
Semi-supervised Neural Networks solve an inverse problem for modeling Covid-19 spread

Alessandro Patocchio, Tommaso Scarlatti, Marco Brambilla

Politecnico di Milano, Milan, Italy

{alessandrosaverio.patocchio, tommaso.scarlatti}@mail.polimi.it
marco.brambilla@polimi.it

Marios Mattheakis, Pavlos Protopapas

John A. Paulson School of Engineering and Applied Sciences, Harvard University
Cambridge, Massachusetts 02138, United States

{mariosmat, pavlos}@seas.harvard.edu

Abstract

Studying the dynamics of COVID-19 is of paramount importance to understanding the efficiency of restrictive measures and develop strategies to defend against upcoming contagion waves. In this work, we study the spread of COVID-19 using a semi-supervised neural network and assuming a passive part of the population remains isolated from the virus dynamics. We start with an unsupervised neural network that learns solutions of differential equations for different modeling parameters and initial conditions. A supervised method then solves the inverse problem by estimating the optimal conditions that generate functions to fit the data for those infected by, recovered from, and deceased due to COVID-19. This semi-supervised approach incorporates real data to determine the evolution of the spread, the passive population, and the basic reproduction number for different countries.

1 Introduction

COVID-19 has had an enormous global impact, resulting in a broad spectrum of crises across multiple sectors, including public health, social structure, economic stability, and access to education. Countries have been affected at different times, and almost all have reacted by imposing strict lockdown measures to contain the pandemic's effects. Studying the evolution of these procedures is vital to evaluating the effectiveness of the adopted measures, formulating new strategies to improve the response for upcoming waves of contagion, and forecasting the virus's spread to allow for policies of early lockdown or re-opening. The spread of a virus is a time-dependent phenomenon that can be described by differential equations (DEs). A fundamental approach used in epidemiological modeling, which consists of a set of DEs, is the Susceptible-Infectious-Removed (SIR) dynamical model [18] that describes how individuals in a population become infected and removed (recovered or died) by a virus. Recent studies that focus on the COVID-19 pandemic propose analyses of the disease dynamics based on the SIR model [1, 9] and its extensions [7, 8, 10, 12, 17, 19].

We introduce the novel application of a semi-supervised neural network (NN) to study the spread of COVID-19. This method consists of unsupervised and supervised parts and is capable of solving inverse problems formulated by DEs. We also propose an extension of the SIR model to include a passive compartment P , which is assumed to be uninvolved in the spread of the pandemic (SIRP), presenting a novel machine learning technique for solving inverse problems and improving disease modeling. We first present our method and use it for studying synthetic data generated by the SIR model. Then, we introduce the SIRP model and study the pandemic's evolution by applying the semi-

supervised approach to real data, capturing the populations infected and removed by COVID-19 in Switzerland, Spain, and Italy. We conclude with a summary of the key ideas and the most significant results presented in this study.

2 Methodology

We developed a semi-supervised method to determine the optimal parameters and initial conditions of a specific DE system, yielding solutions that best fit a given dataset. The unsupervised part consists of a data-free NN that is trained to discover solutions for a DE system in a high-dimensional parametric space that consists of the modeling-parameters and initial conditions [6]. The loss function solely depends on the network predictions providing an unsupervised learning method. The NN solutions are given in a closed differentiable form [13, 11, 14]. Once a NN is optimized for a particular model formed by DEs, and consideration of the differentiability of solutions, a supervised approach employs a gradient descent optimization method to determine the model parameters and initial conditions that best describe ground truth observations. Automatic differentiation [16] computes the derivatives in gradient descent. An advantage of our approach over standard regression methods is that the predictions respect any underlying constraints embedded in the DE system.

The first part of the proposed method is unsupervised where a feed-forward fully connected neural network [11, 14] is employed to learn solutions of a DE system of the form:

$$\frac{d\mathbf{z}}{dt} = g(\mathbf{z}), \quad \text{with the initial condition, } \mathbf{z}(t=0) = \mathbf{z}_0, \quad (1)$$

where t denotes time, $\mathbf{z} = \mathbf{z}(t, \mathbf{z}_0, \theta)$ is a vector that contains the variables, \mathbf{z}_0 holds the initial values for \mathbf{z} , and θ includes the modeling parameters. The NN takes the inputs $(t, \mathbf{z}_0, \theta)$ and is trained in a certain time range and over predefined intervals of \mathbf{z}_0 and θ (called bundles) [6]. The network returns an output vector \mathbf{z}_{NN} of the same dimensions as the target solutions \mathbf{z} . The learned solutions $\hat{\mathbf{z}}$ satisfy the initial conditions identically by considering parametric solutions of the form:

