



System inference for the spatio-temporal evolution of infectious diseases: Michigan in the time of COVID-19

Z. Wang¹ · X. Zhang¹ · G. H. Teichert¹ · M. Carrasco-Teja¹ · K. Garikipati¹

Received: 1 July 2020 / Accepted: 12 July 2020 / Published online: 12 August 2020
© Springer-Verlag GmbH Germany, part of Springer Nature 2020, corrected publication 2020

Abstract

We extend the classical SIR model of infectious disease spread to account for time dependence in the parameters, which also include diffusivities. The temporal dependence accounts for the changing characteristics of testing, quarantine and treatment protocols, while diffusivity incorporates a mobile population. This model has been applied to data on the evolution of the COVID-19 pandemic in the US state of Michigan. For system inference, we use recent advances; specifically our framework for Variational System Identification (Wang et al. in *Comput Methods Appl Mech Eng* 356:44–74, 2019; [arXiv:2001.04816](https://arxiv.org/abs/2001.04816) [cs.CE]) as well as Bayesian machine learning methods.

Keywords Inverse problems · Epidemiology · Compartmental models · Optimization · Neural networks

1 Background

Starting from their origins in the the work of Kermack and McKendrick [1], the use of differential equation models of the course of infectious diseases has grown to become one of the more accessible instances of the reach of mathematics. The current COVID-19 Pandemic has brought them into the common parlance. Even before this, however, the baseline Susceptible-Infected-Recovered (SIR) model had been extended to include Exposed (E) and Deceased (D) compartments and applied with considerable success to influenza, ebola, malaria, cholera, tuberculosis and several other infectious diseases [2–5]. (Some of this literature also includes agent-based models, which we do not consider here.) During the COVID-19 Pandemic, the widespread availability of data in the public domain [6–11] has served to attract methods of mathematics, computation and data science to analyzing this information, inferring the disease’s dynamics and making projections. The present communication is in this spirit, and brings our recent work in large scale computations of partial differential equations (PDEs), system inference and machine learning to this problem [12–17].

Of particular interest to us are two lines of enquiry: The first is that for a rapidly evolving disease such as COVID-19, with its public health, population-based, political, travel and economic manifestations, the classical SIR model of ordinary differential equations (ODEs) with constant coefficients seems inadequate. Driven by data that extends the compartments to the deceased (D), we have adopted the SIRD model. This choice is based entirely on the nature of the data available on the epidemic in the state of Michigan, where the numbers of deceased are reported on a daily basis. The classical SIR model can be extended to compartments additional to the exposed and deceased ones. These models are typically designated as SIS (Susceptible-Infected-Susceptible again, such as in the common cold); MSIR (Maternal-Susceptible-Infected-Recovered, where immunity is derived from the mother in the M compartment); SEIS (Susceptible-Exposed-Infected-Susceptible again, also typical of the common cold); MSEIR and MSEIRD, which combine more of the compartments. However, data are not available to us on the exposed and maternally immune-protected sub-populations in the state of Michigan, and the Maternal compartment is not known to be relevant to COVID-19. We have therefore worked with the SIRD model. The interested reader is directed to Ref. [18] for details of the other SIR model variants.

The first extension that we have undertaken is to allow the ODE coefficients to vary in time to reflect the evolving contours of testing, quarantine and treatment protocols.

✉ K. Garikipati
krishna@umich.edu

¹ Mechanical Engineering, Mathematics and the Michigan Institute for Computational, Discovery and Engineering, University of Michigan, Ann Arbor, USA

This is not necessarily novel, and has been addressed in other work [2,19], although perhaps not with the inference approach of Variational System Identification (VSI) and ODE-constrained optimization that we have adopted.

The second is the fact of a mobile population. Population mobility has been addressed through metapopulation models that characterize how diseases move between population hubs, across countries, or even intercontinentally. The most widely known are gravity models (e.g. [20]), and network and agent based models [21]. Given the prominence that quarantine protocols—adorned with the current-day euphemism of “social distancing”—have played in the COVID-19 Pandemic, it appears natural to seek an extension of the SIRD model to a spatio-temporal PDE model. As the world went into lockdown, but at different rates and degrees of rigor, and then began to emerge from it, the detection of patterns of mobility in space and time presents a compelling avenue for investigation. Such an extension also has been considered—chiefly in the setting of the mathematical analysis of reaction–diffusion systems [22–24]. Our contribution to this aspect of the mathematical treatment is to also allow the diffusivity of the S, I and R sub-populations to vary with time.

To these tasks we have brought the abundance of high-quality, public domain, data on the evolution of the various compartment pertaining to the SIRD model in the US state of Michigan. The temporal resolution by days and spatial resolution by the 85 counties of Michigan has allowed us to apply our methods of Variational System Identification [12,13], PDE-constrained optimization and machine learning [14–17] to these data.

In Sect. 2 we review the foundational SIRD ODE model. Section 3 is on data preparation. The application of system identification and machine learning to the ODE system are, respectively, in Sects. 4 and 5. The results for inferred parameters and forward prediction are presented in Sect. 6. The extension to inferring mobility via reaction–diffusion systems is in Sect. 7. Our conclusions appear in Sect. 8.

2 The compartmental model of infectious disease dynamics

We use the SIRD version of compartmental epidemiology models. The population, taken to remain constant at N , is divided into four disjoint compartments with time-dependent sub-populations: $S(t)$ for susceptible, $I(t)$ for infected, $R(t)$ for recovered and $D(t)$ for deceased individuals. The governing ODEs are:

$$\frac{dS}{dt} = -\frac{\beta}{N}SI + \gamma R \quad (1)$$

$$\frac{dI}{dt} = \frac{\beta}{N}SI - \mu I - \alpha I \quad (2)$$

$$\frac{dR}{dt} = \mu I - \gamma R \quad (3)$$

$$\frac{dD}{dt} = \alpha I \quad (4)$$

$$N = S(t) + I(t) + R(t) + D(t). \quad (5)$$

This is the canonical form of the model where the sub-populations are assumed to be well-mixed so that spatial variations can be ignored over the domain of interest. Here $\beta(t)$ is the infection rate, $\mu(t)$ is the recovery rate, $\gamma(t)$ is the rate of immunity loss, and $\alpha(t)$ is the death rate—all allowed to vary with time. Using the natural temporal unit of one day, we note that $1/\mu(t)$ is also the number of days an individual remains infectious. It follows that $\beta(t)/\mu(t)$ is the effective reproduction number: the total number of the susceptible population that an infectious individual passes the disease to. This quantity is commonly denoted by R_0 , but we use $r_0(t) = \beta(t)/\mu(t)$, to distinguish it from the recovered population, and emphasizing that it, too, varies with time.

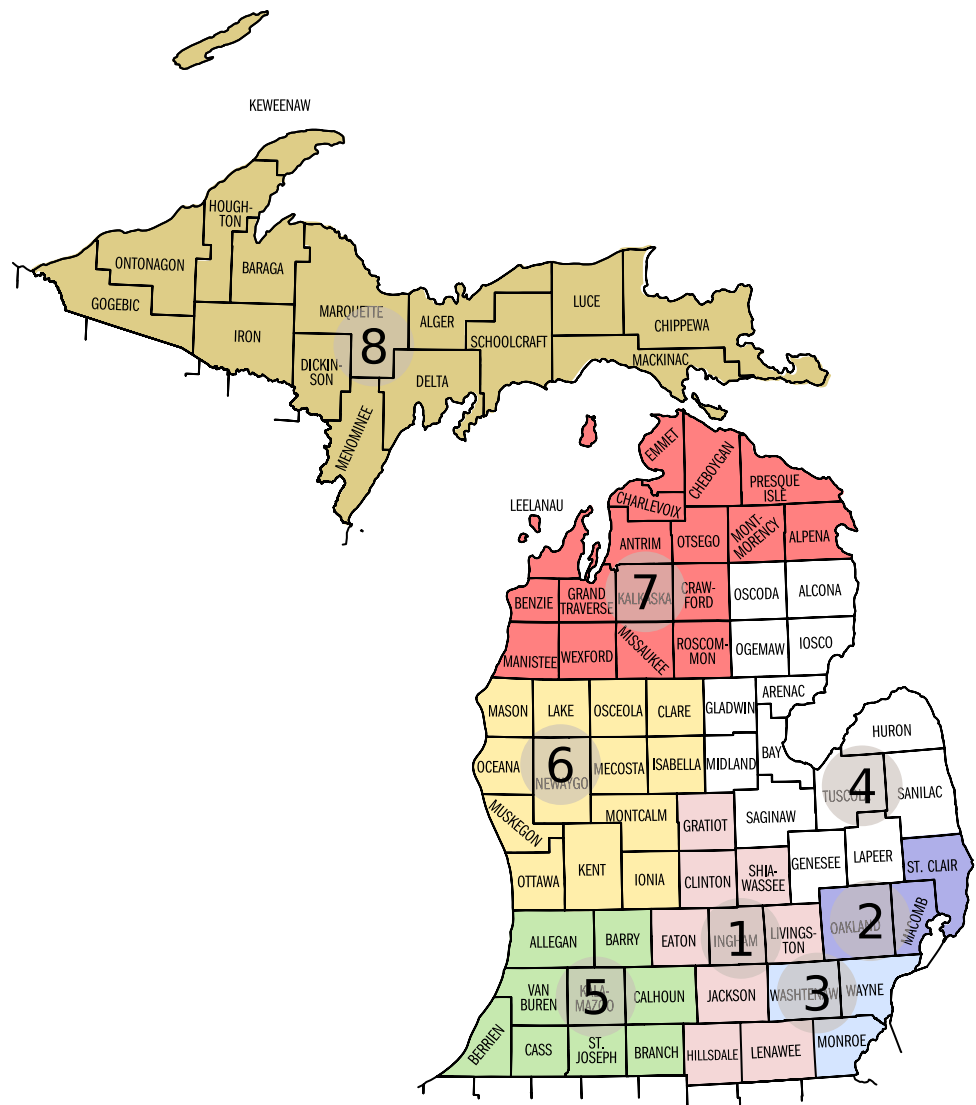
We reiterate what we have outlined in the Background (Sect. 1). Given the rapidly varying nature of testing, reporting, treatment protocols and quarantine conditions over the course of an epidemic, it is natural to allow the coefficients in the SIRD model, Eqs. (1–4) to vary with time. Such variation is evident in epidemiological data. The reader may be familiar with the time varying nature of such factors over the course of the COVID-19 Pandemic. It is a central feature of data preparation in the following section.

3 Data preparation

Counts of new confirmed infected cases $I(t)$ and deaths $D(t)$ were reported in the public domain on a daily basis by the state of Michigan for each county [26], while total recovered cases in the state $R(t)$ were reported weekly [27]. See Fig. 1 for the counties and regions that Michigan is partitioned into. Since county specific recovery data was not reported, the distribution of recovered cases across counties was approximated to be the same as the distribution of cumulative infected cases, $\int_0^t I(\tau)d\tau$ across counties. Estimates for the populations of Michigan’s counties [28] were used to determine the susceptible population, $S(t)$, from Eq. (5).

Some amount of data smoothing was necessary, particularly to account for the weekly instead of daily reporting of the number of recovered cases. To compare the effect of the smoothing method on the data, a moving average filter was applied using 7, 11, and 15-day windows, guided by the week-long period of oscillation in the raw data for daily new infections $I(t) - I(t - 1)$. The 7-day window was applied one, two, and three times. As seen in Fig. 2, the method of

Fig. 1 The map of Michigan delineating the counties and regions (modified from [25])



smoothing has little effect on the trends of the data. However, and as expected, there is a strong effect on the numerical time derivatives (see Fig. 3). It is clear that multiple passes of the filter are required to remove jumps in dR/dt and dI/dt . Since the additional smoothing is helpful for system inference in Sect. 4 and does not negatively affect the data, the 7-day moving average filter applied three times was used for data smoothing.

The lockdown in Michigan began on March 23, 2020. For brevity, we use C for the date when the outdoor construction industry was allowed to resume on May 1, 2020, M for the restart of some manufacturing on May 7, 2020, R for reopening of research laboratories on May 15, 2020, O for broader opening of most other activities and lifting of the stay at home order on June 1, 2020 (albeit with distancing guidelines in place), and E for the end of the data period that we considered (June 28, 2020). This notation is used for the rest of this communication.

4 System identification and ODE-constrained optimization

The SIRD model, Eqs. (1–4) was time-discretized using the Backward Euler method and written as:

$$\frac{S_m^d - S_{m-1}^d}{\Delta t} + \frac{\beta}{N} S_m^d I_m^d - \gamma R_m^d = 0 \tag{6}$$

$$\frac{I_m^d - I_{m-1}^d}{\Delta t} - \frac{\beta}{N} S_m^d I_m^d + \mu I_m^d + \alpha I_m^d = 0 \tag{7}$$

$$\frac{R_m^d - R_{m-1}^d}{\Delta t} - \mu I_m^d + \gamma R_m^d = 0 \tag{8}$$

$$\frac{D_m^d - D_{m-1}^d}{\Delta t} - \alpha I_m^d = 0 \tag{9}$$

where $S_m^d, I_m^d, R_m^d, D_m^d$ are the corresponding data (smoothed as in Sect. 3) at time t_m (the end of the m th day), $\Delta t = 1$ day and Eq. (5) holds: $S_m^d = N - I_m^d - R_m^d - D_m^d$.

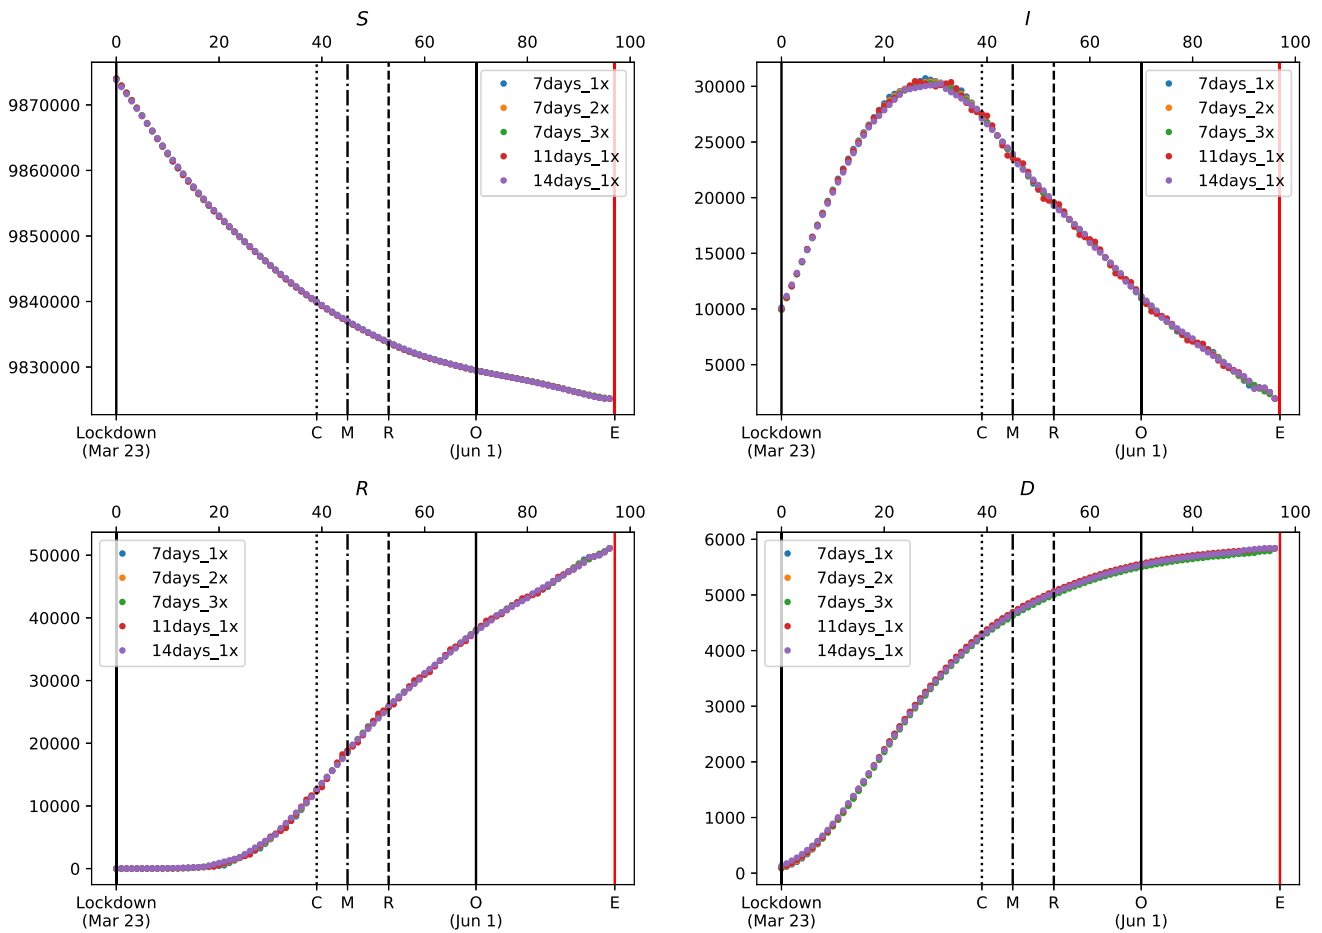


Fig. 2 Cumulative data with different kernel widths and multiples of application of the smoothing filter: 7days_1x represents a 7-day filter applied once. Important dates are marked with the lockdown on March 23, reopening of construction and real estate sites (C) on May 1,

reopening of manufacturing sites (M) on May 7, permission to restart laboratory research (R) on May 15, lifting of the stay-at-home order (O) on June 1, and the end of our data collection (E) on June 28

