AI & Multi-agent Systems for Data-centric Epidemic Forecasting

Doctoral Consortium

Alexander Rodríguez
Georgia Institute of Technology
Atlanta, GA, USA
arodriguezc@gatech.edu

ABSTRACT

Epidemic forecasting is a crucial tool for public health decision making and planning. There is, however, a limited understanding of how epidemics spread, largely due to other complex dynamics, most notably social and pathogen dynamics. With the increasing availability of real-time multimodal data, a new opportunity has emerged for capturing previously unobservable facets of the spatiotemporal dynamics of epidemics. In this regard, my work brings a data-centric perspective to public health via methodological advances in AI at the intersection of time series analysis, spatiotemporal mining, scientific ML, and multi-agent systems. This extended abstract focuses on our new techniques for end-to-end learning with mechanistic epidemiological models—based on differential equations and agent-based models—that bridge ML advances and traditional domain knowledge to leverage individual merits. I finalize discussing some future directions for my work.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; **Agent / discrete models**; • **Applied computing** → Epidemiology.

KEYWORDS

Differentiable Agent-based Modeling; Scientific Machine Learning; Computational Epidemiology; Deep Neural Networks

ACM Reference Format:

Alexander Rodríguez. 2023. AI & Multi-agent Systems for Data-centric Epidemic Forecasting : Doctoral Consortium. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

As the recent COVID-19 experience has shown, preventing epidemics and pandemics is one of the major challenges of our time, with far reaching impacts on health, economy, and broad social well-being. One of the key prevention tasks is prediction of the future spread of epidemics in the population (forecasting). These predictions are used for various decision-making purposes, from resource allocation to individual risk assessment.

AI and public health are becoming increasingly intertwined due to several accelerating trends. (a) There is an increasing availability of multimodal data that provides real-time information about the spatiotemporal dynamics of epidemics. These include mobility, symptomatic searches, behavioral online surveys, and health monitoring from wearable devices. (b) Advances in AI, especially in deep learning, have shown impressive results in capturing complex relationships across data modalities resulting in greater generalization power. Hence, the time is ripe to leverage AI in developing public health solutions.

My work focuses on the following challenges to realize the potential of AI, within the context of a data-centric pipeline for public health [16]. First, data is scarce, sources may disagree, and domains or distributions may unexpectedly shift over time. Second, real-time deployment for response to disease outbreaks faces many issues not present in carefully controlled environments such as reporting delays, missing values, and data revisions. Third, prior epidemiological (scientific) knowledge on the mechanisms of epidemic spread is necessary to predict how epidemics will unfold over long time horizons and to answer counterfactual questions.

To address these challenges, my research proposes: (a) Novel robust and modular deep learning architectures for response to disease outbreaks focused on real-time deployment. For instance, an operational deep learning pipeline for explainable real-time COVID-19 forecasting [15]. (b) New techniques for end-to-end learning with mechanistic epidemiological models (based on differential equations and agent-based models), that bridge ML advances and traditional domain knowledge to leverage individual merits. For example, designing differentiable agent-based models that are seamlessly coupled with neural networks and can leverage gradient-based learning via automatic differentiation [2]. In this extended abstract, my focus is on the second set of methods validated in real-world data for Influenza and COVID-19.

2 BRIDGING NEURAL NETWORKS AND THEORETICALLY-GROUNDED ODE MODELS

My work introduces epidemiologically-informed neural networks (EINNs) [13]. Our neural forecasting models have been successful in short-term forecasting, which typically is up to four weeks ahead–e.g., see [3] for results of our operational deep learning framework GT-DeepCOVID [15]. Nevertheless, long-term predictions well-correlated with epidemic trends remains an open challenge. Epidemiological ODE models–such as the SIR (susceptible-infected-recovered) model [6]–contain mechanisms that can guide us in this task, however, they have limited capability of ingesting data sources and modeling composite signals. Thus, we propose to incorporate epidemic dynamics from an epidemiological ODE model into a neural framework for forecasting, which enables seamless integration of multimodal data, greater representation power, and inclusion of composable neural modules of learned representations. We leverage the rapidly growing literature in physics-informed neural networks (PINNs) [10] that integrate neural networks and

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

ODEs. PINNs, however, cannot incorporate datasets that are not explicitly described in the differential equations and their architecture is the multi-layer perceptron, which is a subpar inductive bias for sequential data. We developed an end-to-end framework to learn the epidemic mechanistic dynamics via a PINN and transfer its representations to a deep sequential neural network which ingests novel data sources. Unlike previous work, we do not assume complete dynamics can be observed and do not need to integrate over the ODE during training. Our approach leads to significant improvements in long-term forecasting and trend correlation.

