

Geospatial forecasting of COVID-19 spread and risk of reaching hospital capacity

Georgiy Bobashev¹, Ignacio Segovia-Dominguez², Yulia R. Gel²,
James Rineer¹, Sarah Rhea¹, Hui Sui³

¹Center for Data Science, RTI International, USA

²Department of Mathematical Sciences, University of Texas at Dallas, USA

³University of North Carolina at Chapel Hill, USA

{bobashev,jrin,srhea,hsui.contractor}@rti.org, {ignacio.segoviadominguez,ygl}@utdallas.edu

Abstract

Prompt surveillance and forecasting of COVID-19 spread are of critical importance for slowing down the pandemic and for the success of any public mitigation efforts. However, as with any infectious disease with rapid transmission and high virulence, lack of COVID-19 observations for near-real-time forecasting is still the key challenge obstructing operational disease prediction and control. In this context, we can follow the two approaches to forecasting COVID-19 dynamics: based on mechanistic models and based on machine learning. Mechanistic models are better in capturing an epidemiological curve, using a low amount of data, and describing the overall trajectory of the disease dynamics, hence, providing long-term insights into where the disease might go. Machine learning, in turn, can provide more precise data-driven forecasts especially in the short-term horizons, while suffering from limited interpretability and usually requiring backlog history on the infectious disease. We propose a unified reinforcement learning framework that combines the two approaches. That is, long-term trajectory forecasts are used in machine learning techniques to forecast local variability which is not captured by the mechanistic model.

1 Introduction

Spatio-temporal forecasts of infectious diseases rapidly move to the forefront of policy and public health response because of their key role in risk mitigation strategies. During the COVID-19 pandemic, this has become especially important in areas with high demographic, economic, and political variability. For example, in North Carolina decisions on opening and closing businesses because of COVID-19 are made at the state, county, and local levels. When Wake County leadership announced the lifting of some restrictions, the town of Apex issued an order to continue keeping the strongest restrictions. Similarly, areas that are either remote or used for seasonal vacations (e.g., mountain or beach counties) might exhibit different disease dynamics than those produced by the rest of the state. Because of high spatial heterogeneity, COVID-19 forecasting is important at the local level. Local outbreaks could overwhelm public health systems, hospitals, and emergency rooms.

Hence, to facilitate hospital preparedness, it is critically important to forecast hospital capacity and probabilities that capacity could be overtaken by emerging patients. This relies on two interconnecting models: one predicts how many people are likely to become infected in each of the areas and another

describes how infected individuals move to and between appropriate health care centers. This second model provides the eventual answer to health officials, while the first model stays in the background until needed. To provide the solution at the local level we use an explicit synthetic population, a database that represents the entire population of U.S. residents (over 300 million synthetic individuals). For each of the synthesized individuals (agents), this database contains demographic characteristics and geographic location (Figure 1). More details and a working viewer are available at <http://synthpopviewer.rti.org/> and elsewhere [3, 13]. Multiple layers can be added to the database to make it relevant to a specific question. The COVID-19 layers could include school and work assignment, hospital and emergency department, etc.




Figure 1: Snapshot of the synthetic population for Wilmington, NC.

For the purpose of forecasting hospital load in North Carolina, a team of Response to Intervention (RTI) researchers developed a spatially explicit agent-based model (ABM) that forecasts to which hospital sick patients are likely to move and where they might be transferred if the hospital is over capacity or doesn't have the proper equipment. This model was developed for North Carolina and considers a synthetic population, where synthesized agents represent over 10 million North Carolina residents. The model also uses 110 short-term acute care hospitals (STACHs), 421 nursing homes, and 10 long-term acute care hospitals (LTACHs). At each day timestep, individual health status and location are updated. Figure 2 shows a map with marked locations of health care providers.

The other model provides a forecast of how many individuals are going to be infected and how many will be sick to the point of going to the hospital in the future. Such a model can have different levels of granularity. One level is a county-level system dynamics model that assumes homogeneous mixing (each individual has the same chance to meet any other individual), which leads to the mass action principle, where the risk of infection is proportional to the prevalence of infectious individuals and the proportion of susceptible individuals. Under these assumptions, individuals are equally likely to get infected and thus randomly spread infection through the population. More sophisticated models can include age and social structure (some people have more and closer contacts than others) and geographic locations and allow for disease transmission to occur in clusters.




