

# Using Agent-Based Simulation to Investigate Behavioral Interventions in a Pandemic

Simulating Behavioral Interventions in a Pandemic

Jan de Mooij<sup>1</sup>, Davide Dell'Anna<sup>2</sup>, Parantapa Bhattacharya<sup>3</sup>, Mehdi Dastani<sup>1</sup>, Brian Logan<sup>1,4</sup> and Samarth Swarup<sup>3</sup>

<sup>1</sup>Intelligent Systems, Information and Computing Sciences, Utrecht University

<sup>2</sup>Control and Operations, Delft University of Technology

<sup>3</sup>Biocomplexity Institute and Initiative, University of Virginia

<sup>4</sup>University of Aberdeen

## Abstract

Simulation is a useful tool for evaluating behavioral interventions when the adoption rate among a population is uncertain. Individual agent models are often prohibitively expensive, but, unlike stochastic models, allow studying compliance heterogeneity. In this paper we demonstrate the feasibility of evaluating behavioral intervention policies using large-scale data-driven agent-based simulations. We explain how the simulation is calibrated with respect to real-world data, and demonstrate the utility of our approach by studying the effectiveness of interventions used in Virginia in early 2020 through counterfactual simulations.

## Keywords

Agent-based Computational Epidemiology, Agent-based Modeling, Policy Evaluation, Normative Reasoning, Complex Social Simulation, Belief-Desire-Intention, Synthetic Population, Multi-agent Simulation

## Introduction

In the absence of pharmaceutical interventions, behavioral interventions are often the first line of defense early in an epidemic. However, the effect of behavioral interventions can be unpredictable, especially when the adoption rate of the population is unknown. While stochastic models are often used to evaluate the costs and benefits of such interventions [1, 2] due to their low computational complexity, agent-based models allow the effects of compliance heterogeneity to be studied [3].

In this work, we demonstrate the feasibility of using large-scale data-driven agent-based simulation for evaluating the effectiveness of behavioral interventions. In addition, we address the concern of generalization of a population to the ‘*standard human*’ [4] by instantiating the simulated agents from a detailed and representative synthetic population [5]. Each agent in the population is associated with an *activity schedule* that is representative of their socio-demographic

---

AMPM'21: First Workshop in Agent-based Modeling & Policy-Making, December 8, 2021, Vilnius, Lithuania

✉ A.J.deMooij@uu.nl (J. de Mooij); d.dellanna@tudelft.nl (D. Dell'Anna); parantapa@virginia.edu (P. Bhattacharya); m.m.dastani@uu.nl (M. Dastani); b.s.logan@uu.nl (B. Logan); swarup@virginia.edu (S. Swarup)

🆔 0000-0003-4129-6074 (J. de Mooij); 0000-0002-1162-8341 (D. Dell'Anna); 0000-0002-3626-9939 (P. Bhattacharya); 0000-0003-0648-7107 (B. Logan); 0000-0003-3615-1663 (S. Swarup)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

background, and which is used to drive their default behavior. By calibrating mobility patterns and disease progression, the framework may be applied to other regions where interventions or compliance rates may differ, provided the required data is available.

## 1. A Data-Driven Agent-Based Simulation

Between March and June 2020, nine subsets of the interventions shown in Table 1 were implemented in the form of Executive Orders (EOs) issued by the Governor of Virginia to mitigate the spread of COVID-19. The interventions are parameterized by the individuals affected and circumstances under which they apply. The timeline of these interventions is shown in Figure 1.

**Table 1**


A brief explanation of the behavioral interventions enforced in our simulation and their parameters.

