Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

ESTIMATING TREATMENT EFFECTS FROM SIMULATION SAMPLES OF POPULATION-SCALE MODELS

Abdulrahman A. Ahmed

Department of Industrial Engineering
University of Pittsburgh
3700 O'Hara Street
Pittsburgh, PA 15261, USA

ABSTRACT

Large-scale models require an exhaustive amount of computational power to simulate, especially when there are multiple treatment conditions to be evaluated across large geographical regions. Therefore, developing an efficient method to distribute computational resources efficiently is essential for conducting large-scale simulations. Agent-based modeling can generate accurate simulation samples, and our goal is to use them for estimating treatment effects to optimize potential interventions with as few simulation samples as possible. In this abstract, I will show methods that perform better than benchmarks by taking into account the uncertainty in the estimation of treatment effects dynamically and discuss our next steps for improving them.

1 INTRODUCTION

Estimating treatment effects is challenging, especially for large-scale, agent-based simulation models which require exhaustive computational power and time. In particular, in this study, we used simulation samples from a software named FRED. FRED is an agent-based simulation software that is developed by the Public Health Dynamics Laboratory (PHDL) in the University of Pittsburgh School of Public Health (based on U.S. census data) to simulate spatial and temporal behaviors of epidemics (Guclu et al. 2016). PHDL has developed a model in FRED software to understand the dynamics of the Opioid Use Disorder (OUD) epidemic, calibrated with real data in US counties, and to evaluate various interventions such as allocation of Medications for Opioid Use Disorder (MOUD) and overdose reversal medication (e.g., Naloxone).

2 CONTRIBUTIONS ACCOMPLISHED

The first method that one can think of for allocating simulation samples to estimate treatment effects is to conduct an equal number of simulations in each condition until we get a reasonable accuracy (i.e., confidence interval width). This is called the brute-force method, however, due to the computational requirements it could become infeasible as the number of treatment conditions grows. A greedy allocation can improve the brute-force method by allocating samples to the conditions with the widest confidence intervals sequentially until all estimates for the treatment effects are below a predefined width. However, in the greedy method, every additional sample only improves a single treatment condition. Alternatively, we model treatment conditions such by sampling each treatment condition we can learn and update a model that informs all treatment conditions. We experimented with two simple regression methods (called model-based greedy with and without interaction), resulting in a better performance than the brute-force and greedy methods. Table 1 shows the results for the first ten treatment effects in our OUD model. The results show enhancement in the sample size required to estimate the treatment effects compared to the greedy method, where achieving the same confidence levels across all treatment conditions requires 23,350 simulations (Ahmed, Rahimian, and Roberts 2023a).

	Model-based greedy			Model-based without interaction		
TC	Mean	CI width	# runs	Mean	CI width	# runs
Aa	2390.12	3.96	1200	2388.44	4.5	650
Ab	2375.02	2.82	900	2374.1	3.6	500
Ac	2359.9	2.4	900	2359.52	3.2	500
Ad	2344.82	3.01	900	2344.95	3.4	500
Ae	2329.72	4.22	1550	2330.37	4.1	500
Ba	2384.63	2.58	900	2383.25	3.3	500
Bb	2369.89	1.83	900	2368.68	2.4	500
Bc	2354.78	1.53	900	2354.1	2.1	500
Bd	2339.67	1.9	900	2339.53	2.7	1000
Be	2323.71	2.3	900	2324.95	2.7	500
		Total=	9950		Total=	5650

Table 1: Model-based greedy mean results (using a linear regression model with and without interaction terms) for the first ten treatment effects and corresponding CI widths where the third column shows the number of runs required to reach that estimate. This is compared to the greedy method, which requires 23,350 simulation runs to get the same level of confidence for each treatment effect.

3 REMAINING WORK

The ultimate goal of our research is to develop a method that scales quickly with a large number of treatment conditions. We only experimented with two simple models to improve sample size requirements for estimating treatment effects. We plan first to experiment with other parameters of the two simple models. For example, in the regression model, we plan to add quadratic terms to the model to represent the dose-response curve better. Furthermore, we will explore more complicated models such as the Gaussian Process (GP), which shows potential in understanding the dynamics of different epidemics and better resolution of the bias-variance trade-offs in our model selection (Ahmed, Rahimian, and Roberts 2023b). We also plan to model treatment conditions across different counties which could further improve sample size requirements when simulating the entire US population. This is especially important when interventions in a region can affect neighboring regions. This structure makes the problem even more complex, and it is an area where spatial regression using Gaussian processes can be especially powerful as a scalable model to estimate treatment effects across the entire US population.

4 SUMMARY

Estimating treatment effects for large-scale models is time-consuming and computationally exhaustive. It requires the development of models that distribute the computational resources economically as necessary for treatment effect estimation. In this abstract, I showed that we can estimate treatment effects with fewer samples using two model-based greedy methods, compared to the simple greedy or brute-force allocation baselines. We plan to experiment with other parameters and explore different models to achieve the best bias-variance trade-off for scalable estimations of treatment effects for the entire US population.

REFERENCES

Ahmed, A. A., M. A. Rahimian, and M. S. Roberts. 2023a. "Estimating Treatment Effects Using Costly Simulation Samples from a Population-Scale Model of Opioid Use Disorder". In *Proceedings of IEEE International Conference on Biomedical and Health Informatics*. In Press.

Ahmed, A. A., M. A. Rahimian, and M. S. Roberts. 2023b. "Inferring Epidemic Dynamics Using Gaussian Process Emulation of Agent-based Simulations". *arXiv preprint arXiv:2307.12186*.

Guclu, H., S. Kumar, D. Galloway, M. Krauland, R. Sood, A. Bocour, T. B. Hershey, E. van Nostrand, and M. Potter. 2016. "An Agent-Based Model for Addressing the Impact of a Disaster on Access to Primary Care Services". *Disaster Medicine and Public Health Preparedness* 10(3):386–393.