



OPEN

## Explaining COVID-19 outbreaks with reactive SEIRD models

Kunal Menda<sup>1</sup>✉, Lucas Laird<sup>2</sup>, Mykel J. Kochenderfer<sup>1</sup> & Rajmonda S. Caceres<sup>2</sup>

COVID-19 epidemics have varied dramatically in nature across the United States, where some counties have clear peaks in infections, and others have had a multitude of unpredictable and non-distinct peaks. Our lack of understanding of how the pandemic has evolved leads to increasing errors in our ability to predict the spread of the disease. This work seeks to explain this diversity in epidemic progressions by considering an extension to the compartmental SEIRD model. The model we propose uses a neural network to predict the infection rate as a function of both time and the disease's prevalence. We provide a methodology for fitting this model to available county-level data describing aggregate cases and deaths. Our method uses Expectation-Maximization to overcome the challenge of partial observability, due to the fact that the system's state is only partially reflected in available data. We fit a single model to data from multiple counties in the United States exhibiting different behavior. By simulating the model, we show that it can exhibit both single peak and multi-peak behavior, reproducing behavior observed in counties both in and out of the training set. We then compare the error of simulations from our model with a standard SEIRD model, and show that ours substantially reduces errors. We also use simulated data to compare our methodology for handling partial observability with a standard approach, showing that ours is significantly better at estimating the values of unobserved quantities.

Having an accurate understanding of the spread of COVID-19 is essential to containing the virus effectively, and necessary for the deployment and allocation of resources. In order to understand why the spread appears to differ between communities, we require a mathematical model of disease spread capable of expressing the differences. A standard model of disease spread is the SEIRD model, in which each individual is either susceptible (S), exposed (E), infected (I), recovered (R), or dead (D)<sup>1,2</sup>. *Compartmental* SEIRD models consider only the aggregate number of individuals with each disease state, and specify a set of differential equations that govern how the compartmental populations change with time.

The well-studied compartmental SEIRD model is popular because of its simplicity<sup>2</sup>. This standard model, however, only predicts a single, clear peak in infections. While certain counties in the northeast of the United States appeared to exhibit this behavior for a while, most other counties do not.

In this work, we seek to extend the model and fit it to available data that can account for the diversity in COVID-19 outbreaks across the US. Specifically, we seek to learn a model that predicts clear peaks for counties in the northeast that had them, and predicts multiple, flatter peaks for those that did not.

In order to build models capable of expressing more realistic behavior, we relax two assumptions made by the standard compartmental SEIRD model:

- (A1) Stationarity—the disease parameters remain constant over time, and,
- (A2) Non-reactivity—the disease parameters remain constant regardless of the prevalence, i.e., the fraction of the population infected by the disease.

We hypothesize that relaxing these assumptions allows us to better explain the diversity of behavior seen in reality. Specifically, allowing disease parameters to vary with time allows us to account for implicit dependencies on virus mutations, weather, and changing cultural norms. Allowing them to depend on prevalence allows us to account for the reactive nature of a population's behavior to the perceived threat of the virus, and heterogeneous patterns of interactions within populations in different geographical areas, which reflect in the peak prevalence levels reached.

<sup>1</sup>Department of Aeronautics & Astronautics, Stanford University, Stanford, CA, USA. <sup>2</sup>MIT Lincoln Laboratory, Lexington, MA, USA. ✉email: kmenda@alumni.stanford.edu

To relax these assumptions, we consider a “reactive” compartmental SEIRD model (R-SEIRD) in which the transmission rate is a function of both time and the number of infected individuals. To avoid imposing an incorrect prior on the functional form of this relationship, we use a neural network to model it.

A common difficulty in fitting such models to available data is that of *partial observability*. SEIRD models have *states* that vary with time according to their dynamic parameters. Because the available data are not time-series of the models’ states, but only partial observations of it (daily new cases and deaths), we employ tools from system-identification<sup>3</sup> in order to fit these models. Specifically, we use a technique called Certainty-Equivalent Expectation-Maximization (CE-EM)<sup>4</sup>, which was recently shown to be a reliable and low-variance method for learning the parameters of partially observable dynamical systems. A key benefit of this methodology is that we do not assume knowledge of the initial state of the population, an assumption typically made by related approaches<sup>5–8</sup>.

In this work, we present a methodology for fitting R-SEIRD models to available data, and validate our hypotheses by showing that learned models can explain the diversity of behavior seen across the United States, and that is more robust to partial observability. Through our experiments, we show that:

1. Learned R-SEIRD models produce behavior consistent with what was observed in counties across the United States,
2. The simulation error of R-SEIRD models is lower than that of standard SEIRD models when compared against trajectories from outbreaks across the United States, and,
3. CE-EM is better than standard methodologies for fitting SEIRD models to partially observed data.

The contributions of this work are, therefore, to:

1. Propose a novel model that relaxes the assumptions of the standard SEIRD model,
2. Provide a methodology that is novel for fitting such a model to available data, and,
3. Experimentally demonstrate that this model is capable of expressing the diversity of behavior observed across the United States, and that our methodology outperforms standard approaches in partially observed settings.

This paper is organized as follows. In the “Background” section, we formalize compartmental SEIRD models and review the CE-EM algorithm, as well as related approaches to fitting such models to data. In the “Methodology” section, we introduce the R-SEIRD model and describe our methodology for fitting it to available data. In the “Fitting COVID-19 epidemics across The United States” section, we fit the model to data from a selection of representative counties across the United States and show that it is able to reflect the observed behavior. In the “Inferring versus Assuming Initial Conditions” section, we compare CE-EM and a baseline methodology in their ability to fit partially observed disease data. Finally, in the “Conclusion” section, we discuss possible limitations to the scope of our work, as well as directions for further study.

## Background

In this section, we review compartmental SEIRD models, as well as CE-EM, an algorithm for fitting partially observed state-space models to data. We also discuss the literature addressing related problems.

**SEIRD models.** SEIRD models are mathematical models of the spread of an infectious disease. Every individual in a population is in one of five states—they are either susceptible (S) to the disease, exposed (E) to the disease, infected (I) by the disease, or who have recovered (R) or died (D) from the disease<sup>2</sup>. In this context, we assume an exposed individual is ‘pre-symptomatic’, i.e., they can spread the disease but have not yet tested positive for it, while infected individuals are symptomatic and have tested positive for the disease. There are numerous extensions to this model<sup>2,6,8</sup>. For example, a SEIR model typically groups those who have recovered and died from the disease into a single state. SEIRD models may also differ in their modeling of reinfection, limited testing, or asymptomatic and quarantined individuals.

