

Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Signature:

Yujia Hao

April 3, 2023

Modeling the COVID-19 Pandemic: A Model-Driven and a Data-Driven Analyses

By

Yujia Hao

Alessandro Veneziani, Ph.D.
Advisor

Mathematics

Alessandro Veneziani, Ph.D.
Advisor

Elizabeth Newman, Ph.D.
Committee Member

Ruoxuan Xiong, Ph.D.
Committee Member

2023

Modeling the COVID-19 Pandemic: A Model-Driven and a Data-Driven Analyses

By

Yujia Hao

Alessandro Veneziani, Ph.D.
Advisor

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences of
Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Mathematics

2023

Abstract

Modeling the COVID-19 Pandemic: A Model-Driven and a Data-Driven Analyses
By Yujia Hao

The COVID-19 pandemic has sparked a plethora of mathematical models, each constructed with varying assumptions and methodologies. However, the nature of infectious disease dynamics, which are inherently complex and influenced by numerous factors, poses challenges to the development of accurate models. In this project, we propose a space-dependent compartmental model based on the classical Susceptible, Infected, Recovered (SIR) epidemiological model as well as a data-driven approach that combines the SIR model and data assimilation through physics-informed neural networks (PINN) using generated training data. The space-dependent model allows for control over each geographical unit (GU) in a global domain, making the problem scalable, while time-dependent parameters are included to simulate the effect of interventions like lockdowns and vaccination campaigns. Following the space-dependent model, we investigate the most influential parameters of the basic SIR model through a first-order local sensitivity analysis. The results show that the reproduction/death rate has the most impact on all compartments of the SIR model. This analysis provides insights into which parameters should be prioritized in future studies and can help in developing more effective interventions. Finally, the PINN approach based on SIR models exhibits satisfactory predictive capability for parameter estimation and dynamics simulation, even with limited and noisy data. This project is a promising start to the modelling of intricate and multifaceted dynamics pertaining to the transmission of infectious diseases.

Modeling the COVID-19 Pandemic: A Model-Driven and a Data-Driven Analyses

By

Yujia Hao

Alessandro Veneziani, Ph.D.
Advisor

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Mathematics

2023

Acknowledgments

I would like to express my deep gratitude to Dr. Alessandro Veneziani for his mentorship since my sophomore year. I am inspired by his passion and insights towards mathematics as well as his kindness and patience towards an undergraduate student. Working with him has been a true pleasure that has led to my aspiration of becoming a full-fledged researcher.

I would also like to thank the members of my research teams. In the multi-regional COVID project, I worked with Yuting Hou, Siwei Xu, Kai Chang, and Zhen Wu. We had a great time constructing the solver and cross-validating the results. The machine learning part of the project is inspired by my research in the summer of 2022 with Kishore Basu, Delphine Hintz, and Dev Shah. Together we learned about physics-informed machine learning as well as its great potentials in applications.

Finally, I would like to thank my family and friends for their unwavering support and encouragement throughout my undergraduate journey. Their love and support have been essential in keeping me motivated and focused on achieving my goals.

This research would not have been possible without the support and contributions of these individuals and many others who have supported me along the way. Thank you all for your invaluable assistance and support.

Contents

1	Introduction	1
2	Mono-region SIR-Like Models for the Outbreak	3
2.1	The basic SIR model	3
2.2	Asymptotic behaviors of the SIR model	5
3	Inter-regional Mobility	6
3.1	The mobility matrix	6
3.2	The numerical solver	8
3.3	Design of the objected-oriented solver	8
3.3.1	The local and global models	9
3.4	Benchmarks and results	10
3.4.1	Consistency test with theorems for one-region models	10
3.4.2	Cross-validation	11
3.4.3	Epidemic vs. endemic	12
3.4.4	Effects of lockdown	13
4	Sensitivity Analysis for Mono-region Models	17
4.1	First-order local sensitivity analysis on the endemic model	18
5	Parameter Estimation and Prediction with Physics-Informed Neural Networks	22

5.1	Introduction to physics-informed neural networks (PINN)	23
5.2	Using PINN for parameter estimation and prediction of the SIR en- demic model	27
5.2.1	Using PINN for noisy data	28
5.2.2	Using PINN for limited data	30
6	Conclusions: Limitations and Perspectives	33
	Bibliography	37

List of Figures

2.1	SIR Compartmental Model	4
3.1	top: simulation results of the endemic SIR model for the GU “Alabama”; bottom: MATLAB output of comparison between simulated and real asymptotic values.	11
3.2	left: Results for our multi-regional MATLAB solver in “Georgia” and “Alabama”; right: results for the same case with the Python solver.	12
3.3	Results in “Georgia”: left: epidemic model ($\mu = 0$); right: endemic model.	13
3.4	Local evolution in “Georgia” without lockdown (left) and with lockdown (right).	14
3.5	left: without lockdown; right: with lockdown.	14
3.6	Multi-regional case with a lockdown in Georgia: left: Florida; right: New York.	15
3.7	Measurements for fighting the pandemic in a multi-regional solver, with focus on “Georgia”: left, lockdown; right, vaccination.	16
4.1	Sensitivities of the S , I , R compartments on: top left: β , the transmission rate; top right: γ , the recovery rate; bottom: μ , the reproduction/death rate.	20
5.1	An example of a feedforward neural network	23

5.2	Plots of commonly used activation functions.	25
5.3	left: training data for SIR endemic model; right: PINN prediction for SIR endemic model.	28
5.4	PINN predictions from noisy data.	30
5.5	PINN predictions from limited data.	32

Chapter 1

Introduction

Modeling the dynamics of an infectious disease like the COVID-19 outbreak, researchers need to not only incorporate and improve upon existing mathematical models but also utilize data-driven methods to link theoretical concepts to reality. In this project, we consider different aspects of the description of the outbreak, whose importance was recently defined by the pandemic and its implications on other infectious diseases. We first introduce a multi-regional model addressing the space-dependence of the outbreak by splitting a domain of interest (the “Domain”) into subdomains (the “Geographical Units”) connected by matrices representing the inter-regional mobility. With a robust model, policymakers may directly or indirectly alter the disease parameters within the model in order to control the pandemic. In this project, we alter the parameters by adding time dependence to mimic the realistic situations of lockdowns and vaccination campaigns. Therefore, it is crucial to analyze the sensitivity of the compartmental model on each parameter. For this reason, we perform a first-order local sensitivity analysis on the mono-regional endemic SIR model. Finally, as an alternative approach, we introduce physics-informed neural networks that combines the mono-regional SIR model with a standard neural network to perform parameter estimation and prediction from simulated limited or noisy data.

In **Chapter 2**, we consider the classical compartmental modeling, based on the Susceptible, Infected, and Recovered (SIR) populations. In **Chapter 3**, we introduce the construction of inter-regional mobility based on mono-regional models and discuss the corresponding design of an object-oriented solver. We present some benchmarks and results of the simulations. Note that the majority of Chapters 2 to 4 is an excerpt from a collaborative conference proceeding [10] for the CILAMCE-PANACM 2021 Conference with four other Emory students under the mentorship of Dr. Alessandro Veneziani. In **Chapter 4**, we present the sensitivity analysis for the mono-regional endemic SIR model. In **Chapter 5**, we introduce physics-informed neural networks and present results on parameter estimation and model predictions on the mono-regional model. In **Chapter 6**, we conclude this project and discuss future perspectives.

Chapter 2

Mono-region SIR-Like Models for the Outbreak

2.1 The basic SIR model

In a classical SIR model for a specific region, the total population, denoted with $N(t)$ individuals function of time t , is divided into three compartments: Susceptible $S(t)$, Infected $I(t)$, and Recovered $R(t)$. SIR models are categorized into two models: the *epidemic* model that describes a rapid outbreak of the infectious disease and the *endemic* model that portrays the long-term dynamics by including regional birth and death rates (see, e.g., [11]). Similarly to a predator-prey model, the infection occurs proportionally to the product SI . More precisely, the endemic model considered here reads

$$\begin{cases} d_t S = \mu N - \beta IS/N - (\mu + \rho)S \\ d_t I = \beta IS/N - \gamma I - \mu I \\ d_t R = \rho S + \gamma I - \mu R \end{cases} \quad (2.1)$$

Here, β is the number of the average adequate contacts for disease transmission per person per unit of time, and γ is the rate of recovery; μ is the reproduction/death

rate assumed to be constant for simplicity, and ρ is the vaccination rate which is proportional to S [17]. Should we set $\mu = 0$, dropping the reproduction/death rate (as reasonable in a short time range), and $\rho = 0$, we obtain the so-called “*epidemic*” model. The inclusion of vaccination is essential to its application in the real-world scenario. We are considering an idealized case, where vaccinated individuals are perfectly protected. More complex scenarios will be considered elsewhere. The vaccination coefficient ρ can be time-dependent as multiple factors can affect the vaccine distribution. We envision that the vaccination rate may range from a minimal value ρ_m , and increases to a maximum ρ_M as the organization and the advertising reach out for more people. We model this aspect by a logistic dynamics [4],

$$\rho(t) = \frac{\rho_m e^{a(t-t_V)}}{1 + \rho_m/\rho_M(e^{a(t-t_V)} - 1)} \quad (2.2)$$

where t_V is the first day of vaccinations. Notice that in this model, the total number of individuals N is constant (summing up the three equations, for $N = S + I + R$ we obtain $d_t N = 0$). No space dependence is indicated in this system; thus, it refers to an isolated geographical unit (GU). This system is intuitively illustrated in Figure 2.1.




Figure 2.1: SIR Compartmental Model

For the sake of notation, we set $\mathbf{A} \equiv [S, I, R]^T$ and the shorthand notation for the

model (2.1) reads: $d_t \mathbf{A} = \mathbf{F}(\mathbf{A})$ where we omit for brevity the dependence of \mathbf{F} on the parameters μ, β, γ, ρ .

