Workshop 2: Modeling the Impact of Intervention Policies for Disease Prediction

Publication Note: This document is a "white paper" created as background for and a summary of discussion by participants at a hybrid workshop held on January 27, 2023. The workshop was supported by the National Science Foundation Predictive Intelligence for Pandemic Prevention (PIPP) Initiative. The contents of this document are the opinions of the authors and participants.

Introduction and Background

For decades, the study of infectious diseases has been of great interest and importance, given their potential to cause significant harm to human civilization. One commonly used approach to modeling infectious diseases is the compartmental epidemic model, which assigns members of the population to labeled compartments, such as Susceptible, Infected, and Removed (SIR model).

Compartmental models have gained wide adoption due to their simplicity and ease of implementation. However, as these models often assume homogeneous mixing of the population, they are often unable to capture the complexities of real-world populations, particularly as the world becomes more connected and populations become more diverse.

To address this, researchers have turned to agent-based model/ing (ABM). Unlike traditional compartmental models, ABMs aim to model the population at a more fine-grained individual level, incorporating realistic population characteristics, social contact networks, and a parameterized disease model. Within ABMs, microscopic pandemic simulators have been proposed [1], [2], building on top of an SIR-style epidemic model to simulate agents' behavior at the most granular level (e.g., the varying frequency of store visits and social events among each age group, the household supply level of each person on each day, and so on). With the help of pandemic simulators, ABMs can predict the effects of government intervention policies, such as the impact of different physical distancing measures and stay-at-home orders [3].

Like ABMs, game theory also focuses on agents and their interactions and thus is often adopted for modeling decision-making processes during epidemics. Game-theoretic models can operate at various levels, from individual to population, and can account for different types of decision-makers, including individuals and government authorities. These models can also focus on the microscopic interaction dynamics between individuals or the interactions between regions, states, or even countries at the population level.

Broadly, this workshop explored recently proposed methods for prioritizing and predicting the impact of intervention policies for limiting outbreak threats at both the individual and population levels. The workshop was organized into four sessions, with each session addressing a specific topic within this broader context. In the first three sessions, an expert speaker presented recent methods related to the following three topics: (1) AI and multiagent systems for social impact, (2) developing game-theoretic epidemic models for multi-region intervention planning, and (3) microscopic epidemic models in agent-based modeling. These speaker presentations were then followed by a discussion amongst all attendees about potential new methods that could be developed for predicting the impact of intervention policies during epidemics. Then, the fourth and final session consisted of a breakout group discussion and an all-onboard discussion, which served as a platform for the attendees to delve deeper into the insights gained from the previous sessions.

AI and Multiagent Systems for Social Impact

The use of AI and multiagent systems has gained popularity in optimizing the allocation of limited resources in public health settings. These models have been applied to various interventions, including tuberculosis treatment adherence [4], maternal and child health [5], and general adherence dynamics in the context of educational training programs and medication compliance [8]. As the world responds to future and ongoing infectious disease crises with limited resources, understanding how AI and multiagent techniques have previously been used in the public health domain is essential.

The guest speaker of the session presented a talk that focused on the Restless Multi-Armed Bandits (RMABs) model. This model has been successfully employed in public health interventions, particularly in the maternal and child care sector in India, where the mortality rate is considerably higher than the UN target. For example, a program that sends weekly automated calls to provide health information to Indian mothers has demonstrated a 30% reduction in maternal mortality; however, a significant percentage (30-40%) of the recipients did not engage with the program actively (low listeners). To address this issue, AI and multi-agent systems have been utilized to deliver personalized calls to those in danger of becoming low listeners. Due to limited resources, only a small fraction of beneficiaries could be contacted. Thus, to identify the mothers who were most likely to disengage from the program, the RMAB model was used to rank all the participants, and the top 1,000 were selected for outreach from a pool of nearly 200,000. This resulted in a significant increase in health message uptake. The implementation of this RMAB model, named "SAHELI" and launched in April 2022, has increased content exposure by 130% among the bottom 25% of beneficiaries with respect to listenership [9].