**Compartmental models.** In reality, each individual interacts with only a subset of individuals in the population. Thus the spread of the disease ought to be considered as propagating over a social interaction network of sparsely connected nodes. However, by introducing this fidelity into the model, the system’s state grows exponentially with the number of individuals, and the task of fitting the model to data becomes difficult. A compartmental SEIRD model ignores the network structure of a population by assuming that the population is *homogeneously mixed*. That is, we assume that every individual in the population is equally likely to interact with any other individual in the population on a given day. By making this assumption, we can dramatically simplify the state of the system. We need only track the number of individuals in each disease state, referred to as the *compartmental populations*, and not specific individuals.

We can model the dynamics of the state  $x_t = [S_t, E_t, I_t, (RD)_t]^T$  as the following deterministic system of differential equations:

$$\frac{d}{dt} \begin{bmatrix} S \\ E \\ I \\ RD \end{bmatrix} = \begin{bmatrix} -\left(\beta_E \frac{S_t E_t}{N} + \beta_I \frac{S_t I_t}{N}\right) \\ \beta_E \frac{S_t E_t}{N} + \beta_I \frac{S_t I_t}{N} - \gamma E \\ \gamma E - \lambda I \\ \lambda I \end{bmatrix}. \quad (1)$$

Here, the parameters  $\beta_E$  and  $\beta_I$  (referred to as the *infection rates*) are interpreted as the average number of individuals that an exposed or infected individual comes in contact with per unit time, multiplied by the probability that the contact results in disease transmission, respectively. The parameter  $\gamma$  specifies the average number of exposed individuals who transition to the infected state, per unit time, and  $\lambda$  specifies the average number of infected individuals who recover or die, per unit time. We define infected individuals to be those who have *tested positive* for the disease, and thus  $\gamma^{-1}$  is the average amount of time it takes for someone who contracted the disease to test positive for it. The *recovered-deceased* (RD) compartment includes both individuals who have truly recovered from the disease and those who have died from it. Hence,  $\lambda^{-1}$  is the average amount of time it takes for someone to no longer be infectious after having tested positive, as a result of recovery or death. If we assume a fixed mortality rate  $\mu$  for the disease, then we can compute the number of individuals who have died from the disease as  $D_t = \mu(RD)_t$ , and those who have recovered from it as  $R_t = (1 - \mu)(RD)_t$ .

The above framing describes a deterministic dynamical system. There are many methods for framing compartmental SEIRD systems using stochastic differential equations<sup>9</sup>, and one such formulation will be discussed in the “Methodology” section.

It is worth noting the various effects *not* modeled here. In this model, we do not account for individuals who never develop symptoms. Furthermore, this model assumes an individual’s infectiousness changes when they test positive—the truth of which may depend on the delays experienced with RT-PCR testing. SEIRD models can be extended to account for these effects, though such extensions may introduce more parameters and not be identifiable from the available data.

It would be straightforward to learn the parameters  $\theta = [\beta_E, \beta_I, \gamma, \lambda, \mu]$  from data if we could directly observe the compartmental populations. However, the available data is typically only of the aggregate transitions from the *E* compartment to the *I* compartment, when individuals test positive for the disease, and aggregate transitions from *I* to *R*, when individuals recover from or die due to the disease. As a result, estimation of  $\theta$  must be performed under partial observability. Methods for doing so, including CE-EM, will be discussed next.

**Certainty-equivalent expectation maximization.** Suppose we seek to find the parameters of the following state-space dynamical system:

$$\begin{aligned} x_{t+1} &= f_\theta(t, x_t) + w_t; & w_t &\sim p_w(\cdot), \\ y_t &= h_\theta(t, x_t) + v_t; & v_t &\sim p_v(\cdot). \end{aligned} \quad (2)$$

Here,  $x_t$  is the state of the system,  $f_\theta(t, x_t)$  is a parameterized model of the dynamics,  $w_t$  is an additive *process noise* term, and  $p_w$  is the distribution from which  $w_t$  is sampled. Furthermore,  $y_t$  is an observation of  $x_t$ ,  $h_\theta(t, x_t)$  is a parameterized observation model,  $v_t$  is the *observation noise* term, and  $p_v$  is the distribution from which  $v_t$  is sampled. In a *nonlinear Gaussian* system,  $p_w$  and  $p_v$  are multivariate Gaussian distributions.

Methods for parameter estimation typically attempt to quantify the likelihood  $p(y_{1:T} | \theta)$  of a time-series of observations  $y_{1:T}$  given some choice of parameters  $\theta$ . A set of methods called approximate Bayesian computing (ABC) use this estimate to characterize the Bayesian posterior  $p(\theta | y_{1:T})$  using approximate methods such as Markov Chain Monte Carlo<sup>10,11</sup>. Maximum-likelihood (MLE) methods attempt to find the likelihood maximizing parameters  $\theta_{ML} = \arg \max_{\theta} p(y_{1:T} | \theta)$ . Both methods rely on being able to estimate the data likelihood  $p(y_{1:T} | \theta)$  with low-variance.

In order to estimate the data likelihood, we must marginalize over the unobserved states of dynamical system:

$$p(y_{1:T} | \theta) = \int p(y_{1:T}, x_{1:T} | \theta) dx_{1:T}. \quad (3)$$

Many approaches to fitting SEIRD models to data assume known initial conditions  $x_1$  when estimating  $p(y_{1:T} | \theta)$ <sup>5,6,12</sup>. Key drawbacks to such approaches are their sensitivity to the choice of  $x_1$ , as well as a degradation in their ability to find global optima if being fit to long time-series. Since computing this expectation is generally intractable without knowing initial conditions and for long time-series, many approaches instead find the posterior distribution over states, i.e.  $p(x_{1:T} | y_{1:T}, \theta)$ , and then approximate quantities of interest using Monte-Carlo samples<sup>13</sup>.

In a subset of MLE methods, the *smoothing distribution*  $p(x_{1:T} | y_{1:T}, \theta)$  is used to estimate the joint data log-likelihood:

$$Q(\theta, \theta^k) = \mathbf{E}_{x_{1:T} \sim p(\cdot | y_{1:T}, \theta^k)} [\log p(y_{1:T}, x_{1:T} | \theta)]. \quad (4)$$

The function  $Q(\theta, \theta^k)$  is then maximized to yield  $\theta^{k+1}$ , and this two-step procedure is repeated until convergence in an algorithm called Expectation-Maximization (EM)<sup>14</sup>. Such approaches overcome the need to precisely know  $x_1$  and can straightforwardly handle long time-series<sup>4</sup>.

Various algorithms differ in how they compute the smoothing distribution. ParticleEM uses particle smoothing, an approach that uses sequential Monte-Carlo to approximately sample from the smoothing distribution<sup>13,15</sup>.

Though this approach makes very few assumptions, it can require a prohibitive number of Monte-Carlo samples to yield sufficiently low variance estimates of  $Q(\theta, \theta^k)$ .