The system identification problem is to infer the time-dependent coefficients $\beta(t)$, $\gamma(t)$, $\mu(t)$, $\alpha(t)$, which we choose to expand in a polynomial basis (other choices of bases are admissible).

$$\beta(t) = \theta_0 + \theta_1 t + \theta_2 t^2 + \theta_3 t^3 \tag{10}$$

$$\gamma(t) = \theta_4 + \theta_5 t + \theta_6 t^2 + \theta_7 t^3 \tag{11}$$

$$\mu(t) = \theta_8 + \theta_9 t + \theta_{10} t^2 + \theta_{11} t^3 \tag{12}$$

$$\alpha(t) = \theta_{12} + \theta_{13} t + \theta_{14} t^2 + \theta_{15} t^3 \tag{13}$$

The parameters to be inferred are collected into a vector $\theta = \langle \theta_0, \dots, \theta_{15} \rangle^T$. Since the data are known the label vector can be constructed as:

$$y_m = \begin{Bmatrix} \frac{S_m^d - S_{m-1}^d}{\Delta t} \\ \frac{I_m^d - I_{m-1}^d}{\Delta t} \\ \frac{R_m^d - R_{m-1}^d}{\Delta t} \\ \frac{D_m^d - D_{m-1}^d}{\Delta t} \end{Bmatrix} \tag{14}$$

and a matrix can be assembled from the reaction terms in the time-discretized SIRD equations (6–9):

$$\Xi_m = \begin{bmatrix} \frac{S_m^d I_m^d}{N} \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & -R_m^d \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle \\ -\frac{S_m^d I_m^d}{N} \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle & I_m^d \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & I_m \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle \\ \langle 0 \ 0 \ 0 \ 0 \rangle & R_m^d \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & -I_m^d \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle \\ \langle 0 \ 0 \ 0 \ 0 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle & \langle 0 \ 0 \ 0 \ 0 \rangle & -I_m^d \langle 1 \ t_m \ t_m^2 \ t_m^3 \rangle \end{bmatrix} \tag{15}$$

The columns of Ξ_m can be regarded as discretized versions of the basis operators that appear as reaction terms on the right hand-side of the SIRD model (1–4). The label vectors and

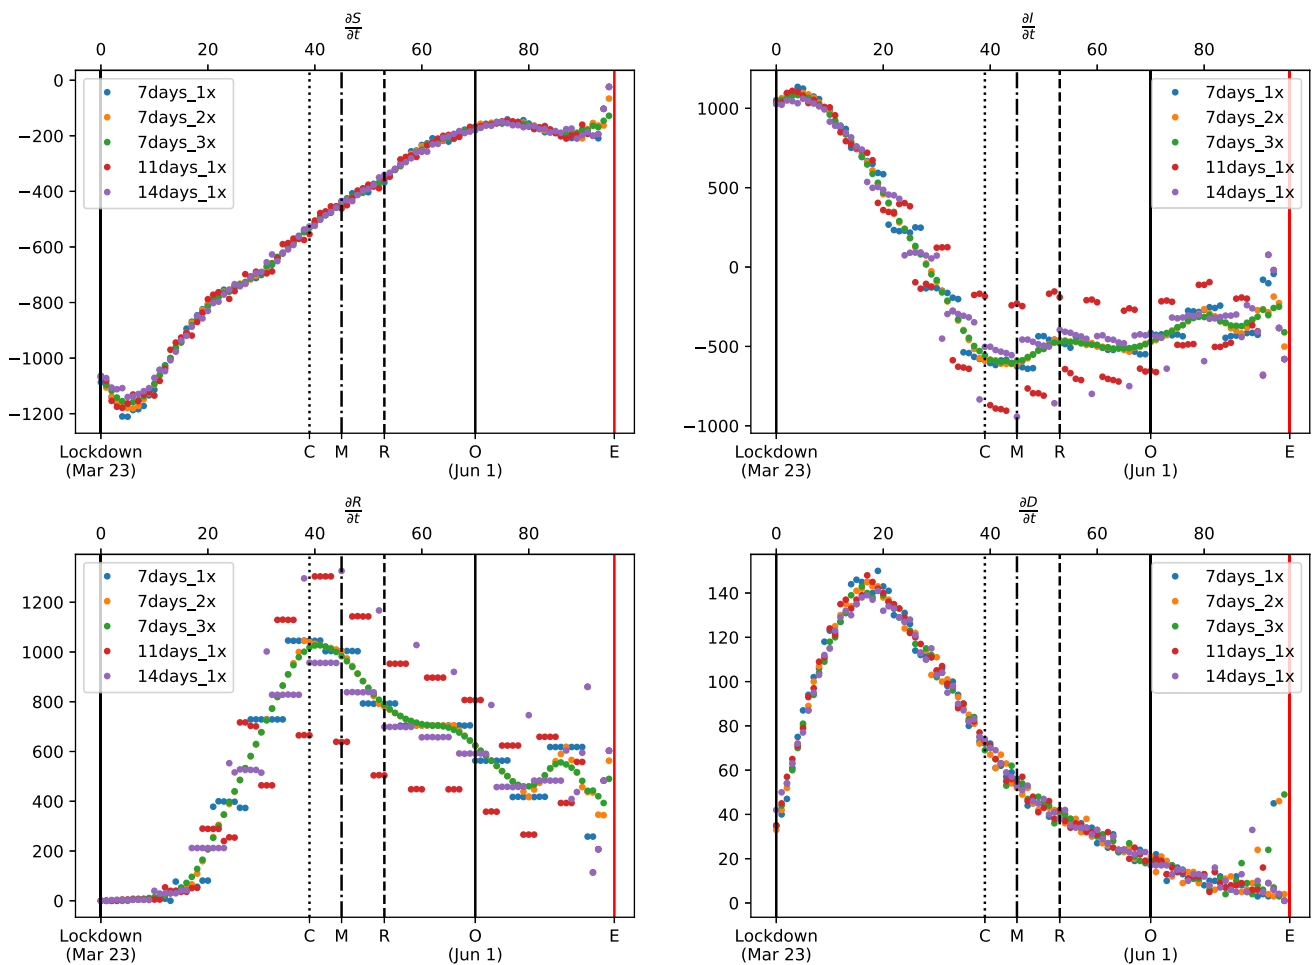


Fig. 3 Time derivatives (daily change) of sub-population data with different kernel widths and multiples of application of the smoothing filter. Important dates are marked with the lockdown on March 23, reopening of construction and real estate sites (C) on May 1, reopening of manu-

facturing sites (M) on May 7, permission to restart laboratory research (R) on May 15, lifting of the stay-at-home order (O) on June 1, and the end of our data collection (E) on June 28

matrices of basis operators at times t_0, \dots, t_M are collected into

$$\mathbf{y} = \begin{Bmatrix} \mathbf{y}_0 \\ \vdots \\ \mathbf{y}_M \end{Bmatrix}, \quad \mathbf{\Xi} = \begin{bmatrix} \mathbf{\Xi}_0 \\ \vdots \\ \mathbf{\Xi}_M \end{bmatrix} \tag{16}$$

$4(M+1) \times 1$ $4(M+1) \times 16$

and the residual vector is defined:

$$\mathcal{R}(\boldsymbol{\theta}) = \mathbf{y} - \mathbf{\Xi}\boldsymbol{\theta} \tag{17}$$

Our approach to inference combines system identification by stepwise regression [12,13] and ODE-constrained optimization using adjoints. We define a loss function that incorporates penalization on $\boldsymbol{\theta}$ (leading to ridge regression

below):

$$\ell(\boldsymbol{\theta}) = |\mathcal{R}(\boldsymbol{\theta})|^2 + \frac{1}{2}\lambda|\boldsymbol{\theta}|^2 \tag{18}$$

Our stepwise regression techniques incorporate two algorithms listed next:

Algorithm 1: Model selection by Stepwise regression:

```

j = 0, Q = 0,
While j ≤ P - 1 - Q do
Step 0:
    Establish target vector y and matrix of bases E.
Step 1:
    Solve for θj by ridge regression:

θj = arg minθ̃ ℓ(θ̃)
      = ( ETE + λ1 )-1 ETy
      (19)

    Calculate the loss function at this iteration, ℓj.
Step 2:
    Apply the F-test introduced below.
    IF F-test eliminates an operator
    THEN Set Q = Q + 1
        Set to zero the corresponding component of θ.
        GOTO Step 1. %Loss function remains small (ℓj ~ ℓj-1); solution may be
overfit.
    ENDIF
Step 3:
    Stop if the F-test does not allow elimination of any more basis
operators.
    % Beyond this, the loss function increases dramatically for any further
reduction.

```

There are several possible criteria for eliminating basis terms. Here, we adopt a widely used statistical criterion called the *F*-test, also used by us previously [12,13]. The significance of the change between the model at iterations j and $j - 1$ is evaluated by:

$$F = \frac{\frac{\ell_j - \ell_{j-1}}{p_{j-1} - p_j}}{\frac{\ell_{j-1}}{P - p_{j-1}}} \quad (20)$$

where p_j is the number of bases at iteration j and $P = 16$ is the total number of operator bases. The *F*-test is achieved through the application of Algorithm 2:

Model selection thus finds θ consisting of a minimal set of non-zero components, ensuring that the coefficients $\beta(t), \dots, \alpha(t)$ admit a parsimonious representation as polynomials in t . For clarity, we collect this set of *non-zero* coefficients into another vector, ϑ_0 . Using $\dim(\bullet)$ to represent the dimension of a Euclidean vector, we have $\dim(\vartheta_0) \leq \dim(\theta)$.

The next step is to further refine the values of the non-zero polynomial coefficients using ODE-constrained optimization starting from the initial guess ϑ_0 , and regarding $S_m(\tilde{\vartheta}), I_m(\tilde{\vartheta}), R_m(\tilde{\vartheta}), D_m(\tilde{\vartheta})$ as the forward solution to the discretized SIRD model (22–25) with coefficient $\beta(t), \dots, \alpha(t)$ values drawn from $\tilde{\vartheta}$:

$$\vartheta = \arg \min_{\tilde{\vartheta}} \sum_{m=0}^M \left(\frac{S_m(\tilde{\vartheta}) - S_m^d}{W_1} \right)^2 + \left(\frac{I_m(\tilde{\vartheta}) - I_m^d}{W_2} \right)^2 + \left(\frac{R_m(\tilde{\vartheta}) - R_m^d}{W_3} \right)^2 + \left(\frac{D_m(\tilde{\vartheta}) - D_m^d}{W_4} \right)^2 \quad (21)$$

Subject to the discretized SIRD model:

$$\forall m \in \{0, \dots, M\}$$

$$\frac{S_m(\tilde{\vartheta}) - S_{m-1}(\tilde{\vartheta})}{\Delta t} + \frac{\beta}{N} S_m(\tilde{\vartheta}) I_m(\tilde{\vartheta}) - \gamma R_m(\tilde{\vartheta}) = 0 \quad (22)$$

$$\frac{I_m(\tilde{\vartheta}) - I_{m-1}(\tilde{\vartheta})}{\Delta t} - \frac{\beta}{N} S_m(\tilde{\vartheta}) I_m(\tilde{\vartheta}) + \mu I_m(\tilde{\vartheta})$$

$$+ \alpha I_m(\tilde{\vartheta}) = 0 \tag{23}$$

$$\frac{R_m(\tilde{\vartheta}) - R_{m-1}(\tilde{\vartheta})}{\Delta t} - \mu I_m(\tilde{\vartheta}) + \gamma R_m(\tilde{\vartheta}) = 0 \tag{24}$$

$$\frac{D_m(\tilde{\vartheta}) - D_{m-1}(\tilde{\vartheta})}{\Delta t} - \alpha I_m(\tilde{\vartheta}) = 0 \tag{25}$$

dictions on case numbers at day $m+k$ based on case numbers reported at day m . In this work, k is chosen to vary from 1 to $M - m$, where $M = 97$ is the number of days that we used data for. In both types of NNs, S_m^d and S_{m+k}^d are computed based on the constraint Eq. (5).

Algorithm 2: Application of the F -test:

```

Step 1: %Find the least significant basis
i = 0
While i ≤ P - 1 - Q      Tentatively eliminate each basis corresponding to
non-zero coefficients in θ
    Set the corresponding coefficient to zero in θ.
    Evaluate the loss function followed by ridge regression on the reduced
basis set.
    Compute the  $F$ -value on the reduced basis set with smallest loss
function. Label this coefficient  $\theta_k$ 
Step 2:
    IF  $F < \alpha$ 
    THEN formally eliminate the term corresponding to  $\theta_k$  in matrix  $\Xi$ , by
deleting the corresponding column.
    ENDIF
    
```

where

$$W_1 = \max_m S_m^d - \min_m S_m^d$$

$$W_2 = \max_m I_m^d - \min_m I_m^d$$

$$W_3 = \max_m R_m^d - \min_m R_m^d$$

$$W_4 = \max_m D_m^d - \min_m D_m^d$$

The ODE-constrained optimization problem is solved iteratively, and requires the gradient of the ODE constraint (22–25) with respect to $\tilde{\vartheta}$. We adopt the classical approach requiring a single solution of the adjoint equation of the original ODE-constraint in each iteration. In this work we use the L-BFGS-B optimization algorithm from SciPy [29] and the dolfin-adjoint software library [30] to compute the gradient.

5 Deep and Bayesian neural networks

We also explore multilayer feedforward neural networks (NNs), which are universal function approximators [31], to learn the disease’s dynamics via the data $S_m^d, I_m^d, R_m^d, D_m^d$ at discrete times and to infer the coefficients in Eqs. (1–4), as an alternative to the approach presented in Sect. 4. Specifically, we construct two NNs to represent the data, with one as a deterministic model and the other being a probabilistic model.