The speaker then presented the next steps of their research agenda for the use of RMABs for Social Impact, particularly in the context of public health. First, for improving the quality of decision-making, there is an important effort to incorporate decision-focused learning in restless bandits. Decision-focused learning is an approach to training AI algorithms that prioritizes the quality of decision-making over intermediate goals, such as accurately predicting the behavior of every individual. For instance, in the maternal and child care program mentioned above, an RMAB model trained using decision-focused learning would aim to allocate resources more efficiently by having better predictive abilities over a specific subgroup of beneficiaries who are more likely to disengage from the program, rather than striving for high prediction accuracy over all beneficiaries [7]. Second, the speaker mentioned that the researchers working on this project are creating a strategy called Robust Restless Bandits via MinMax Regret. This approach considers the uncertainty in the transition between states that represent the level of engagement of beneficiaries in the program. Additionally, it was also mentioned that adherence monitoring for preventing tuberculosis through calls to remind patients to take their medication is also being explored using Partially Observable Markov Decision Processes (POMDPs), namely because the state of the patient is not known ahead of time. Finally, the speaker mentioned that flexible budget RMABs, which enable the ability to have varying amounts of available resources at each decision step, are also being investigated. The speaker and their team is particularly interested in exploring whether a flexible budget, along with the previously mentioned extensions to RMABs, could be useful in epidemic settings.

The discussion following the speaker's presentation focused on how RMABs can be used in pandemics and which characteristics of pandemics can be modeled using an RMAB framework. It was suggested that flexible budgets, non-stationary transition probabilities, and interventions like awareness campaigns and pharmaceutical (i.e., prophylactic, diagnostic, and therapeutic) allocation are important in pandemic settings. The group also explored how RMABs can be used in hospital settings for labor allocation, emergency room triage, and bed allocation, as well as in awareness campaigns for population-level impact. During the discussion, ethical and fairness concerns related to the use of these models for allocating life-saving interventions were raised. Participants stressed the significance of having well-evaluated fairness metrics in place to ensure that resources are allocated in a manner that is equitable, particularly for disadvantaged and marginalized subpopulations during infectious disease crises.

Developing Game-Theoretic Epidemic Model for Multi-Region Intervention Planning

During a pandemic, it is essential to develop effective interventions to control and prevent the spread of infectious diseases. However, designing interventions is a complex task, as the spread of diseases is influenced by various factors, including social, economic, and health policies. To address this challenge, the session speaker presented the innovative introduction of game-theoretic approaches into traditional multi-region epidemic models. By incorporating the interactions between regions and the impact of policies issued by multiple region-specific decision-makers, these models provide a more comprehensive understanding of the spread of infectious diseases across regional boundaries and enable the exploration of a wide range of intervention strategies. As such, they have the potential to inform policymakers and public health officials on the most effective measures to mitigate the spread of pandemics.

The speaker presented a novel game-theoretic multi-region SEIR model that takes into account how social and health policies issued by multiple region-specific decision-makers affect the progress of infectious diseases. However, dealing with this complexity poses a significant computational challenge due to the large number of possible decisions to take and the interdependence of such decisions. To overcome this challenge, the model presented uses an enhanced deep fictitious play algorithm.

The presented game-theoretic multi-region SEIR model was tested on a COVID-19 dataset from three states, New York, New Jersey, and Pennsylvania. The study showed how travel ban policies affected the spread of SARS-CoV-2 for each state and demonstrated how the public's willingness to comply and decision-makers' attitudes towards death influenced equilibrium strategies. However, policies instituted during a pandemic are usually short-term and frequently adjusted, and thus repeated games can be used to model this phenomenon. Decision-makers can also infer costs incurred by other regions' from past game outcomes. Further research is needed to refine the existing model by relaxing certain assumptions, such as assuming parameter homogeneity across regions. Additionally, it is crucial to incorporate the dynamics of previously implemented policies that are still in effect, as these may impact the estimated efficacy of new policies and contribute to more accurate simulations of health outcomes.