The Certainty-Equivalent EM (CE-EM) algorithm makes the approximation that  $p(x_{1:T} | y_{1:T}, \theta)$  can be modeled by a Dirac-delta function located at the smoothing distribution's mode<sup>4</sup>. In the case where  $p_w(\cdot)$  and  $p_v(\cdot)$  are Gaussian distributions, the assumptions made allow CE-EM to find an estimate for  $\theta^{\text{ML}}$  using block-coordinate ascent. By doing so, CE-EM finds solutions with significantly higher efficiency and lower variance than other estimation procedures<sup>4</sup>. However, by making this approximation, the algorithm is known to be biased in the presence of large process noise.

In this work, we propose using CE-EM to fit SEIRD models, including our proposed extension to available county-level data. As one might expect, such an approach should yield inaccurate results if the disease progression is highly stochastic, or poorly modeled by a compartmental model. However, we will show that even when fitting a single model to data from six counties with diverse epidemics, simulations from the model appear to reflect reality.

**Related work.** Since the dawn of the COVID-19 pandemic, many efforts have been directed at forecasting its evolution, as well as fitting modified SEIR models to specific outbreaks<sup>16,17</sup>. Approaches vary in their characterization of the model, assumptions made, and their methodology for parameter estimation under partial observability. Korolev<sup>5</sup> demonstrates issues with the identifiability of the SEIRD model and presents an estimation technique for the basic reproduction number  $R_0$ . They fit a standard SEIRD model to COVID-19 data, but assume known initial conditions to address partial observability—an assumption that can only be made in restricted settings. He et al.<sup>6</sup> use particle swarm optimization to optimize the parameters of a SEIR model extended to account for hospitalized and quarantined individuals. Sun and Wang<sup>7</sup> fit an extended SEIR model that accounts for a threshold in recovery rate, as well as asymptomatic patients. Both of these studies also address partial observability by assuming known initial conditions, but only fit to a single, short trajectory of data from the Hubei province and Heilongjiang province in China, respectively. Arik et al.<sup>8</sup> introduce additional data sources such as mobility as covariates to expand the explanatory power of SEIR models, and specify a distribution over initial conditions to overcome partial observability. One of the more successful approaches to forecasting the early pandemic progression also fit SEIR models with a subset of variables that vary with time<sup>12</sup>. Their approach also assumes a known initial condition and finds parameters by minimizing simulation error. In contrast, our methodology does not need to assume knowledge of initial conditions, and instead infers unobserved states such that they maximize the likelihood of the data.

Recent works<sup>18–20</sup> have also attempted to use neural networks to forecast the spread of COVID-19. While Wieczorek et al.<sup>19</sup> and Melin et al.<sup>20</sup> do not use neural-networks to model relationships within a SEIR model, Dandekar and Barbastathis<sup>18</sup> use them to model a quarantine control function. Similar to other discussed methodologies, they assume initial conditions to overcome partial observability. Yang et al.<sup>21</sup> fit both a SEIR model with additional compartments, as well as a black-box LSTM to a short time-series of COVID-19 infection data from Hubei, China, and compare their forecasting ability. Unlike our model and most other approaches, their SEIR model is not fit to the data by attempting to reproduce a time-series, but by linearizing the model around certain set points and assuming values for unknown quantities.

The methods mentioned above attempt to fit disease models that assume the population to be in one of many disease states, and specify differential equations that govern their rate-of-change. Once model parameters are estimated with a dataset, the methods perform forecasts by simulating the estimated disease states forward in time. Another approach commonly taken to forecasting the near-term future of a disease outbreak is to use statistical methods that directly attempt to predict a given quantity at some future point in time. For example, Castillo and Melin<sup>22,23</sup> construct a fuzzy model that uses linear and nonlinear fractal dimensions of a time-series in order to make forecasts of future cases and deaths. Though such statistical models are useful for forecasting the evolution of the disease, they often lack interpretability and behavioral guarantees provided by disease models using differential equations to govern transitions between disease states.

To our knowledge, our work is the first to use a neural network to model the relationship between time, prevalence, and infection rate, the first to use Certainty-Equivalent EM for the purpose of parameter estimation under partial observability, and also the first to fit a single model to multiple time-series from across the United States.

## Methodology

This section describes data sources used in this work, the proposed modification to SEIRD models yielding R-SEIRD models, the formulation of R-SEIRD models as nonlinear Gaussian systems, and a procedure for fitting these models to available data.

**Data sources.** In this work, the primary data source considered is the Novel Coronavirus (COVID-19) Cases Dataset, provided by JHU CSSE<sup>24</sup>, which provides a time-series of daily reported COVID-19 cases and deaths for every county in the United States.

**R-SEIRD models.** As stated in the “Introduction” section, we propose to relax the assumption that the infection rate  $\beta_E$  is constant with time. Instead, we model it to be a function of time and the *observed prevalence* of the disease. Specifically, we let:

$$\beta_E(t) = \tilde{\beta}_E \cdot \sigma \left( \text{NN}_{\tilde{\theta}} \left( t, \frac{I_t}{N_t} \right) \right). \quad (5)$$

Here,  $N_t = S_t + E_t + I_t + (RD)_t$  is the *effective* population size at time  $t$ , a quantity that we allow to vary with time<sup>25</sup> and be dynamically inferred during learning. Additionally,  $I_t/N_t$  is the observed prevalence,  $\text{NN}_{\hat{\theta}}$  is a neural network mapping  $\mathbf{R}^2 \rightarrow \mathbf{R}$ ,  $\sigma(\cdot)$  is the sigmoid function, and  $\beta_E$  is a learned coefficient. For simplicity, we assume  $\beta_I = 0$ , which implies that the number of infections caused by an individual after they test positive for COVID-19 is negligible compared to the number of infections prior to them knowing they have the disease. Furthermore, since much is now known about the typical duration between exposure and symptom onset, and symptom onset and death, we rely on literature to provide estimates of  $\gamma$  and  $\lambda$ <sup>26,27</sup>. Our methodology does, however, allow us to treat them as learned parameters. We treat the mortality rate  $\mu$  as learned, and thus  $\theta = [\beta_E, \hat{\theta}, \mu]$  are the learned parameters.

With the modification that  $\beta_E$  depends on time, we allow the model to reflect changes in the infection rate that may depend on changing behaviors over time, such as climate, and the dominant strain of the virus. With the modification that it depends on prevalence, we allow for changes in infection rate in response to the level of infection in a county, as well as account for heterogeneous interaction patterns between geographies, that may affect the peak infection rates reached in those geographies.

**R-SEIRD as a nonlinear Gaussian system.** In order to effectively learn the parameters of an R-SEIRD model using CE-EM, we represent it as a nonlinear Gaussian system. To do so, we must specify the state  $x_t$ , discrete-time dynamics function  $f_{\theta}(t, x_t)$ , observation  $y_t$ , and observation function  $h_{\theta}(t, x_t)$ .

