

PARAMETER ESTIMATION FOR THE SEIR MODEL USING RECURRENT NETS

Anonymous authors

Paper under double-blind review

ABSTRACT

The standard way to estimate the parameters Θ_{SEIR} (e.g., the transmission rate β) of an SEIR model is to use grid search, where simulations are performed on each set of parameters, and the parameter set leading to the least L_2 distance between predicted number of infections and observed infections is selected. This brute-force strategy is not only time consuming, as simulations are slow when the population is large, but also inaccurate, since it is impossible to enumerate all parameter combinations. To address these issues, in this paper, we propose to transform the non-differentiable problem of finding optimal Θ_{SEIR} to a differentiable one, where we first train a recurrent net to fit a small number of simulation data. Next, based on this recurrent net that is able to generalize SEIR simulations, we are able to transform the objective to a differentiable one with respect to Θ_{SEIR} , and straightforwardly obtain its optimal value. The proposed strategy is both time efficient as it only relies on a small number of SEIR simulations, and accurate as we are able to find the optimal Θ_{SEIR} based on the differentiable objective. On two COVID-19 datasets, we observe that the proposed strategy leads to significantly better parameter estimations with a smaller number of simulations.

1 INTRODUCTION

The SEIR model (Kermack & McKendrick, 1927; Bjørnstad et al., 2002; Harko et al., 2014; Kendall, 2020) is a widely-used epidemiological model to predict the macroscopic behavior of disease spread through a population, e.g., the spread of COVID-19 among different populations at different locations over time (Estrada, 2020; Prem et al., 2020; Godio et al., 2020; Aleta et al., 2020; Chang et al., 2021; Grimm et al., 2021). A typical SEIR model models the spread dynamics of disease using four states of populations: susceptible (S), exposed (E), infectious (I) and recovered (R), where susceptible individuals can be transformed into the exposed, and later to the infectious and finally to the recovered. Each transition is associated with corresponding parameter(s), forming the parameter set Θ_{SEIR} for the model: *transmission rate* for $S \rightarrow E$, *infection rate* for $E \rightarrow I$ and *recovery rate* for $I \rightarrow R$. Θ_{SEIR} captures the spreading pattern of the disease, the learning of which is thus crucial for understanding the disease spread, and developing social-distancing or interventions policies.

Existing SEIR models relies on grid search to estimate Θ_{SEIR} via simulations. Model simulations are performed on different sets of parameters, and the set that has the smallest L_2 distance between the predicted number of infections and the observed number of infections is selected as the best parameter combination. This is because of the non-differentiable nature of the L_2 objective with respect to Θ_{SEIR} : predicted number of infections are obtained from discrete simulations. This learning process is (1) *high time-intensive*: the simulation process can be slow when the population is large, which gives a time complexity of $\mathcal{O}(MNT)$ where M is the number of simulations performed, N is total number of individuals and T is maximum time steps; besides, the search space will grow exponentially with respect to the number of parameters we need to estimate; and (2) *inaccurate*: because it is impractical to enumerate all possible combinations of parameters in the continuous space, the resulting parameter combination would be sub-optimal given a limited amount of trials.

To address these issues, in this paper, we propose to transform the original non-differentiable simulation problem into a differentiable one using neural recurrent nets for parameter estimation in SEIR models. The basic idea is that neural recurrent nets are differentiable with respect to the input learnable parameters, so that they are able to automatically learn these parameters via gradient

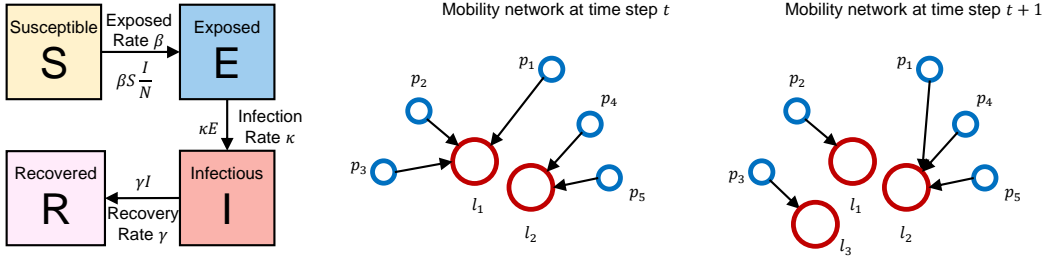


Figure 1: An overview of a standard SEIR model (left) the mobility network (right: from time t to time $t + 1$, person p_3 moved from location l_1 to l_3 , and p_1 moved from l_1 to l_2).

descent and backpropagation (Rumelhart et al., 1986). Another advantage neural recurrent nets offer is that they intrinsically model the temporal changes of observations, and have the potentials to make accurate predictions about the number of individuals at each time step. To train the recurrent net, we first harvest the training data by running a few SEIR simulations, leading to a collection of simulated data under different parameter combinations. Next, a recurrent net is trained to fit these simulations. Then with model parameters of the recurrent net Θ_{LSTM} fixed, the observation data are used to train the net with respect to the SEIR parameters Θ_{SEIR} via gradient descent, leading to the optimal parameter values Θ_{SEIR}^* .