During the discussion following the speaker's presentation, various points related to the aforementioned modeling framework were raised. The assumption of a fixed transition probability between regions was discussed, which led to the introduction of a paper that explores the usefulness of cellular phone data for mobility estimation in the context of pandemic prediction and prevention. Workshop participants also discussed characteristics that could enhance the proposed model. These include differentiating the types of interventions being modeled, such as pharmaceutical and non-pharmaceutical policies, as well as considering the individual's willingness to comply with the policy, which can provide valuable insights on the optimal use of more or less strict policies. The discussion then shifted towards exploring the relationship between the strictness of a simulated intervention policy and the level of compliance it generates, as well as the trade-off between the level of compliance and the effectiveness of the intervention. Participants acknowledged that while stricter policies can lead to effective outcomes within shorter time durations, such policies may experience compliance challenges in the long term.

Microscopic Epidemic Models in Agent-Based Modeling

In agent-based modeling, building a microscopic pandemic model allows for simulations to be performed at the most granular (i.e., individual) level. This fine-grained control brings about multiple advantages, including the ability to simulate and learn human behavior, as well as the ability to evaluate government policies for decision-makers.

When building a microscopic model, it is important to incorporate a human behavior model that captures the decision-making process of individuals. The session speaker highlighted a model that takes into account two primary objectives of individuals during the COVID-19 pandemic: minimizing the chances of getting infected and maintaining a certain level of social activity. However, these objectives conflict with each other, as limiting activity minimizes the chance of infection, but maintaining activity levels requires higher social interaction and thus increases the chance of infection. Therefore, a reasonable general objective would be a weighted sum of the two. Additionally, when modeling the decision-making process of infected individuals, the objective should also prioritize limiting disease transmission; otherwise, the rational decision for infected individuals would be to maintain full social activity, given that they cannot (at present) be infected again. Under these objectives, a multi-agent reinforcement learning algorithm may be used to learn the human behavior model.

Modeling social activities between humans was also discussed by the session speaker and the participants, necessitating the consideration of different facilities, such as households, workplaces, hospitals, and schools. In the presented microscopic model, each facility has a unique level of social interaction and a corresponding probabilistic infection sub-model. For instance, the infection probability in hospitals is much higher than in households due to the higher number of contagious and infected patients. The participants also discussed misidentification of COVID-19 status, whereby individuals interact with others while unknowingly infected.

To fully develop a microscopic model, the speaker noted that the available actions humans can take should be taken into consideration. At the most basic level, each agent needs to decide their activity level for each day, which is crucial for determining the probability of infection. The session speaker and the participants also discussed extensions of this basic action, such as how people can choose to wear masks and how people need to decide when to shop for everyday necessities. Furthermore, the participants discussed pro-public-health actions individuals can take, such as proactively screening themselves for infection and notifying others about one’s infection promptly.

The participants then discussed how the presented microscopic model could serve as a testbed for evaluating governmental policies. The session speaker provided two examples of governmental policies that have thus far been evaluated: (1) information disclosure efforts (i.e., pandemic-related information disclosed by the government to the public) and (2) quarantine strategies. Simulated results on both strategies indicate their effectiveness through different stages of the COVID-19 pandemic and shed light on the policy-making process.

Despite the benefits of microscopic models, they also face notable challenges. Among others, a notable challenge is the lack of a principled approach to designing feedback systems (i.e., rewards) that guide the training of the underlying human behavior model. One way to design the reward function is to combine the objectives described above with parameterized scaling factors, as discussed by the session speaker. The parameters are then estimated by a calibration process that attempts to fit simulated data to data collected in the real world. Workshop participants discussed that this calibration process might be susceptible to overfitting such that the fitted parameters may not be accurate outside the training data. Participants also discussed the possibility that, in real life, individuals may not necessarily optimize for the objectives that are defined by the model and that humans can have multiple objective functions that are challenging to summarize into a single scalar reward.