Though the state of a SEIRD system is described by the number of individuals in each compartment, we note two facts about these quantities:

- The number of individuals in any compartment is a positive quantity, and,
- These quantities scale various orders of magnitude<sup>24</sup>.

For this reason, it is sensible to let the state  $x_t$  of the system correspond to the logarithm of each compartment's population as opposed to their absolute value. As a result, applying Gaussian process or observation noise to these quantities loosely corresponds to assuming that noise is proportional to the absolute value of the quantity it is applied to. Specifically, let the state of the system be  $x_t = \log [S_t, E_t, I_t, (RD)_t]$ , and the observation be  $y_t = \log [\Delta C_t, \Delta D_t]$ . Here,  $\Delta C_t$  corresponds to the number of new confirmed cases on day  $t$  and tracks the total number of individuals that have transitioned from the  $E$  to the  $I$  compartment, and  $\Delta D_t$  corresponds to the number of new deaths on day  $t$ .

The dynamics and observation models for the R-SEIRD model are then specified as follows:

$$\Delta \log \begin{bmatrix} S \\ E \\ I \\ RD \end{bmatrix}_t = \int_t^{t+1} \begin{bmatrix} -\beta_E(\tau)SE \\ \frac{\beta_E(\tau)SE}{N} - \gamma E \\ \gamma E - \lambda I \\ \lambda I \end{bmatrix}_{\tau} \circ \begin{bmatrix} 1/S \\ 1/E \\ 1/I \\ 1/RD \end{bmatrix}_{\tau} d\tau + w_t; \quad w_t \sim \mathcal{N}(0, \Sigma_w), \quad (6)$$

$$\log \begin{bmatrix} \Delta C_t \\ \Delta D_t \end{bmatrix} = \log \begin{bmatrix} \gamma E \\ \mu \lambda I \end{bmatrix} + v_t; \quad v_t \sim \mathcal{N}(0, \Sigma_v). \quad (7)$$

where  $\beta_E(t)$  is computed according to Equation 5. The process noise covariance  $\Sigma_w$  and observation noise covariance  $\Sigma_v$  are hyperparameters that can be chosen to be identity or diagonal matrices scaled by  $\sigma_w^2$  and  $\sigma_v^2$  respectively. Integration is performed using a Runge-Kutta method<sup>28</sup>.

Since the parameters  $\beta_E$  and  $\mu$  are positive, we optimize their logarithms as opposed to their absolute values. Framed as a nonlinear Gaussian system, the R-SEIRD model can straightforwardly fit a batch of observation time-series from multiple counties simultaneously by using CE-EM<sup>4</sup>. To improve fit reliability and eliminate periodic drops from weekends, we apply a 7-day moving average filter to daily case and death data before fitting to them.

In our experiments, we fit the R-SEIRD model to data from six counties in the United States exhibiting diversity in how COVID-19 has spread, and show that it is capable of expressing this diversity.

## Experiments

In this section, we first perform an experiment to fit SEIRD and R-SEIRD using CE-EM to COVID-19 data from counties across the United States. In a second experiment, we generate simulated data to compare the use of CE-EM and a baseline methodology in their ability to fit partially observed disease data.

**Fitting COVID-19 epidemics across The United States.** Our experiment aims to achieve the following goals:

1. Fit a single R-SEIRD model to data from a variety of counties across the United States,
2. Show that the learned model, when simulated with realistic initial conditions, can reproduce observed behavior, and,
3. Justify the modifications to the SEIRD model by showing that the mean squared error (MSE) of simulations under the R-SEIRD model is much lower than that of a fit SEIRD model.

To achieve these goals, we fit models to data from three counties in the northeast of the United States (Middlesex, MA, Kings, NY, and Fairfield, CT), as well as three counties across the United States that have had diverse epidemics (Los Angeles, CA, Miami-Dade, FL, and Cook, IL). We consider data from February 22, 2020, to September 27, 2020, because during this period, the counties in the northeast exhibit single, clear peaks in infection, though the remaining counties exhibit multiple peaks. The remaining worst-hit counties in each of the United States compose a test-set to evaluate the model.

We fit both an R-SEIRD model and a standard SEIRD model (where we learn just the constant parameters  $\beta_E$  and  $\mu$ ) to data from these six counties, and then attempt to see if the learned models can reproduce the behavior of all six counties, as well as of counties not trained on, from appropriate initial conditions. We visually compare simulations on select counties to show that deterministic simulations of the R-SEIRD system can express multiple peaks while those of a SEIRD system cannot. We also visualize the learned relationship between time, prevalence, and infection rate, to provide an intuition for how the model can express such behavior. We then compare the MSE of simulations from both learned models on the worst-hit county of each of the United States (not in the training set), showing that simulation error is much lower when using the R-SEIRD model. The codebase for running these experiments can be found at <https://github.com/sisl/rseird>.

**Experimental setup.** Here we detail the hyperparameters used when training the R-SEIRD and SEIRD models, the methodology for selecting appropriate initial conditions for states when evaluating the models, as well as metrics for evaluation.

*Hyperparameters.* When learning the R-SEIRD model or SEIRD model, we specify the values of  $\gamma$  and  $\lambda$ . The literature suggests that the median time from exposure to developing symptoms is five days<sup>27</sup> and that the median time between symptom onset and recovery is 21 days<sup>26</sup>. As a result, we let  $\gamma = 1/5$  and  $\lambda = 1/21$ . We let  $\sigma_w = \sigma_v/10$ , assuming that corruptions to observations are an order of magnitude noisier than corruptions to the process. Furthermore, we let the observation noise on deaths be double that on cases, since reporting is typically noisier when the number of deaths is small. The neural network used in the R-SEIRD model has three hidden layers with 32 hidden units each and tanh activation functions. Both models are optimized in the CE-EM learning step using an Adam optimizer with a learning rate of  $5 \times 10^{-4}$ , which is harmonically decayed over time, and CE-EM trust-region parameters of  $\rho_x = 0.5$  and  $\rho_\theta = 0.01^4$ . We additionally smooth the observed data using a 7-day moving average filter to smooth high-frequency noise, including the expected drop in cases over weekends.

*Initial condition selection.* When evaluating learned models, partial observability makes it such that we do not know the initial conditions (i.e., the number of individuals in each disease compartment) to simulate the models from. Related work often assumes the population is almost completely unexposed at  $t = 1$ , and assumes the susceptible population is a community's true population<sup>6</sup>. However, as stated in the "Methodology" section, we allow the population size to be determined by the sum of the compartmental population, and therefore be a free parameter. Doing so is necessary because, for trajectories in our dataset, it is unreasonable to assume the population is initially unexposed.