Both NNs take $\{I_m^d, R_m^d, D_m^d \Delta t\}$ as features and $\{I_{m+k}^d, R_{m+k}^d, D_{m+k}^d\}$ as labels. Thus, the two NNs make pre-

The deterministic model is a deep neural network (DNN) that consists of multiple fully connected layers, whose model parameters (i.e. weights and bias) can be obtained in a straightforward manner by minimizing the loss function

$$\mathcal{L}_{DNN} = \text{MSE} \tag{26}$$

through an optimization algorithm, such as stochastic gradient descent, via backpropagation. The probabilistic model is a Bayesian neural network (BNN), which also consists of multiple fully connected layers, but with its model parameters (i.e. weights and bias) being sampled from a posterior distribution $P(\theta|\mathcal{D})$ that is computed based on Bayes’ theorem

$$P(\theta|\mathcal{D}) = \frac{P(\mathcal{D}|\theta)P(\theta)}{P(\mathcal{D})}, \tag{27}$$

where \mathcal{D} denote the i.i.d. observations (training data) and P represents the probability density function. In Eq. (27), $P(\mathcal{D}|\theta)$ is the likelihood, $P(\theta)$ is the prior probability, and $P(\mathcal{D})$ is the evidence, respectively. The posterior distribution of θ is computed based on variational inference (VI), which approximates the exact posterior distribution $P(\theta|\mathcal{D})$ with a more tractable distribution $Q(\theta)$ by minimizing the Kullback–Leibler (KL) divergence [32–34]

$$Q^* = \arg \min \text{KL}(Q(\theta)||P(\theta|\mathcal{D})). \tag{28}$$

The KL divergence is computed as

$$\text{KL}(Q(\theta)||P(\theta|\mathcal{D})) = \mathbb{E}[\log Q(\theta)] - \mathbb{E}[\log P(\theta, \mathcal{D})] + \log P(\mathcal{D}), \quad (29)$$

which requires computing the logarithm of the evidence, $\log P(\mathcal{D})$ in Eq. (27) [33]. Since $P(\mathcal{D})$ is hard to compute, it is challenging to directly evaluate the objective function in Eq. (28). Alternately, we can optimize the evidence lower bound (ELBO) defined as

$$\text{ELBO}(Q) = \mathbb{E}[\log P(\theta, \mathcal{D})] - \mathbb{E}[\log Q(\theta)], \quad (30)$$

which is equivalent to the KL-divergence up to an additive constant coming from the evidence. Thus, maximizing the ELBO is equivalent to minimizing the KL-divergence. The loss function for the BNN has the following form:

$$\mathcal{L}_{\text{BNN}} = \omega_1 \text{MSE} + \omega_2 \text{ELBO}, \quad (31)$$

where ω_1 and ω_2 are weighting parameters, with $\omega_1 = 50$ and $\omega_2 = 1$ being chosen in this work. A specific weight perturbation method, known as Flipout [35], is followed to infer $Q(\theta)$ by minimizing Eq. (31) through mini-batch training via backpropagation with stochastic optimization algorithms. Flipout has been implemented in the TensorFlow Probability Library. The architectures of both NNs are summarized in Table 1. Both NNs were trained by using the Adam optimizer following an exponentially decaying learning rate

$$\text{lr} = \text{lr}_0 \cdot \text{pow}\left(v_{\text{decay}}, \frac{N_{\text{total}}}{N_{\text{decay}}}\right) \quad (32)$$

with an initial learning rate $\text{lr}_0 = 0.001$, a decay rate $v_{\text{decay}} = 0.91$, a decay step $N_{\text{decay}} = 100$, and a final $N_{\text{total}} = 10,000$ epochs.

6 Results

Because of the extremely nonuniform distribution of the population of Michigan, we first studied the SIRD model for the entire state consisting of the lower and upper peninsulas (Fig. 1). Following this, the SIRD models were inferred for the eight Regions (also shown in Fig. 1) individually as one direct approach to study the effect of spatial variations in the populations and sub-populations corresponding to the model's compartments.

6.1 System identification and ODE-constrained optimization

Figure 4 shows the progression of stepwise regression to infer the active time-dependent terms in Eqs. (10–13) via Algorithms 1 and 2. The stem-and-leaf plots on the left illustrate

the fate of the terms $\theta_0 - \theta_{15}t^3$ over eight iterations of stepwise regression. Each stem-and-leaf represents one term out of $\theta_0 - \theta_{15}t^3$ and the values are scaled to 1 (active) or 0 (inactive) for each iteration. On the right is the loss, which remains low until Iteration 10 and increases dramatically in Iteration 11, if any further terms are eliminated. Following the F -test used in Algorithm 2, the large increase in loss after Iteration 10 exceeds the threshold for acceptable model error. Thus system identification converges to the inferred model in ten iterations.

Figure 5 shows, on the left, the evolution of SIRD model parameters and, on the right, a comparison of the predictions of the inferred model versus the data after ODE-constrained optimization that follows the system identification step. It is important to recall that these results are representative of the population of the entire state of Michigan. The SIRD model, having only four compartments, and applied to data that are the outcome of changing characteristics of testing, quarantine and treatment protocols, does not resolve many details of the public health aspects of the epidemic. The immunological characteristics of the disease itself are accounted for only in a very aggregated sense.

In Fig. 5, the important dates when the lockdown was imposed, and its gradual lifting are indicated by vertical lines to aid an understanding of the results. We first draw attention to the conclusion that $\gamma(t) = 0$; the inference indicates that recovery from COVID-19 confers permanent immunity—an important conclusion, that remains to be confirmed by immunologists. As may be expected, the population's infection rate, $\beta(t)$, declined as the initially higher rates of positive diagnoses fell with fewer infected individuals. However, it began to rise again upon the opening of construction activities (C), and continued to do so through the lifting of stay at home orders (O). The recovery rate, $\mu(t)$, showed a long initial increase as growing numbers of infected individuals recovered. Our interpretation of the initially high death rate, $\alpha(t)$, is that many of the early cases already had advanced progression of the disease. Its rapid decline can be attributed to the ramp up of the public health campaign, hospitalization and emergency response of the medical system. The success that the state of Michigan gained by mandating an aggressive lockdown of nearly all societal, educational, commercial and industrial activity is best reflected in the rapid decline of the effective reproduction number, $r_0(t)$. According to the inference presented here, $r_0(t) < 1.0$ for $m > 32$ (April 24, 2020), after which the typical infected individual passed the disease on to less than one other person. The death rate increased over the last few days for which data were obtained, perhaps as some number of individuals who had been infected for a longer time failed to recover. This affected the recovery rate as well, which fell. The close match between the simulations with the inferred ODE SIRD model and data (Fig. 5, right plot) validate the systems inferred. Such validation

Table 1 Model architecture for the DNN and BNN. Dense and DenseFlipout refer to specific NN architectures (layers) used in the TensorFlow Library

Layer type	Description
<i>DNN</i>	
Input layer (features)	$I_m^d, R_m^d, D_m^d \Delta t$
Dense layer	neurons = 40 (Sigmoid)
Dense layer	neurons = 40 (Sigmoid)
Output Dense Layer (labels)	$I_{m+k}^d, R_{m+k}^d, D_{m+k}^d$ (Softplus)
<i>BNN</i>	
Input layer (features)	$I_m^d, R_m^d, D_m^d \Delta t$
DenseFlipout layer	neurons = 40 (Sigmoid)
DenseFlipout layer	neurons = 40 (Sigmoid)
Output DenseFlipout Layer (labels)	$I_{m+k}^d, R_{m+k}^d, D_{m+k}^d$ (Softplus)

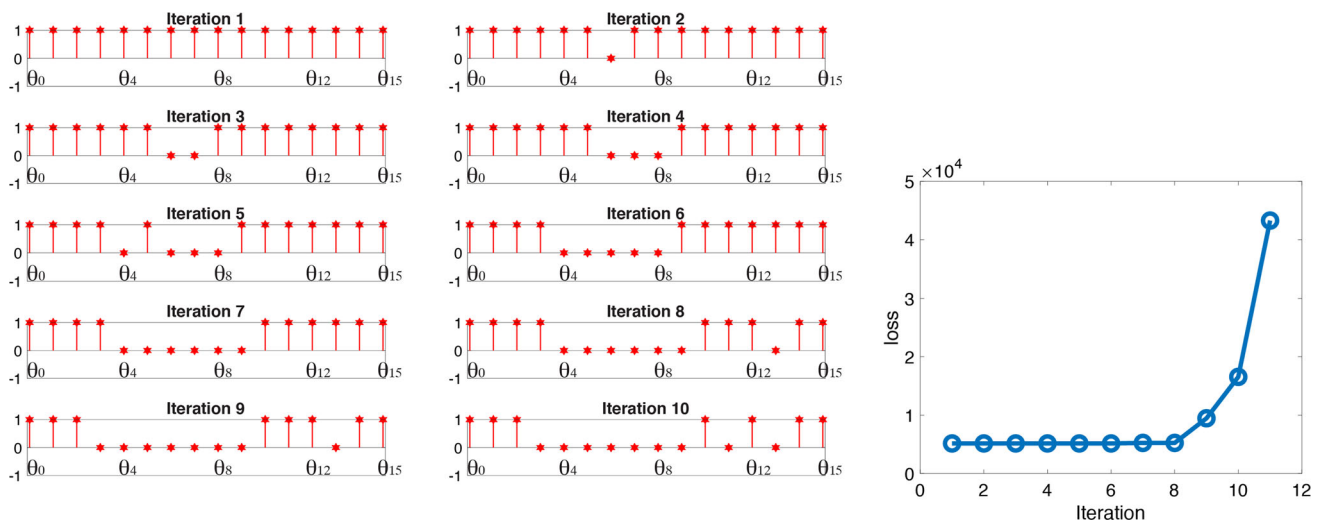


Fig. 4 Left: Stem-and-leaf plot illustrating system identification of active time-dependent SIRD parameters using data for the entire state of Michigan. Each stem-and-leaf represents one term of $\theta_0, \dots, \theta_{15}t^3$, scaled to 1 (active) or 0 (inactive). Right: The changing loss as terms are

eliminated from the set of time-dependent coefficients. System identification converges at Iteration 10 as the loss increases dramatically for further elimination of terms

against the data holds for all the inferred results presented in this communication, although the non-uniqueness of inverse problems does not preclude the existence of multiple sets of inferred coefficients.

Figures 6 and 7, respectively, illustrate the time-dependent SIRD coefficients and comparison between data and simulation using the inferred model (with the inferred ODE SIRD model) of the disease for Regions 1–8 delineated in Fig. 1. This is an important step toward a more fine-grained understanding of the geographical distribution of the disease in the state. The Southeastern part of the state is more heavily populated, especially Regions 2 and 3, which also bore the greatest burden of the disease. The city of Detroit, at the Western tip of Region 3, was the worst affected, reflecting its well-known socio-economic challenges. By contrast, Washtenaw County, about 50 km to the West, but also in Region 3, bore among the lowest burdens, per capita. At the risk of stating the obvious, we note that Regions 1–4, which account

for nearly 80% of the state’s population displayed very similar characteristics in the evolution of the data, as well as in SIRD coefficients and forward simulation results. We do not enter a more detailed analysis of these results here, deferring a different approach to spatial aspects of the spread of the disease to Sect. 7.

6.2 Deep and Bayesian neural networks

To infer the coefficients $\beta(t), \gamma(t), \mu(t), \alpha(t)$, we first compute the time derivatives of $S(t), I(t), R(t), D(t)$ by using the automatic differentiation API from TensorFlow. The coefficients are then computed by inverting Eqs. (1–4) at each time instant. For DNNs, we obtained deterministic results for all the coefficients. With BNNs, a Monte Carlo Sampling is performed to compute the mean and the standard deviation of the coefficients.

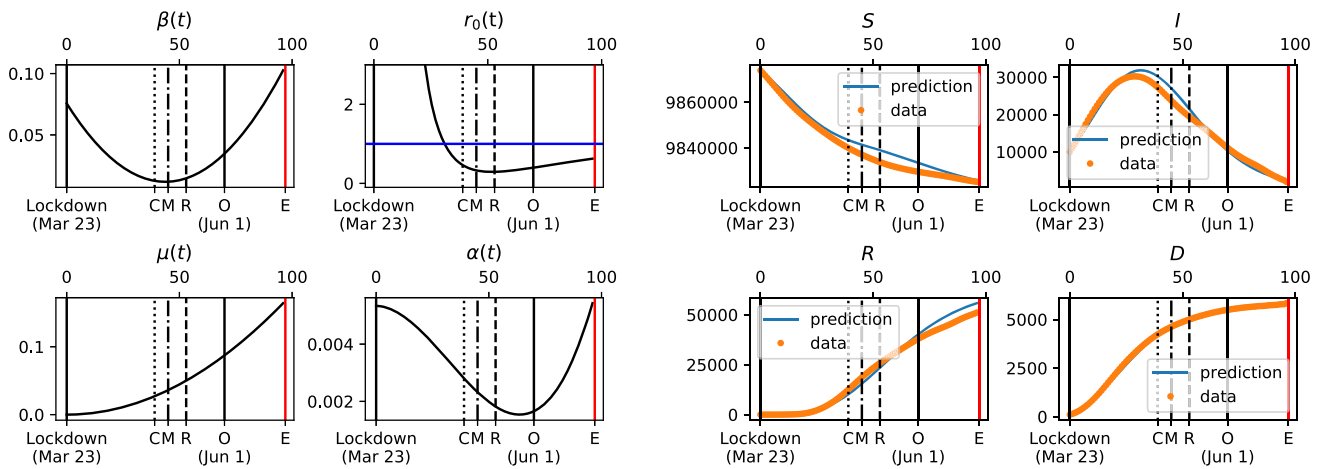


Fig. 5 Left: The time-dependent SIRD parameters after tuning by ODE-constrained optimization following system identification are: $\beta(t) = 0.0756 - 0.0029t + 3.33 \times 10^{-5}t^3$, $\gamma(t) = 0$, $\mu(t) =$

$1.78 \times 10^{-5}t^2$, $\alpha(t) = 0.0053 - 2.8 \times 10^{-6}t^2 + 2.93 \times 10^{-8}t^3$. Right: Simulation of the four compartments using the inferred ODE SIRD model, in comparison with the data

The constraint Eq. (5) is used to obtain S_m^d from I_m^d, R_m^d, D_m^d . This ensures that the discrete time derivatives in Eqs. (22–25) satisfy

$$\frac{S_m^d - S_{m-1}^d}{\Delta t} = -\frac{I_m^d - I_{m-1}^d}{\Delta t} - \frac{R_m^d - R_{m-1}^d}{\Delta t} - \frac{D_m^d - D_{m-1}^d}{\Delta t} \tag{33}$$

The constraint Eq. (5) also has been imposed in the DNN and BNN representations by training networks for I, R and D and then defining the network for S by this conservation of total population. Therefore, in using Eqs. (1–4) to invert the DNN/BNN representations for $\beta(t), \gamma(t), \mu(t), \alpha(t)$ at each time instant, a linear dependence is encountered: The summed left and right hand-sides of (2–4) exactly equal the left and right hand-side of (1), respectively. A unique solution for $\beta(t), \gamma(t), \mu(t), \alpha(t)$ is not possible due to linear dependence introduced by the population constraint. To circumvent this indeterminacy, we endow the system with additional information by requiring that $\gamma(t) = 0$. This represents the conferral of immunity on the recovered population, and importantly, is detected by our inference results using system identification and ODE-constrained minimization, as discussed in Sect. 6.

The inferred values, extended to a 30-day prediction (until July 28, 2020) for $\beta(t), \mu(t), \alpha(t), r_0(t)$ and $S(t), I(t), R(t), D(t)$ obtained from both DNNs and BNNs for Michigan are presented in Figs. 8 and 9, while the results for the eight Regions are given in Appendices “DNN results for different regions” and “BNN results for different regions”. One can observe that these time-dependent coefficients in Figs. 8a and 9a have a similar *initial* trend as those inferred by the system inference approach in Fig. 5. The effective repro-

duction number $r_0(t) < 1$ for $m > 30$ (April 23), in good agreement with its value obtained via system inference in Fig. 5. As polynomial approximation is used by the system inference approach, the inferred coefficients in Fig. 5 are very smooth, whereas inversion using the NN approach captures the detailed fluctuation of these coefficients, particularly, the rising infection rate after the open of the lockdown on June 1, 2020. In Fig. 9, the band around the inferred coefficients and the NN predictions shows the mean \pm one standard deviation of the corresponding results. Note the high standard deviation in parameters at early times, due to the noise in the data at small numbers. The regional results in Appendices “DNN results for different regions” and “BNN results for different regions” indicate that an accelerating infection rate for all the regions after the open of the lockdown. In particular, Region 7 and 8 have a predicted $r_0(t)$ value that is greater than 1. In addition, we observed that the BNN inferred coefficients in the regional results have a narrower range compared to those from the DNN.