2.2 Asymptotic behaviors of the SIR model

The SIR system is well-studied, and it is known to exhibit some analytical properties as demonstrated in [11]. We state the two theorems

Theorem 1 (Epidemic SIR system). *Let $\mu = \rho = 0$, $N = 1$ in Equation (2.1) and define $\sigma \equiv \beta/\gamma$. Let $(S(t), I(t))$ be a solution of this system. If $\sigma S(0) \leq 1$, then $I(t)$ decreases to zero as $t \rightarrow \infty$. If $\sigma S(0) > 1$, then $I(t)$ first increases up to a maximum value $I_{max} = I(0) + S(0) - 1/\sigma - [\ln(\sigma S(0))]/\sigma$ and then decreases to zero as $t \rightarrow \infty$. The susceptible fraction $S(t)$ is a decreasing function and the limiting value S_∞ is the unique root in $(0, 1/\sigma)$ of the equation*

$$I(0) + S(0) - S_\infty + \ln(S_\infty/S(0))/\sigma = 0.$$

Theorem 2 (Endemic SIR system). *Let $\rho = 0$, $N = 1$ in Equation (2.1) and define $\vartheta \equiv \beta/(\gamma + \mu)$. Let $(S(t), I(t))$ be a solution of this system. If $\vartheta \leq 1$, the solutions approach disease-free equilibrium given by $S = 1$ and $I = 0$; if $\vartheta > 1$, the all solutions approach the endemic equilibrium given by $S_e = 1/\vartheta$, $I_e = \mu(\vartheta - 1)/\beta$.*

Theorems 1 and 2 are verified numerically with results presented in Section 3.4.1.

Chapter 3

Inter-regional Mobility

One of the main limitations of our basic SIR model is the absence of space dependence and, in general, of the mobility of individuals across different geographical regions. The inclusion of such mobility would be critical, for instance, for assessing the impact of transportation and restrictions on the passengers mobility on the outbreak [2]. Space dependence can be included as systems of partial differential equations [22] or fractional differential operators (to incorporate preferential directions like highways or train lines). In this work, we resort to a simpler approach. We consider a region of interest as a set of geographical units (GU) where we assume an SIR dynamic. Then, we describe migration among the different regions in the form of “mobility matrices” that quantify the number of people leaving one region for another. In this way, we have a simplified description of the space dependence, yet potentially able to consider the impact of the mobility on the outbreak.

3.1 The mobility matrix

Let us consider now a set of n areas or GU denoted by the compartment vectors \mathbf{A}_i , $i = 1, 2, \dots, n$. Each area will have an intrinsic dynamic corresponding to the vector $\mathbf{F}_i(\mathbf{A}_i)$ as specified in Section 2.1. Then, we consider the mobility among the regions

as follows. As a simplifying assumption, we assume that *people of a compartment in an area do not change compartment when moving*: a susceptible in the GU j will keep the status of susceptible also in the target region i . The rate of susceptible moving from j to i will be $\sigma_{ij}S_j$ where σ_{ij} are mobility coefficients. In this way, the incoming rate of susceptibles in region i is given by $\sum_{j=1, j \neq i}^n \sigma_{ij}S_j$. Correspondingly, the outgoing rate of susceptible from the region i is $(-\sum_{j=1, j \neq i}^n \sigma_{ji})S_i$. We can introduce the S-mobility matrix

$$M^S \equiv [m_{i,j}^S] = \begin{cases} \sigma_{ij} & \text{for } i \neq j \\ -\sum_{k=1, k \neq i}^n \sigma_{ki} & \text{for } i = j, \end{cases} \quad (3.1)$$

where the off-diagonal entries $m_{i,j}^S, i \neq j, i, j \in \{1, \dots, n\}$ are the incoming rates of susceptible from each region j to each region i , and the diagonal entries $m_{i,i}^S, i \in \{1, \dots, n\}$ are the outgoing rates of susceptible from each region i to all other regions as discussed above.

Similarly, we introduce an I-mobility and an R-mobility matrices M^I and M^R , and finally the block-diagonal $(n_c \times n)^2$ matrix $M = \text{diag}(M^S, M^I, M^R)$, where n_c is the number of compartments per region (3 in our case). This matrix is promptly reordered consistently with the vector \mathbf{A} , oriented by geographical areas as opposed to a compartment-wise ordering. We keep calling M the rearranged matrix. In this way, the Domain, i.e. the set of GU, denoted by the vector $\mathbf{C} \equiv [\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_n]$ will obey the ordinary differential system

$$d_t \mathbf{C} = \mathbf{D}(\mathbf{C}) + M \mathbf{C} \quad (3.2)$$

where $\mathbf{D}(\mathbf{C}) \equiv [\mathbf{F}_1(\mathbf{A}_1), \dots, \mathbf{F}_n(\mathbf{A}_n)]^T$. Should we admit that a population can change compartment while moving across regions, the matrix M would have a full pattern (not block-diagonal).

3.2 The numerical solver

As analytical solutions for this problem are not available, we opted for a numerical solver. In particular, we choose an explicit solver like Heun’s method [20] to avoid the solution of non-linear equations. This method solves the system with $\mathcal{O}(\Delta t^2)$ accuracy. Specifically, let us denote by Δt a user-selected time-step, and by t_k the collocation instants when we compute the solution. The Heun’s scheme then reads: given the solution \mathbf{C}^k at time t_k , compute \mathbf{C}^{k+1} as:

$$\begin{aligned} \text{Predictor : } \mathbf{C}^* &= \mathbf{C}^k + \Delta t (\mathbf{D}(\mathbf{C}^k) + M\mathbf{C}^k) \\ \text{Corrector : } \mathbf{C}^{k+1} &= \mathbf{C}^k + \frac{\Delta t}{2} (\mathbf{D}(\mathbf{C}^k) + \mathbf{D}(\mathbf{C}^*) + M(\mathbf{C}^k + \mathbf{C}^*)). \end{aligned} \tag{3.3}$$

At time t_0 the initial condition \mathbf{C}^0 is provided. The scheme is only conditionally stable: we empirically select our time-step to ensure numerical stability.

3.3 Design of the objected-oriented solver

One of the design specifics of our solver is to be flexible to both *horizontal* and *vertical* extensions. With this, we mean that we want a solver easily scalable to any given number of GU, so to manage the 20 regions in Italy (one of the countries most affected by the virus [7]), the 50 states of the United States or even the counties [3] in one or more of the US (horizontal expansion). Meanwhile, we want to manage possible refinements of the models for the single area; while the SIR model is a good starting point, more compartments may be needed for a better representation of the outbreak (e.g. SEIR models); in an even more sophisticated modeling, we can consider Partial Differential Equations models, where the local mobility is assimilated to a Brownian motion [22] (*vertical flexibility*). For this reason, *Object-Oriented* programming provides the right framework to our solver, particularly for the features of encapsulation and inheritance.

3.3.1 The local and global models

The GU class (local model) The core class of our solver is the GU Class implementing the model for the single area represented by the generic equation

$$d_t \mathbf{A} = \mathbf{F}(\mathbf{A}) \quad (3.4)$$

with the corresponding initial conditions. At the basic level, the class implements the SIR model in (2.1). As members of the class we include:

- the vector \mathbf{A} (by default given by S, I, R , with the number of compartments $n_c (= 3)$);
- the vector \mathbf{F} with the associated parameters β, μ, γ, ρ , possibly function of time;
- the initial conditions.

More sophisticated models are inherited by this class. For instance, to implement a PDE-based model like [22] we can add as members the mesh of the region of interest. Similarly, to include the structure of age, we can extend the number of compartments and parameters to describe Susceptible etc. at the different ages, by exploiting the inheritance from the basic class.

The domain class (global model) To cover the entire domain of interest, we introduce a class with the following members:

- an array of GU;
- the mobility matrix M ;
- a solver for the problem (3.2) with the initial conditions of the GU, for instance (3.3).

The solver was implemented both in Python and Matlab and it is available at the GitHub repository github.com/HowardYutingHou/CovidProject.

3.4 Benchmarks and results

We present a number of tests we run to check the correctness of our solver. Specifically, we test the accuracy of our simulations against analytical results, when available. Then, we use our two different and independent implementations to cross-validate the results on a fictitious benchmark consisting of seven regions (having the nickname of seven states of the U.S). Successively, we test the endemic vs. the epidemic models, to stress the importance of the parameter μ in the final results. In the second part of this Section, we illustrate the results we obtain when we want to mimic the presence of lockdown policies and a vaccination campaign.

3.4.1 Consistency test with theorems for one-region models

When cross-regional interactions are disregarded, each local dynamics is represented locally by the classical epidemic or endemic SIR models. In this case, analytical results are available as mentioned in Section 2.2. Using the numerical method described in Section 3.2, we verified Theorems 1 and 2 with various parameter values in an unnormalized population, i.e. N is the total population and does not equal to 1.

For example, in the case of an endemic model for the GU “Alabama” with $N = 78,000$ we assign the parameter values $\beta = 0.3$, $\gamma = 0.1$, and $\mu = 0.01$, resulting in $\vartheta = \beta/(\gamma + \mu) = 0.3/(0.1 + 0.01) > 1$. By Theorem 2, all solutions of the system should approach the endemic equilibrium given by $S_e = N/\vartheta \approx 2.882 \times 10^4$ and $I_e = \mu N(\vartheta - 1)/\beta \approx 4.525 \times 10^3$. The numerical solver successfully simulates this asymptotic behavior as shown in Figure 3.1.