For these reasons, we optimize for the initial conditions from which deterministically simulated trajectories minimize the MSE of the observed number of daily cases. We optimize the initial conditions by using the Cross-Entropy Method<sup>29</sup>, selecting the top 100 candidates at each epoch, perturbing each candidate to generate 1000 candidates for the next epoch, and running the optimizer for 10 epochs.

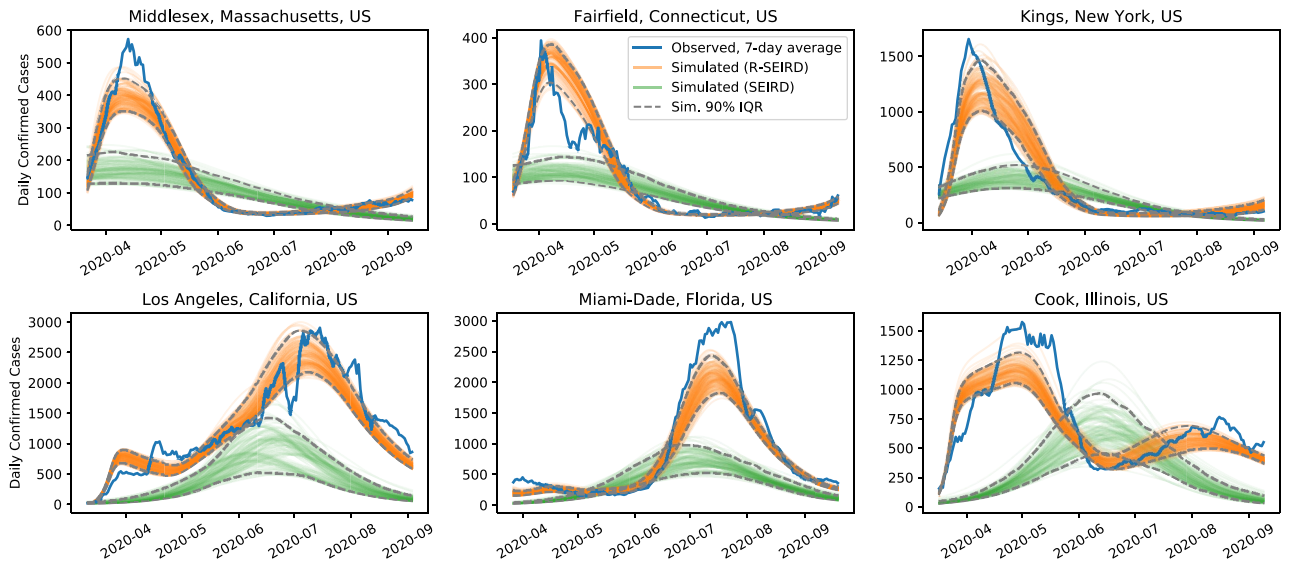
*Evaluation metrics.* When selecting initial conditions, we measure the MSE between the log of the daily cases in simulation (i.e.,  $\log(\gamma E_t)$ ), and reality (i.e.,  $\log \Delta C_t$ ), and use the distribution over this metric to compare the R-SEIRD and standard SEIRD models. We compare histograms of these errors on data from counties not in the training set.

*Results.* In Fig. 1, we see the simulated number of cases from optimized initial conditions for the six counties on which the R-SEIRD model was trained. As is clearly seen, the model can express both a single clear peak present in the northeastern counties and the multiple irregular peaks present in the other counties. Since the simulations are deterministic, they are smooth and only approximately reflect the trends in the trajectories they are supposed to match.

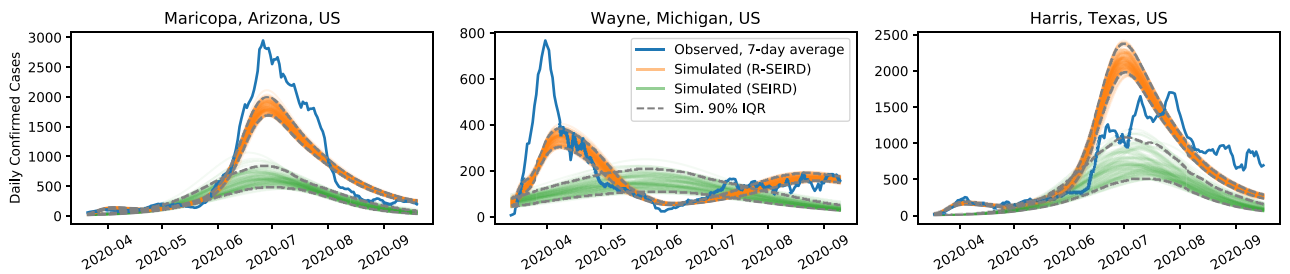
In contrast, and consistent with expectation, a SEIRD model, even when fit to the same data and simulated from optimized initial conditions, cannot reflect the multiple peaks in the training data. Furthermore, in an attempt to fit the multi-peak behavior in three counties, the fit is severely compromised in counties with a single, clear peak.

In Fig. 2, we compare the learned R-SEIRD and SEIRD models on three hard-hit counties not in the training set. Again, the SEIRD model cannot express more than a single peak, and thus cannot capture the behavior reflected in reality. Despite not being trained on data from these counties, the R-SEIRD model is, however, able to reflect the double-peaks seen in the data, which vary in onset and magnitude.

We quantify the difference in simulation quality by measuring the MSE between the simulations (from optimized initial conditions) with the observed trajectories. In Fig. 3, we show a histogram of these errors for simulations on the worst-hit county of each of the United States (not in the training set). Consistent with the improvement visible in Fig. 2, we see that the MSE of simulations by the learned R-SEIRD model is significantly lower than that of the SEIRD model.



**Figure 1.** Simulations of the learned R-SEIRD and learned SEIRD systems on the six counties the models are trained on. The trajectory in blue is the 7-day moving average of observed daily confirm cases. These results demonstrate the expressiveness of our model. In particular, our model (orange) is able to track observed disease progression (blue) much closer than the canonical SEIRD model (green). In addition, it is able to identify the observed multiple peak behavior for which SEIR is not designed to capture. These results are consistent across diverse geographic areas of US.