The proposed strategy is both time efficient in that it only relies on a small number of SEIR simulations, and accurate because we are able to find the optimal parameters Θ_{SEIR}^* based on the differentiable objective. Through experiments on two COVID-19 datasets, we observe that the proposed strategy leads to more accurate parameter estimates with significantly better time efficiency.

2 BACKGROUND AND PROBLEM STATEMENT

This work focuses on estimating the parameters Θ_{SEIR} in the SEIR mode. A standard SEIR model is formulated in terms of 4 populations of individuals: the susceptible population (S), the exposed population (E), the infectious population (I) and the recovered population (R), which are respectively comprised of all individuals susceptible to the infection, all individuals that have contacted infected patients but are currently during their incubation period, infected individuals that can transmit the disease to susceptible population, and recovered individuals that cannot become infected again and cannot transmit the disease to others. Given a mobility network $\mathcal{G}_t = (\{\mathcal{P}, \mathcal{L}\}, \mathcal{E})_t$ at time step $t (t = 1, 2, \dots, T)$ where \mathcal{P} is the set of person nodes, \mathcal{L} is the set of location nodes and \mathcal{E} is the set of edges linking persons to locations, the SEIR model needs to predict, what the numbers of individuals for these populations are in different locations at time step t , according to the preceding populations progress, the current mobility network and a set of parameters controlling the probability of transferring from one state of population to the next state.

Formally, we use the superscript l to represent a specific location l and use the subscript t to represent the time point t , e.g., S_t^l is the number of individuals for state S in location l at time step t . The number for all populations in location l at time step t is denoted by $N_t^l = S_t^l + E_t^l + I_t^l + R_t^l$, and $S_t = \sum_l S_t^l$ is the population for all locations at time t (the same for E_t, I_t, R_t). We further assume a constant overall population $N = N_t = S_t + E_t + I_t + R_t, \forall t$. There are three sets of parameters β, κ and γ controlling how likely that a person currently at one particular state would transfer to the next state. β controls the probability of transferring from S to E , and κ and γ are respectively responsible for transferring from E to I and I to R . β can be a set of values, representing the transmission rate for different location categories or cities. κ and γ can also be a set of values, representing the infection and recovery rates for different population groups. With a fixed set of parameters and the mobility network, the SEIR model can simulate the disease spread process and predict the numbers of these four populations at each time step for a location. The probability of an individual transferring from

one state to the next at location l for time step t can be formulated as:

$$\begin{cases} p_t^l(S \rightarrow E) = \beta \frac{I_t^l}{N_t^l} \\ p_t^l(E \rightarrow I) = \kappa \\ p_t^l(I \rightarrow R) = \gamma \end{cases} \quad (1)$$

During simulations, the state for each individual is sampled. An illustration is shown in Figure 1. The sum of the L_2 distances between the observed infections and the predicted infections over all time steps is the simulation error for a particular set of parameters, and the best parameter combination is selected to minimize the simulation error:

$$\Theta_{\text{SEIR}}^* = \beta^*, \kappa^*, \gamma^* = \arg \min_{\beta, \kappa, \gamma} \sum_t \|I_t - \hat{I}_t\|^2 \quad (2)$$

We use the hat notation $\hat{\cdot}$ to represent the prediction made by the model. Towards finding the (near) optimal parameter combination in Eq.(2), the SEIR models needs to run Eq.(1) over all individuals for T time steps, which gives a total time complexity of $\mathcal{O}(NT)$.

3 MODELING THE SEIR PROCESS USING RECURRENT NETS

3.1 OVERVIEW

The core idea of the proposed framework is that, instead of using the brute force grid search strategy to obtain the parameter set that leads to the least L_2 distance between predictions and observations, we train a recurrent net to fit the time-series data of SEIR based on a small number of simulations using different sets of parameters. The trained recurrent net can then be used to directly find the optimal value of Θ_{SEIR}^* , which denotes the parameter set for SEIR that leads to the optimal predictions. We employ the widely used Long Short-Term Memory network (LSTM) (Hochreiter & Schmidhuber, 1997) as the recurrent net model backbone.

3.2 SIMULATION DATASET CONSTRUCTION

The first step of the proposed framework is to generate simulation data, which will be used to train the recurrent LSTM net. We first use different sets of Θ_{SEIR} to perform simulations on the predefined network, as in Eq.(1). We perform M' simulations with M' different sets of Θ_{SEIR} , denoted by $\{\Theta_{\text{SEIR}}^{(m')}\}_{m'=1}^{M'}$. At each time step t , we iterate over all individuals, assign the previous state (i.e., S , E , I or R) of an individual, and sample its current state based on the SEIR model. Then, we sum up all individuals belonging to the same states, and obtain $\{S_t^{(m')}, E_t^{(m')}, I_t^{(m')}, R_t^{(m')}\}$ for time step t .