$$\hat{\mathbf{z}} = \mathbf{z}_0 + f(t)(\mathbf{z}_{\text{NN}} - \mathbf{z}_0), \quad (2)$$

where $f(t) = 1 - e^{-t}$ [14]. The loss function used in the NN optimization is defined by Eq. (1) as:

$$L = \left\langle \left(\frac{d\hat{\mathbf{z}}}{dt} - g(\hat{\mathbf{z}}) \right)^2 \right\rangle_t, \quad (3)$$

where $\langle \cdot \rangle_t$ denotes averaging with respect the time. The auto-differentiation technique [16] is used for the calculation of time derivatives. The proposed architecture is outlined by Fig. 1. Once the NN is trained to provide solutions for the system of Eq. (1), it is used to develop a supervised pipeline for the estimation of \mathbf{z}_0 and θ , leading to solutions that fit given observations denoted by $\tilde{\mathbf{z}} = \tilde{\mathbf{z}}(t)$. This procedure is illustrated in blue in Fig. 1. Starting from random \mathbf{z}_0 and θ , a solution $\hat{\mathbf{z}}(t)$ is generated, then a gradient descent optimizer adjusts \mathbf{z}_0 and θ in order to minimize the loss function:

$$L_{\text{inv}} = \left\langle (\hat{\mathbf{z}} - \tilde{\mathbf{z}})^2 \right\rangle_t. \quad (4)$$

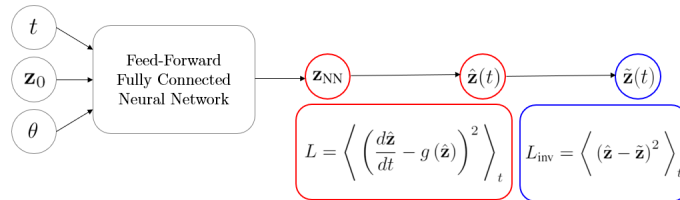


Figure 1: Semi-supervised neural network architecture. Red and blue indicate, respectively, the unsupervised and supervised learning parts.

We first assessed the performance of the proposed method by studying synthetic data generated by the SIR model. The SIR model is a system of non-linear DEs given by:

$$\frac{dS}{dt} = -\frac{\beta SI}{N}, \quad \frac{dI}{dt} = \frac{\beta SI}{N} - \gamma I, \quad \frac{dR}{dt} = \gamma I, \quad (5)$$

where N is the time-invariant total population, $N = S + I + R$. We use $S = S(t)$, $I = I(t)$, $R = R(t)$ to keep the notation elegant. The flow from S to I is regulated by the *infection rate* parameter β , while the flow from I to R is determined by the *recovery rate* parameter γ . An important assumption in the SIR model is that the population in R does not flow either to S or to I . A high-level description of the dynamics of epidemic phenomena is given by *basic reproduction number* \mathfrak{R}_0 that estimates how many new contagions are generated by a single infected person in a population composed only by susceptible people [5, 18]. In the context of the SIR model, we read $\mathfrak{R}_0 = \beta/\gamma$.

For the SIR model, $\theta = (\beta, \gamma)$, $\hat{\mathbf{z}} = (\hat{S}, \hat{I}, \hat{R})$, and $\mathbf{z}_0 = (S_0, I_0, R_0)$, where $S_0 = N - I_0 - R_0$, subsequently, \mathbf{z}_0 is determined by I_0 and R_0 . We work with relative values of $\hat{\mathbf{z}}$ and \mathbf{z}_0 that represent a probability of a compartment. This is achieved by dividing all the compartments by N , yielding a normalized total population equal to one and thus, the constraint $S + I + R = 1$ dictates the quantities \mathbf{z}_0 and $\hat{\mathbf{z}}$ to be bounded between 0 and 1. We use a softmax activation in the output layer of the NN forcing \mathbf{z}_{NN} to take values in $[0, 1]$. Considering that $f(t) \in [0, 1]$, Eq. (2) yields $\hat{\mathbf{z}} \in [0, 1]$, which is the accepted range. According to Eq. (3), the loss function for the SIR model of Eqs. (5) reads:

$$L = \left\langle \left(\frac{d\hat{S}}{dt} + \beta\hat{S}\hat{I} \right)^2 + \left(\frac{d\hat{I}}{dt} - \beta\hat{S}\hat{I} + \gamma\hat{I} \right)^2 + \left(\frac{d\hat{R}}{dt} - \gamma\hat{I} \right)^2 \right\rangle_t. \quad (6)$$

In the unsupervised training process, $2 \cdot 10^3$ equally-spaced time points are sampled from the range $[0, 20]$. The points are perturbed in each iteration, improving the NN predictability [14]. The network architecture is composed by 4 hidden layers with 50 neurons per hidden layer. For the training, we use Adam optimizer with learning rate $8 \cdot 10^{-4}$, and the network is trained for $2 \cdot 10^4$ epochs. In the supervised optimization part of our method we split the training set in 4 mini-batches. We also use Adam optimizer with learning rate 10^{-3} , and the model is trained for 10^3 epochs. For each experiment, we have trained 10 networks and chose the model with the lowest MSE.