Figure 2: Locations of health care providers

ABM approaches have been widely used in a large number of areas including health care, epidemiology, economics and finance, and social sciences. Epidemiology of infectious diseases is one of the most natural areas to apply it. For example, an ABM was developed to forecast pandemic flu in North Carolina and New York, where the model explicitly described the household structure and people’s movements and activities through the day including going to work and school, being in a household, or in community spaces and transportation [4, 5]. These models allowed us to identify the areas at the highest risk and estimate the contribution of these areas to the epidemic. For example, for New York, it was shown that public transportation (including subways) contributed less than 10% to disease incidence, while households and schools provided the majority of new cases and thus acted as transmission clusters [10]. Recently we considered synthetic populations to describe seasonal influenza in Russia [8]. The study had an additional challenge of estimating background susceptibility.

The length of the forecasting time horizon is critical for preparedness, but it is also challenging. The longer the time horizon, the higher the uncertainty, especially when policy and environment change in unpredictable ways. Nevertheless, one could foresee at least some impacts such as weekends, major holidays, and scheduled public announcements such as on business and school closures and openings.

2 Methodology

There are at least two approaches to forecasting COVID-19 dynamics: based on mechanistic models and based on machine learning. Mechanistic models are better at capturing an epidemiological curve and describing the overall trajectory of disease dynamics, hence providing long-term insights into where the disease might go. Machine learning, in turn, can provide more precise data-driven forecasts, especially in the short-term horizons, while suffering from limited interpretability. We propose a unified reinforcement learning (RL) framework combining the two approaches. That is, long-term trajectory forecasts are used in machine learning techniques to forecast local variability, which is not captured by a mechanistic model.

The choice of the “best” model generally balances model fidelity, explanatory features, data availability, and computational requirements. A summary of these features is described in [2] and is presented in Figure 3. In [11] we conducted a comparison of the simpler system dynamics model and an ABM of pandemic flu with a number of interventions. Not surprisingly, higher granularity brings higher

fidelity but also increases uncertainty because of the variability of model parameters and structural assumptions.

In the current study, we start mechanistic modeling with a system dynamics approach. System dynamics models are quick to execute and thus are easy to calibrate. Future ABMs will start with the average parameter values of the system dynamics model and will expand around those values. Our system dynamics model divides the population into Susceptible, Exposed, Infectious, and Recovered individuals and considers the movements of individuals between these compartments. Thus, the model is commonly called an SEIR model. Specifically, for COVID-19 we also consider whether individuals are symptomatic or asymptomatic, which in turn requires estimation of disease transmission from an asymptomatic person. Assuming homogeneous mixing, the initial SEIR model can be described in a differential equation form:

$$\begin{aligned}
\frac{dS^m}{dt} &= - \sum_v (\beta_a^{mv} I_a^v + \beta_s^{mv} I_s^v) S^m, \\
\frac{dE_k^m}{dt} &= \sum_v p_k^m [\beta_a^{mv} (I_a^v - \xi_a I_a^v) + \beta_s^{mv} (I_s^v - \xi_s I_s^v)] S^m - \mu_k E_k^m, \\
\frac{dI_k^m}{dt} &= \mu_k E_k^m - \gamma_k I_k^m - \delta_k I_k^m \quad \frac{dR_k^m}{dt} = \gamma_k I_k^m
\end{aligned} \tag{1}$$

where S , E , I , and R denote susceptible, exposed, infected, and recovered, respectively, and for simplicity total immunity is assumed. In turn, $k = a, s$ denote symptomatic and asymptomatic subgroups, respectively; p_k is a fraction of symptomatic or asymptomatic cases; and β , μ , γ and δ are transmission rate, infectivity period, recovery, and excess mortality rates, respectively. Quarantine of identified symptomatic and asymptomatic cases is denoted by ξ . We assume that there exists no excess mortality among asymptomatic cases (i.e., $\delta_a = 0$). Upper indexes correspond to age groups m and v . Finally, we assume homogeneity of the daily number of contacts within an age group and heterogeneity between




Figure 3: Hierarchy of simulation models in increasing levels of complexity. At the higher level Statistical, Markov, and System dynamics models do not distinguish between individuals in the populations and describe populations (or subpopulations) as a whole. As the names suggest, microsimulation and agent-based models describe each individual in the population and thus could be averaged across specific characteristics to obtain population-level estimates. In the current project, we consider a challenging question. Do these modeling approaches have to be mutually exclusive? Each brings something to the forecast, and perhaps we can benefit from combining at least two of them.

groups.