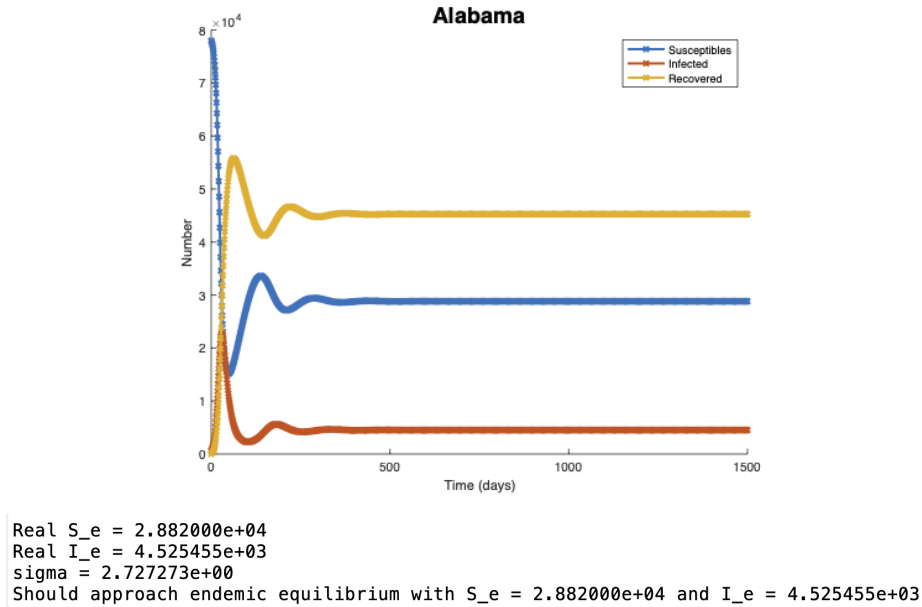


Figure 3.1: top: simulation results of the endemic SIR model for the GU “Alabama”; bottom: MATLAB output of comparison between simulated and real asymptotic values.

We run similar tests on both epidemic and endemic models, and the consistency of our simulations with these results was carefully assessed and verified. Thus, this verification with existing analytical results gives confidence to the numerical solver.

3.4.2 Cross-validation

Another test was performed by comparing the results of the two different codes developed with two different languages (MATLAB and Python). The Python version of the solver is constructed by my teammates in [10]. To this aim, we introduce a fictitious data-set, with seven GUs. For the sake of notation, we give these units the name of seven States of the U.S. The parameters and the mobility matrices are, however, selected without a specific tuning of the data, so the names are to be intended as nicknames, with no particular reference to the real states. Fine tuning of the parameters is, in fact, an important follow-up of the present work. For the present proof-of-concept, we calibrate the results arbitrarily, based on the available

literature. Figure 3.2 shows the computational solutions of the GUs “Georgia” and “Alabama” with the endemic model from both versions of the solver.

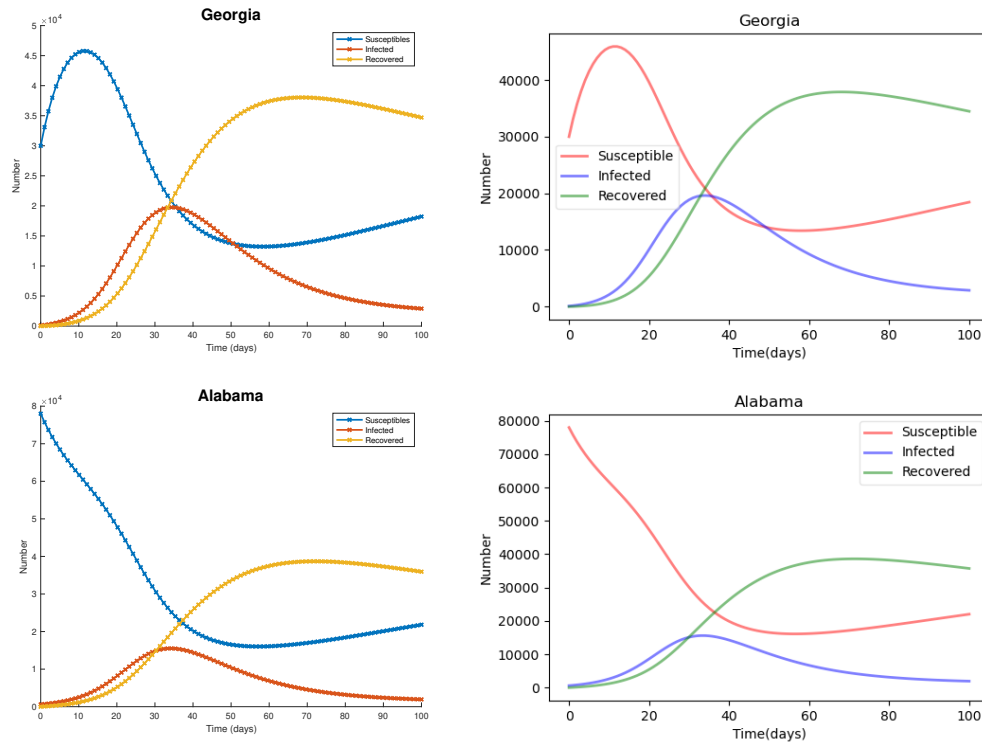


Figure 3.2: left: Results for our multi-regional MATLAB solver in “Georgia” and “Alabama”; right: results for the same case with the Python solver.

By visual and quantitative inspections, the results do coincide. The two solvers were developed independently with just the same code-specifications. The comparison provides therefore an excellent cross-validation of our results, giving confidence in the correctness of the solvers.

3.4.3 Epidemic vs. endemic

Figure 3.3 shows the comparison of the multi-regional solver in “Georgia”, when we drop the local reproduction rate ($\mu = 0$), i.e. the “epidemic model”, on the left, and when we include it (“endemic model”) on the right. The increment of S in the first 10 days is due to the mobility (particularly, in this case, from “California” and “Texas”,

with a mobility rate of 0.01 and 0.02 respectively). The comparison between the two plots highlights the importance of the coefficient μ on the long time range, as the endemic model exhibits an additional increment of S , absent in the epidemic one.

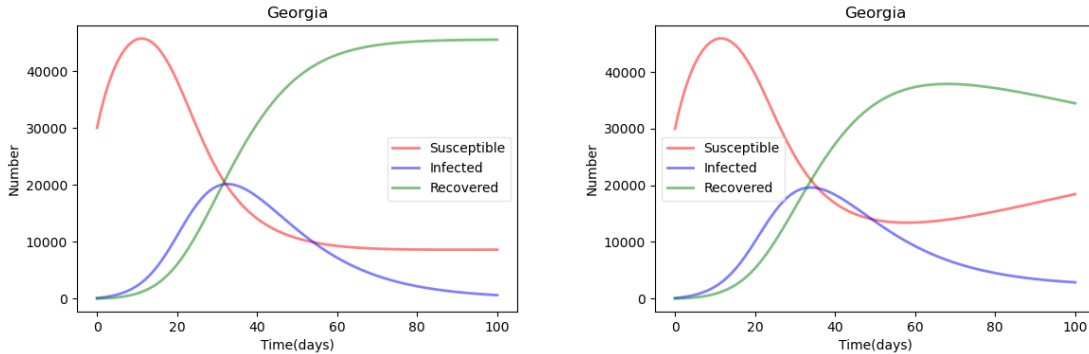


Figure 3.3: Results in “Georgia”: left: epidemic model ($\mu = 0$); right: endemic model.

3.4.4 Effects of lockdown

We simulate the effects of lockdown restrictions by reducing the infection rate β by a factor of 10 as a result of restrictive measures like mask-enforcement, curfew, etc.

Local lockdown We consider first the effects of lockdown in an isolated region or GU. In Fig. 3.4 we show the results in “Georgia” with no mobility: on the left the no-lockdown results for 200 days, on the right the results with a lockdown enforced from day 20 to 60. The impact of the lockdown is evident, particularly on the I population. At the end of the restrictive measures, there is a partial comeback, however the lockdown significantly helps containing the number of infected (from 10,000 to 5,000). At the end of the lockdown there is a spike of infected, as the number of susceptible is considerably higher than in the no-lockdown case, yet the peak of I is halved.




Figure 3.4: Local evolution in “Georgia” without lockdown (left) and with lockdown (right).

Multi-regional lockdown We consider now the multi-regional case: during lockdown, β is reduced by a factor of 10 and the borders are closed, i.e. the mobility coefficients are set to 0. In particular, we assume that in the 7-state case, only Georgia enforces a lockdown (Fig. 3.5). As expected, with the restrictions on the mobility, the results are similar to the isolated case, and the lockdown reduces the number of infected to about 10,000 vs. the 20,000 of the no-lockdown case.

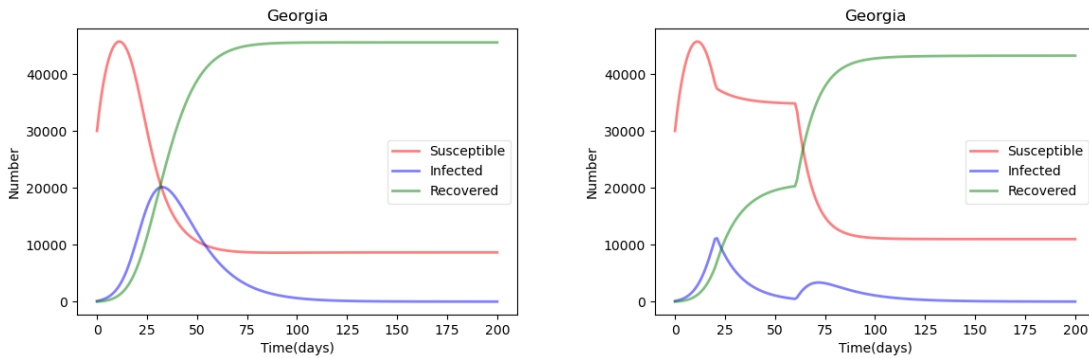


Figure 3.5: left: without lockdown; right: with lockdown.

We also plot in Figure 3.6 Florida and New York from day 0 to day 100 to show how a lockdown in Georgia can slightly affect other regions. After day 60, Florida immediately experiences an increase in S because once the lockdown is dismissed, 0.03 of Georgia’s susceptible (incremented by the lockdown) move into Florida (only 0.02

of Florida's S is assumed to move into Georgia). As for New York, there is a slower increase in R after day 60. As 0.01 of Georgia's recovered population is moving into New York, and the recovered population of Georgia is lower than without lockdown, while 0.015 of New York recovered is moving into Georgia. In conclusion, if the lockdown is enforced only on one state, the effects for the other states appear to be minor.