3 DIFFERENTIABLE PROGRAMMING WITH ABMS AND NEURAL NETWORKS

My work introduces GradABM, a differentiable design for agent-based epidemiological models that can be seamlessly coupled with neural networks [2]. Agent-based models (ABMs) are simulators where agents act and interact within a computational environment. In epidemiology, ABMs have become increasingly popular since they are more flexible and able to model more details than ODEs. These models are, however, traditionally slow [1], difficult to calibrate (fit) to real-world data [4, 17], and require additional layers of complexity for incorporating novel datasets [7]. We could use neural modules to seamlessly incorporate data sources, but ABMs are based on non-differentiable and discrete operations, which also prevents us from using a framework like EINNs. Thus, we propose to redesign ABMs to be tensorized and fully differentiable simulators, which enables end-to-end learning with deep neural networks to incorporate novel data sources into calibration and prediction of ABMs.

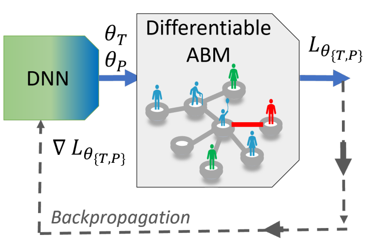

We take advantage of existing invariances in disease transmission and reformulate it as message passing operations over sparse networks. This results in quick and highly parallelized forward simulations. To make our ABM fully differentiable, we reparametrized discrete distributions used in disease transmission and progression with continuous relaxations. This allows our method to take advantage of the practical benefits of gradient-based optimization for calibration and forecasting.

4 DISCUSSION AND FUTURE WORK

This extended abstract outlines two techniques that lay the groundwork for facilitating a closer integration between AI advancements and mechanistic models/simulators in epidemiology. However, incorporating more intricate mechanistic epidemiological models presents certain challenges that need to be addressed in future work. For example, as the number of compartments in ODE models increases, the resulting complexity leads to greater unobservability, necessitating additional engineering efforts to constrain the optimization problem. Additionally, to introduce more features such as vaccination strategies into our differentiable ABM, we must find a way to make such interventions a differentiable operation. It appears unlikely that a single solution that suits all situations can be found in these efforts. On another front, we have demonstrated that our techniques are applicable to various airborne infectious diseases, namely Influenza and COVID-19. In future research, it would be valuable to explore the constraints of these approaches in modeling other infectious diseases that have distinct transmission modes, such as vector-borne diseases and sexually transmitted diseases.

Additionally, I am interested in building on our successes to continue to push the boundaries of ML methods for modeling epidemics and other spatiotemporal processes in social and technological systems. The following are some promising directions I am excited about. Despite focusing on public health as their motivation, I am eager to examine these research directions in other social impact areas, such as socioeconomic development, humanitarian crises, sustainability, and community resilience. For the latter, I plan to leverage my previous work on disaster resilience of critical infrastructure networks [11, 14].

AI-augmented differentiable simulators. My current research addresses end-to-end learning with fully detailed epidemiological models. I am eager to study cases where domain or scientific knowledge is incomplete. For instance, the effects of genetic mutations and climate on disease infectiousness are scientifically understood, yet existing models of disease transmission (such as those used in my research) do not take these factors into account, and instead largely center around human interactions. To bring together these two pieces of scientific knowledge in a unified simulator, I will leverage my work in differentiable ABMs and create differentiable simulators augmented with ML modules that fill in the gaps in scientific knowledge. These new simulators can also be coupled with differentiable solvers for discrete optimization [18] and trained end-to-end, which allows us to optimize our predictions to be most useful for downstream decisions to be taken. I believe this novel class of simulators has a strong potential to amplify the impact of AI across a variety of scientific disciplines.

Principled AI methods for bias and equity issues. The results of my research on systematically addressing missing values [15] and data revisions [8] show that developing methods to tackle data-related issues is a promising direction. I am interested in addressing other urgent issues such as accounting for sampling biases and inequity in data collection. For instance, the CDC reports flu-like illness patients based on data from a limited number of voluntary healthcare providers, which accounts for less than 5% of the population and is often biased towards wealthier individuals. Further, datasets from digital sources also present biases due to costs associated with access and their opt-in nature. By utilizing these datasets to monitor and predict disease activity, there is a risk of underestimating the healthcare needs of low-income communities, thus further diverting resources away from them. Currently, it is challenging to create equitable ML methods, as it requires taking into account how people from different subgroups of the population interact with one another. I plan to investigate ML methods that can systematically account for these bias and inequity issues using my deep learning frameworks and my previous studies on using AI safety techniques [12]. In my work, I will provide an alternative viewpoint on addressing data biases by utilizing theoretically-grounded models and granular representations, which will complement existing related research [5, 9].

ACKNOWLEDGEMENTS

This work was supported in part by the NSF (Expeditions CCF-1918770, CAREER IIS-2028586, PIPP CCF-2200269) and funds/computing resources from Georgia.