Building off this challenge, participants discussed paths forward toward a more principled approach to designing rewards that incentivize and guide agents' behaviors. For example, workshop participants explored the possibility of having agents determine internalized rewards for themselves while being guided only by broader, higher-level objectives. The emergence of altruism, which was observed in the real world during the COVID-19 pandemic [6], would also be interesting to observe in the model.

In addition, participants discussed modeling the human adaptation process, as human responses to pandemics differ during different stages due to increased knowledge of the situation. While the current reinforcement learning algorithm attempts to learn a human behavior model, the way that it learns is vastly different from how humans learn. Therefore, modeling how humans gradually reach an equilibrium state instead of solving for a converged equilibrium using machine learning algorithms would be an interesting direction to explore.

Finally, workshop participants discussed several future directions for microscopic models in agent-based modeling. One possibility is to design models that are more interpretable to policymakers and the public. Specifically, participants discussed how, in certain situations, a model with the highest accuracy may not be desirable if it would lead to a significantly more complex and less interpretable model. Therefore, to make the best use of simulation results, policymakers may prefer models that are more understandable, even if they are less accurate in some instances.

References

[1] Z. Tang, K. Yan, L. Sun, W. Zhan, and C. Liu, "A Microscopic Pandemic Simulator for Pandemic Prediction Using Scalable Million-Agent Reinforcement Learning." arXiv, Aug. 14, 2021. doi: 10.48550/arXiv.2108.06589.

[2] R. Anirudh, J. J. Thiagarajan, P.-T. Bremer, T. Germann, S. D. Valle, and F. Streitz, "Accurate Calibration of Agent-based Epidemiological Models with Neural Network Surrogates," in *Proceedings of the 1st Workshop on Healthcare AI and COVID-19, ICML 2022*, PMLR, Jul. 2022, pp. 54–62. Accessed: Dec. 12, 2022. [Online]. Available: https://proceedings.mlr.press/v184/anirudh22a.html

[3] B. Wilder *et al.*, "Modeling between-population variation in COVID-19 dynamics in Hubei, Lombardy, and New York City," *Proc. Natl. Acad. Sci.*, vol. 117, no. 41, pp. 25904–25910, Oct. 2020, doi: 10.1073/pnas.2010651117.

[4] J. A. Killian, B. Wilder, A. Sharma, V. Choudhary, B. Dilkina, M. Tambe. "Learning to Prescribe Interventions for Tuberculosis Patients using Digital Adherence Data". *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2019. doi: 10.1145/3292500.3330777

[5] A. Mate, L. Madaan, A. Taneja, N. Madhiwalla, S. Verma, G. Singh, A. Hegde, P. Varakantham and M. Tambe. "Field Study in Deploying Restless Multi-Armed Bandits: Assisting Non-Profits in Improving Maternal and Child Health" in *Proceedings of the AAAI Conference on Artificial Intelligence.* 2022

[6] K. K. Cheng, T. H. Lam, and C. C. Leung, "Wearing face masks in the community during the COVID-19 pandemic: altruism and solidarity," The Lancet, vol. 399, no. 10336, pp. e39–e40, Apr. 2022, doi: 10.1016/S0140-6736(20)30918-1.

[7] A. Mate, A. Taneja, K. Wang, M. Tambe, and S. Verma. "Case Study: Applying Decision Focused Learning in the Real World" in *Proceedings of the NeurIPS Workshop on Trustworthy and Socially Responsible Machine Learning*. 2022

[8] J. A. Killian, A. Lalan, A. Mate, M. Jain, A. Taneja, and M. Tambe."Adherence Bandits" in *Proceedings of the AI for Social Good (AI4SG) Workshop at the AAAI Conference on Artificial Intelligence*. 2023

[9] S. Verma, G. Singh, A. Mate, N. Madhiwalla, A. Hegde, D. Thakkar, M. Jain, M. Tambe, and A. Taneja. "SAHELI for Mobile Health Programs in Maternal and Child Care: FurtherAnalysis" in *Proceedings of the AI for Social Good (AI4SG) Workshop at the AAAI Conference on Artificial Intelligence*. 2023