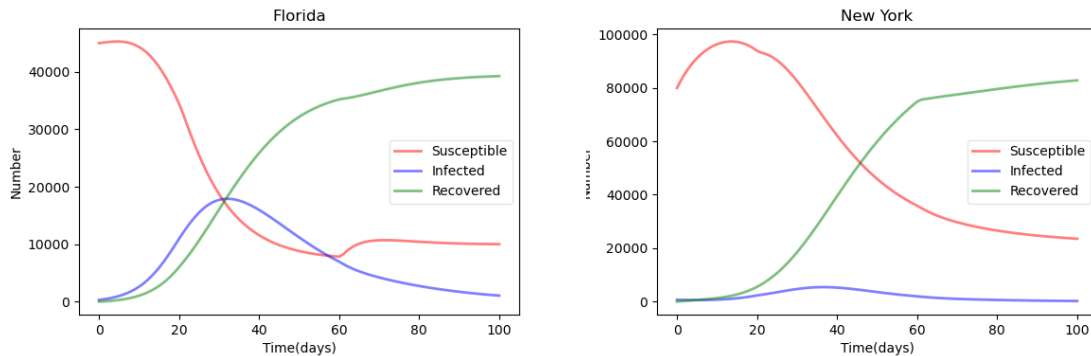


Figure 3.6: Multi-regional case with a lockdown in Georgia: left: Florida; right: New York.

Effects of vaccination

To test the effects of vaccination (in absence of lockdown), we introduce the vaccine at day 80, where the vaccination rate is modeled by (2.2). We refer again to “Georgia.” We observe that at day 80, the development pattern of the recovered is interrupted and R increases at a much higher rate. At the same time, the population of the susceptible drop significantly, decreasing faster than before.

Notice that the total number of recovered reaches a greater value than that without vaccination. Also, the susceptible population decreases to 0 at the end, which does not happen where vaccination is not implemented, reaching herd immunity. Therefore, our model suggests that the introduction of vaccination can end the outbreak. We compare the effects of a temporary measure like the lockdown and the vaccination in Fig. 3.7.

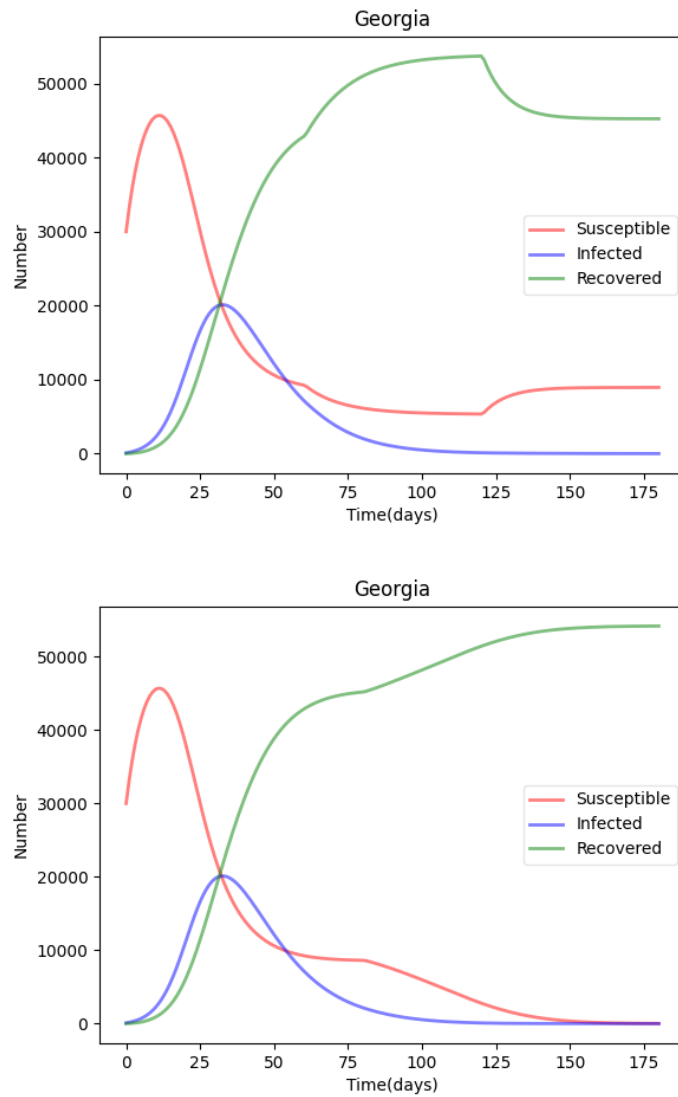


Figure 3.7: Measurements for fighting the pandemic in a multi-regional solver, with focus on “Georgia”: left, lockdown; right, vaccination.

Chapter 4

Sensitivity Analysis for Mono-region Models

COVID-19 modeling is subject to several uncertainties, including unknown disease parameters and limited data availability. Sensitivity analysis is a valuable tool that can help researchers assess the impact of these uncertainties on COVID-19 model predictions. By systematically varying the disease parameters within plausible ranges, sensitivity analysis can provide insights into which parameters have the most significant impact on model outcomes and which are less critical. This information can guide researchers and policymakers in prioritizing data collection efforts and developing effective interventions to control the spread of COVID-19. In this chapter, we perform first-order local sensitivity analysis on the endemic model without consideration for vaccination, i.e. $\rho = 0$.

4.1 First-order local sensitivity analysis on the endemic model

To perform first-order local sensitivity analysis the SIR endemic model, we investigate the effect of perturbations in each disease parameter, i.e., β (the number of average adequate contacts per person per unit of time), γ (the recovery rate) and μ (the reproduction/death rate), on each S , I , R compartment. More specifically, we perturb β , γ and μ by δ_β , δ_γ and δ_μ respectively; then the corresponding perturbed system dynamics are derived as

$$\begin{cases} d_t \hat{S}_\beta = \mu N - (\beta + \delta_\beta) \hat{I}_\beta \hat{S}_\beta / N - \mu \hat{S}_\beta \\ d_t \hat{I}_\beta = (\beta + \delta_\beta) \hat{I}_\beta \hat{S}_\beta / N - \gamma \hat{I}_\beta - \mu \hat{I}_\beta \\ d_t \hat{R}_\beta = \gamma \hat{I}_\beta - \mu \hat{R}_\beta \end{cases} \quad (4.1)$$

$$\begin{cases} d_t \hat{S}_\gamma = \mu N - \beta \hat{I}_\gamma \hat{S}_\gamma / N - \mu \hat{S}_\gamma \\ d_t \hat{I}_\gamma = \beta \hat{I}_\gamma \hat{S}_\gamma / N - (\gamma + \delta_\gamma) \hat{I}_\gamma - \mu \hat{I}_\gamma \\ d_t \hat{R}_\gamma = (\gamma + \delta_\gamma) \hat{I}_\gamma - \mu \hat{R}_\gamma \end{cases} \quad (4.2)$$

$$\begin{cases} d_t \hat{S}_\mu = (\mu + \delta_\mu) N - \beta \hat{I}_\mu \hat{S}_\mu / N - (\mu + \delta_\mu) \hat{S}_\mu \\ d_t \hat{I}_\mu = \beta \hat{I}_\mu \hat{S}_\mu / N - \gamma \hat{I}_\mu - (\mu + \delta_\mu) \hat{I}_\mu \\ d_t \hat{R}_\mu = \gamma \hat{I}_\mu - (\mu + \delta_\mu) \hat{R}_\mu \end{cases} \quad (4.3)$$

where the hat indicates the compartmental population after perturbation of the parameter in the subscript; for example, \hat{S}_β is the perturbed susceptible population after perturbing β by δ_β .

To obtain $d_t(S - \hat{S})$, $d_t(I - \hat{I})$, $d_t(\hat{S}_\mu - \hat{R})$, we subtract systems (4.1), (4.2), and (4.3) from the original system (2.1). Since we are not considering the effect of vaccination, we set ρ , the vaccination rate in the original system, to zero. We also normalize the total population to $N = 1$. For concision, we will exemplify below the

computation of $d_t(S - \hat{S}_\beta)$ and omit those of the other compartments. We have

$$\begin{aligned}
d_t(\hat{S}_\beta - S) &= d_t\hat{S}_\beta - d_tS \\
&= (\mu N - (\beta + \delta_\beta)\hat{I}_\beta\hat{S}_\beta/N - \mu\hat{S}_\beta) - (\mu N - \beta IS/N - \mu S) \\
&= -\mu(\hat{S}_\beta - S) - \beta(\hat{S}_\beta - S)(\hat{I}_\beta - I) - \beta S(\hat{I}_\beta - I) - \beta(\hat{S}_\beta - S)I - \delta_\beta\hat{S}_\beta\hat{I}_\beta
\end{aligned}$$

We then compute

$$d_t \frac{\hat{S}_\beta - S}{\delta_\beta} = -\mu \frac{(\hat{S}_\beta - S)}{\delta_\beta} - \beta \frac{(\hat{S}_\beta - S)(\hat{I}_\beta - I)}{\delta_\beta} - \beta S \frac{(\hat{I}_\beta - I)}{\delta_\beta} - \beta \frac{(\hat{S}_\beta - S)}{\delta_\beta} I - \hat{S}_\beta \hat{I}_\beta \quad (4.4)$$

Since we assume the response of the susceptible and infected populations from a small perturbation of β is small, we assume $\lim_{\delta_\beta \rightarrow 0} (\hat{S}_\beta - S)(\hat{I}_\beta - I) = 0$. With this assumption, we set the second term of equation (4.4) to be zero. Also, note that $\lim_{\delta_\beta \rightarrow 0} \hat{S}_\beta = S$, $\lim_{\delta_\beta \rightarrow 0} \hat{I}_\beta = I$, etc. In this way, we arrive at the relation

$$\begin{aligned}
d_t \sigma_S^\beta &= d_t \lim_{\delta_\beta \rightarrow 0} \frac{\hat{S}_\beta - S}{\delta_\beta} = -\mu \frac{(\hat{S}_\beta - S)}{\delta_\beta} - \beta S \frac{(\hat{I}_\beta - I)}{\delta_\beta} - \beta \frac{(\hat{S}_\beta - S)}{\delta_\beta} I - SI \\
&= -\mu \sigma_S^\beta - \beta S \sigma_I^\beta - \beta \sigma_S^\beta I - SI,
\end{aligned}$$

where $\sigma_S^\beta \equiv \lim_{\delta_\beta \rightarrow 0} (\hat{S}_\beta - S)/\delta_\beta$ and $\sigma_I^\beta \equiv \lim_{\delta_\beta \rightarrow 0} (\hat{I}_\beta - I)/\delta_\beta$ are defined respectively as the sensitivities of the S and I compartments on β .