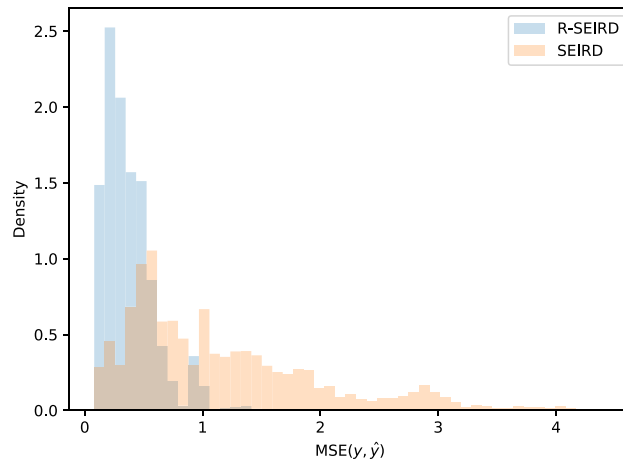
**Figure 2.** Simulations of the learned R-SEIRD and learned SEIRD systems on the three counties the models are not trained on. The trajectory in blue is the 7-day moving average of observed daily confirm cases. These results demonstrate the generalizability of our model against unseen data. We show representative results across different geographic regions of how our model can more accurately predict the different signatures of disease progression. Note that the SEIRD model consistently under-estimates the prevalence of disease spread, which has significant implications for policy-making decisions on the severity of intervention.

To gain an intuition for how the R-SEIRD model is able to express multiple irregular peaks, we visualize the learned mapping between time, the prevalence of the disease ( $I_t/N_t$ ) and the basic reproduction number  $R_0 = \beta_E(t, I_t/N_t)/\gamma$ , which indicates a growing epidemic for values greater than 1. In Fig. 4, we see the learned mapping, which suggests that  $\beta_E$  decreases initially as prevalence increases, up until a point, after which it begins to surge again before decreasing. This behavior would be consistent with a two-tiered response, in which weak restrictions were adopted until they lost effect, and then stronger restrictions were imposed. Over time, we see that  $\beta_E$  begins to require larger prevalence levels to decay and bring  $R_0$  below one, suggesting that communities are growing more reluctant to adopt the weak restrictions when infection rates are low.

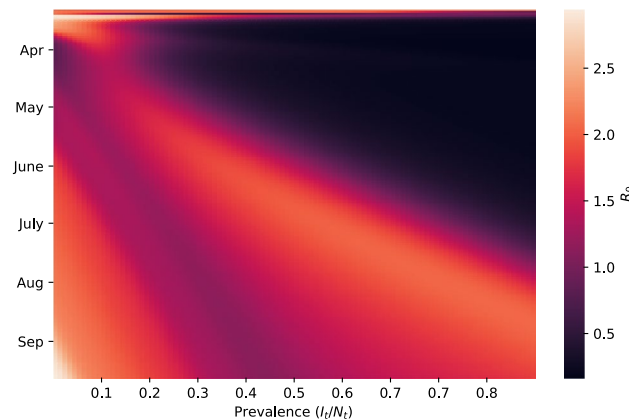
**Inferring versus assuming initial conditions.** Our second experiment aims to compare CE-EM with a baseline methodology that assumes known initial conditions, for fitting compartmental models to partially observed data. Specifically, we aim to show that:

1. CE-EM more accurately estimates the values of unobserved quantities, such as the number of exposed and infected individuals at a given time, and,
2. Estimating the effective population size with CE-EM leads to better fits to the data than by assuming its value.

To do so, we simulate a disease spreading through a population of sparsely connected individuals, but only partially observe the aggregated disease states of the population. We then attempt to fit a compartmental SEIRD



**Figure 3.** Comparison of prediction errors between R-SEIRD and SEIRD for the epidemic in the worst-hit county of each of the United States. We observe the errors of our model are substantially smaller than the canonical SEIRD model.



**Figure 4.** The learned mapping between time, prevalence, and the infection rate. The infection rate is represented by the  $R_0$  value and represented by the intensity of the color, with higher values reflecting more a infectious disease. This analysis aids to the interpretability of the deep learning model. It can lead to meaningful insights about potential intervention strategies.

model to the observed data using both CE-EM and a baseline methodology. We demonstrate that CE-EM is better at recovering unobserved states than the baseline methodology, and produces a better quality fit to the data.

*Baseline methodology.* We baseline against a common approach<sup>5,6,8,12</sup> that assumes the initial conditions of unobserved state-variables are known, which we refer to as the Fixed Initial Condition (FIC) baseline. Specifically, we let  $\tilde{x}_1$  be the assumed initial condition, and let  $\theta$  be the parameters of a deterministic compartmental disease model  $x_{t+1} = f_\theta(x_t), y_t = g_\theta(x_t)$ . We can forward-simulate the initial conditions according to the disease model to yield a state-trajectory  $\tilde{x}_{1:T}^\theta$ , and observation-trajectory  $\tilde{y}_{1:T}^\theta = g_\theta(\tilde{x}_{1:T}^\theta)$ . We can then define a loss:

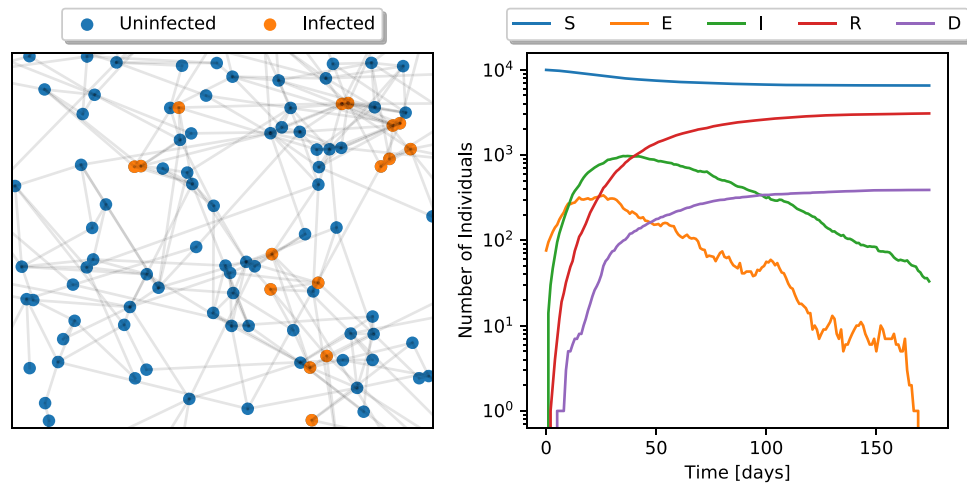
$$\mathcal{L}(\theta) = \|\tilde{y}_{1:T}^\theta - y_{1:T}\|_2^2 \tag{8}$$

and minimize the loss with respect to  $\theta$ .

To overestimate the performance of the FIC baseline methodology, we provide the baseline methodology with the exact initial condition of the unobserved disease state. Furthermore,  $\mathcal{L}(\theta)$  is nonconvex, and thus we optimize the objective using Differential Evolution to explore more of the space for a global optimum.. It is possible to use a global optimizer on this problem because the dimensionality of  $\theta$  is low.