3.3 TRAINING THE RECURRENT NETS

Next, we train an recurrent net to fit $\{S_t^{(m')}, E_t^{(m')}, I_t^{(m')}, R_t^{(m')}\}$. Our goal is to use LSTM to predict the values of $\{S_t^{(m')}, E_t^{(m')}, I_t^{(m')}, R_t^{(m')}\}$ for each time step t for each parameter combination in $\{\Theta_{\text{SEIR}}^{(m')}\}_{m'=1}^{M'}$. Specifically, for each time step t , the input to the LSTM model is denoted by $\mathbf{x}_t^{(m')} = \{S_{t-1}^{(m')}, E_{t-1}^{(m')}, I_{t-1}^{(m')}, R_{t-1}^{(m')}\}$, the populations of the four states at the previous time step. The model should also be aware of the value of $\Theta_{\text{SEIR}}^{(m')}$ and the mobility network structure at time step t because they directly decide the number of individuals for the four states for the next time step. To serve this propose, we also feed $\Theta_{\text{SEIR}}^{(m')}$ as input to the LSTM model at time step t . For the time-varying mobility network \mathcal{G}_t , we map it to a vector representation $\mathbf{h}_{\mathcal{G}_t}$, which is able to capture its mobility structure and be conveniently fed to the LSTM network.

Mapping \mathcal{G}_t to $\mathbf{h}_{\mathcal{G}_t}$ We use the DIFFPOOL model (Ying et al., 2018) to map the mobility network \mathcal{G}_t at time step t to its high-dimensional representation $\mathbf{h}_{\mathcal{G}_t}$. DIFFPOOL is a differentiable pooling model that can generate hierarchical representations of graphs by progressively clustering nodes into a coarser graph. Specifically, at each layer l , DIFFPOOL learns an assignment matrix $\mathbf{S}^{(l)} \in \mathbb{R}^{n_l \times n_{l+1}}$ to assign each node at layer l to a cluster in the next layer $l + 1$, where n_l is the number of nodes

(clusters) at layer l . Given node embeddings $\mathbf{Z}^{(l)} \in \mathbb{R}^{n_l \times d}$ and the adjacency matrix $\mathbf{A}^{(l)} \in \mathbb{R}^{n_l \times n_l}$ at layer l , DIFFPOOL generates the new node embeddings $\mathbf{Z}^{(l+1)} \in \mathbb{R}^{n_{l+1} \times d}$, the new adjacency matrix $\mathbf{A}^{(l+1)} \in \mathbb{R}^{n_{l+1} \times n_{l+1}}$ and the new assignment matrix $\mathbf{S}^{(l+1)} \in \mathbb{R}^{n_{l+1} \times n_{l+2}}$ by applying the following equations:

$$\begin{aligned} \mathbf{Z}^{(l+1)} &= \text{GNN}_{l,\text{embed}}(\mathbf{A}^{(l+1)}, \mathbf{X}^{(l+1)}), \mathbf{S}^{(l+1)} = \text{softmax}\left(\text{GNN}_{l,\text{pool}}(\mathbf{A}^{(l+1)}, \mathbf{X}^{(l+1)})\right) \\ \mathbf{X}^{(l+1)} &= \mathbf{S}^{(l)\top} \mathbf{Z}^{(l)}, \mathbf{A}^{(l+1)} = \mathbf{S}^{(l)\top} \mathbf{A}^{(l)} \mathbf{S}^{(l)} \end{aligned} \quad (3)$$

$\text{GNN}_{l,\text{embed}}$ and $\text{GNN}_{l,\text{pool}}$ are two distinctly parameterized GNN, and are respectively responsible for generating new embeddings and producing a distribution over next-layer clusters. Setting the number of clusters at the last layer L to 1, DIFFPOOL outputs a single high-dimensional graph representation. We train DIFFPOOL to classify the graph \mathcal{G}_t at time step t to the category label t , and use the extracted graph representation $\mathbf{h}_{\mathcal{G}_t}$ as input to the LSTM.

Training LSTMs Θ_{LSTM} Given $\mathbf{h}_{t-1}^{(m')}$, $\mathbf{x}_t^{(m')}$, $\Theta_{\text{SEIR}}^{(m')}$ and $\mathbf{h}_{\mathcal{G}_t}$, we are able to obtain the hidden vector representation $\mathbf{h}_t^{(m')}$ for the time step t :

$$\begin{bmatrix} \mathbf{i}_t \\ \mathbf{f}_t \\ \mathbf{o}_t \\ \mathbf{l}_t \end{bmatrix} = \begin{bmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{bmatrix} \mathbf{W} \cdot \begin{bmatrix} \mathbf{h}_{t-1}^{(m')} \\ \mathbf{x}_t^{(m')} \\ \mathbf{h}_{\mathcal{G}_t} \\ \Theta_{\text{SEIR}}^{(m')} \end{bmatrix} \quad (4)$$

$$\mathbf{c}_t = \mathbf{f}_t \cdot \mathbf{c}_{t-1} + \mathbf{i}_t \cdot \mathbf{l}_t \quad (5)$$