Using similar computations, we obtain sensitivities of every compartment on every parameter as follows:

$$\left\{ \begin{array}{l} d_t \sigma_S^\beta = -\mu \sigma_S^\beta - \beta S \sigma_I^\beta - \beta \sigma_S^\beta I - SI \\ d_t \sigma_I^\beta = \beta S \sigma_I^\beta + \beta I \sigma_S^\beta + SI - \gamma \sigma_I^\beta - \mu \sigma_I^\beta \\ d_t \sigma_R^\beta = \gamma \sigma_I^\beta - \mu \sigma_R^\beta \end{array} \right. \quad (4.5)$$

$$\begin{cases} d_t \sigma_S^\gamma = -\mu \sigma_S^\gamma - \beta S \sigma_I^\gamma - \beta I \sigma_S^\gamma \\ d_t \sigma_I^\gamma = \beta S \sigma_I^\gamma + \beta I \sigma_S^\gamma - I - \gamma \sigma_I^\gamma - \mu \sigma_I^\gamma \\ d_t \sigma_R^\gamma = \gamma \sigma_I^\gamma - \mu \sigma_R^\gamma + I \end{cases} \quad (4.6)$$

$$\begin{cases} d_t \sigma_S^\mu = 1 - S - \mu \sigma_S^\mu - \beta S \sigma_I^\mu - \beta I \sigma_S^\mu \\ d_t \sigma_I^\mu = \beta S \sigma_I^\mu + \beta I \sigma_S^\mu - \gamma \sigma_I^\mu - \mu \sigma_I^\mu - I \\ d_t \sigma_R^\mu = \gamma \sigma_I^\mu - \mu \sigma_R^\mu - R \end{cases} \quad (4.7)$$

With these systems of ODEs, we can solve for these sensitivities through numerical methods. We use the `ode45()` method from MATLAB and obtain the following results shown in Figure 4.1.

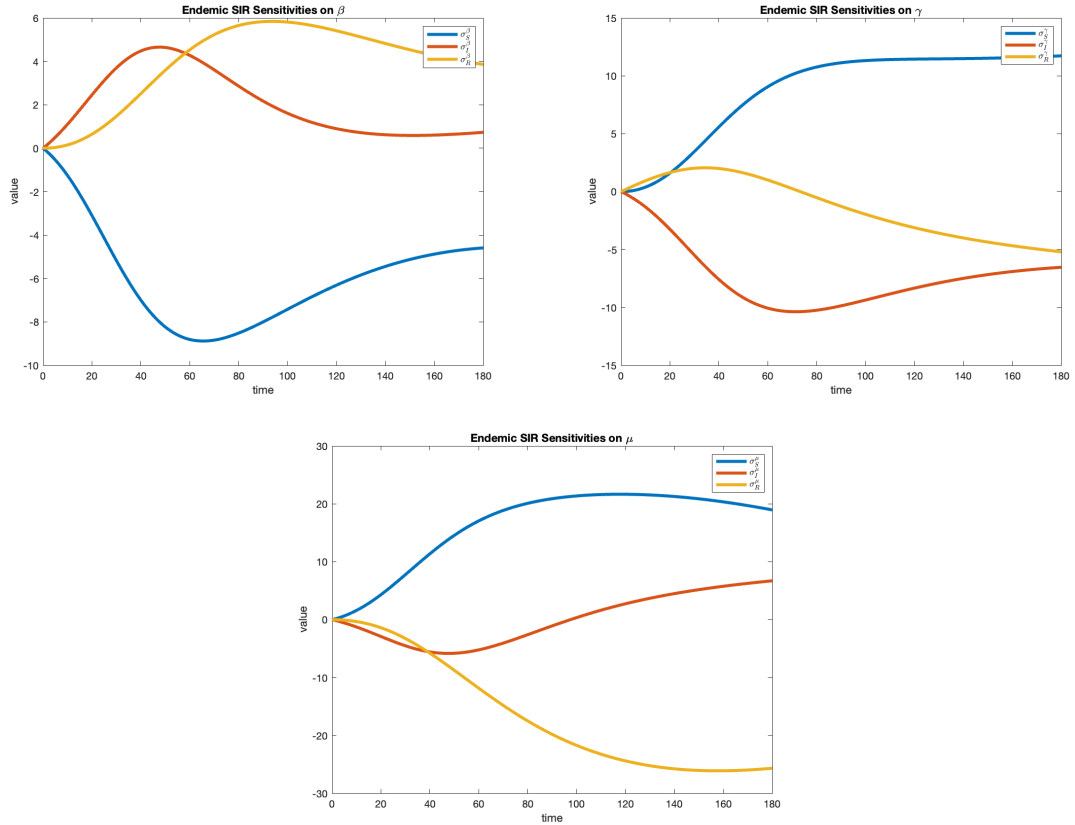


Figure 4.1: Sensitivities of the S , I , R compartments on: top left: β , the transmission rate; top right: γ , the recovery rate; bottom: μ , the reproduction/death rate.

We observe that the transmission rate β is negatively associated with the suscepti-

ble population and positively associated with the infected and recovered populations. Regulation measures such as lockdown and mask mandates that aim to reduce β can effectively reduce the infected population. From the sensitivity on γ , the recovery rate, we observe the recovered population is positively associated with γ until around day 70; then, the relationship becomes negative. We reason that an increased recovery rate induces reduction in infected population, which ultimately leads to a reduction of the recovered population over time. From sensitivity on μ , we observe that the reproduction/death rate, μ , has the highest magnitude of sensitivities. More interestingly, the sensitivity of the infected population on μ is negative before day 90 but positive afterwards. This result indicates that if the policymakers are able to control μ in this short period of time, increasing μ in the first 90 days reduces the infected population.

Chapter 5

Parameter Estimation and Prediction with Physics-Informed Neural Networks

During the COVID-19 pandemic, governments around the world collected a vast amount of data to be analyzed for disease control. In this chapter, we utilize the versatility of neural networks to estimate the disease parameters and make predictions of the dynamics. Neural networks have shown great promise and remarkable accuracy on image processing and other classification problems [14][15]. However, standard neural networks produce unsatisfactory results on domains beyond the training data; thus, we take advantage of physics-informed neural networks (PINN) [13], which enforce physical constraints to eliminate unreasonable results. In Section 5.1, we provide a brief introduction to PINN. In Section 5.2, we discuss results from using PINNs to estimate disease parameters and make predictions, from noisy or limited data, for the endemic mono-regional SIR model.

5.1 Introduction to physics-informed neural networks (PINN)

Neural networks consist of multiple layers of neurons that are connected to each other in a systematic way, where outputs from neurons of a given layer become inputs for the next layer. Figure 5.1 displays an example of a feedforward neural network [1].




Figure 5.1: An example of a feedforward neural network. The neurons are partitioned into four layers. The first layer is referred to as the input layer and the last layer is referred to as the output layer. Each layer between these two layers is a hidden layer.

Mathematically, each data point in the dataset is trained by a neural network, $\hat{f}_{\mathbf{W}, \mathbf{b}}$, defined as

$$\begin{aligned} \mathbf{a}_0 &= \mathbf{x}, \\ \mathbf{a}_i &= \sigma(\mathbf{W}_i \mathbf{a}_{i-1} + \mathbf{b}_i), \quad \forall i \in \{1, 2, \dots, L\}, \\ \hat{f}_{\mathbf{W}, \mathbf{b}}(\mathbf{x}) &= \mathbf{a}_L, \end{aligned}$$

where $\mathbf{x} \in \mathbb{R}^{M_0}$ denotes a vector of M_0 input variables. The weight matrix connecting the $(i-1)^{\text{th}}$ layer to the i^{th} layer is denoted by $\mathbf{W}_i \in \mathbb{R}^{M_i \times M_{i-1}}$, where M_i denotes the number of neurons in the i^{th} layer of the network with $i \in \{1, 2, \dots, L\}$, and the bias, the other set of learning parameters that are added to the weighted sums of inputs to each neuron in the network, is denoted by $\mathbf{b}_i \in \mathbb{R}^{M_i}$. The nonlinear differentiable

activation function, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, is applied element-wise to the weighted input vector $(\mathbf{W}_i \mathbf{a}_{i-1} + \mathbf{b}_i)$. The output of the i^{th} layer is denoted \mathbf{a}_i and is passed to the following layer as input. The final layer's output, \mathbf{a}_L , is the output of the neural network.

To train a standard neural network, the weights and biases are adjusted to fit the training data well. This is performed by minimizing the loss function \mathcal{L} . In practice, the full dataset is divided into batches, and the losses from all batches are summed to compute the total loss. A typical example of a loss function is the mean squared error. With $\hat{f}_{\mathbf{w},\mathbf{b}}(\mathbf{x})$ denoting the neural network output from a batch of S input data points and $f(\mathbf{x})$ the actual values, the loss function is defined as

$$\mathcal{L} = \frac{1}{S} \left\| \hat{f}_{\mathbf{w},\mathbf{b}}(\mathbf{x}) - f(\mathbf{x}) \right\|_2^2, \quad (5.1)$$

where $\| \cdot \|_2$ denotes the L_2 norm. Then, expressing the loss function of the i th batch as \mathcal{L}_i , we find the total loss function as

$$\mathcal{L}_{\text{tot}} = \sum_{i=1}^{N_b} \mathcal{L}_i,$$

where N_b is the total number of batches.