*Simulated data.* To generate simulated diseases data, we consider a population in which individuals are sparsely connected to one another. We generate this graph by sampling the ‘locations’ of 10,000 individuals as random points on a unit square. We then assign an edge between two individuals  $i$  and  $j$  with a probability that





**Figure 5.** (Left) Depiction of the graph connecting the population, and the individuals infected or left uninfected over the course of the simulation. (Right) Aggregated numbers of individuals in each disease state over the course of the simulation. We see that due to the network's sparsity, the disease only propagates through a subset of the population, affecting just  $\sim 30\%$  of individuals.

exponentially decays with the distance between them. In this experiment, individuals have an average of 6.126 connections between them.

Individuals in the population can each be in one of five states: susceptible (S), exposed (E), infected (I), recovered (R), or dead (D). At a given instant in time, a susceptible individual becomes exposed with probability:

$$p_i^{S \rightarrow E} = 1 - (1 - \beta_E)^{\sum_{j \in \mathcal{N}(i)} E(j)}, \quad (9)$$

where  $\mathcal{N}(i)$  is the set of individuals that individual  $i$  is connected to, and  $E(j)$  is 1 if individual  $j$  is exposed, and otherwise 0. Furthermore, every exposed individual transitions to the infected state with probability  $\gamma$ , and every infected individual to the recovered state with probability  $\lambda(1 - \mu)$ , and to the dead state with probability  $\lambda\mu$ . We simulate an outbreak initialized with 1% of randomly chosen individuals in the exposed state, and the remainder in the susceptible state. We let  $\beta_E = 0.06$ ,  $\gamma = 1/5$ ,  $\lambda = 1/21$  and  $\mu = 0.1$ , and we simulate 175 days of the outbreak. At every time-step, we observe only the number of individuals who have transitioned from the  $E$  to  $I$  compartment, and the number who have transitioned from the  $I$  to  $D$  compartment.

In Fig. 5, we see both a depiction of the graph connecting the population, and the aggregated number of individuals in each disease state over the course of the simulation. Since the population is sparsely connected (which is a realistic property of observed interaction networks), only  $\sim 30\%$  is ever infected with the disease.

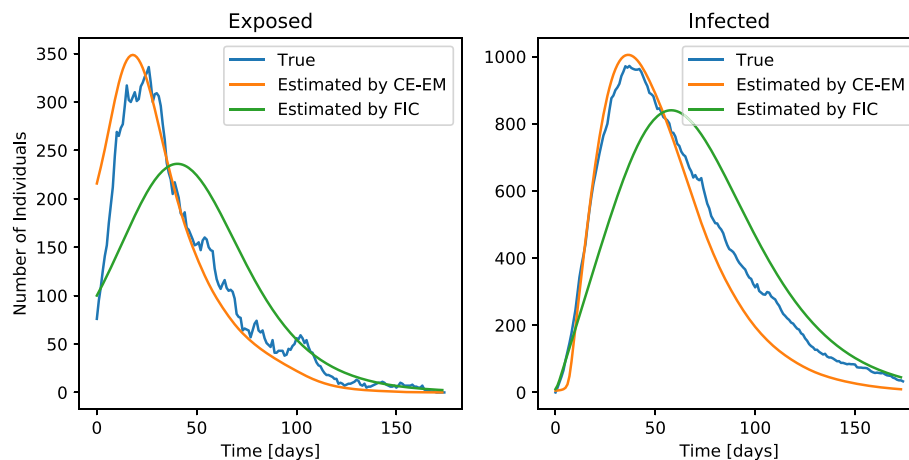
**Results.** In Fig. 6, we show the aggregated number of exposed and infected individuals over the course of the simulation. These two quantities are not directly observed during training, and are thus estimated by the compared algorithms. Overlaid on these trajectories are the trajectories estimated by CE-EM and the FIC baseline methodology. We see that the trajectories estimated by CE-EM almost perfectly align with the ground truth in their shape, peak location, and peak magnitude. Those estimated by the FIC baseline methodology, however, align poorly with the ground truth.

We quantitatively compare the mean-square-error (MSE) of predicted observations to those observed during the simulation. We predict observations using the models learned by each method and the states estimated by each learning methodology. We find that the MSE of predicted observations is 0.046 when using CE-EM and 0.241 when using the FIC baseline methodology.

The qualitative and quantitative comparisons between these methods show that CE-EM more accurately estimates the values of unobserved quantities, such as the number of exposed and infected individuals at a given time. They also show that estimating the effective population size instead of assuming its value, CE-EM leads to better fits to the data than by using the FIC baseline methodology.

## Conclusion

In this work, we proposed an extension to the compartmental SEIRD model that relaxes the assumptions of static model parameters and non-reactivity to disease prevalence called the R-SEIRD model. We did so by training a neural network to map the time and the prevalence of the disease to the infection rate. In order to fit available data, we employed Certainty-Equivalent Expectation-Maximization (CE-EM), which is a technique suited to fitting nonlinear Gaussian state-space models to data without direct observation of the system's state variables, and does not assume knowledge of the system's initial conditions. We provided a methodology for framing the R-SEIRD model as a nonlinear Gaussian system, and for fitting it to available data on daily confirmed cases and deaths.



**Figure 6.** Estimation of unobserved quantities by CE-EM and FIC baseline methods when fitting a SEIRD model to simulated data. (Left) Estimation of the number of exposed individuals. (Right) Estimation of the number of infected individuals. The estimates made by CE-EM align almost perfectly with ground truth in terms of peak intensity, peak height, and curve shape, while those of the FIC baseline do not.

Our experiments fit both the R-SEIRD and standard compartmental SEIRD models to data from six counties across the United States. We showed that the R-SEIRD model learned is capable of expressing the range of multi-peak behavior exhibited not only in the training data, but also on counties not trained on. We showed quantitatively that the simulation error when trying to reproduce the behavior of the worst-hit counties in the United States is much lower when using the R-SEIRD model compared to the standard SEIRD model. We further justify the use of CE-EM as a methodology for fitting compartmental disease models in partially observed settings by showing it achieves better fits to simulated data than baseline methodologies that assume initial conditions are known.

This work showed that by allowing the infection rate to be time-varying and reactive, that the much more complex behaviors exhibited by the epidemic across the United States can be recovered. However, we do not suggest that the R-SEIRD model definitively explains why a given outbreak progressed the way it did, but instead proposes a hypothesis that is consistent with observations. To come closer to a definitive explanation, we must depart from another assumption made by both standard compartmental SEIRD and R-SEIRD models—that of *homogeneous mixing*. By explicitly modeling the observed heterogeneous social interaction patterns in a county's population, we gain the capacity to account for super-spreading individuals, as well as non-uniform population densities.