$$\mathbf{h}_t^{(m')} = \mathbf{o}_t \cdot \tanh(\mathbf{c}_t) \quad (6)$$

where $\mathbf{W}_i, \mathbf{W}_f, \mathbf{W}_o, \mathbf{W}_l \in \mathbb{R}^{K \times (2K+4+|\Theta_{\text{SEIR}}|)}$ where K is the dimensionality of $\mathbf{h}_t^{(m')}$ and $\mathbf{h}_{\mathcal{G}_t}$. $\mathbf{h}_t^{(m')}$ is then passed to a fully connected layer to obtain $\hat{\mathbf{h}}_t^{(m')}$, which is mapped to scalars to predict $\{S_t^{(m')}, E_t^{(m')}, I_t^{(m')}, R_t^{(m')}\}$:

$$\begin{aligned} \hat{S}_t^{(m')} &= \mathbf{h}_S^\top \times \hat{\mathbf{h}}_t^{(m')} \\ \hat{E}_t^{(m')} &= \mathbf{h}_E^\top \times \hat{\mathbf{h}}_t^{(m')} \\ \hat{I}_t^{(m')} &= \mathbf{h}_I^\top \times \hat{\mathbf{h}}_t^{(m')} \\ \hat{R}_t^{(m')} &= \mathbf{h}_R^\top \times \hat{\mathbf{h}}_t^{(m')} \end{aligned} \quad (7)$$

where $\mathbf{h}_S, \mathbf{h}_E, \mathbf{h}_I, \mathbf{h}_R \in \mathbb{R}^{K \times 1}$. The training objective is minimizing the distance between simulation outputs $\{S_t^{(m')}, E_t^{(m')}, I_t^{(m')}, R_t^{(m')}\}$ and LSTM predictions $\{\hat{S}_t^{(m')}, \hat{E}_t^{(m')}, \hat{I}_t^{(m')}, \hat{R}_t^{(m')}\}$:

$$\Theta_{\text{LSTM}}^* = \arg \min_{\Theta_{\text{LSTM}}} \sum_{t,m'} [\|S_t^{(m')} - \hat{S}_t^{(m')}\|^2 + \|E_t^{(m')} - \hat{E}_t^{(m')}\|^2 + \|I_t^{(m')} - \hat{I}_t^{(m')}\|^2 + \|R_t^{(m')} - \hat{R}_t^{(m')}\|^2] \quad (8)$$

Eq. 8 can be trained in an end-to-end fashion to obtain optimal Θ_{LSTM} .

3.4 FINDING OPTIMAL Θ_{SEIR}

The LSTM model with Θ_{LSTM}^* is able to generalize the behavior of the SEIR model with a specific value of Θ_{SEIR} . When training the LSTM to learn Θ_{LSTM}^* , Θ_{SEIR} is set to a fixed value of $\Theta_{\text{SEIR}}^{(m')}$ and fed as input to LSTMs at each time step. Due to the fact that Θ_{SEIR} can also be viewed as parameters in LSTMs, i.e., the input to each time step, we can fix Θ_{LSTM}^* and relax Θ_{SEIR} , treating Θ_{SEIR} as learnable parameters, to minimize Eq.(2) which we write down here for reference:

$$\Theta_{\text{SEIR}}^* = \beta^*, \kappa^*, \gamma^* = \arg \min_{\beta, \kappa, \gamma} \frac{1}{T} \sum_t \|I_t - \hat{I}_t\|^2 \quad (2)$$

where \hat{I}_t is the output from LSTM, and I_t is the observation data rather than the simulation data used to train the LSTM. Eq.(2) is differentiable with respect to Θ_{SEIR} and can be trained in an end-to-end

fashion based on SGD (Kiefer et al., 1952; Rumelhart et al., 1986). To this end, we learn the optimal values of Θ_{SEIR} that minimize the L_2 distance between predicted infections and observations.

In the case where we have prior knowledge about the values Θ_{SEIR} , e.g., all values in Θ_{SEIR} should be larger than 0, the value of β is usually smaller than 0.1 based on clinical observations for COVID-19 (Chang et al., 2021), we can incorporate regularizers as side objectives:

$$\Theta_{SEIR}^* = \arg \min_{\Theta_{SEIR}} = \sum_t \frac{1}{T} \|I_t - \hat{I}_t\|^2 + \lambda \|\Theta_{SEIR} - \text{prior}(\Theta_{SEIR})\|^2 \quad (9)$$

where $\text{prior}(\Theta_{SEIR})$ denotes the human prior knowledge regarding the values of Θ_{SEIR} , and λ controls the trade-off. We will explore the effects of $\text{prior}(\Theta_{SEIR})$ and λ in experiments.

4 EXPERIMENTS

4.1 DATASETS AND CORRESPONDING SEIR MODELS

We use two public datasets for evaluations: the infection network of Covid-19 in China (Covid-China) (Sun et al., 2021; Liu et al., 2021), and the infection network of Covid-19 in the US (Covid-US) (Chang et al., 2021).