REFERENCES

[1] Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Arnau Quera-Bofarull, Miguel Icaza-Lizaola, Aidan Sedgewick, Henry Truong, Aoife Curran, Edward Elliott, Tristan Caulfield, Kevin Fong, et al. 2021. June: open-source individual-based epidemiology simulation. *Royal Society open science* 8, 7 (2021), 210506.

[2] Ayush Chopra*, A. Rodríguez*, Jayakumar Subramanianikari, Balaji Krishnamurthy, B Aditya Prakash, and Ramesh Raskar. 2023. Differentiable Agent-based Epidemiology. In *Proceedings of the 22th International Conference on Autonomous Agents and MultiAgent Systems.*

[3] Estee Y Cramer, Evan L Ray, Velma K Lopez, Johannes Bracher, Andrea Brennen, Alvaro J Castro Rivadeneira, Aaron Gerding, Tilmann Gneiting, Katie H House, Yuxin Huang, et al. 2022. Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the United States. *Proceedings of the National Academy of Sciences* 119, 15 (2022), e2113561119.

[4] Wouter Edeling, Hamid Arabnejad, Robert C Sinclair, Diana Suleimenova, Krishnakumar Gopalakrishnan, Bartosz Bosak, Derek Groen, Imran Mahmood, Qureshi Hashmi, Daan Crommelin, et al. 2020. Model uncertainty and decision making: Predicting the Impact of COVID-19 Using the CovidSim Epidemiological Code. (2020).

[5] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 259–268.

[6] Herbert W Hethcote. 2000. The mathematics of infectious diseases. *SIAM review* 42, 4 (2000), 599–653.

[7] Robert Hinch, William JM Probert, Anel Nurtay, Michelle Kendall, Chris Wymant, Matthew Hall, Katrina Lythgoe, Ana Bulas Cruz, Lele Zhao, Andrea Stewart, et al. 2021. OpenABM-Covid19—An agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing. *PLoS computational biology* 17, 7 (2021), e1009146.

[8] Harshavardhan Kamarthi, A. Rodríguez, and B Aditya Prakash. 2022. Back2Future: Leveraging Backfill Dynamics for Improving Real-time Predictions in Future. *International Conference on Learning Representations* (2022).

[9] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems* 33, 1 (2012), 1–33.

[10] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. 2021. Physics-informed machine learning. *Nature Reviews Physics* 3, 6 (2021), 422–440.

[11] A. Rodríguez, Bijaya Adhikari, Andres D. Gonzalez, Charles Nicholson, Anil Vullikanti, and B. Aditya Prakash. 2020. Mapping Network States using Connectivity Queries. In *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, Atlanta, GA, USA, 778–787. https://doi.org/10.1109/BigData50022.2020.9378355

[12] A. Rodríguez, Bijaya Adhikari, Naren Ramakrishnan, and B Aditya Prakash. 2020. Incorporating Expert Guidance in Epidemic Forecasting. In *ACM SIGKDD 2020 Epidemiology Meets Data Mining and Knowledge Discovery (epiDAMIK)*.

[13] A. Rodríguez, Jiaming Cui, Naren Ramakrishnan, Bijaya Adhikari, and B. Aditya Prakash. 2023. EINNs: Epidemiologically-Informed Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37.

[14] A. Rodríguez and Charles D Nicholson. 2018. Data-based Methodology for Mitigation Decision-making in Critical Infrastructure. In *IISE Annual Conference. Proceedings*. Institute of Industrial and Systems Engineers (IISE), 182–187.

[15] A. Rodríguez, Anika Tabassum, Jiaming Cui, Jiajia Xie, Javen Ho, Pulak Agarwal, Bijaya Adhikari, and B Aditya Prakash. 2021. DeepCOVID: An Operational Deep Learning-driven Framework for Explainable Real-time COVID-19 Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 15393–15400. Issue: 17.

[16] Alexander Rodríguez, Harshavardhan Kamarthi, Pulak Agarwal, Javen Ho, Mira Patel, Suchet Sapre, and B. Aditya Prakash. 2022. Data-Centric Epidemic Forecasting: A Survey. *arXiv preprint arXiv:2207.09370* (2022).

[17] Santiago Romero-Brufau, Ayush Chopra, Alex J Ryu, Esma Gel, Ramesh Raskar, Walter Kremers, Karen S Anderson, Jayakumar Subramanian, Balaji Krishnamurthy, Abhishek Singh, et al. 2021. Public health impact of delaying second dose of BNT162b2 or mRNA-1273 covid-19 vaccine: simulation agent based modeling study. *bmj* 373 (2021).

[18] Bryan Wilder, Bistra Dilkina, and Milind Tambe. 2019. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1658–1665.