Id	Interpretation	Parameters
$n_1$	Mask wearing is allowed and encouraged	-
$n_2$	Businesses of type $type$ are closed	$type \in \{NEB\}$ : the type of business, $NEB = \text{Non Essential Business}$
$n_3$	Employees working in retail must wear a mask during work activities	-
$n_4$	Telework is encouraged	-
$n_5$	Physical distance of 1.5m should be maintained	-
$n_6$	Capacity of business should be reduced to $perc$	$perc$ : percentage of business capacity
$n_7$	Schools of type $type$ are closed	$type \in \{K12, HE, K12 \text{ or } HE\}$ : the type of school, $K12 = \text{primary and secondary education}$ $HE = \text{Higher Education (HE)}$
$n_8$	The maximum allowed size of groups of type $type$ is $size$	$type \in \{public, private, all\}$ : the target settings, either public, private or both ( $all$ ); $size \in \mathbb{N}$ : maximum size of groups
$n_9$	Stay at home if belong to category $appl$	$appl \in \{sick \text{ or } age \geq 65, all\}$ : the group of agents to which the norm applies, either people sick or older than 65 ( $sick \text{ or } age \geq 65$ ), or everyone ( $all$ )
$n_{10}$	Only take away allowed for restaurants	-
$n_{11}$	A mask must be worn in public indoor settings	-

In previous work, we developed a large-scale data-driven normative agent-based simulation to study the effects of compliance with these interventions [6, 7]<sup>1</sup>. The simulation was developed using an extension of the agent programming language 2APL [8, 9] called *Sim-2APL*<sup>2</sup>, and consisted of 86 thousand agents drawn from a detailed synthetic population [5]. Each agent adopts goals to perform actions from a representative activity schedule. When no norms are active,

<sup>1</sup>The version of the framework described in this paper is archived at <https://github.com/A-Practical-Agent-Programming-Language/Normative-COVID-19-Simulation/releases/tag/v1.0.0>.

<sup>2</sup>Archived at <https://github.com/A-Practical-Agent-Programming-Language/Sim-2APL/releases/tag/v1.0.0>



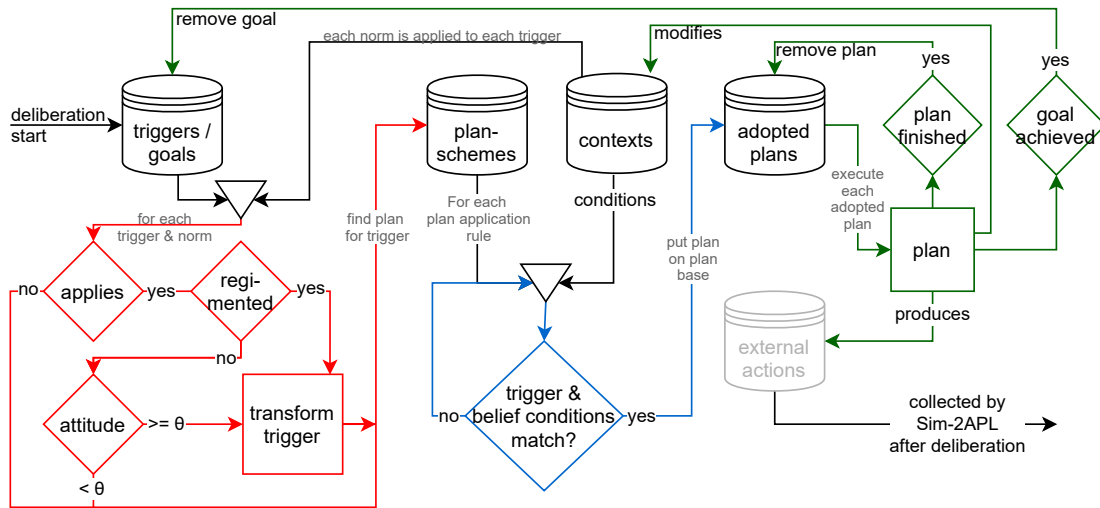
**Figure 1:** Timeline of the 9 major Executive Orders (EOs) implemented in Virginia, USA, between March and June, 2020. The column on the left indicates the norms belonging to the 9 EOs. For each norm in an EO (e.g.,  $n_1$  in EO<sub>1</sub>) a black line indicates the period during which the norm was enforced (e.g., the first two black lines indicate that both  $n_1$  and  $n_2$  from EO<sub>1</sub> were enforced from March 12 to July 1).