More broadly, we note the difference in trends between the inferred time-dependent coefficients with the DNNs and BNNs in Figs. 8 and 9 in comparison with those in Fig. 5. This is due to the local inversion at each data point to infer the coefficients with the DNNs and BNNs versus the global optimization of losses for system inference in Sect. 6. As was referred to above, inverse problems allow non-unique solutions. It will be instructive to compare the predictions made by the DNN and BNN representations with the data when they become available.

A result that is consistent across all inference methods: system identification with ODE-constrained optimization, DNNs and BNNs, and for the state as a whole as well as its Regions is the following: The infection rate, $\beta(t)$, initially

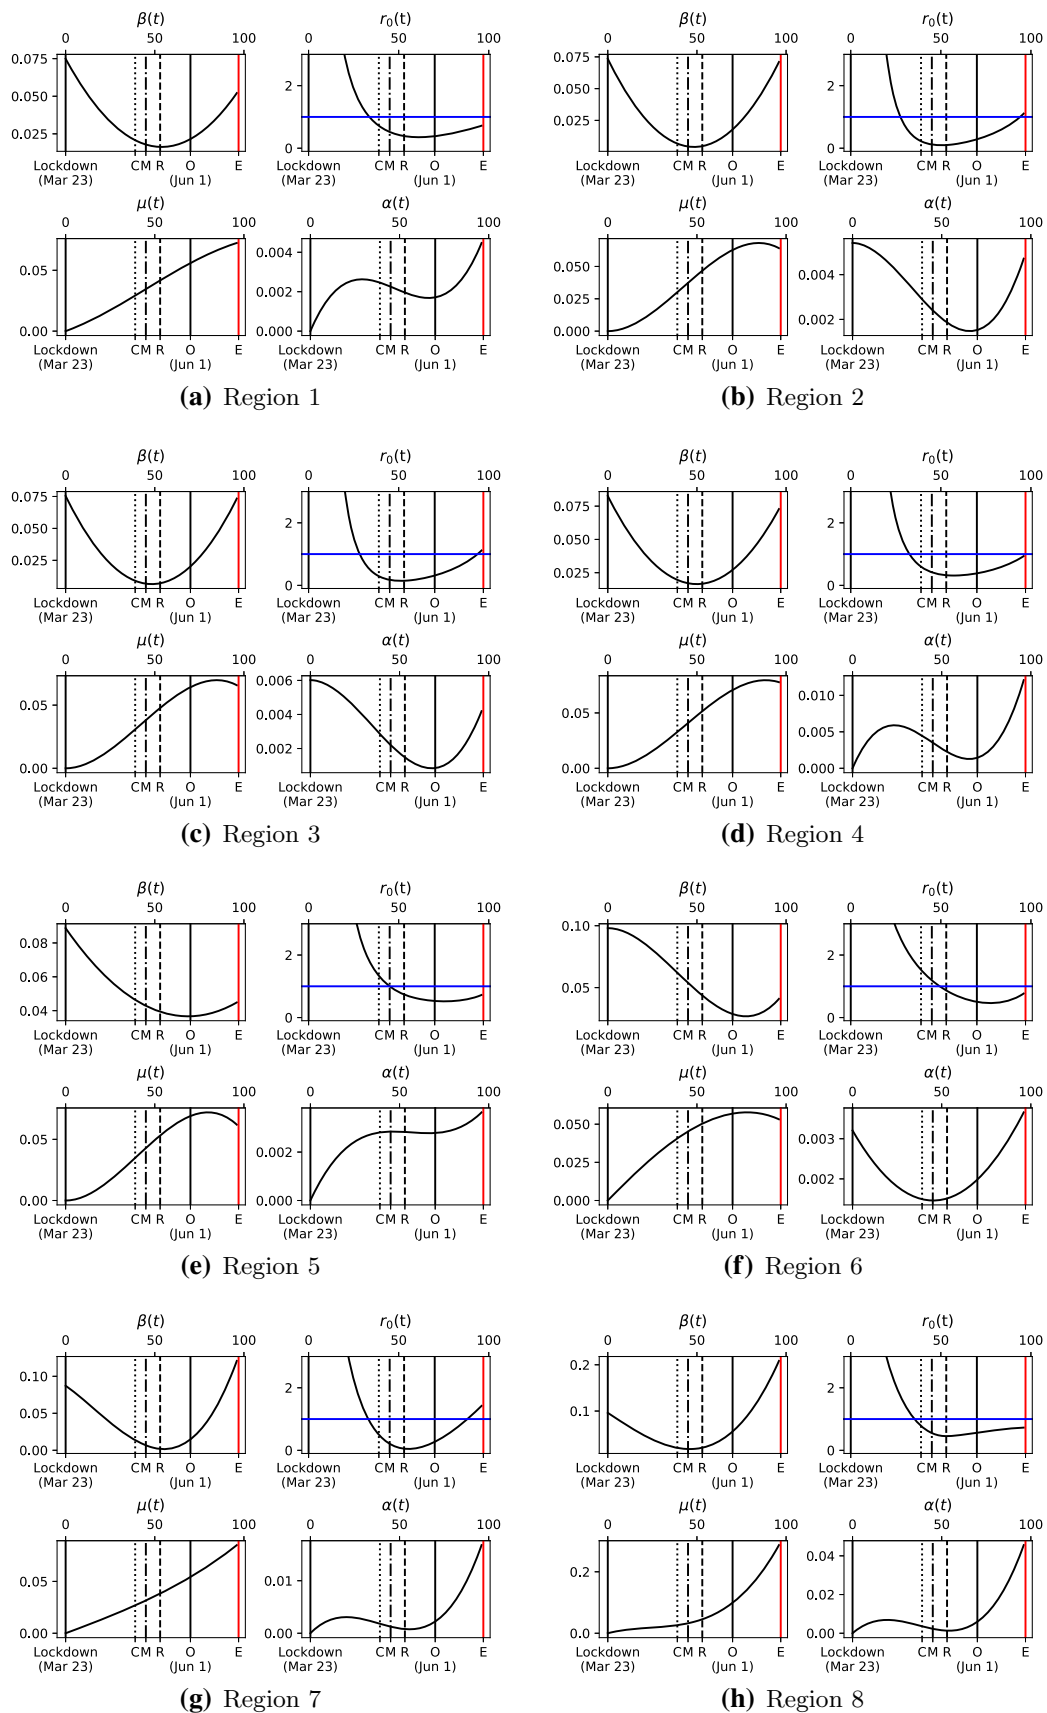


Fig. 6 Parameters of time-dependent SIRD coefficients, $\beta(t)$, $\mu(t)$, $\alpha(t)$, and the effective reproduction number, $r_0(t)$, for Regions 1–8 (see Fig. 1) of Michigan

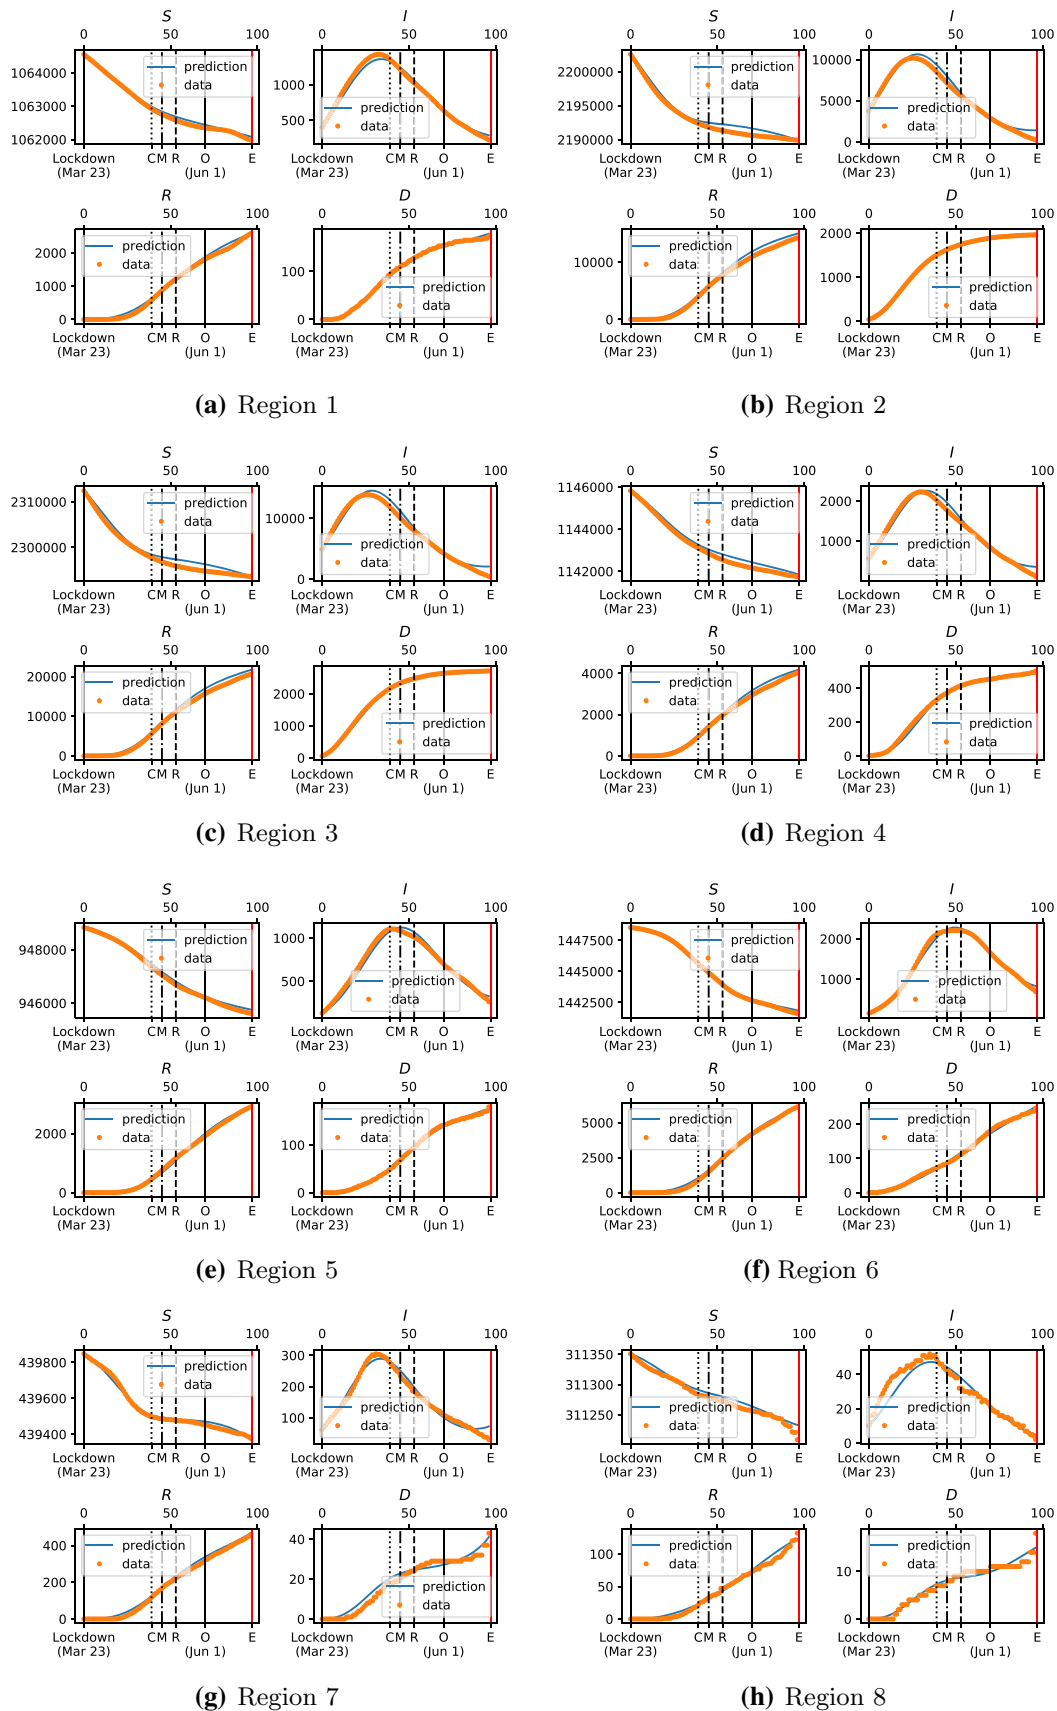


Fig. 7 Comparison of the simulation using inferred SIRD parameters (Fig. 6) for Regions 1–8 of Michigan

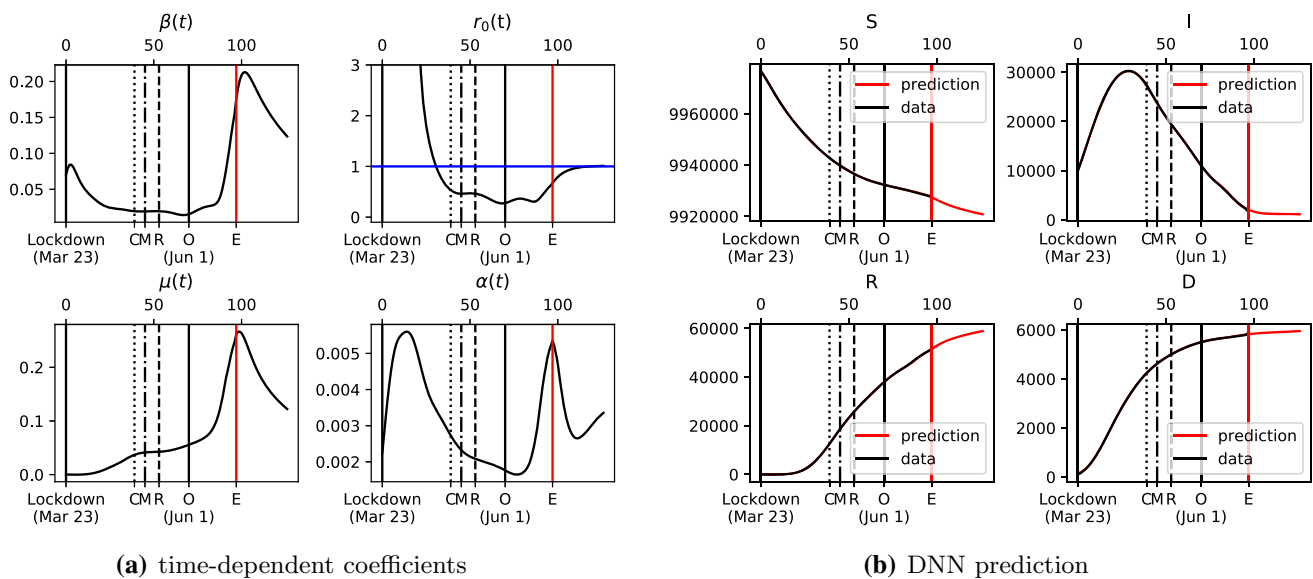


Fig. 8 **a** Time-dependent coefficients identified by DNNs, where an increased infection rate after the opening (O) of the lockdown on June 1st is observed. **b** DNNs learned $S(t)$, $I(t)$, $R(t)$, $D(t)$ based on the full extent of data points, and made a 30-day prediction