In follow-on work, we propose a SEIRD model that accounts for the network structure of a community and an estimation procedure for fitting it to data. With such a model, we gain another explanatory tool to analyze the diversity in outbreaks and a model that can be used to evaluate localized containment strategies at a community level. We intend to expand on the ideas behind CE-EM by using probabilistic programming to perform inference in network-based epidemiological simulations, thereby explicitly modeling the dependence on known interaction patterns. Additionally, we intend to apply these methodologies to outbreaks in other parts of the world, which might differ substantially from the United States in climate, cultural norms, disease variants, and containment measures. While R-SEIRD is flexible enough to be directly applied to aggregated data from any locality, certain regions may warrant the relaxing of assumptions beyond that which is done in this work. In a future study, we seek to understand which regions of the world sufficiently share characteristics with the United States, so that we may use a single model to explain behavior across regions.

### Data availability

The datasets generated and/or analyzed during the current study are available in the COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University (<https://github.com/CSSEGISandData/COVID-19>)<sup>24</sup>.

Received: 17 March 2021; Accepted: 20 August 2021

Published online: 09 September 2021

### References

1. William, O. K. & Anderson, G. M. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A* **115**(772), 700–721 (1927).
2. Herbert, W. H. The mathematics of infectious diseases. *SIAM Rev.* **42**(4), 599–653 (2000).
3. Lennart L. System identification. In *Wiley Encyclopedia of Electrical and Electronics Engineering*, pp. 1–19 (1999).
4. Kunal M., Jean de B., Jayesh K.G., Ilan K., Mykel J.K., Zachary M. Scalable identification of partially observed systems using Certainty-Equivalent EM. In *International conference on machine learning (ICML)*, 2020.
5. Korolev, I. Identification and estimation of the SEIRD epidemic model for COVID-19. *J. Econom.* **220**(1), 63–85 (2021).

6. He, S., Peng, Y. & Sun, K. SEIR modeling of the COVID-19 and its dynamics. *Nonlinear Dyn.* **101**(3), 1667–1680 (2020).
7. Sun, T. & Wang, Y. Modeling COVID-19 epidemic in Heilongjiang province, China. *Chaos, Solitons Fractals* **138**, 109949 (2020).
8. Sercan O.A., Chun-Liang L., Jinsung Y., Rajarishi S., Arkady E., Long T.L., Vikas M., Shashank S., Leyou Z., Nate Y., et al. Interpretable sequence learning for COVID-19 forecasting. In *Advances in neural information processing systems (NIPS)*, 2020.
9. Priscilla E.G., Luis F.G. Stochastic epidemic modeling. *Math. Stat. Estim. Approach. Epidemiol.*, pp. 31–52, 2009.
10. Sunnåker, M. *et al.* Approximate Bayesian computation. *PLoS Comput. Biol.* **9**(1), e1002803 (2013).
11. Grant, D. B., Aaron, T. P., Jacob, J. O. & Jessica, A. H. Approximate Bayesian computation for spatial SEIR (S) epidemic models. *Spatial Spatio-Temporal Epidemiol.* **24**, 27–37 (2018).
12. Youyang G. COVID-19 projections using machine learning. <https://covid19-projections.com>, 2020.
13. Nikolas, K., Arnaud, D., Sumeetpal, S. S., Jan, M. & Nicolas, C. On particle methods for parameter estimation in state-space models. *Stat. Sci.* **30**(3), 328–351 (2015).
14. Arthur, P. D., Nan, M. L. & Donald, B. R. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B (Methodol.)* **39**(1), 1–22 (1977).
15. Thomas, B. S., Adrian, W. & Brett, N. System identification of nonlinear state-space models. *Automatica* **47**(1), 39–49 (2011).
16. Wim N. Artificial intelligence against COVID-19: An early review. *IZA Discussion Paper No. 13110*, 2020. <https://ssrn.com/abstract=3568314>.
17. Amir, A., Sunita, G., Santosh K.R., Gagan K., Sharaf J.M., Omar M.B. The number of confirmed cases of COVID-19 by using machine learning: Methods and challenges. *Archives of Computational Methods in Engineering*, pp. 1–9, 2020.
18. Raj, D., George, B. Neural network aided quarantine control model estimation of global COVID-19 spread. *arXiv preprint arXiv:2004.02752*, 2020.
19. Wiczołek, M., Silka, J. & Woźniak, M. Neural network powered COVID-19 spread forecasting model. *Chaos, Solitons Fractals* **140**, 110203 (2020).
20. Patricia, M., Julio C.M., Daniela, S., Oscar, C. Multiple ensemble neural network models with fuzzy response aggregation for predicting COVID-19 time series: the case of Mexico. In *Healthcare*, volume 8, page 181. Multidisciplinary Digital Publishing Institute, 2020.
21. Yang, Z. *et al.* Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. *J. Thorac. Dis.* **12**(3), 165 (2020).
22. Castillo, O. & Melin, P. Forecasting of COVID-19 time series for countries in the world based on a hybrid approach combining the fractal dimension and fuzzy logic. *Chaos, Solitons Fractals* **140**, 110242 (2020).
23. Oscar, C., Patricia, M. A novel method for a COVID-19 classification of countries based on an intelligent fuzzy fractal approach. In *Healthcare*, volume 9, page 196. Multidisciplinary Digital Publishing Institute, 2021.
24. Ensheng, D., Hongru, D., and Lauren, G. An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infectious Dis.*, 20(5):533–534, 2020. URL <https://github.com/CSSEGISandData/COVID-19>.
25. Michael, Y. L., John, R. G., Liancheng, W. & János, K. Global dynamics of a SEIR model with varying total population size. *Math. Biosci.* **160**(2), 191–213 (1999).
26. Qifang, B., Yongsheng, W., Shujiang, M., Chenfei, Y., Xuan, Z., Zhen, Z., Xiaojian, L., Lan, W., Shaun A.T., Tong Z., et al. Epidemiology and transmission of COVID-19 in Shenzhen China: analysis of 391 cases and 1,286 of their close contacts. *MedRxiv*, 2020.
27. Stephen, A. L. *et al.* The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. *Ann. Internal Med.* **172**(9), 577–582 (2020).
28. Runge, C. Ueber die numerische auflösung von differentialgleichungen. *Math. Ann.* **46**(2), 167–178 (1895).
29. Rubinstein, R. The cross-entropy method for combinatorial and continuous optimization. *Methodol. Comput. Appl. Probab.* **1**(2), 127–190 (1999).

## Acknowledgements

The authors would like to acknowledge Neela Kaushik, Ross Alexander, Christian Vanderloo, and Shushman Choudhury for their valuable insights.

## Author contributions

K.M. proposed the methodology for this manuscript and conducted the experiments. K.M., L.L., M.J.K, R.S.C. contributed to the writing of the manuscript and the review of relevant literature. All authors reviewed the manuscript.

## Funding

This material is supported by the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Under Secretary of Defense for Research and Engineering.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to K.M.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

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



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021