We consider a total population $N = 10^7$ and the normalized bundles are: $I_0 = [0.2, 0.4]$, $R_0 = [0.1, 0.3]$, $\beta = [0.4, 0.8]$, and $\gamma = [0.3, 0.7]$. The loss function (6) during the training is represented by the left graph in Fig. 2, where softmax (green) and identity activation functions (red) are used in the output layer. We observe that lower loss value is obtained when softmax activation is used. We implemented the proposed NN in PyTorch [16] and published the code in github ¹.

We employed the semi-supervised pipeline to explore two datasets generated by the SIR model to be considered as the ground truth; these sets are denoted as $\tilde{\mathbf{z}} = (\tilde{S}, \tilde{I}, \tilde{R})$. The aim is to determine which \mathbf{z}_0 and θ generate the $\tilde{\mathbf{z}}$. Indeed, minimizing Eq. (4) yields the \mathbf{z}_0 and θ and the associated SIR solutions that fit $\tilde{\mathbf{z}}$. The middle and right graphs in Fig. 2 present the results of the supervised pipeline. The solid lines show the predictions and the points indicate $\tilde{\mathbf{z}}$. Specifically, we sample 20 equally-spaced points from SIR solutions where 16 points (green points) are used for training, and 4 points (red) are used for validation. Only the infected curves are displayed for simplicity, but we obtained equally accurate predictions for the other compartments. We point out that the predicted fitting curves ensure the conservation of the total population since Eqs. (5) are embedded in the NN architecture, establishing this as an epidemiology-informed model. We proceed by applying the method in a realistic model that is able to describe real data for COVID-19 dynamics.

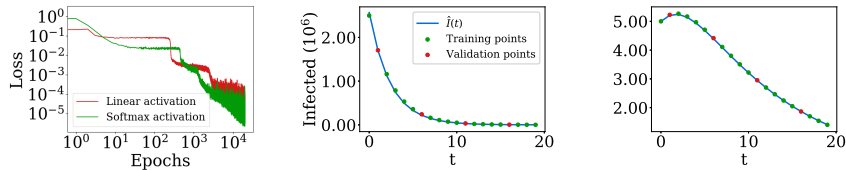


Figure 2: Left: Loss function during the training with softmax (green) and identity (red) activation functions in the last layer. The rapid oscillations appear due to the random perturbation of the training points and are visually amplified by the logarithmic scale of the y-axis. Middle and right: Predictions (solid lines) of the infected population for two different experiments. Green and red points indicate the training and validation sets.

¹https://github.com/alessandropaticchio/semi_supervised_nn_for_covid19

3 COVID-19: Real data

The complexity of the virus spread and the partial quality of data make the simple SIR model incapable of capturing the dynamics of COVID-19. Previous studies used the SIR model to fit only the accumulated infected population [2, 9, 15], while other, more complex, models have been proposed to fit both infected and removed populations [7], we present a simple extension of the SIR model, called SIRP, which can closely fit the data for infected and removed individuals. The model assumes a passive compartment that is not involved in the pandemics' whole dynamics. We examined the effectiveness of the SIRP model and the semi-supervised method by fitting data obtained during the COVID-19 pandemic for three countries: Switzerland, Spain, and Italy [3].

The passive population does not interact with the active compartments S , I , R and thus, P is not considered as susceptible and remains constant in time. Mathematically speaking, we introduce the fourth equation $dP/dt = 0$, with solution $P(t) = P_0$, where P_0 is the initial passive population. The total population in the SIRP model reads $N = S + I + R + P$. We modify the network architecture used to solve Eqs. (5), supplementing an additional input P_0 , resulting in the loss function:

$$L = \left\langle \left(\frac{d\hat{S}}{dt} + \beta\hat{S}\hat{I} \right)^2 + \left(\frac{d\hat{I}}{dt} - \beta\hat{S}\hat{I} + \gamma\hat{I} \right)^2 + \left(\frac{d\hat{R}}{dt} - \gamma\hat{I} \right)^2 + \left(\frac{d\hat{P}}{dt} \right)^2 \right\rangle_t. \quad (7)$$

