Int. J. Infect. Dis.*, vol. 94, pp. 96–102, May 2020. [Online]. Available: <https://doi.org/10.1016/j.ijid.2020.03.076>
- [26] J. S. Jia, X. Lu, Y. Yuan, G. Xu, J. Jia, and N. A. Christakis, "Population flow drives spatio-temporal distribution of COVID-19 in China," *Nature*, vol. 582, no. 7812, pp. 389–394, Apr. 2020. [Online]. Available: <https://doi.org/10.1038/s41586-020-2284-y>
- [27] C. C. F. D. C. Epidemiology Working Group for NCIP Epidemic Response and Prevention, "The epidemiological characteristics of an outbreak of 2019 novel coronavirus diseases (COVID-19) in China," *Zhonghua Liu Xing Bing Xue Za Zhi = Zhonghua Liuxingbingxue Zazhi*, vol. 41, no. 2, pp. 145–151, Feb. 2020. [Online]. Available: <https://doi.org/10.3760/cma.j.issn.0254-6450.2020.02.003>

- [28] W.-H. Chiang, X. Liu, and G. Mohler, "Hawkes process modeling of COVID-19 with mobility leading indicators and spatial covariates," *Int. J. Forecast.*, vol. 38, no. 2, pp. 505–520, Apr.–Jun. 2022, doi: [10.1016/j.ijforecast.2021.07.001](https://doi.org/10.1016/j.ijforecast.2021.07.001).
- [29] N. Altieri *et al.*, "Curating a COVID-19 data repository and forecasting county-level death counts in the United States," *Harvard Data Sci. Rev.*, vol. 1, Feb. 2021, doi: [10.1162/99608f92.1d4e0dae](https://doi.org/10.1162/99608f92.1d4e0dae).
- [30] A. Kapoor *et al.*, "Examining COVID-19 forecasting using spatio-temporal graph neural networks," 2020, *arXiv:2007.03113*.
- [31] T. Varsavsky *et al.*, "Detecting COVID-19 infection hotspots in England using large-scale self-reported data from a mobile application: A prospective, observational study," *Lancet Public Health*, vol. 6, no. 1, pp. e21–e29, 2021.
- [32] M. Shariati, T. Mesgari, M. Kasraee, and M. Jahangiri-Rad, "Spatiotemporal analysis and hotspots detection of COVID-19 using geographic information system (march and april, 2020)," *J. Environ. Health Sci. Eng.*, vol. 18, no. 2, pp. 1499–1507, 2020.
- [33] A. Getis and J. K. Ord, "The analysis of spatial association by use of distance statistics," in *Perspectives on Spatial Data Analysis*, Berlin, Germany: Springer, 2010, pp. 127–145.
- [34] M. Feng, A. Hickok, and M. A. Porter, "Topological data analysis of spatial systems," 2021, *arXiv:2104.00720*.
- [35] S. Dhamodharavadhani and R. Rathipriya, "COVID-19 mortality rate prediction for India using statistical neural networks and gaussian process regression model," *Afr. Health Sci.*, vol. 21, no. 1, pp. 194–206, 2021.
- [36] R. M. A. Velásquez and J. V. M. Lara, "Forecast and evaluation of COVID-19 spreading in USA with reduced-space Gaussian process regression," *Chaos, Solitons Fractals*, vol. 136, 2020, Art. no. 109924.
- [37] S. Ketu and P. K. Mishra, "Enhanced Gaussian process regression-based forecasting model for COVID-19 outbreak and significance of IoT for its detection," *Appl. Intell.*, vol. 51, no. 3, pp. 1492–1512, 2021.
- [38] T. N. Y. Times, "We're sharing coronavirus case data for every U.S. County," 2020. [Online]. Available: <https://www.nytimes.com/article/coronavirus-county-data-us.html>
- [39] A. M. Oster, E. Caruso, J. DeVies, K. P. Hartnett, and T. K. Boehmer, "Transmission dynamics by age group in COVID-19 hotspot counties—United States, April–September 2020," *Morbidity Mortality Weekly Rep.*, vol. 69, no. 41, pp. 1494–1496, 2020.
- [40] Google, "COVID-19 community mobility reports," 2020. [Online]. Available: <https://www.google.com/covid19/mobility/>
- [41] J. Bernardo *et al.*, "Non-stationary spatial modeling," 1998.
- [42] S. Zhu, S. Li, Z. Peng, and Y. Xie, "Imitation learning of neural spatio-temporal point processes," *IEEE Trans. Knowl. Data Eng.*, early access, Jan. 26, 2021, doi: [10.1109/TKDE.2021.3054787](https://doi.org/10.1109/TKDE.2021.3054787).
- [43] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*, Berlin, Germany: Springer, 2003, pp. 63–71.
- [44] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Proc. 12th Int. Conf. Artif. Intell. Statist.*, 2009, pp. 567–574.
- [45] J. Hensman, N. Fusi, and N. D. Lawrence, "Gaussian processes for Big Data," in *Proc. 29th Conf. Uncertainty Artif. Intell.*, 2013, pp. 282–290.
- [46] J. Hensman, A. Matthews, and Z. Ghahramani, "Scalable variational gaussian process classification," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2015, pp. 351–360.
- [47] E. Snelson and Z. Ghahramani, "Sparse Gaussian processes using pseudo-inputs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2005, pp. 1257–1264.
- [48] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 1951.
- [49] Q. Liu and D. A. Pierce, "A note on gauss–hermite quadrature," *Biometrika*, vol. 81, no. 3, pp. 624–629, 1994.
- [50] J. Martens, "New insights and perspectives on the natural gradient method," *J. Mach. Learn. Res.*, vol. 21, pp. 146:1–146:76, 2020. [Online]. Available: <http://jmlr.org/papers/v21/17-678.html>
- [51] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge U.K.: Cambridge Univ. Press, 2014.