Covid-China consists of infection networks for 31 provinces in China from Apr 2020 to Feb 2021, extracted from action tracking reports of Covid-19 patients. The network of Covid-China consists of two types of nodes: patient and location. Time-varying edges are constructed between a patient node and a location node if the patient visited the location at time t . Each location takes an attribute from 11 categories of locations: households, workplaces, hotels, supermarkets, banks, restaurants, parks, barber shops/hairdressers, trains, buses, and airplanes, and the economic city tier (first, second or third) that it belongs to. A patient node is characterized by features of age (taking the value of children, youths, adults or seniors) and gender (taking the value of male and female). Each attribute for gender, age, city-tier and location type is associated with a specific transmission rate β . The transmission rate for a certain person node of gender s , age a in location of type c in a city of tier t is the additive combination of corresponding β :

$$\beta(s, a, t, c) = \beta_s + \beta_a + \beta_t + \beta_c \quad (10)$$

$\beta = \{\beta_s, \beta_a, \beta_t, \beta_c\}$, along with γ and κ are parameters to learn. The network for each city is sliced into consecutive time snippets, with the size of stride set to two weeks. For each city, we have daily gold number of infections, public by Chinese CDC. These gold numbers of infections are used to learn Θ_{SEIR} . Snippets without any infection are removed. Each time step of each city is labeled with gold number of infections, which is used to train the SEIR parameters. Snippets are divided to 80%/10%/10% for training, dev and test.

Covid-US consists of networks that capture hourly visits from each population group to each location in 10 metro areas in the US. The network is extracted from mobility data provided by the SafeGraph application. The network consists of two types of nodes: population group and location. A time-varying edge with weight $w_{i,j}$ is constructed, if at time t , the number of people from population group i visiting location j is $w_{i,j}$. $w_{i,j}$ is column-normalized. Each location is associated with a location category (e.g., full-service restaurants, grocery stores, etc), and each population group is associated with its race and median income. For the SEIR model, transmissions can happen within groups or across population groups when two people from two groups visit the same location. Simulations are performed at the population group level. Each category location c is associated with a specific transmission rate β_c , which captures the inter-group transmissions across groups in the locations. For each population group with race r and income decile (i), the intra-group transmission is set to $\beta = \beta_r + \beta_i$. Each area is associated with gold number of infections published by the The New York Times¹. Each snippet consists of the network for a single city and lasts a week. We divided snippets to 80%/10%/10% for training, dev and test.

4.2 EXPERIMENTAL DETAILS

Generating Simulated Data We first need to sample Θ_{SEIR} . We limit the value of each β to the scope of $[0, 0.1]$, and we randomly sample its value within the scope. For κ and γ , based on previous

¹<https://github.com/nytimes/covid-19-data>

# Simulations	Vanilla	LSTM	BayesOpt	Regression-ABC	ABC-MCMC
<i>Covid-China</i>					
20	34.2	22.2	27.5	25.4	29.1
100	30.1	13.1	17.0	16.2	18.3
500	18.7	9.9	13.2	12.9	14.0
1000	15.2	9.2	11.7	11.1	12.4
5000	13.5	8.6	9.5	9.3	10.4
<i>Covid-US</i>					
20	1870	1250	1491	1371	1538
100	1530	1130	1289	1275	1438
500	1320	930	1035	1004	1175
1000	1120	824	965	932	1013
5000	970	674	780	776	812

Table 1: Average square L_2 distances for vanilla grid search simulation, recurrent LSTMs, Bayesian optimization, regression-ABC and ABC-MCMC on Covid-China and Covid-US.

clinical observations (Kucharski et al., 2020) where κ^{-1} is around 96 hours and γ^{-1} is around 84 hours, we sample κ and γ from a normal distribution with expectation set to 96 and 84. Given a sampled set of Θ_{SEIR} , we run simulations on the training datasets to obtain the simulation data, i.e., the number of individuals for all four states for each time step. For each episode, we take K samples of Θ_{SEIR} , leading to a total number of $K * |\text{train}|$ training sequences, where $|\text{train}|$ denotes the number of training episodes.

Learning LSTMs Θ_{LSTM} to Fit Simulated Data We split the simulated data to 90/10 for training and validation. We train a three-layer LSTM with residual connections (He et al., 2016; Kim et al., 2017) to fit the simulated time-series data using for training, based on Eq.(8). Then size of hidden states is set to 128. The value of batch size is set to 256, and SGD is used for optimization. LSTM parameters and embeddings are initialized from a uniform distribution in $[-0.08, 0.08]$. Gradient clipping is adopted by scaling gradients when the norm exceeds a threshold of 1. Dropout rate, learning rate and the number of training epochs are treated as hyper-parameters to be tuned on the dev set.

Learning Θ_{SEIR} We optimize Θ_{SEIR} based on Eq.(9) on the number of gold daily infections, using the daily reported infections as labels. We use AdaGrad (Duchi et al., 2011) for optimization. Dropout rate, learning rate, the number of training epochs, and the hyper-parameter λ are tuned on the dev set.