The goal of neural networks is to minimize the loss function with respect to the weights and biases so that \hat{f} approximates f . In practice, the loss is computed over each batch by using an algorithm called stochastic gradient descent with backpropagation [16]. This algorithm computes the gradient of the loss function over the batch, and moves in the direction of steepest descent to minimize the loss function. As different activation functions render different derivative behaviors, the choice of activation function plays an important role in the model's efficiency. Commonly used activation functions include Sigmoid, tanh, rectified linear unit (ReLU), and swish (see Table 5.1 and Figure 5.2).

	Mathematical formulation
Sigmoid	$\sigma(x) = \frac{1}{1 + e^{-x}}$
tanh	$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$
ReLU	$R(x) = \max\{0, x\}$
swish	$S(x) = x \cdot \sigma(\beta x), \beta \in \mathbb{R}$

Table 5.1: Mathematical formulations of common activation functions.

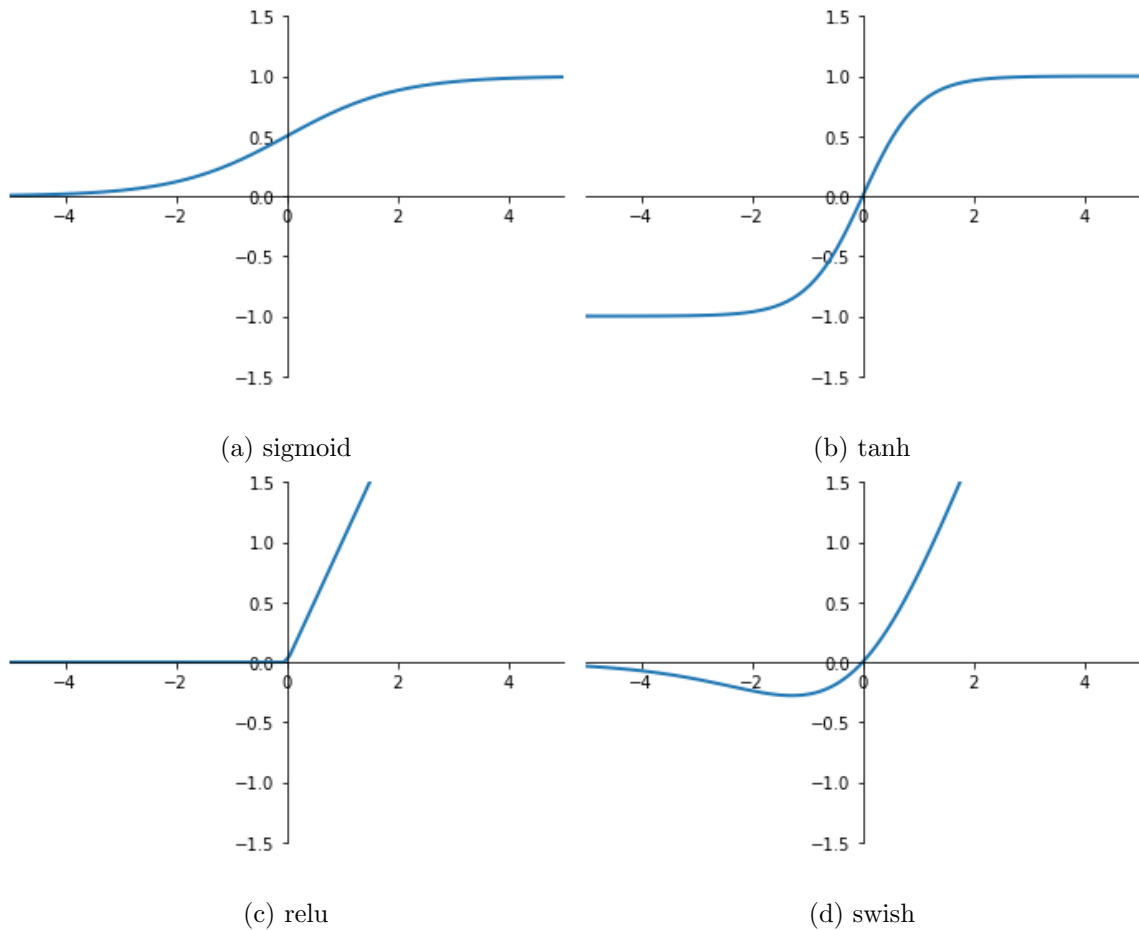


Figure 5.2: Plots of commonly used activation functions.

By the Universal Approximation Theorems [19], neural networks are exceptional at approximating any arbitrary function. However, neural networks can easily fall susceptible to overfitting the noise in the training data while failing to approximate

any meaningful trend. This drawback reduces the ability of neural networks to extrapolate beyond the training data domain [12]. To minimize this error, one can train directly on data in the extrapolation region, but generating additional data is usually inefficient.

The extrapolation limitation of standard neural networks can be overcome with physics-informed neural networks (PINN). By incorporating constraints based on physical, empirical, or mathematical consistency, a standard neural network is endowed with physical intuition. Consequently, PINNs yield more interpretable results that are accurate with noisy or a lack of data [13]. The incorporation of physical constraints is accomplished through augmenting the loss function with a regularizer. Ideally, regularization increases generalization accuracy without greatly decreasing accuracy over the training data and decreases variance of the model predictions, $\mathbb{V}(\hat{f}(\mathbf{x}))$ without greatly increasing the bias of the model output, $(\mathbb{E}[\hat{f}(\mathbf{x})] - f(\mathbf{x}))^2$ [9].

In this project, we simply use the ODE residuals of the SIR model as the regularizer. For example, the regularizers attributed to each of the compartment in the SIR epidemic model are

$$\begin{aligned}\mathcal{N}_S[\hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}), \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t})] &= d_t \hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) + \beta \hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) - \mu + \mu \hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) \\ \mathcal{N}_I[\hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}), \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t})] &= d_t \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) - \beta \hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) + (\mu + \gamma) \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) \\ \mathcal{N}_R[\hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}), \hat{R}_{\mathbf{W},\mathbf{b}}(\mathbf{t})] &= d_t \hat{R}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) - \gamma \hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t}) + \mu \hat{R}_{\mathbf{W},\mathbf{b}}(\mathbf{t}),\end{aligned}$$

where the \mathcal{N}_S , \mathcal{N}_I , and \mathcal{N}_R represent the differential operator for the S , I , and R compartments respectively, and $\hat{S}_{\mathbf{W},\mathbf{b}}(\mathbf{t})$, $\hat{I}_{\mathbf{W},\mathbf{b}}(\mathbf{t})$, and $\hat{R}_{\mathbf{W},\mathbf{b}}(\mathbf{t})$ are the predicted compartmental populations with parameters \mathbf{W} and \mathbf{b} . In practice, the first-order derivatives of the model predictions are approximated with automatic differentiation.

Taking the susceptible compartment for example, this augmentation leads to a

new loss function:

$$\mathcal{L}_S = \frac{1}{S} \left\| \hat{S}_{\mathbf{w},b}(\mathbf{t}) - S(\mathbf{t}) \right\|_2^2 + \left| \mathcal{N}_S[\hat{S}_{\mathbf{w},b}(\mathbf{t}), \hat{I}_{\mathbf{w},b}(\mathbf{t})] \right|, \quad (5.2)$$

Trained with augmented loss functions for each compartment, PINNs are able to not only fit the training data well but also to comply with the underlying physical principles, rendering more confidence in the predictions.

5.2 Using PINN for parameter estimation and prediction of the SIR endemic model

One of the most prominent challenges of COVID modeling is parameter estimation. Accurate parameter estimation facilitates effective policy making and pandemic control. However, in most cases, these parameters are difficult to obtain analytically and thus must be approximated with data-driven methods. In this Section, we simulate data of the SIR system using predetermined parameters; then, we use the `deepxde` package developed by Lu et al. [16] to solve this inverse problem for the endemic SIR system as shown in equation (2.1), with $\rho = 0$. We construct a fully connected deep neural network with 1 input layer of 1 neuron for the time input, 3 hidden layers of 30 neurons each, and 1 layer of 3 neurons for the S , I , R outputs. This construction of the neural network is relatively small, as the size of the dataset is small, and the dynamics is not too complex. The network is trained for 30,000 iterations with a learning rate of 0.001. Since the solution for each compartment strictly lies in the interval $[0, 1]$, an activation function with a smaller range usually facilitates fast convergence. Therefore, we choose the tanh function as the activation function.

First, we experiment with the SIR endemic model with $\rho = 0$, $N = 1$, $I_0 = 0.1$, $\beta = 0.08$, $\gamma = 0.03$, and $\mu = 0.01$. The estimated parameters are β , γ , and μ , which

are initialized with values 0.01, 0.01, and 0 respectively. The training data comprises of 237 data points over the span of 180 days and is unperturbed in this case.

The neural network successfully recovers the parameters as $\hat{\beta} = 7.98 \times 10^{-2}$, $\hat{\gamma} = 3.00 \times 10^{-2}$, $\hat{\mu} = 9.97 \times 10^{-3}$. Using these approximations, the model accurately replicates the SIR dynamics with a test loss (using L2 relative error metric) of 1.05×10^{-6} :

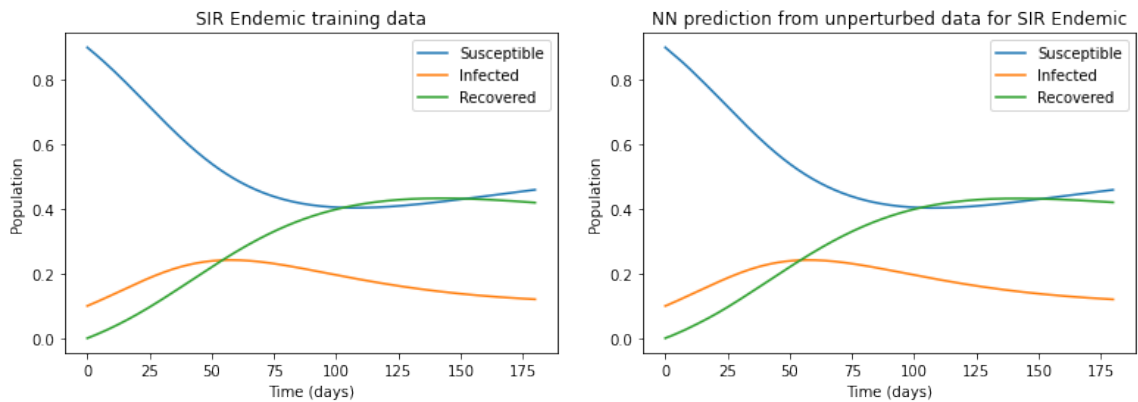


Figure 5.3: left: training data for SIR endemic model; right: PINN prediction for SIR endemic model.