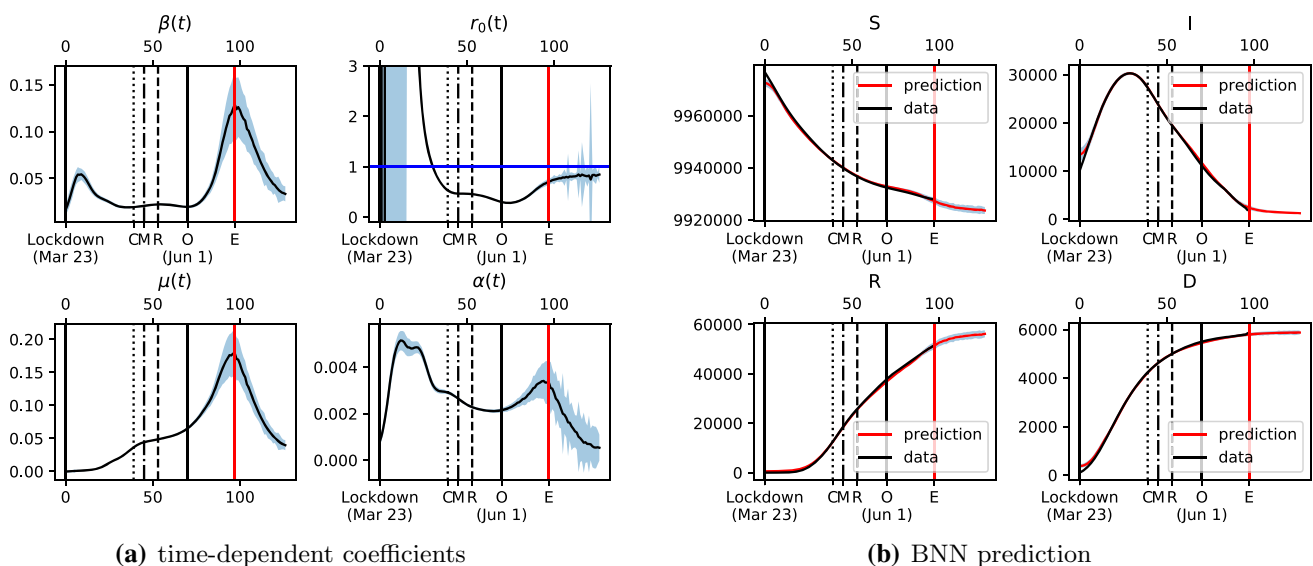


Fig. 9 **a** Time-dependent coefficients identified by BNNs, where an increased infection rate after the opening (O) on June 1st is observed. **b** BNNs learned $S(t)$, $I(t)$, $R(t)$, $D(t)$ based on the full extent of data points, and made a 30-day prediction. Bands correspond to \pm standard deviation over the mean

fell with the public health campaign, especially driven by the lockdown orders. However, it began to rise with the first step of opening (C), and even accelerated as more aspects of public, recreational, commercial and industrial activities were relaxed (M, R, O). Yet, every one of the versions of forward simulations with corresponding and consistently inferred systems matched very well with the data, which confirm that the state has largely controlled the pandemic, and continues to do so. As the number of remaining infected individuals, $I(t)$, has fallen steeply, there are fewer conveyors of infection, and even the higher $\beta(t)$ has not yet led to another explosion of

infection. This also can be seen by the sharply rising recovery rate, $\mu(t)$, and is verified by the effective reproduction number $r_0(t)$ falling below 1.0 after April 20 or 23 (the later date according to the DNN and BNN inference methods). A warning bell, however, must be rung as the results also indicate that $r_0(t) \rightarrow 1.0$ from below as we approach the end of our data and the time of writing. Michigan’s numbers for $I(t)$ are rising, although not yet exponentially. See Figs. 5, 6, 7, 8, 9, and Appendix sections “DNN results for different regions”, “BNN results for different regions”.

7 Two dimensional SIRD model with diffusion

Classical epidemiological models hold in the well-mixed limit, which is reflected in the compartments and sub-populations, S, I, R, D being total numbers over some geographical region. Spatial effects have been introduced by simply resolving smaller regions and treating them individually, as demonstrated here with our inference of SIRD coefficients over the regions of Michigan’s lower peninsula (Figs. 6, 7). However, while affording a spatially finer-grained treatment, this approach cannot, of course, address the mobility of the population. This is an important consideration, especially in light of the imposition and lifting of quarantines. In the COVID-19 Pandemic, the effects of social distancing, and the possibility of surges with their lifting revolve on the question of the time (and spatially) varying mobility of the population. At the finest resolution, this must be approached via agent-based models refined to resolve individuals. However, an intriguing question to explore is whether simple reaction–diffusion models can detect the evidence of mobility in these data. With our approach to model inference, we have access to methods of identifying mechanisms from data in which their action, while weak, may hold the key to important insights to the system. In this section, we embark down such a path, while noting that reaction–diffusion models of epidemiology have been considered previously from the perspective of analysis of the corresponding PDEs [22–24].

We now extend the SIRD model to PDEs in two spatial dimensions using the same compartments. However, the population variables are now replaced with spatio-temporally varying densities, $\widehat{S}(x, t), \widehat{I}(x, t), \widehat{R}(x, t), \widehat{D}(x, t)$ defined as numbers per unit area.

$$\frac{\partial \widehat{S}}{\partial t} = \mathcal{D}_S \nabla^2 \widehat{S} - \frac{\beta}{N} \widehat{S} \widehat{I} + \gamma \widehat{R} \tag{34}$$

$$\frac{\partial \widehat{I}}{\partial t} = \mathcal{D}_I \nabla^2 \widehat{I} + \frac{\beta}{N} \widehat{S} \widehat{I} - \mu \widehat{I} - \alpha \widehat{I} \tag{35}$$

$$\frac{\partial \widehat{R}}{\partial t} = \mathcal{D}_R \nabla^2 \widehat{R} + \mu \widehat{I} - \gamma \widehat{R} \tag{36}$$

$$\frac{\partial \widehat{D}}{\partial t} = \alpha \widehat{I} \tag{37}$$

Where $\mathcal{D}_S, \mathcal{D}_I, \mathcal{D}_R$ are diffusivities of the corresponding compartments, and represent the mobility of the population via random walks. We define $\widehat{(\bullet)} = (\bullet) / \int_{\Omega} dA$ where Ω is the domain of the lower peninsula of Michigan, to which we restrict our PDE SIRD studies. Furthermore the population constraint holds: $\int_{\Omega} \widehat{N} dA = \int_{\Omega} \widehat{S}(t) dA + \int_{\Omega} \widehat{I}(t) dA + \int_{\Omega} \widehat{R}(t) dA + \int_{\Omega} \widehat{D}(t) dA$.

7.1 Inference on the PDE form of the SIRD model

We adopt the weak form, and specifically, the finite element framework for inference on the above system of PDEs. For a generic, finite-dimensional field u^h , the problem is stated as follows: Find $u^h \in \mathcal{S}^h \subset \mathcal{S}$, where $\mathcal{S}^h = \{u^h \in \mathcal{H}^1(\Omega) \mid u^h = \bar{u} \text{ on } \Gamma^u\}$, such that $\forall w^h \in \mathcal{V}^h \subset \mathcal{V}$, where $\mathcal{V}^h = \{w^h \in \mathcal{H}^1(\Omega) \mid w^h = 0 \text{ on } \Gamma^u\}$, the finite-dimensional (Galerkin) weak form of the problem is satisfied. The variations w^h and trial solutions u^h are defined component-wise using a finite number of basis functions,

$$w^h = \sum_{a=1}^{n_b} c^a N^a, \quad u^h = \sum_{a=1}^{n_b} d^a N^a, \tag{38}$$

where n_b is the dimensionality of the function spaces \mathcal{S}^h and \mathcal{V}^h , and N^a represents the basis functions. To obtain the Galerkin weak forms, we multiply each strong form by the corresponding weighting function, use Backward Euler method for time-discretization, integrate by parts and apply boundary conditions appropriately, leading to:

$$\int_{\Omega} w_1^h \frac{\widehat{S}_m^h - \widehat{S}_{m-1}^h}{\Delta t} ds = - \int_{\Omega} \mathcal{D}_S \nabla w_1^h \cdot \nabla \widehat{S}_m^h ds - \int_{\Omega} w_1^h \left(\frac{\beta}{N} \widehat{S}_m^h \widehat{I}_m^h - \gamma \widehat{R}_m^h \right) ds \tag{39}$$

$$\int_{\Omega} w_2^h \frac{\widehat{I}_m^h - \widehat{I}_{m-1}^h}{\Delta t} ds = - \int_{\Omega} \mathcal{D}_I \nabla w_2^h \cdot \nabla \widehat{I}_m^h ds + \int_{\Omega} w_2^h \left(\frac{\beta}{N} \widehat{S}_m^h \widehat{I}_m^h - \mu \widehat{I}_m^h - \alpha \widehat{I}_m^h \right) ds \tag{40}$$

$$\int_{\Omega} w_3^h \frac{\widehat{R}_m^h - \widehat{R}_{m-1}^h}{\Delta t} ds = - \int_{\Omega} \mathcal{D}_R \nabla w_3^h \cdot \nabla \widehat{R}_m^h ds + \int_{\Omega} w_3^h \left(\mu \widehat{I}_m^h - \gamma \widehat{R}_m^h \right) ds \tag{41}$$

$$\int_{\Omega} w_4^h \frac{\widehat{D}_m - \widehat{D}_{m-1}}{\Delta t} ds = \int_{\Omega} w_4^h \alpha \widehat{I}_m^h ds \tag{42}$$

Where, boundary terms disappear because we assume that the populations do not cross the state boundary, or into the upper peninsula. The system identification problem is to infer the time-dependent coefficients $\mathcal{D}_S(t), \mathcal{D}_I(t), \mathcal{D}_R(t)$, and we also choose to expand them in a polynomial basis

$$\mathcal{D}_s(t) = \theta_{16} + \theta_{17}t + \theta_{18}t^2 + \theta_{19}t^3 \tag{43}$$

$$\mathcal{D}_i(t) = \theta_{20} + \theta_{21}t + \theta_{22}t^2 + \theta_{23}t^3 \tag{44}$$

$$\mathcal{D}_r(t) = \theta_{24} + \theta_{25}t + \theta_{26}t^2 + \theta_{27}t^3 \tag{45}$$

along with the time-dependent coefficients $\beta(t), \gamma(t), \mu(t), \alpha(t)$ shown in Eq. (10–13). We expect that the effect of mobility on the evolution of population densities is small over the

course of the COVID-19 Pandemic. However, our interest is in inferring the presence of this effect in the data following the relaxation of lockdown orders. In order to identify the diffusivities despite the expected dominance of the reaction terms in the data obeying Eqs. (42), we adopt two stage Variational System Identification [13].

We define Stage 1 by choosing $w_i^h = 1, i = 1, \dots, 4$, yielding:

$$\int_{\Omega} \frac{\widehat{S}_m^h - \widehat{S}_{m-1}^h}{\Delta t} dA = -\beta \int_{\Omega} \frac{1}{\widehat{N}} \widehat{S}^h \widehat{I}^h ds - \gamma \int_{\Omega} \widehat{R}^h dA \tag{46}$$

$$\int_{\Omega} \frac{\widehat{I}_m^h - \widehat{I}_{m-1}^h}{\Delta t} ds = \beta \int_{\Omega} \frac{1}{\widehat{N}} \widehat{S}^h \widehat{I}^h ds - \int_{\Omega} \mu \widehat{I}^h ds - \alpha \int_{\Omega} \widehat{I}^h dA \tag{47}$$

$$\int_{\Omega} \frac{\widehat{R}_m^h - \widehat{R}_{m-1}^h}{\Delta t} ds = \mu \int_{\Omega} \widehat{I}^h ds - \gamma \int_{\Omega} \widehat{R}^h ds \tag{48}$$

$$\int_{\Omega} \frac{\widehat{D}_m^h - \widehat{D}_{m-1}^h}{\Delta t} ds = \alpha \int_{\Omega} \widehat{I}^h dA \tag{49}$$

The diffusion operators vanish since, for a constant weighting function, $\nabla w = 0$. In order to avoid a proliferation of superscripts and subscripts, we simply denote the data interpolated over the finite element mesh at time m by $(\bullet)_m^d$, dispensing with the superscript $(\bullet)^h$ for the finite-dimensional fields. The label vector and matrix of bases can be constructed as:

$$y_m = \begin{Bmatrix} \int_{\Omega} \frac{\widehat{S}_m^d - \widehat{S}_{m-1}^d}{\Delta t} dA \\ \int_{\Omega} \frac{\widehat{I}_m^d - \widehat{I}_{m-1}^d}{\Delta t} dA \\ \int_{\Omega} \frac{\widehat{R}_m^d - \widehat{R}_{m-1}^d}{\Delta t} dA \\ \int_{\Omega} \frac{D_m^d - D_{m-1}^d}{\Delta t} dA \end{Bmatrix} \tag{50}$$

$$\Xi_m = \begin{bmatrix} \int_{\Omega} \frac{\widehat{S}_m^d \widehat{I}_m^d}{N} ds \langle 1 t_m t_m^2 t_m^3 \rangle & \int_{\Omega} -\widehat{R}_m^d ds \langle 1 t_m t_m^2 t_m^3 \rangle & \langle 0 0 0 0 \rangle & \langle 0 0 0 0 \rangle \\ \int_{\Omega} -\frac{\widehat{S}_m^d \widehat{I}_m^d}{N} ds \langle 1 t_m t_m^2 t_m^3 \rangle & \langle 0 0 0 0 \rangle & \int_{\Omega} \widehat{I}_m^d ds \langle 1 t_m t_m^2 t_m^3 \rangle & \int_{\Omega} I_m ds \langle 1 t_m t_m^2 t_m^3 \rangle \\ \langle 0 0 0 0 \rangle & \int_{\Omega} \widehat{R}_m^d ds \langle 1 t_m t_m^2 t_m^3 \rangle & \int_{\Omega} -\widehat{I}_m^d ds \langle 1 t_m t_m^2 t_m^3 \rangle & \langle 0 0 0 0 \rangle \\ \langle 0 0 0 0 \rangle & \langle 0 0 0 0 \rangle & \langle 0 0 0 0 \rangle & \int_{\Omega} -\widehat{I}_m^d ds \langle 1 t_m t_m^2 t_m^3 \rangle \end{bmatrix} \tag{51}$$

Once the reaction terms are identified, we return to the original weak forms Eqs. (39–42). Accounting for the arbitrariness of w^h in \mathcal{V}^h , the finite-dimensionality leads to a system of residual equations for each degree of freedom (DOF):

$$\mathcal{R}_i = \mathcal{F}_i \left(S_{m-1}^d, S_m^d, \nabla S_m^d, \dots, D_s \langle 1 t_m t_m^2 t_m^3 \rangle, \dots, N, \nabla N \dots \right), \tag{52}$$

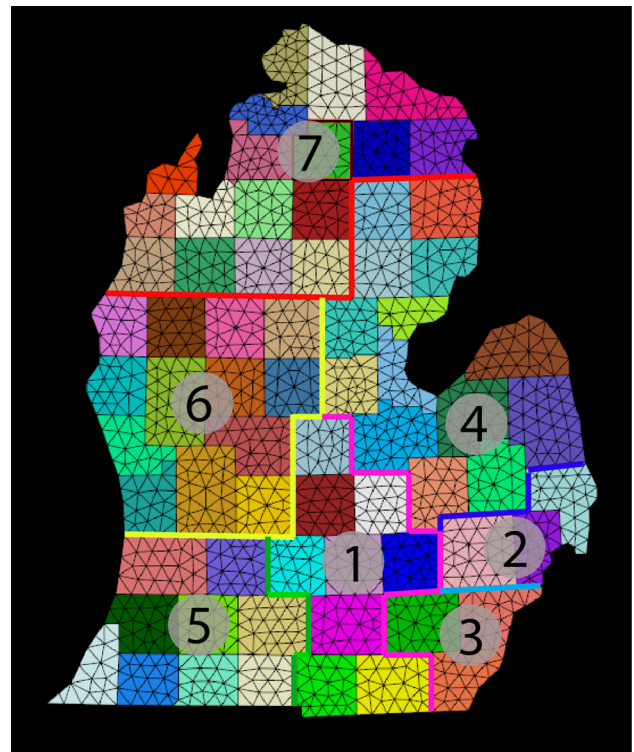


Fig. 10 A finite element mesh of the map of Michigan delineating the counties. Only Regions 1–7 were used in the PDE inference problem

where \mathcal{R}_i is the i th component of the residual vector. The diffusion terms can then be identified by the two stage approach to Variational System Identification detailed in [12].