4.3 RESULTS

For the first baseline, we search the optimal Θ_{SEIR} using vanilla grid search SEIR simulations. For each set of Θ_{SEIR} , simulations are performed on all training episodes, and the parameter set that leads to the minimum L_2 loss is selected as the final value. Suppose that we conduct K explorations for Θ_{SEIR} . This means we need to perform K simulations on each training episode. The K is here thus comparable to and the same as the K for simulation data generation. We also implement the BayesOpt algorithm using the open-sourced package,² along with the regression-ABC model described in (Saulnier et al., 2017) and the ABC-MCMC model (Sunnåker et al., 2013).

We report the average of the square of L_2 distance between the predicted number and reported number of infections on the test episodes with varying number of simulations K on the test set. Lower values indicate superior models. Results are shown in Table 1. As can be seen, the proposed LSTM model performs better than BayesOpt, regression-ABC and ABC-MCMC, which in turn performs better than the vanilla model, showing the superiority of the proposed model over BayesOpt. The explanation is that (1) the neural network-based LSTM provides with better expressivity than the Bayesian baselines; and (2) the neural model has better capacity of capturing the pattern in the time-series data of the SEIR model, leading to better performances in general. Comparing the vanilla simulation and recurrent LSTM, we have the following observations: (1) as the number of simulations

²<https://github.com/rmcantin/bayesopt>




Figure 2: Daily reported cases, predictions made by the vanilla SEIR model, and the predictions made by LSTM for three cities in China – Shijiazhuang, Harbin and Suihua. Curves are smoothed by 5-day average.

$\log_{10}(\lambda)$	0	1	2	3	4	5	6	7
Covid-China	10.4	9.9	9.6	9.4	10.1	11.4	13	23.5
Covid-US	890	866	861	842	832	824	899	1423

Table 2: The effect of different values of λ on Covid-China and Covid-US. We show the logarithm with base 10 for λ . The average square L_2 distances are reported.

K increases, the performances for both the vanilla model and the proposed model improve. This is in accord with our expectations: for the vanilla model, a larger number of simulations means that the model is able to explore the search space more thoroughly to obtain the optimal value; for the proposed LSTM model, the model learns better with more training data and avoids overfitting; (2) with the same number of K , the proposed LSTM model performs significantly better than the vanilla brute-force search model. This is due to the generalization ability of proposed framework: the vanilla model can only select the optimal parameters from the set it tries, while the proposed framework can generalize to the un-tried parameter set; and (3) notably, the proposed framework is able to achieve comparable performance to the vanilla search model with significantly smaller number of simulations. Specifically, for Covid-China, the performance obtained with 100 simulations (13.1) is comparable to the vanilla model with 5,000 simulations (13.5); for Covid-US, the performance obtained with 500 simulations (930) is comparable to the vanilla model with 5,000 simulations (970). This further illustrates the superiority of the proposed framework.

Figure 2 shows simulations performed on the four test episodes in Covid-China dataset using parameters learned from the proposed model and the vanilla SEIR model. As can be seen, the proposed framework offers more accurate predictions.

4.4 ABLATION STUDIES

In this subsection, we explore the effect of different modules, along with hyper-parameters in the proposed framework to explore their influence.

The effect of λ The hyper-parameter controls the trade-off between observations and external prior knowledge. The effect of λ is shown in Table 2. As can be seen, model performance first improves and then declines as the value of λ grows. Finding the sweep spot for the balance between observations and prior knowledge leads to the best performance.

The effect of Recurrent Structures We also conduct experiments using other recurrent net structures, including vanilla recurrent net (RNN), Gated Recurrent Unit (GRU) (Cho et al., 2014), and Simple Recurrent Units (SRU) (Lei et al., 2017). Results for different recurrent structures are shown in Table 3. As can be seen, the LSTM structure performs comparable to GRU (slightly worse than GRU on Covid-China and slightly better on Covid-US), better than the vanilla recurrent net and SRU.

How to incorporate Θ_{SEIR} into LSTMs We explore the effects of different ways to incorporate Θ_{SEIR} in the LSTM model, including the current strategy of (1) Θ_{SEIR} being concatenated with the input x_t at each time step (**Each**); (2) Θ_{SEIR} being incorporated only at the first time step (**First**); and (3) Θ_{SEIR} element-wise multiplies (Hadamard product) the input for each time step (**Hadamard**).

	CN	US
LSTM	9.4	824
GRU	9.1	831
RNN	12.1	890
SRU	11.5	872

Table 3: The effect of recurrent structures.

	CN	US
Each	9.4	824
First	13.5	1007
Hadamard	9.1	821

Table 4: The effect of using different strategies to incorporate SEIR parameters.

	CN	US
w/o Graph	24.2	2531
Constant	18	2013
Varying	9.4	824

Table 5: The effect of graph representations.