Similar accuracy is found in the prediction of the SIR epidemic model and is omitted for concision.

5.2.1 Using PINN for noisy data

However, during this pandemic, the inadequate reporting systems in most countries have resulted in noisy data. Therefore, we use simulated noisy data to examine the model performance under this limitation. We compare the model accuracy under 5%, 10%, and 20% Gaussian noise as demonstrated in Table 5.2 and Figure 5.4. We observe that with 5% noise, the PINN demonstrates success in parameter estimation, with highly accurate predictions that closely align with true values. As anticipated, the introduction of higher levels of noise at 10% and 20% result in a degradation of performance with regard to both parameter estimation and prediction, yet the PINN

still captures the general dynamics of system with respect to peak of infection and asymptotic behaviors. However, intriguingly, the parameter estimation outcomes for the 20% noise scenario are actually superior to those obtained with 10% noise, but the predictions for the 10% noise scenario are closer to truth, as reflected by the test losses. These findings highlight the complexity of the interplay between noise levels and the performance of the PINN, and underscore the importance of careful consideration of these factors in optimizing the accuracy of the model.

	$\hat{\beta}$	$\hat{\gamma}$	$\hat{\mu}$	Test loss
True value	8×10^{-2}	3×10^{-2}	1×10^{-2}	
5% noise	8.07×10^{-2}	3.01×10^{-2}	1.02×10^{-2}	1.094×10^{-2}
10% noise	7.47×10^{-2}	2.84×10^{-2}	9.29×10^{-3}	2.536×10^{-2}
20% noise	8.47×10^{-2}	3.08×10^{-2}	1.04×10^{-2}	6.289×10^{-2}

Table 5.2: PINN parameter estimations and test losses from noisy data.

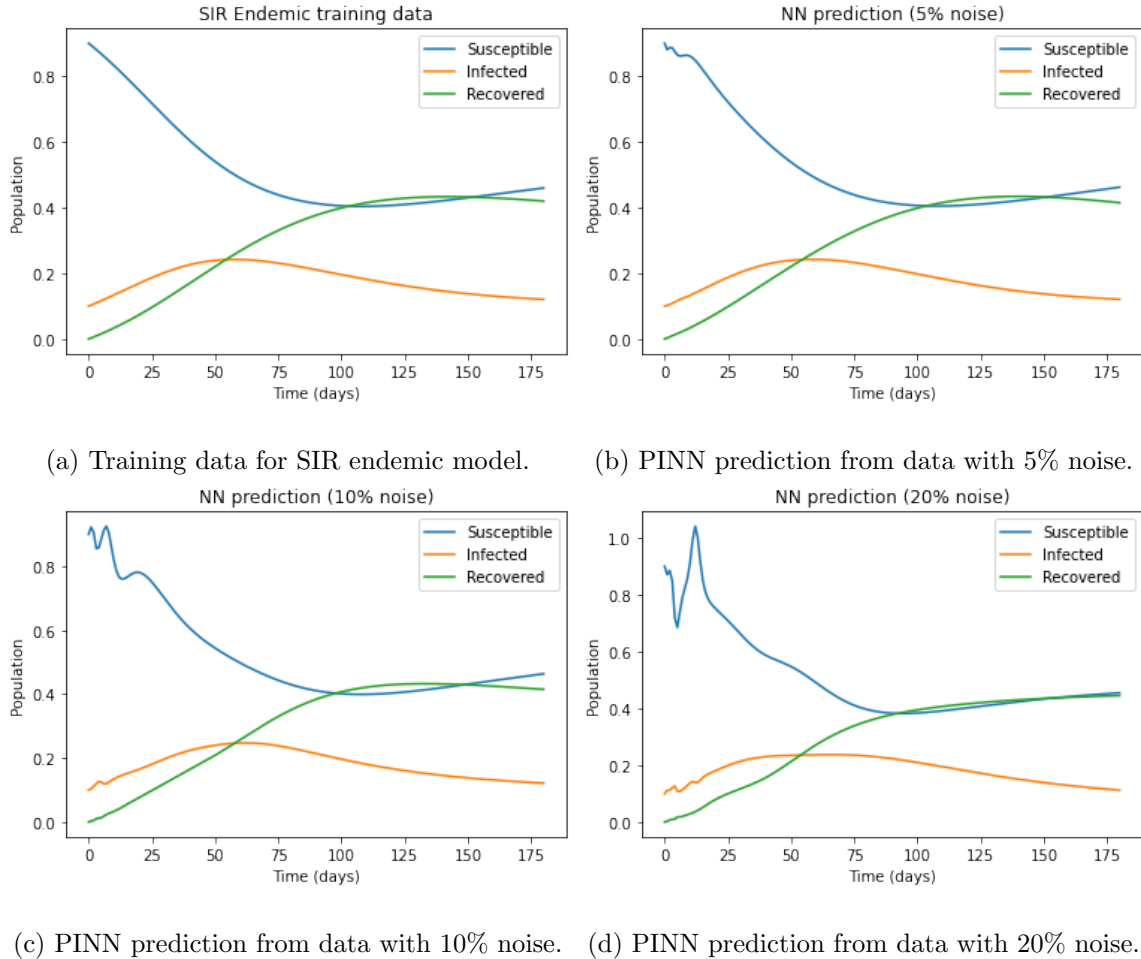


Figure 5.4: PINN predictions from noisy data.

5.2.2 Using PINN for limited data

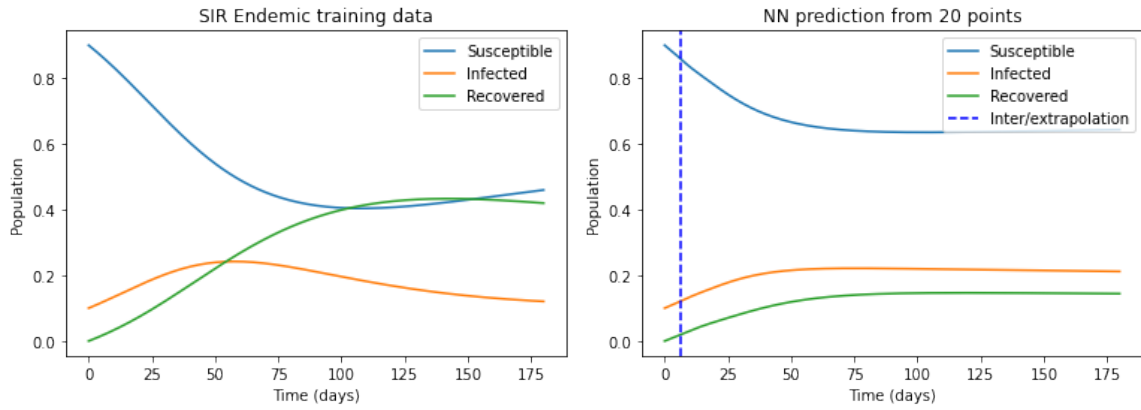
Infectious disease modeling usually takes place while the pandemic is in progress, as policymakers are interested the future dynamics of the pandemic given limited amount of initial data. Therefore, PINN's ability to extrapolate beyond training domain is especially beneficial in this case. We again consider the SIR endemic model for demonstration. By limiting the training data to a certain point in time, we simulate the real situation during a pandemic where we must estimate parameters and predict the future development of the disease with limited data. Table 5.3 and

Figure 5.5 show the SIR parameter estimations, test losses, and predictions by the PINN using the first 20, 50, 100 (out of 237 total) training data points.

As expected, augmenting the volume of training data significantly improves the predictive ability of the PINN, as the test loss consistently decreases as the volume of data increases. However, similarly to findings in Section 5.2, the increase of training data from 50 to 100 points does not correspond with superior parameter estimation as expected. The total training loss are evaluated for each of the three cases examined: 3.74×10^{-8} for the 20-point scenario, 4.37×10^{-8} for the 50-point scenario, and 1.38×10^{-7} for the 100-point scenario. We observe that the training losses are broadly consistent across all scenarios. This result suggests that the PINN may have converged upon alternative sets of parameters that are similarly capable of satisfying the underlying ODE system, despite being distinct from the true values.

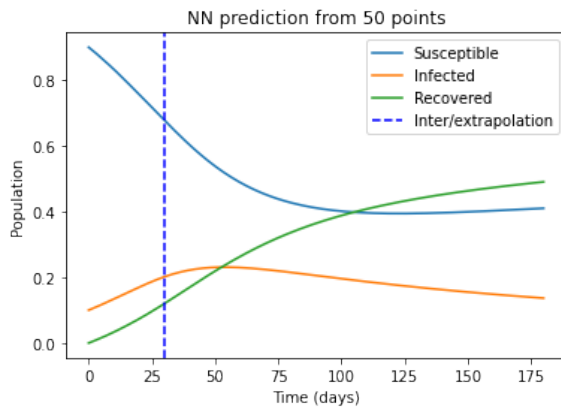
	$\hat{\beta}$	$\hat{\gamma}$	$\hat{\mu}$	Test loss
True value	8×10^{-2}	3×10^{-2}	1×10^{-2}	
20 data points	1.28×10^{-1}	3.31×10^{-2}	4.91×10^{-2}	4.013×10^{-1}
50 data points	8.84×10^{-2}	3.14×10^{-2}	1.50×10^{-2}	4.204×10^{-2}
100 data points	7.08×10^{-2}	2.69×10^{-2}	6.82×10^{-3}	3.043×10^{-2}

Table 5.3: PINN parameter estimations and test losses from limited data.