This model is defined at the county level and could be further expanded to represent rural or urban parts of the county, adding migration of individuals between counties and between urban and rural county components. The model is adequate enough to describe the dynamics of COVID-19 in most counties. For low-density rural counties (e.g., in Appalachia) the assumption of homogeneous mixing is no longer valid, and an explicit ABM might be more adequate. For the sake of simplicity, we still keep the differential equation formalism but add a stochastic transmission component. Future models will be fully agent-based to consistently describe local geographic elements of disease clusters and patient assignments to hospitals. We fit the model to the reported case data where we also need to consider multiple reporting biases such as under-reporting, reporting of certain subpopulations (e.g., age 65+), and the availability of disease test kits.

Disease transmission parameters are key to the understanding of future disease dynamics; therefore, we also consider county vulnerability indices, which we calculated based on multiple sources of socioeconomic and health data. We have developed a county vulnerability dashboard (Figure 4) that is publicly available at <https://RTImerge.org>.




Figure 4: County vulnerability index with respect to COVID-19 dynamics

This mechanistic modeling effort results in describing the dynamics of symptomatic and asymptomatic infected individuals. Figure 5 shows an example of a model fit to Gaston county data. Deterministic SEIR models produce a smooth curve fit that tracks past mean-field disease dynamics. The key component of prediction is forecasting how policy measures will impact the transmission rate. For that purpose, we developed a model linking the dynamics of a SEIR beta parameter with public health actions. Mechanistic models allow one to simulate a variety of scenarios and pre-train a deep learning model on these scenarios. Although the SEIR model allows us to consider “what if” scenarios and provide mechanistic explanations of “why,” they don’t capture all the richness of disease dynamics. The addition of stochastic components leads to consideration of uncertainty and higher boundaries of risk through a family of stochastic realizations.

Mechanistic models are limited in their applications because they only produce results based on a hardwired mechanism which could be miss-specified. Furthermore, such models can neglect a variety




Figure 5: Examples of model fit to Gaston County case reports

of smaller factors that are collectively influential. These factors could be captured by a mechanism-agnostic data-driven model that learns with the data. For example, Long-Short Term Memory (LSTM) algorithms can capture short- and long-term factors that can also change in time.

In this project, we aim to develop a reinforcement learning (RL) framework which systematically combines mechanistic SEIR models with data-driven LSTM algorithms to get the best of both worlds: that is, the interpretability of mechanistic models and predictive capabilities of deep learning methods.

Our approach is described below and is illustrated in Figure 6:

1. Based on past data we fit the SEIR model and train the LSTM model.
2. Based on forecasted subjective beliefs on policy changes in the future use the SEIR model to predict the numbers of infected individuals (smooth curves).
3. Use the SEIR-predicted data as input into the LSTM model to produce an improved forecast.
4. Through reinforcement learning update both the SEIR and LSTM forecasts as new data become available.

Our goal is to build a highly adaptive model capable of readjusting its prediction on the future disease course, in accordance with changes in public response and bio-atmospheric information. As shown in Figure 6, RL plays an important role to control the dynamic of the SEIR beta parameter, which is directly linked to public health actions. Hence, this model-free RL adapts its parameters on the fly (i.e., learning from experience and using the latest updated official data). As a consequence, we will be able to update our predictions in potentially highly uncertain and volatile disease scenarios such as the current coronavirus spread.

The main difference between the proposed RL methodology with respect to other approaches is the ability to summarize multiple policies for outbreak response via parameter adaptation in our mechanistic model. Indeed, existing techniques primarily focus on learning context-dependent policies for complex epidemiological models, in which the conventional approach consists of evaluating the expected




Figure 6: Reinforcement learning of mechanistic (SEIR) and data-driven (LSTM) models.

performance of different potential interventions via stochastic simulations [9, 12]. However, such approaches predominantly model the RL problem from a spatial perspective, leaving the time variable as only part of the evolution of the states [1, 6] and mostly obtaining optimal decisions in fully observable environments [7]. In contrast, our proposed approach accounts for the dynamics of COVID epidemics via mechanistic models and allows us for building highly adaptive models capable of readjusting the delivered forecasts for the future disease course.