Although the model parameters can be time-dependent, in specific periods such as lockdown they can be considered constants [4]. We therefore trained our NN that in the lockdown period it would assume constant modeling parameters. Additionally, it has been reported that the real number of I and R is about ten times larger than what data show. This is due to the pandemic's early stage, where testing was not accurate, and samples were not enough to get accurate statistics. Subsequently, the data obtained by [3] are multiplied by a factor of 10. Data give the I_0 and R_0 and are not therefore determined through the pipeline. The optimization process is employed to determine the parameters β, γ , and the conditions S_0 and P_0 . All the compartments have been normalized for the total population of $N \simeq 8.5 \cdot 10^6$ for Switzerland, $N \simeq 4.7 \cdot 10^7$ for Spain, and $N \simeq 6 \cdot 10^7$ for Italy. For these experiments, the training set is split into 10 mini-batches during the supervised optimization.

For each country, we employ and train a different network to learn the solutions of the SIRP system. For the training of each model, we define the bundles on the knowledge of the initial conditions (given by data) and on the intuition about the regions where the real parameters lied. Although I_0 and R_0 are provided by the data, we did not fix them in the unsupervised part. We have empirically observed that the network generalizes better when it is train being trained on a bundle of initial conditions than fixed values. Table 3 defines the training bundles used to study the three countries:

	I_0	R_0	P_0	β	γ
Switzerland	[0.01, 0.02]	[0.001, 0.006]	[0.9, 0.97]	[0.7, 0.9]	[0.15, 0.3]
Spain	[0.01, 0.02]	[0.004, 0.009]	[0.9, 0.97]	[0.4, 0.6]	[0.1, 0.2]
Italy	[0.01, 0.02]	[0.001, 0.006]	[0.9, 0.97]	[0.4, 0.6]	[0.1, 0.2]

Figure 3 presents real data (color points) and predictions (solid lines) for infected (upper row) and removed (lower row) populations. The left column outlines Switzerland's results, the middle accounts for Spain, and the right column represents Italy. We consider training (green) and validation (red) datasets sampled before the end of lockdown, which occurred on April 27th in Switzerland, and on May 4th in Spain and Italy. The training set consists of the first 80% of the data, while the last 20% are used for validation. The data after the lockdown period (orange) are used to evaluate our method's long-term predictability and not involved in any part of the optimization process. We observe that Italy has been the most impacted country, among the ones considered, reaching $\mathfrak{R}_0 = 4.7$ with a significant portion of the population, $P = 96\%$, in the passive state. Spain follows with $\mathfrak{R}_0 = 3.3$ and $P = 95\%$. Switzerland has also $P = 96\%$, with the smallest \mathfrak{R}_0 , resulting in $\mathfrak{R}_0 = 2.7$.

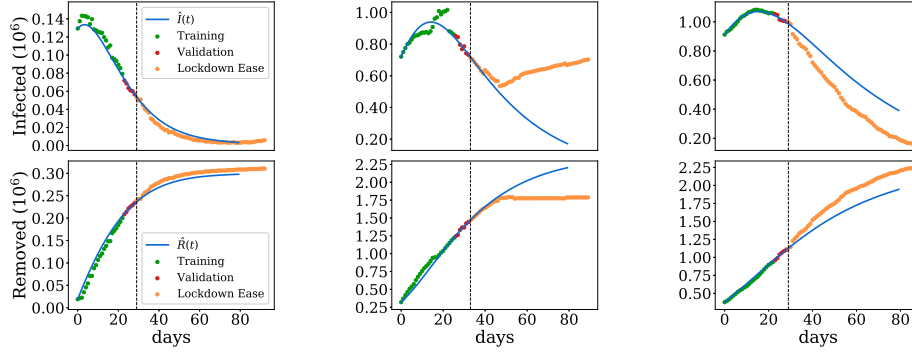


Figure 3: Infected and removed populations for Switzerland (left column), Spain (middle), Italy (right). Points indicate data, solid lines denote predictions, dashed lines show the end of lockdown.

4 Conclusion

We introduced a semi-supervised neural network to solve inverse problems which are formulated by DEs. The method consists of unsupervised and supervised parts. An unsupervised network solves DEs over a range of parameters and initial conditions. A supervised approach incorporates data and uses a gradient descent algorithm to determine the optimal initial conditions and modeling parameters that best fit a given dataset considering a certain model of DEs. We extended the SIR model to include a passive compartment, and showed that the new model, called SIRD, captures the dynamics of COVID-19 spread. We applied the proposed semi-supervised method on real data to study the COVID-19 spread in Switzerland, Spain, and Italy.