For (3), since the dimensionalities of Θ_{SEIR} and x_t are different, x_t is first passed to an FFN, the output of which has the same dimensionality with Θ_{SEIR} . Results for the three strategies are shown in Table 4. As can be seen, Θ_{SEIR} incorporated only at the first time step significantly underperforms the strategy that incorporates Θ_{SEIR} at every time step. This is because of the gradient vanishing effect of recurrent nets: reminding the model of Θ_{SEIR} at each time leads to better performances. The concatenation strategy obtains comparable performances to the Hadamard product strategy.

The Effect of Graph Representation The time-varying location-population network is captured by the graph representations through DIFFPOOL. This is critical since the number of infections highly relies on the spreading network for each time. We explore its necessity by comparing it with other variants: (1) no network (**w/o Graph**): where no time-varying graph embedding is incorporated; (2) constant graph representation (**Constant**): the graph embedding is not time-varying, where we use the graph embedding of the first time step for all time steps; (3) time-varying graph representation (**Varying**): the strategy adopted in this work where time-varying representations are incorporated. Results are shown in Table 5. As can be seen, when no network information is incorporated, the model nearly fails to learn anything. Constant networks perform slightly better than no network, but still significantly worse than the time-vary graph representations. This is in accord with our expectations since the time-varying network decides the number of infections and the disease spread patterns at each time.

5 RELATED WORKS

The classical compartmental models simplify the mathematical modeling of disease spread by simulating the population transitions between different states in the disease spread process. Since the outbreak of Covid-19, the SEIR model has been widely used to model the Covid-19 spread around the world (Estrada, 2020; Godio et al., 2020; Grimm et al., 2021) and provide important insights regarding isolation policies (Endo et al., 2020; Carcione et al., 2020; Chang et al., 2021) and vaccine delivery strategies (Bubar et al., 2021; Ghostine et al., 2021). To achieve a faster simulation speed and better simulation results, recent works have proposed to leverage deep neural networks in place of SEIR models to predict pandemic dynamics over time. For example, (Yang et al., 2020) used the LSTM model to predict the numbers of new infections given the contact statistics and the pre-selected transmission/incubation/recovery/death rates. (Gao et al., 2021) incorporated sequential network structures and graph attention to predict the number of infections upon a temporal and spatial mobility graph. These works aim at taking advantage of SEIR models to inform effective strategies in response to the disease spread.

With regard to the estimation of the parameters in SEIR models, a simple approach is to enumerate parameter combinations, run the SEIR model with each combination and select the one with the smallest error. An alternative to grid search is to use approximate Bayesian computation (ABC) (Sunnåker et al., 2013; Saulnier et al., 2017; Raynal et al., 2019), a technique that maintains a small fraction of simulations that are close to the target statistics in the light of the computed distance. These simulations are treated as posterior distributions of the SEIR parameters, which are then used to infer the optimal parameters. Another approach to estimating parameters is to view the process of population transitions between states as a problem of ordinary differentiable equations (ODEs) (Kermack & McKendrick, 1927; Hethcote, 2000; Harko et al., 2014). However, ODEs can only give approximate numerical solutions, which could be inaccurate for real-world modeling.

The most relevant work is from (Tessmer et al., 2018), who proposed to make direct use of existing neural networks to predict the basic reproduction number (R_0), the number of secondary cases

generated by an infectious individual in a fully susceptible host population. This work is different from (Tessmer et al., 2018) in that (1) they sought to estimate the basic reproduction number R_0 and we estimate the parameters in SEIR (or other SEIR variants) models; and more importantly (2) they propose to directly output the number to estimate given inputs, whereas we propose to automatically learn the parameters through neural network gradient descent and back-propagation. The proposed method can be extended to other fields that require time-consuming simulations to estimate necessary parameters.

6 CONCLUSION

In this work, we propose to transform the original non-differentiable simulation problem of SEIR parameter estimation into a differentiable one by leveraging neural recurrent nets. The recurrent net is first trained to fit a small number of simulation data, and then trained on the observation data to derive the optimal SEIR parameters. This strategy bypasses the needs of time-consuming simulations, and automatically induces the optimal parameters via gradient descent, leading to both accuracy and efficiency gains.