7.2 Data preparation on the 2D map of Michigan

We first construct a two-dimensional mesh that fully resolves the counties as shown in Fig. 10. Recall that only the lower peninsula, consisting of Regions 1–7 was included in the PDE inference problem. The data are available as cumulative sub-population numbers I_m^d, R_m^d, D_m^d at the county level (Michigan’s lower peninsula has 68 counties). We use a uni-

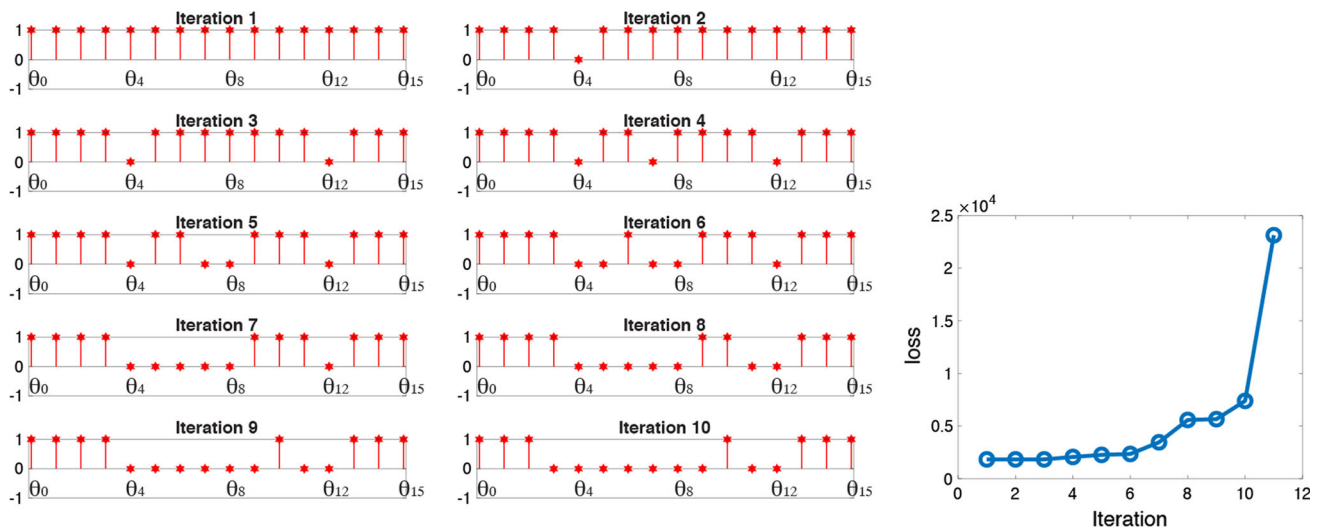


Fig. 11 Left: stem-and-leaf plot illustrating system identification of active reaction parameters in the PDE SIRD model in Stage 1 of Variational System Identification. Each stem and leaf represents one term of $\theta_0, \dots, \theta_{15}t^3$, scaled to 1 (active) or 0 (inactive). Right: The changing loss as terms are eliminated from the set of time-dependent coefficients. System identification converges at Iteration 10 as the loss increases dramatically for further elimination of terms

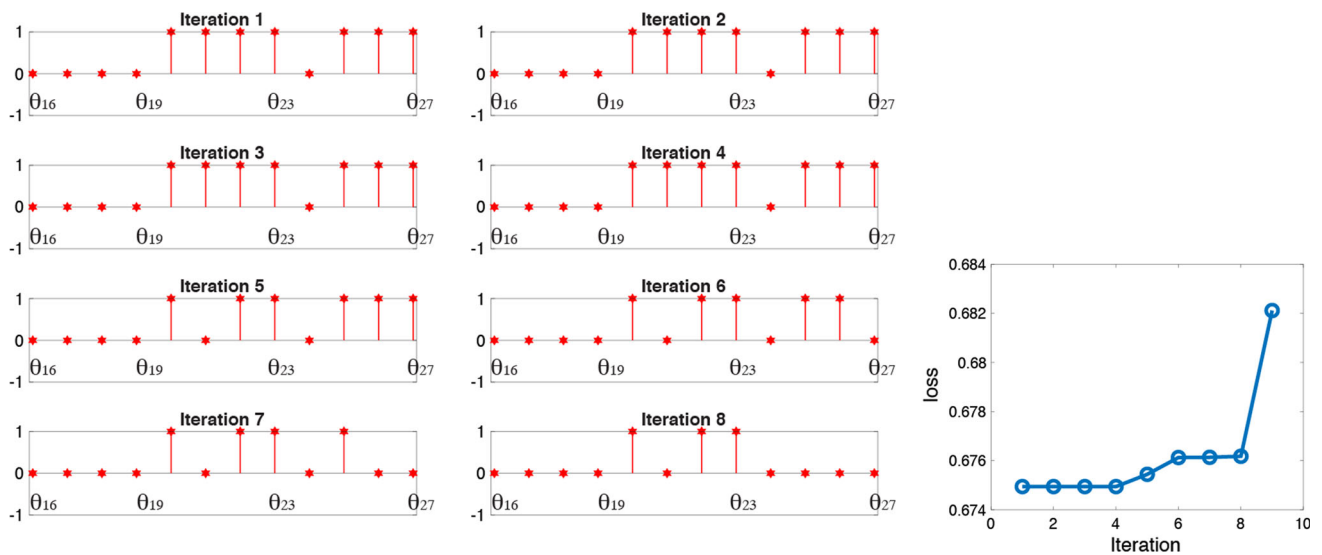


Fig. 12 Left: stem-and-leaf plot illustrating system identification of active diffusion parameters in the PDE SIRD model in Stage 2 of Variational System Identification. Each stem and leaf represents one term of $\theta_{16}, \dots, \theta_{27}t^3$, scaled to 1 (active) or 0 (inactive). Right: The changing loss as terms are eliminated from the set of time-dependent coefficients. System identification converges at Iteration 8 as the loss increases dramatically for further elimination of terms

form density of each sub-population to compute $\widehat{I}_m^d, \widehat{R}_m^d, \widehat{D}_m^d$ within the county, and applied Gaussian filtering to smooth the discontinuities between counties. Note that the discrete Gaussian filter can not be applied in a straightforward manner to unstructured meshes. Here we start with continuous Gaussian filtering over the infinite domain:

$$u(\mathbf{x}_0) = \int_{-\infty}^{\infty} G(\mathbf{x}_0, \mathbf{x}) u_{\text{raw}}(\mathbf{x}) dv \tag{53}$$

$$= \int_{\Omega} G(\mathbf{x}_0, \mathbf{x}) u_{\text{raw}}(\mathbf{x}) dv \tag{54}$$

where u could be any of the four sub-population densities, and $G(\mathbf{x}_0, \mathbf{x}) = \frac{1}{2\pi\sigma^2} e^{-\frac{\|\mathbf{x}\|^2}{2\sigma^2}}$ is the two dimensional Gaussian distribution function. The parameter σ is the standard deviation of the Gaussian distribution which is related to the kernel size in the discrete Gaussian filter. Since $\int_{\Omega} G dA < 1$ we scale up the filtered displacement at each node:

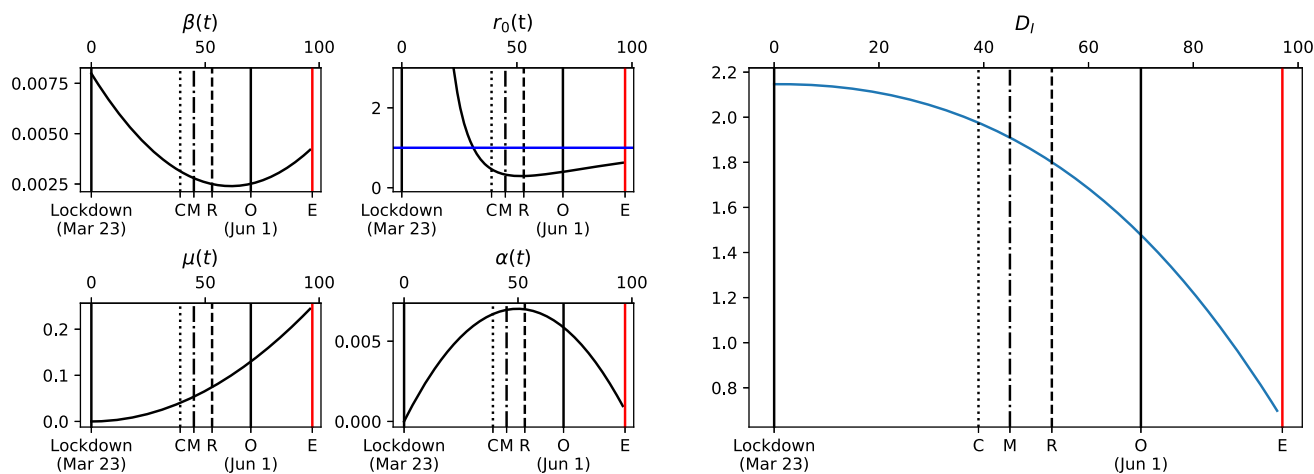


Fig. 13 Left: The time-dependent reaction parameters in the 2D SIRD model after tuning by PDE-constrained optimization: $\beta(t) = 0.00798 - 1.82 \times 10^{-4}t + 1.49 \times 10^{-6}t^2$, $\gamma(t) = 0$, $\mu(t) = 2.65 \times 10^{-5}t^2$, $\alpha(t) = 2.82 \times 10^{-4}t - 2.83 \times 10^{-6}t^2$. Right: The sole time-dependent diffusivity is for the infected sub-population: $\mathcal{D}_I = 2.146 - 8.12 \times 10^{-5}t^2 - 7.914 \times 10^{-7}t^3$

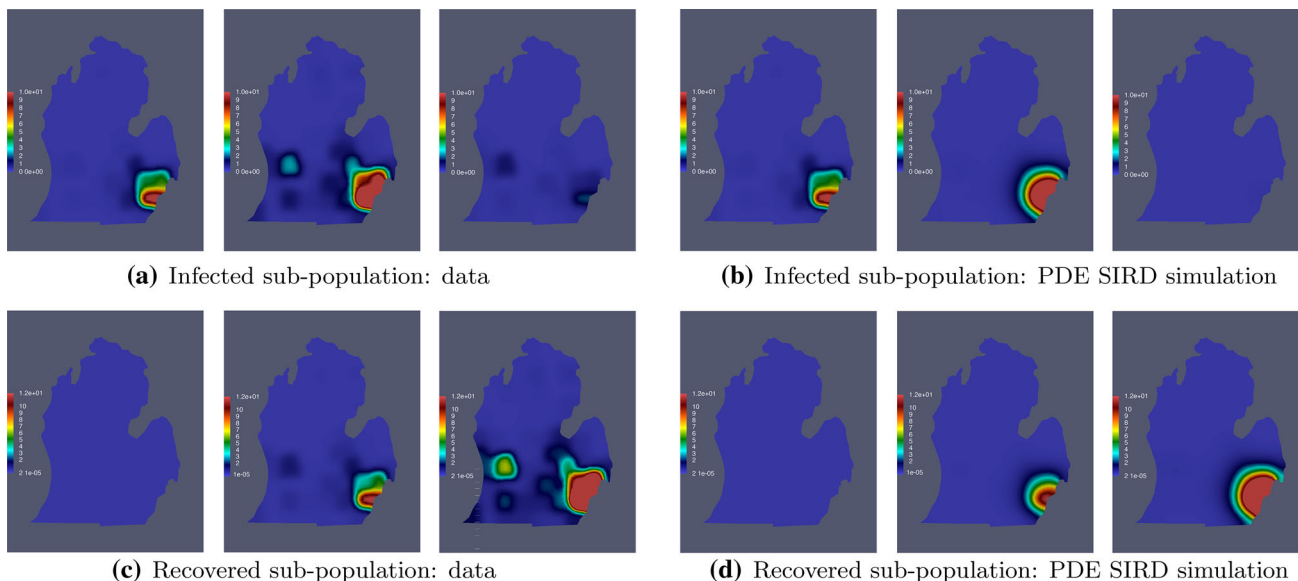


Fig. 14 Comparison of the data on distributions of the infected (a) and recovered (c) sub-populations against forward PDE SIRD simulations with inferred quantities, (b) and (d), respectively. Data and simulation results are shown for Day 0 (lockdown, March 23), Day 40 (May 2, when the infected sub-population had its greatest spread across the state, but still restricted to Southeastern Michigan) and Day 96 (end of our data range, June 28, 2020)

$$u(\mathbf{x}_0) = \frac{\int_{-\infty}^{\infty} G(\mathbf{x}_0, \mathbf{x}) dv}{\int_{\Omega} G(\mathbf{x}_0, \mathbf{x}) dv} \int_{\Omega} G(\mathbf{x}_0, \mathbf{x}) u_{\text{raw}}(\mathbf{x}) dv \quad (55)$$

$$= \frac{1}{\int_{\Omega} G(\mathbf{x}_0, \mathbf{x}) dv} \int_{\Omega} G(\mathbf{x}_0, \mathbf{x}) u_{\text{raw}}(\mathbf{x}) dv \quad (56)$$

The spatio-temporal evolution of these fields was used in PDE inference via two-stage Variational System Identification as described in Sect. 7.1 followed by optimization constrained by the PDEs in (39–42) using adjoints. Stem-and-leaf plots and the losses for Stage 1 of Variational System

Identification appear in Fig. 11. Recall that in this stage only the reaction terms $\beta(t)$, $\gamma(t)$, $\mu(t)$, $\alpha(t)$ are identified. These inference results for active coefficients should be compared with the ODE SIRD model in Fig. 5. This is followed by Stage 2 of Variational System Identification with stem-and-leaf plots and losses appearing in Fig. 12. Note that the diffusivities of the susceptible and recovered populations, $\mathcal{D}_S = 0$ and $\mathcal{D}_R = 0$. However, the infected population has a time-varying diffusivity \mathcal{D}_I that declines.

7.3 Results of system identification of two dimensional SIRD model with diffusion

Figure 13 shows the inference (two stage Variational System Identification followed by PDE-constrained optimization) for the coefficients $\beta(t)$, $\mu(t)$, $\alpha(t)$, the effective reproduction number, $r_0(t)$ as well as the diffusivity $\mathcal{D}_I(t)$ in the PDE SIRD model. On comparing with Fig. 5 some differences are revealed in the time dependence of $\beta(t)$, $\mu(t)$, $r_0(t)$, $\alpha(t)$. This is to be expected in adopting the PDE SIRD model over the ODE form. The inference of time-dependent diffusion in the mobility of the infected sub-population, \mathcal{D}_I , naturally affects the other quantities. While the preliminary nature of these results warrants caution, it is worth noting the inference of decreasing mobility of the infected sub-population in \mathcal{D}_I . Figure 14 compares data and the forward simulation with inferred quantities for the distribution of the infected and recovered sub-populations on days corresponding to the initial lockdown, the maximum spread of the infected sub-population (May 2), and at the end of our data collection. Notably, the restriction of the high density of the infected population, \hat{I} to Southeastern Michigan reflects the success, to date, of the state's public health response. While the correspondence is reasonable, the statewide sub-populations $S(t)$, $I(t)$, $R(t)$, $D(t)$ obtained by integrating the corresponding densities over the lower peninsula, show a poorer match in Fig. 15. While the trends are reproduced, there are notable errors over time. A major improvement is possible in the PDE SIRD model by allowing the coefficients β , γ , μ , α , \mathcal{D}_S , \dots , \mathcal{D}_R to also vary over space. This would allow better representation of the system, in keeping with the inferred difference in $\beta(t)$, $\mu(t)$, $r_0(t)$, $\alpha(t)$ over the eight Regions in Fig. 6, which led to the excellent agreement between data and the forward ODE SIRD simulations in Fig. 7. From a purely data representation standpoint, the greater number of parameters will allow lower errors.

The code used for the inference, machine learning and forward simulations is available in the `mechanoChem` and `mechanoChemML` libraries at <https://github.com/mechanoChem/>.

8 Conclusion

We have brought machine learning inference techniques to bear upon the data on progression of COVID-19 across the state of Michigan by applying three distinct approaches: (a) Our methods of system identification to delineate the operational mechanisms, followed by (b) adjoint-based model-constrained optimization for refinement of the parameters, and (c) deep and Bayesian neural networks. Our interest in this study has been two-fold.