Acknowledgments and Disclosure of Funding

The authors would like to thank the anonymous reviewers for their valuable comments on this manuscript.

5 Broader Impact

The semi-supervised method and the analysis presented in this manuscript contribute to the study of COVID-19. Our method was used to solve the inverse problem for existing and new established disease models by incorporating real data. We believe that our results can be further leveraged for the study of virus spread, especially with rigorous data collection. As countries have significantly improved their testing capacity and tracking strategies, data collection now depicts a more realistic scenario, specifically in regards to the early phases of the pandemic. However, while this work presents an elegant and simple method for improving on epidemiological models, it also has applications for applied sciences where DEs play an important role. It could be useful for elaborating problems such as designing material and metamaterials with specific optical properties which consist of an inverse problem. We do not foresee any way that our study can yield any negative outcome regarding ethical aspects. We believe that our work can help in defending upcoming waves of COVID-19 and consequently, retain the balance in the society and improve the daily living conditions.

References

- [1] G.D. Barmparis and G.P. Tsironis. Estimating the infection horizon of covid-19 in eight countries with a data-driven approach. *Chaos, Solitons & Fractals*, 135:109842, 2020.
- [2] Tanujit Chakraborty and Indrajit Ghosh. Real-time forecasts and risk assessment of novel coronavirus (covid-19) cases: A data-driven analysis. *Chaos, Solitons & Fractals*, page 109850, 2020.
- [3] JHU CSSE. Covid-19 data repository by the center for systems science and engineering (csse) at johns hopkins university. <https://github.com/CSSEGISandData/COVID-19>, 2020.

- [4] Jonas Dehning, Johannes Zierenberg, F. Paul Spitzner, Michael Wibral, Joao Pinheiro Neto, Michael Wilczek, and Viola Priesemann. Inferring change points in the spread of covid-19 reveals the effectiveness of interventions. *Science*, 2020.
- [5] Klaus Dietz. The estimation of the basic reproduction number for infectious diseases. *Statistical methods in medical research*, 2(1):23–41, 1993.
- [6] Cedric Flamant, Pavlos Protopapas, and David Sondak. Solving differential equations using neural network solution bundles, 2020.
- [7] Giulia Giordano, Franco Blanchini, Raffaele Bruno, Patrizio Colaneri, Alessandro Di Filippo, Angela Di Matteo, and Marta Colaneri. Modelling the covid-19 epidemic and implementation of population-wide interventions in italy. *Nature Medicine*, 26(6):855–860, Apr 2020.
- [8] Efthimios Kaxiras and Georgios Neofotistos. Multiple epidemic wave model of the covid-19 pandemic: Modeling study. *J Med Internet Res*, 22(7):e20912, Jul 2020.
- [9] Efthimios Kaxiras, Georgios Neofotistos, and Eleni Angelaki. The first 100 days: Modeling the evolution of the covid-19 pandemic. *Chaos, Solitons & Fractals*, 138:110114, 2020.
- [10] 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*, 2020.
- [11] I.E. Lagaris, A. Likas, and D.I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5):987–1000, 1998.
- [12] Qianying Lin, Shi Zhao, Daozhou Gao, Yijun Lou, Shu Yang, Salihu S Musa, Maggie H Wang, Yongli Cai, Weiming Wang, Lin Yang, and Daihai He. A conceptual model for the coronavirus disease 2019 (covid-19) outbreak in wuhan, china with individual reaction and governmental action. *International journal of infectious diseases : IJID : official publication of the International Society for Infectious Diseases*, 93:211—216, April 2020.
- [13] Martin Magill, Faisal Qureshi, and Hendrick de Haan. Neural networks trained to solve differential equations learn general representations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 4071–4081. Curran Associates, Inc., 2018.
- [14] Marios Mattheakis, David Sondak, Akshunna S. Dogra, and Pavlos Protopapas. Hamiltonian neural networks for solving differential equations. *arXiv*, 2020.
- [15] Faical Ndairou, Ivan Area, Juan J Nieto, and Delfim FM Torres. Mathematical modeling of covid-19 transmission dynamics with a case study of wuhan. *Chaos, Solitons & Fractals*, page 109846, 2020.
- [16] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary Devito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [17] 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*, 2020.
- [18] Pauline [van den Driessche]. Reproduction numbers of infectious disease models. *Infectious Disease Modelling*, 2(3):288 – 303, 2017.
- [19] Shilei Zhao and Hua Chen. Modeling the epidemic dynamics and control of covid-19 outbreak in china. *Quantitative Biology*, pages 1–9, 2020.