3 Conclusion and Future Work

The ultimate goal of this project is to develop a novel methodology for forecasting the COVID-19 spread via synergistic interaction between mechanistic and data-driven models under the RL framework. In particular, the proposed methodology is based on the new idea of using RL as a part of solving a time-series forecasting problem under the assumption of dynamic stability and requires identification of the following main components: states, environment, reward function, and agent interactions. Since RL focuses on learning an optimal policy, we also need to obtain significant feedback between agents and the ongoing behavior of the system and ensure that our adaptive learning still can be formulated as a Markov Decision Process problem (i.e., the challenges that are both largely unsolved in a context of RL for space-time data).

By effectively combining mechanistic models with deep learning tools, the proposed RL approach to epidemiological forecasting can harness the strength of both theoretical and data-based models and deepen our understanding of the hidden mechanisms behind COVID-19 progression. In the near term future, we plan to investigate the utility and limitations of the proposed methodology at the county level and then investigate the transferability of the derived tools to other states and spatial data resolution.

Acknowledgements

This work is supported in part by the grants from the National Science Foundation DMS 2027793/2027802 and the ConTex Postdoctoral Fellowship grant.

References

- [1] D. Bertsekas. *Rollout, Policy Iteration, and Distributed Reinforcement Learning*. Athena Scientific, 1 edition, 8 2020.
- [2] G. Bobashev. Simulation modeling of hiv infection—from individuals to risk groups and entire populations. In C. Chan, M. G. Hudgens, and S.-C. Chow, editors, *Quantitative Methods for HIV/AIDS Research*, chapter 10, pages 201–229. Chapman & Hall/CRC Biostatistics Series, 1 edition, 8 2017.
- [3] J. Cajka, P. Cooley, and W. Wheaton. Attribute assignment to a synthetic population in support of agent-based disease modeling. *Methods report (RTI Press)*, 19:1–14, 09 2010.
- [4] P. Cooley, L. Ganapathi, G. Ghneim, S. Holmberg, W. Wheaton, and C. Hollingsworth. Using influenza-like illness data to reconstruct an influenza outbreak. *Mathematical and Computer Modelling*, 48:929–939, 02 2008.
- [5] N. Ferguson, D. Cummings, C. Fraser, J. Cajka, P. Cooley, and S. Burke. Strategies for mitigating an influenza pandemic. *Nature*, 442:448–52, 08 2006.
- [6] P. Hernandez-Leal, B. Kartal, and M. Taylor. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 33, 10 2019.
- [7] S. Ivanov and A. D’yakonov. Modern deep reinforcement learning algorithms, 2019. Available at <https://arxiv.org/abs/1906.10025>.
- [8] V. Leonenko and G. Bobashev. Analyzing influenza outbreaks in russia using an age-structured dynamic transmission model. *Epidemics*, 29:100358, 2019.
- [9] P. Libin, A. Moonens, T. Verstraeten, F. Perez-Sanjines, N. Hens, P. Lemey, and A. Nowé. *Deep reinforcement learning for large-scale epidemic control*. ALA 2020, Auckland, New Zealand, 5 2020.
- [10] I. M. Longini, A. Nizam, S. Xu, K. Ungchusak, W. Hanshaoworakul, D. A. T. Cummings, and M. E. Halloran. Containing pandemic influenza at the source. *Science*, 309(5737):1083–1087, 2005.
- [11] J. Mathieu, M. Pfaff, G. Klein, J. Drury, M. Geodecke, J. James, P. Mahoney, and G. Bobashev. *Tactical Robust Decision-Making Methodology: Effect of Disease Spread Model Fidelity on Option Awareness*. Seattle, WA, 01 2010.
- [12] W. Probert, S. Lakkur, C. Fonnesebeck, K. Shea, M. Runge, M. Tildesley, and M. Ferrari. Context matters: using reinforcement learning to develop human-readable, state-dependent outbreak response policies. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 374:20180277, 07 2019.
- [13] W. D. Wheaton, J. Cajka, B. M. Chasteen, D. Wagener, P. Cooley, L. Ganapathi, D. J. Roberts, and J. Allpress. Synthesized population databases: A us geospatial database for agent-based models. *Methods report*, 2009 10:905, 2009. Available at <http://www.rti.org/sites/default/files/resources/mr-0010-0905-wheaton.pdf>.