(a) Training data for SIR endemic model.

(b) PINN prediction from 20 points.



(c) PINN prediction from 50 points.



(d) PINN prediction from 100 points.

Figure 5.5: PINN predictions from limited data.

Chapter 6

Conclusions: Limitations and Perspectives

In conclusion, this project aims to simulate the dynamics of the COVID-19 pandemic through both (1) a space-explicit compartmental model approach and (2) a data-driven machine learning approach; a first-order local sensitivity analysis is also used to identify the impact of each parameter on the SIR system. The space dependence of the model-driven approach holds significant potential for policymakers who are seeking to navigate the complex regulatory landscape of pandemics that involve a multitude of regions. We validate our solver against existing theorems and by cross-validation of two solvers developed independently with different languages. This solver is just a starting point for a tool to test in realistic and complex scenarios (e.g. [7]). The physics-informed neural networks approach combines classic compartmental models with observable data to produce enhanced extrapolation ability for parameter estimation and model prediction with data limitations. Our results show that this method captures the general dynamics of an SIR model with very few data points, while significant data limitation hinders accurate parameter estimation.

For the space-dependent model approach, we consider the following improvements:

Local model The SIR model is just a starting point. Class inheritance enables using more complex models like SEIR [7], up to models incorporating a local space dependence like in [22]. In the latter case, we will combine a multi-scale representation of population mobility, with a PDE for local range and matrices for the long range. Inheritance of classes should guarantee a prompt extension of the current solver when all the GUs are modeled by the same system, more challenging will be the implementation of heterogeneous local solvers (e.g. with complex models in some regions and simpler in others). Other improvements refer to (1) the introduction of a structure of age (particularly evident in the COVID-19 pandemic); (2) the introduction of re-infection of vaccinated people and virus variants in the model.

Numerical methods The Heun’s method is explicit, so to avoid the computational burden of nonlinear solvers, consequently it is only conditionally stable. This requires a manual tuning of the time-step, to avoid numerical oscillations. We are considering the introduction of adaptive explicit solvers (like the Runge-Kutta Fehlberg 45), where the time step is selected automatically to guarantee the accuracy. Notice that the region of stability of the multi-regional solver is driven by the smallest stability region of the single GU. A challenging task will be the introduction of a solver-by-subdomains, where the solution of the global domain is achieved by solving iteratively the subdomains with different time-steps selected adaptively. This may be crucial in a multi-regional multi-scale solvers based on PDEs to contain the computational costs.

Parameter tuning The local parameters and the mobility data need to be retrieved from measurements. As suggested by our results of parameter estimation with neural networks and the first-order sensitivity analysis, the reliability of the results strongly depend on these disease parameters. We argue that, even in presence of complete measures, the parameters will be affected by noise. The first-order local sensitivity analysis that was conducted in the current study is a valuable initial step

in identifying the most influential parameters of the model. However, it should be noted that this method is inherently limited in its ability to accurately capture the complex interactions that may occur between the various model parameters. As such, the use of higher-order sensitivity analysis or global sensitivity analysis is needed to better capture the complex dynamics of the model and thus allows a more comprehensive understanding of the underlying mechanisms driving the model. For parameter estimation, we consider other data-driven methods other than the current neural network approach, such as minimal variance techniques for non-linear problems like the Ensemble Kalman-Filter or Bayesian approaches [5][6].

For the data-driven physics-informed neural networks approach, we propose the following improvements:

Uncertainty quantification Given the promise that PINNs have shown in its predictive power, we need to perform uncertainty quantification (UQ) to ensure its reliability and robustness with the lack of actual data. The UQ process involves characterizing the uncertainties that arise from model inputs, parameters, and the inherent stochasticity of the model itself. Thus, it can provide researchers and policymakers with a more complete understanding of the potential sources of error in the model’s predictions and enable more effective decision-making. Furthermore, UQ can facilitate the identification of areas of the model that require further refinement or calibration to improve its predictive accuracy. Some of the existing UQ methods include Bayesian neural networks [8] and various types of ensemble methods [18].

Model complexity The neural network in this project only considers the classic SIR model, which does not capture the full development of the COVID pandemic. For example, the model does not consider waning immunity, for no recovered population re-enters the susceptible. The model also does not take space dependence into account. For future work, we can continue using neural networks for more com-

plex models such as the Susceptible, Exposed, Infected, Recovered, Waning immunity (SEIRW) model with space dependence as we considered earlier. We may also apply this approach to the space-dependent model in this project. In this case, this deep learning approach is especially beneficial, as the increased dimensionality of a complex problem poses minimal computational hindrance.

Time dependence of parameters Moreover, the current use of neural networks for parameter estimation and prediction does not take time-dependent model parameters into consideration. Although the current model shows predictive potentials with static parameters, its reliability is undetermined under the evolving nature of the pandemic. Future research efforts should explore incorporating time-dependent parameters into the model, potentially through the use of recurrent neural networks [21] or other techniques capable of modeling time-series data. By accounting for the dynamic nature of disease transmission, these models can provide a more comprehensive understanding of the pandemic's evolution and subsequently enable more effective decision-making by policymakers and public health officials.

Bibliography

- [1] Neural network diagram. <https://alexlenail.me/NN-SVG/LeNet.html>. Accessed: 2022-07-17.
- [2] Duygu Balcan, Bruno Gonçalves, Hao Hu, José J. Ramasco, Vittoria Colizza, and Alessandro Vespignani. Modeling the spatial spread of infectious diseases: The global epidemic and mobility computational model. *Journal of Computational Science*, 1(3):132–145, Aug 2010.
- [3] Mohamed Aziz Bhourri, Francisco Sahli Costabal, Hanwen Wang, Kevin Linka, Mathias Peirlinck, Ellen Kuhl, and Paris Perdikaris. Covid-19 dynamics across the us: A deep learning study of human mobility and social behavior. *Computer Methods in Applied Mechanics and Engineering*, 382:113891, 2021.
- [4] Alexander Aram Bruckhaus, Aidin Abedi, Sana Salehi, Trevor A. Pickering, Yujia Zhang, Aubrey Martinez, Matthew Lai, Rachael Garner, and Dominique Duncan. Covid-19 vaccination dynamics in the us: Coverage velocity and carrying capacity based on socio-demographic vulnerability indices in california. *Journal of Immigrant and Minority Health*, 24(1):18–30, 2021.
- [5] Daniela Calvetti, Alexander Hoover, Johnie Rose, and Erkki Somersalo. Bayesian dynamical estimation of the parameters of an se (a) ir covid-19 spread model. *arXiv preprint arXiv:2005.04365*, 2020.

- [6] Peng Chen, Keyi Wu, and Omar Ghattas. Bayesian inference of heterogeneous epidemic models: Application to covid-19 spread accounting for long-term care facilities. *Computer Methods in Applied Mechanics and Engineering*, 385:114020, Jul 2021.
- [7] Marino Gatto, Enrico Bertuzzo, Lorenzo Mari, Stefano Miccoli, Luca Carraro, Renato Casagrandi, and Andrea Rinaldo. Spread and dynamics of the covid-19 epidemic in italy: Effects of emergency containment measures. *Proceedings of the National Academy of Sciences*, 117(19):10484–10491, 2020.
- [8] Ethan Goan and Clinton Fookes. Bayesian neural networks: An introduction and survey. In *Case Studies in Applied Bayesian Data Science*, pages 45–87. Springer International Publishing, 2020.
- [9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [10] Yujia Hao, Yuting Hou, Siwei Xu, Kai Chang, and Zhen Wu. An object-oriented solver for modeling the multi-regional covid-19 outbreak. In *Proceedings of the XLII Ibero-Latin-American Congress on Computational Methods in Engineering and III Pan-American Congress on Computational Mechanics*, CILAMCE-PANACM 2021, Rio de Janeiro, Brazil, 2021. CILAMCE-PANACM. Available upon request.
- [11] Herbert W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.
- [12] Laurent Valentin Jospin, Wray L. Buntine, Farid Boussaïd, Hamid Laga, and Mohammed Bennamoun. Hands-on bayesian neural networks - a tutorial for deep learning users. *CoRR*, abs/2007.06823, 2020.

- [13] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 6 2021.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [15] Akshi Kumar, Sukriti Verma, and Himanshu Mangla. A survey of deep learning techniques in speech recognition. *2018 International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*, Jul 2019.
- [16] Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. Deepxde: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021.
- [17] Barbara D MacCluer, Paul S Bourdon, and Thomas L Kriete. *Differential Equations: Techniques, Theory, and Applications*. American Mathematical Soc., 2019.
- [18] Apostolos F. Psaros, Xuhui Meng, Zongren Zou, Ling Guo, and George Em Karniadakis. Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *Journal of Computational Physics*, 477:111902, Mar 2023.
- [19] Anton Maximilian Schäfer and Hans Georg Zimmermann. Recurrent neural networks are universal approximators. volume 17, pages 253–263, 8 2007.
- [20] Endre Süli and David F. Mayers. *An introduction to numerical analysis*. Cambridge University Press, 2003.
- [21] Ahmed Tealab. Time series forecasting using artificial neural networks method-

ologies: A systematic review. *Future Computing and Informatics Journal*, 3(2):334–340, Nov 2018.

- [22] Alex Viguerie, Guillermo Lorenzo, Ferdinando Auricchio, Davide Baroli, Thomas JR Hughes, Alessia Patton, Alessandro Reali, Thomas E Yankeelov, and Alessandro Veneziani. Simulating the spread of covid-19 via a spatially-resolved susceptible–exposed–infected–recovered–deceased (seird) model with heterogeneous diffusion. *Applied Mathematics Letters*, 111:106617, 2021.