REFERENCES

- Alberto Aleta, David Martin-Corral, Ana Pastore y Piontti, Marco Ajelli, Maria Litvinova, Matteo Chinazzi, Natalie E Dean, M Elizabeth Halloran, Ira M Longini Jr, Stefano Merler, et al. Modelling the impact of testing, contact tracing and household quarantine on second waves of covid-19. *Nature Human Behaviour*, 4(9):964–971, 2020.
- Ottar N Bjørnstad, Bärbel F Finkenstädt, and Bryan T Grenfell. Dynamics of measles epidemics: estimating scaling of transmission rates using a time series sir model. *Ecological monographs*, 72(2):169–184, 2002.
- Kate M Bubar, Kyle Reinholt, Stephen M Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H Grad, and Daniel B Larremore. Model-informed covid-19 vaccine prioritization strategies by age and serostatus. *Science*, 371(6532):916–921, 2021.
- José M Carcione, Juan E Santos, Claudio Bagaini, and Jing Ba. A simulation of a covid-19 epidemic based on a deterministic seir model. *Frontiers in public health*, 8:230, 2020.
- Serina Chang, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, and Jure Leskovec. Mobility network models of covid-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, 2021.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- Akira Endo et al. Estimating the overdispersion in covid-19 transmission using outbreak sizes outside china. *Wellcome Open Research*, 5, 2020.
- Ernesto Estrada. Covid-19 and sars-cov-2. modeling the present, looking at the future. *Physics Reports*, 2020.
- Junyi Gao, Rakshith Sharma, Cheng Qian, Lucas M Glass, Jeffrey Spaeder, Justin Romberg, Jimeng Sun, and Cao Xiao. Stan: spatio-temporal attention network for pandemic prediction using real-world evidence. *Journal of the American Medical Informatics Association*, 28(4):733–743, 2021.
- Rabih Ghostine, Mohamad Gharamti, Sally Hassrouny, and Ibrahim Hoteit. An extended seir model with vaccination for forecasting the covid-19 pandemic in saudi arabia using an ensemble kalman filter. *Mathematics*, 9(6):636, 2021.

- Alberto Godio, Francesca Pace, and Andrea Vergnano. Seir modeling of the italian epidemic of sars-cov-2 using computational swarm intelligence. *International Journal of Environmental Research and Public Health*, 17(10):3535, 2020.
- Veronika Grimm, Friederike Mengel, and Martin Schmidt. Extensions of the seir model for the analysis of tailored social distancing and tracing approaches to cope with covid-19. *Scientific Reports*, 11(1):1–16, 2021.
- Tiberiu Harko, Francisco SN Lobo, and MK Mak. Exact analytical solutions of the susceptible-infected-recovered (sir) epidemic model and of the sir model with equal death and birth rates. *Applied Mathematics and Computation*, 236:184–194, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- David G Kendall. Deterministic and stochastic epidemics in closed populations. In *Contributions to Biology and Problems of Health*, pp. 149–166. University of California Press, 2020.
- William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- Jack Kiefer, Jacob Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952.
- Jaeyoung Kim, Mostafa El-Khamy, and Jungwon Lee. Residual lstm: Design of a deep recurrent architecture for distant speech recognition. *arXiv preprint arXiv:1701.03360*, 2017.
- Adam J Kucharski, Timothy W Russell, Charlie Diamond, Yang Liu, John Edmunds, Sebastian Funk, Rosalind M Eggo, Fiona Sun, Mark Jit, James D Munday, et al. Early dynamics of transmission and control of covid-19: a mathematical modelling study. *The lancet infectious diseases*, 20(5):553–558, 2020.
- Tao Lei, Yu Zhang, Sida I Wang, Hui Dai, and Yoav Artzi. Simple recurrent units for highly parallelizable recurrence. *arXiv preprint arXiv:1709.02755*, 2017.
- Xiao Fan Liu, Xiao-Ke Xu, and Ye Wu. Mobility, exposure, and epidemiological timelines of covid-19 infections in china outside hubei province. *Scientific data*, 8(1):1–7, 2021.
- Kiesha Prem, Yang Liu, Timothy W Russell, Adam J Kucharski, Rosalind M Eggo, Nicholas Davies, Stefan Flasche, Samuel Clifford, Carl AB Pearson, James D Munday, et al. The effect of control strategies to reduce social mixing on outcomes of the covid-19 epidemic in wuhan, china: a modelling study. *The Lancet Public Health*, 5(5):e261–e270, 2020.
- Louis Raynal, Jean-Michel Marin, Pierre Pudlo, Mathieu Ribatet, Christian P Robert, and Arnaud Estoup. Abc random forests for bayesian parameter inference. *Bioinformatics*, 35(10):1720–1728, 2019.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- Emma Saulnier, Olivier Gascuel, and Samuel Alizon. Inferring epidemiological parameters from phylogenies using regression-abc: A comparative study. *PLoS computational biology*, 13(3):e1005416, 2017.
- Xiaofei Sun, Tianjia Guan, Tao Xue, Chun Fan, Meng Yang, Yuxian Meng, Tianwei Zhang, Bahabaik Jiangtulu, Fei Wu, and Jiwei Li. Analysis on action tracking reports of covid-19 informs control strategies and vaccine delivery in post-pandemic era. *medRxiv*, 2021.

Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate bayesian computation. *PLoS Comput Biol*, 9(1):e1002803, 2013.

Heidi L Tessmer, Kimihito Ito, and Ryosuke Omori. Can machines learn respiratory virus epidemiology?: A comparative study of likelihood-free methods for the estimation of epidemiological dynamics. *Frontiers in microbiology*, 9:343, 2018.

Zifeng Yang, Zhiqi Zeng, Ke Wang, Sook-San Wong, Wenhua Liang, Mark Zanin, Peng Liu, Xudong Cao, Zhongqiang Gao, Zhitong Mai, et al. Modified seir and ai prediction of the epidemics trend of covid-19 in china under public health interventions. *Journal of thoracic disease*, 12(3):165, 2020.

Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*, 2018.