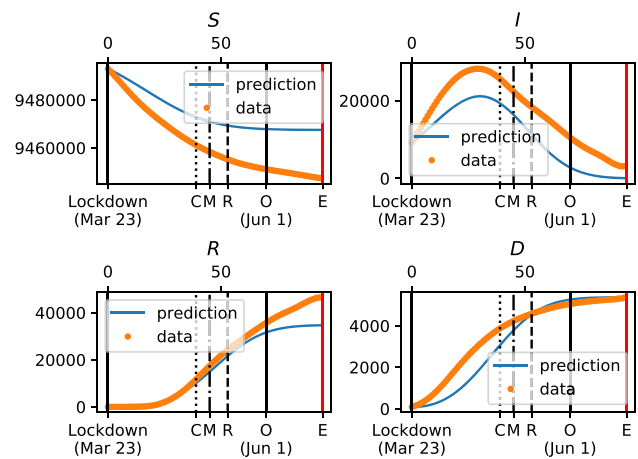


Fig. 15 Simulation of the four compartments using the inferred PDE SIRD model, in comparison with the data

The first has been to seek to infer the time-dependence of the coefficients in the classical ODE SIRD model, motivated by the evolving characteristics of testing, quarantine and treatment protocols over the 97-day course of the pandemic as reflected in the data. As discussed in Sect. 6, our inference methods reveal the course of rates of infection, recovery and death over the state and its eight regions, assuming uniform mixing in each case. Notably, our methods suggest that recovery confers immunity, but we hasten to add that this is a very preliminary conclusion. More detailed and fine-grained studies need to be undertaken to verify it, and of course, immunology will have the final say here. Also of note are our conclusions that while the infection rate has increased after an initial decline, as the state relaxed restrictions, the lower numbers of infectious individuals has meant a lower overall extent of transmission. This is also seen in the effective reproduction rate, which, while below one, has trended dangerously closer to that threshold of exponential growth. The uncertainty in our inference, given the data, is reflected in the results of the Bayesian neural networks in the same section. Of some interest here are the predictions made by BNNs for 30 days beyond the end of the data we have considered; that is until July 28, 2020.

The second facet of our interest is to try and infer spatial dependence by extending the SIRD models to PDEs by incorporating the population's mobility via diffusion. This is a different, and potentially intriguing, approach that complements the resolution of the problem down to the smaller Regions of the state as we did with the ODE SIRD model. On this front, we note that the inference needs to be extended to our methods of two-stage Variational System Identification followed by PDE-constrained optimization. Here, it is of note that the susceptible and recovered populations were found to have vanishing diffusivities (mobilities), while the infected population had a diffusivity that declined over the 97-day extent of the data that we used. This first extension to sys-

tem inference of the PDE SIRD model returned reasonable comparisons with data on distributions of the sub-population densities, although the total numbers integrated over the state were not as well reproduced. As suggested by the notable differences in the ODE SIRD model coefficients for the eight regions of the state, the PDE SIRD model with spatially varying coefficients may be a better representation. This will allow us to make connections with the well-known metapopulation variants [20,21] of the network-based SIR model that account for the mobility of individuals or groups, and thereby represent spatially and group-wise varying diffusivity. Building on these initial results, we see many possibilities for analysis and prediction of the future course and geographical spread of the COVID-19 Pandemic using the PDE SIRD model.

Acknowledgements We acknowledge the support of Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR0011199002, “Artificial Intelligence guided multi-scale multi-physics framework for discovering complex emergent materials phenomena”.

Appendix: Additional regional results

This appendix contains the inferred time-dependent coefficients and the NN prediction results for the eight regions of the Michigan state.

A.1 DNN results for different regions

See Figs. 16 and 17.

A.2 BNN results for different regions

See Figs. 18 and 19.

Fig. 16 Regions 1–8: Time-dependent coefficients identified by DNNs, where an increased infection rate after the opening (O) of lockdown on June 1st is observed

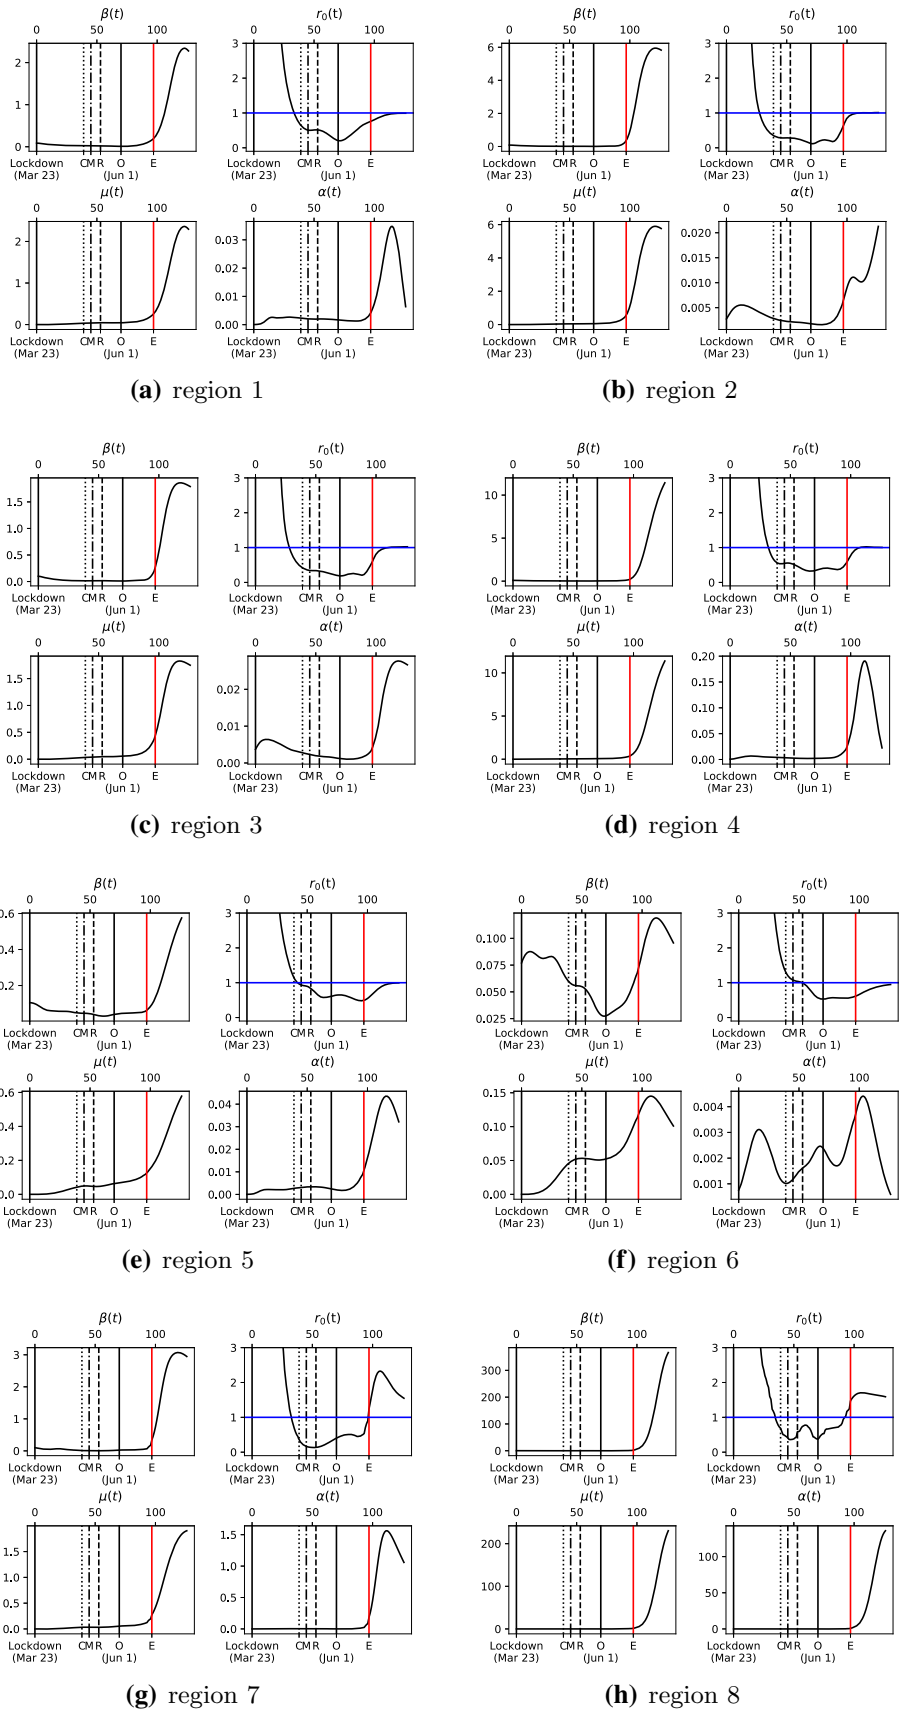


Fig. 17 Regions 1–8: DNNs learned $S(t)$, $I(t)$, $R(t)$, $D(t)$ based on the existing discrete data point, where a 30-day prediction is made by DNNs

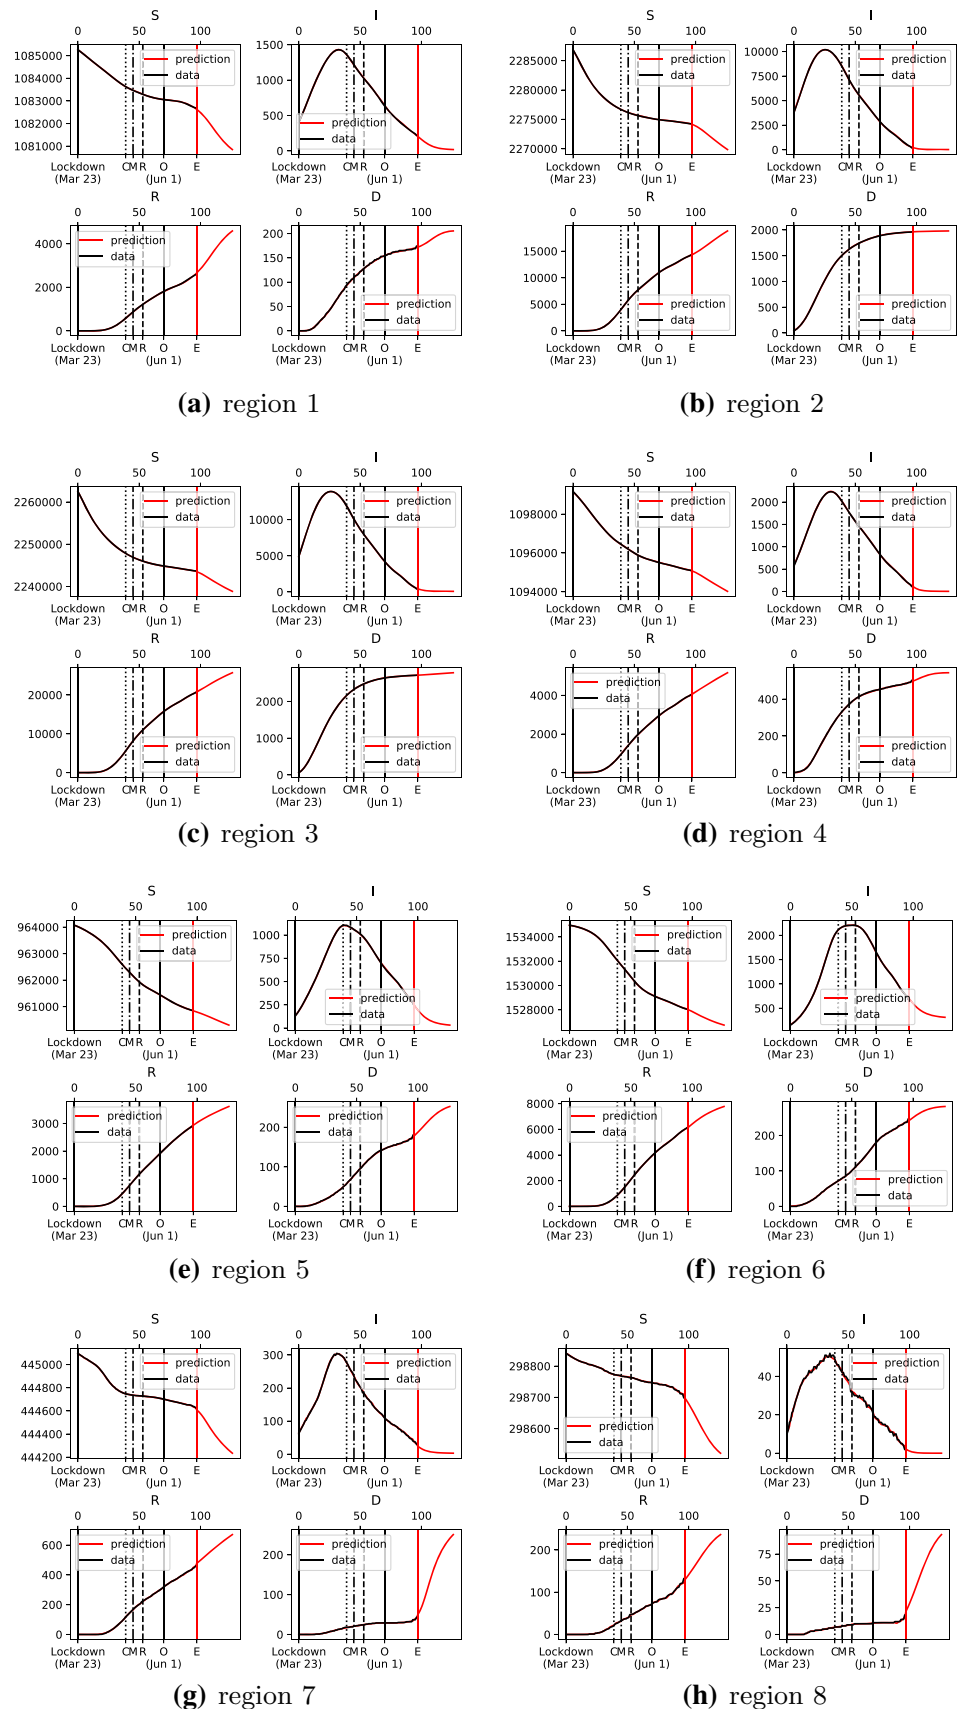


Fig. 18 Regions 1–8: Time-dependent coefficients identified by BNNs, where an increased infection rate after the opening (O) of lockdown on June 1st is observed. Bands correspond to \pm standard deviation over the mean

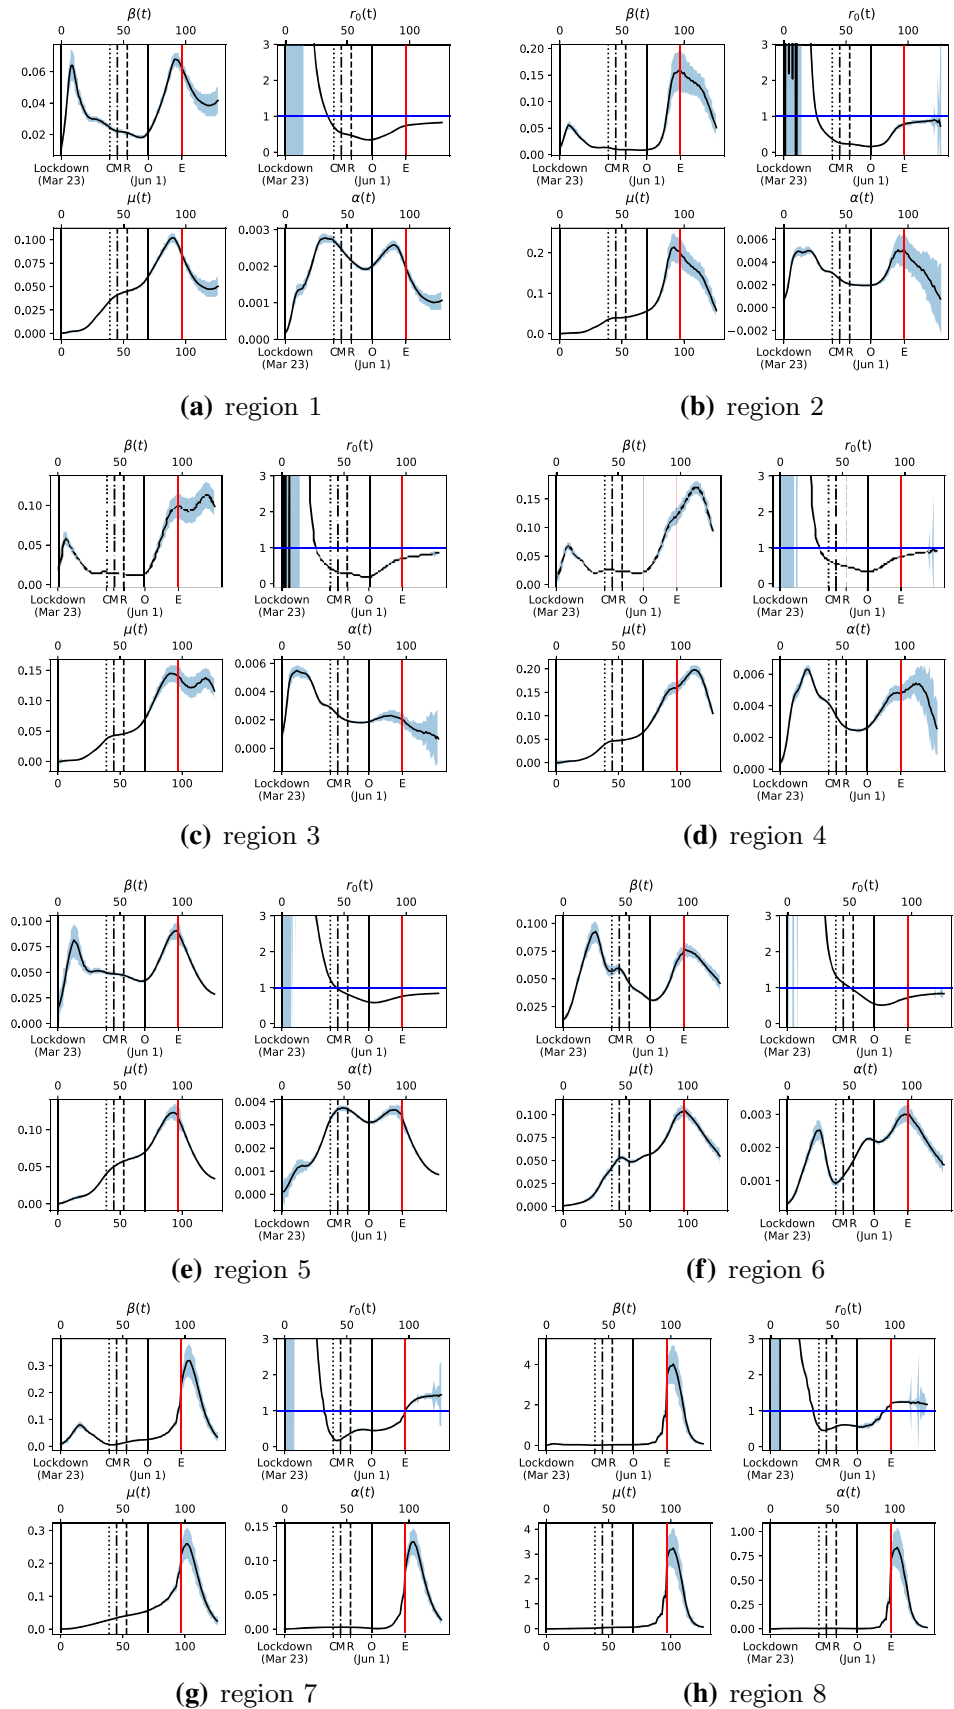
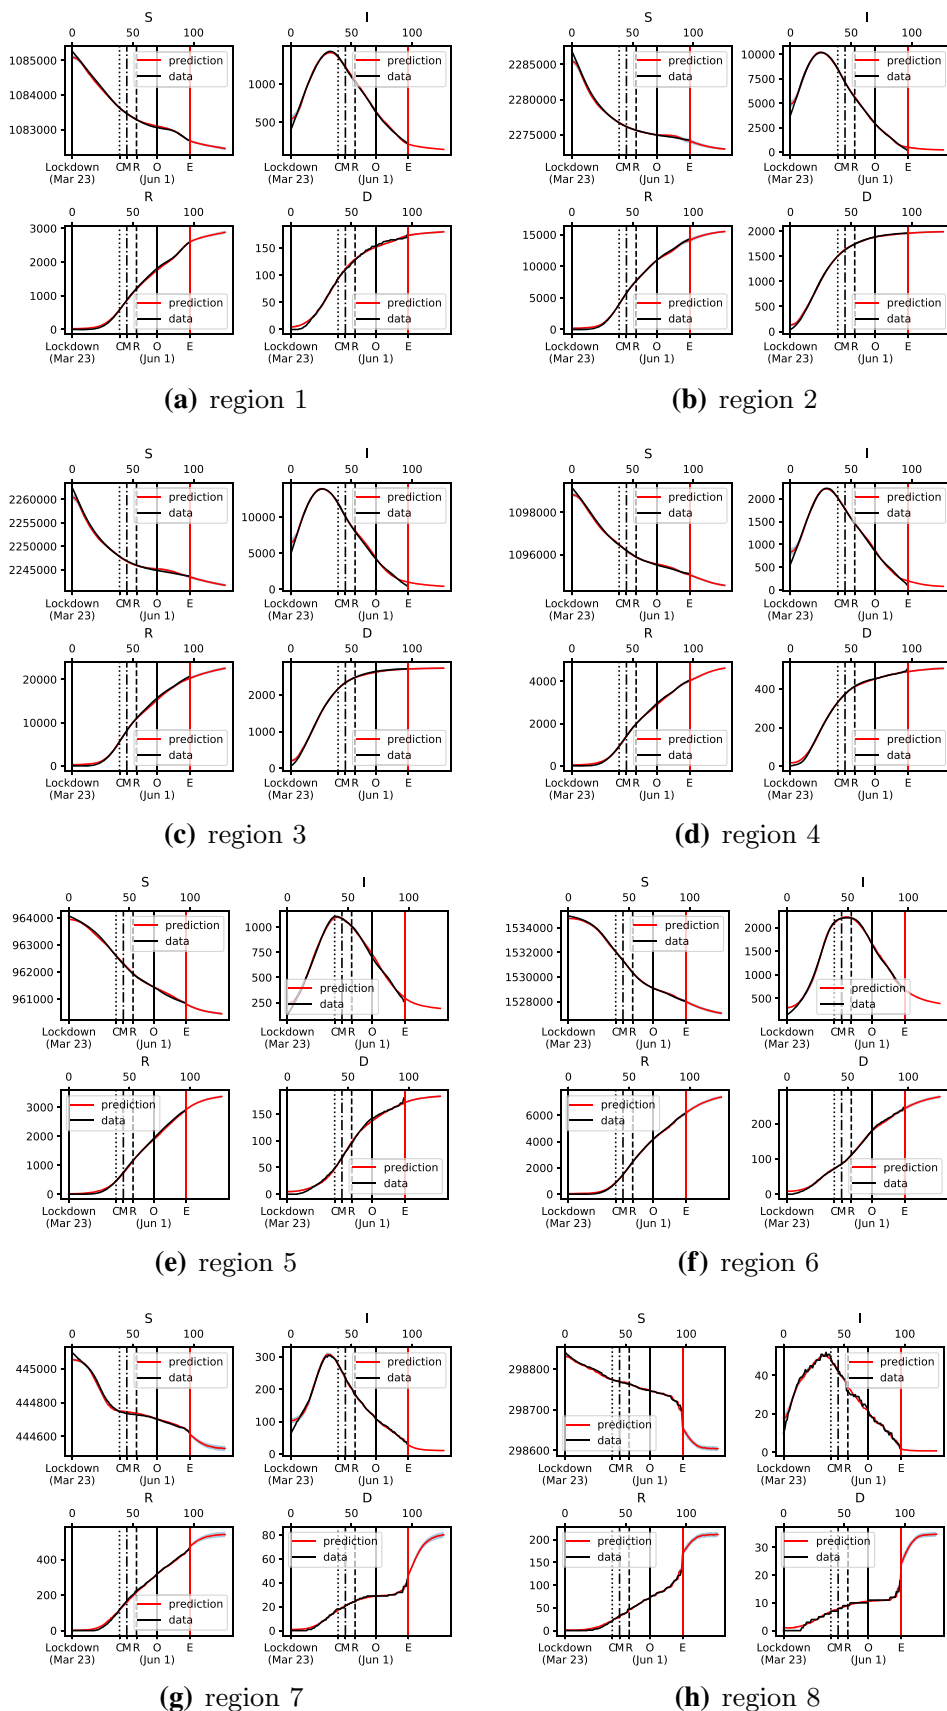


Fig. 19 Region 1–8: BNNs learned $S(t)$, $I(t)$, $R(t)$, $D(t)$ based on the existing discrete data point, where a 30-day prediction is made by BNNs. Bands correspond to \pm standard deviation over the mean



References

- Kermack WO, McKendrick AG (1927) A contribution to the mathematical theory of epidemics. *Proc R Soc Lond Ser A* 115:700–721
- Eisenberg MC, Eisenberg JNS, D'Silva JP, Wells EV, Cherng S, Kao Y-H, Meza R (2015) Forecasting and uncertainty in modeling the 2014–2015 ebola epidemic in West Africa
- Eisenberg M, Kujbida G, Tuite AR, Fisman DN, Tien JH (2013) Examining rainfall and cholera dynamics in haiti using statistical and dynamic modeling approaches. *Epidemics* 5:197–207. <https://doi.org/10.1016/j.epidem.2013.09.004>
- Wesolowski A, Eagle N, Tatem AJ, Smith DL, Noor AM, Snow RW, Buckee CO (2012) Quantifying the impact of human mobility on malaria. *Science* 338:267–270. <https://doi.org/10.1126/science.1223467>
- Colizza V, Barrat A, Barthelemy M, Valleron A-J, Vespignani A (2007) Modeling the worldwide spread of pandemic influenza: baseline case and containment interventions. *PLoS Med* 4:e13. <https://doi.org/10.1371/journal.pmed.0040013>
- 1Point3Acres.com. <https://coronavirus.1point3acres.com/en>
- Yang T, Shen K, He S, Li E, Sun P, Chen P, Zuo L, Hu J, Mo Y, Zhang W, Zhang H, Chen J, Guo Y (2020) Covidnet: To bring data transparency in the era of covid-19
- Johns Hopkins University of Medicine. COVID-19 dashboard by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (JHU). <https://coronavirus.jhu.edu/map.html>
- Michigan State Coronavirus Data. <https://www.michigan.gov/coronavirus/>
- The New York Times. Coronavirus in the U.S.: Latest Map and Case Count—The New York Times. <https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html>
- The Institute for Health Metrics and Evaluation. COVID-19 Projections. <https://covid19.healthdata.org/united-states-of-america>
- Wang Z, Huan X, Garikipati K (2019) Variational system identification of the partial differential equations governing the physics of pattern-formation: Inference under varying fidelity and noise. *Comput Methods Appl Mech Eng* 356:44–74. <https://doi.org/10.1016/j.cma.2019.07.007>
- Wang Z, Huan X, Garikipati K. Identification of the partial differential equations governing microstructure evolution in materials: inference over incomplete, sparse and spatially non-overlapping data. [arXiv:2001.04816](https://arxiv.org/abs/2001.04816)
- Teichert G, Garikipati K (2018) Machine learning materials physics: surrogate optimization and multi-fidelity algorithms predict precipitate morphology in an alternative to phase field dynamics. *Comput Methods Appl Mech Eng* (to appear)
- Teichert GH, Natarajan AR, Van der Ven A, Garikipati K (2019) Machine learning materials physics: integrable deep neural networks enable scale bridging by learning free energy functions. *Comput Methods Appl Mech Eng* 353:201–216. <https://doi.org/10.1016/j.cma.2019.05.019>
- Teichert GH, Natarajan AN, Van der Ven A, Garikipati K (2020) Scale bridging materials physics: active learning workflows and integrable deep neural networks for free energy function representations in alloys. [arXiv:2001.05646](https://arxiv.org/abs/2001.05646)
- Zhang X, Garikipati K (2020) Machine learning materials physics: multi-resolution neural networks learn the free energy and nonlinear elastic response of evolving microstructures. [arXiv:2001.01575](https://arxiv.org/abs/2001.01575)
- Hethcote HW (2000) The mathematics of infectious diseases. *SIAM Rev* 42(4):599–653
- Jo H, Son H, Hwang HJ, Jung SY (2020) Analysis of covid-19 spread in South Korea using the sir model with time-dependent parameters and deep learning. <https://doi.org/10.1101/2020.04.13.20063412>
- Truscott J, Ferguson NM (2012) Evaluating the adequacy of gravity models as a description of human mobility for epidemic modelling. *PLoS Comput Biol* 8(10):e1002699. <https://doi.org/10.1371/journal.pcbi.1002699>
- Hunter E, Namee BM, Kelleher J (2017) A taxonomy for agent-based models in human infectious disease epidemiology. *J Artif Soc Soc Simul* 20(3):2. <https://doi.org/10.18564/jasss.3414>
- Chinviriyasit S, Chinviriyasit W (2010) Numerical modelling of an sir epidemic model with diffusion. *Appl Math Comput* 216(2):395–409. <https://doi.org/10.1016/j.amc.2010.01.028>
- Gai C, Iron D, Kolokolnikov T (2020) Localized outbreaks in an s-i-r model with diffusion. *J Math Biol* 80:1389–1411
- Angulo J, Yu H-L, Langousis A, Kolovos A, Wang J, Madrid AE, Christakos G (2013) Spatiotemporal infectious disease modeling: a bme-sir approach. *PLoS ONE* 8(9):e72168. <https://doi.org/10.1371/journal.pone.0072168>
- U.S. Census Bureau, Census 2000, Michigan Counties map. U.S. Census Bureau, Census 2020. https://www2.census.gov/geo/maps/general_ref/stco_outline/cen2k_pgsz/stco_MI.pdf
- Michigan Data: Cases by County by Date. <https://www.michigan.gov/coronavirus>, 2020
- The COVID Tracking Project: Daily Michigan Data. <https://covidtracking.com/data/state/michigan>, 2020
- Annual Estimates of the Resident Population for Counties in Michigan: April 1 (2010) to July 1, 2019 (CO-EST2019-ANNRES-26). U.S. Census Bureau, Population Division, March, p 2020
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Jarrod Millman K, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, Polat İ, Feng Y, Moore EW, Vand erPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P (2020) and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat Methods* 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Mitusch SK, Funke SW, Dokken JS (2019) dolfin-adjoint 2018.1: automated adjoints for fenics and firedrake. *J Open Source Softw* 4(38):1292. <https://doi.org/10.21105/joss.01292>
- Hornik K, Stinchcombe M, White H et al (1989) Multilayer feedforward networks are universal approximators. *Neural Netw* 2(5):359–366
- Liu Q, Wang D (2016) Stein variational gradient descent: a general purpose Bayesian inference algorithm. In: *Advances in neural information processing systems*, pp 2378–2386
- Blei DM, Kucukelbir A, McAuliffe JD (2017) Variational inference: a review for statisticians. *J Am Stat Assoc* 112:859–877
- Graves A (2011) Practical variational inference for neural networks. In: *Advances in neural information processing systems 24: 25th annual conference on neural information processing systems 2011, NIPS 2011*, pp 1–9
- Wen Y, Vicol P, Ba J, Tran D, Grosse R (2018) Flipout: efficient pseudo-independent weight perturbations on mini-batches. In: *6th International conference on learning representations, ICLR 2018—conference track proceedings*, pp 1–16

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.