these are exactly the activities they perform. However, throughout the simulation, we activate or deactivate norms according to the timeline from Figure 1. When a scheduled activity is prohibited by an intervention modeled as a *norm*, the agents use previous observations and intrinsic values to determine whether to comply with the norm by altering or canceling the activity, or to violate the norm by performing the activity as scheduled. Two important factors used in this reasoning are the agent’s trust in institutions, and the behavior of other agents. In particular, agents tend towards the compliance behavior of other agents they have encountered in the past. However, the magnitude of this tendency is determined by the trust the agent holds in institutions, which is sampled from a Beta distribution at the start of the simulation. The normative deliberation process followed by the Sim-2APL agents each simulated day is shown in Figure 2. Agent deliberation essentially follows the Belief-Desire-Intention (BDI) approach: each agent pursues a set of *goals* (in the current work, these are the activities from their activity schedules), and responds to a set of *triggers*, such as messages, or events from the environment. At each stage in the reasoning process, each agent also has access to a *context* that encodes its beliefs and other persistent data. *Plan Schemes* map triggers to plans. Each *plan* encodes a particular form of agent behavior as a sequence of zero or more actions. When the reasoning process of all agents is complete, the actions generated by the plans selected by each agent are processed by the environment to move the global simulation state to the next time step.

A contagion simulator called PanSim simulates the progression of COVID-19 based on agents’ collocation [7]. PanSim also allows the distribution and synchronization of a simulation across multiple compute nodes, which allows scaling up beyond the initial 86 thousand agents.

## 2. Calibration

To ground the model in real-world data, both the behavior model and the disease model were calibrated. For the behavior model, the parameters that were calibrated were the mean of the Beta distribution from which the agents’ trust is sampled, a discount factor termed *fatigue* and the *start time* at which fatigue comes into effect. Fatigue reduces the trust of all agents linearly starting from the *start time*, and models the gradual increase in mobility after the first intervention which is not associated with the deactivation of norms. For the disease model, the parameters that were



**Figure 2:** The normative deliberation process of a Sim-2APL agent

calibrated were the probability of symptomatic and asymptomatic infected agents infecting their susceptible contacts.


The behavior interventions had—by design—significant impact on mobility patterns. This was used to calibrate the behavior model using the *mobility index*: the percentage change in Cuebiq mobility data for each county compared to pre-pandemic levels. The mobility index of each of the four simulated counties in our calibrated simulation is plotted against those reported by Cuebiq in Figure 3a.

The disease model was calibrated using the cumulative number of reported cases. To account for under-testing at the start of the pandemic, we multiplied the number of reported cases by 30. While this scale factor is arbitrary, it can be changed without affecting the methodology of our work. The simulated case counts are plotted against the reported case count in Figure 3b.

For both calibration processes, Nelder-Mead [10] was used to minimize the Root Mean Square Error (RMSE) between the real-world data and the simulated data across 5 simulations for each proposed configuration of parameters.<sup>3</sup> Each simulation spans the time period 1 March to 29 June, and in each simulation, the activation and deactivation of norms follows the timeline from Figure 1. For the behavior model, the RMSE was calculated as the error between the real-world and simulated mobility index for each county individually. The final obtained RMSE was 17.65. For the disease model, where numbers are significantly lower, the cumulative value across all simulated counties was used. The final obtained RMSE was 2052 on a population of 86150 agents.

As can be seen in Figure 3b, the disease model was not able to accurately reproduce either our estimate of the total number of infections or the time at which infections start to increase significantly. Other approaches to estimating the real number of cases, such as the number of

<sup>3</sup>Initial simulations indicated that the standard deviation is sufficiently small for a reasonable level of confidence; the purpose of using more than one simulation is to avoid a statistical outlier as a global minimum




**Figure 3:** The mobility index observed in the simulation plotted against that recorded by Cuebiq in each simulated county (3a), and percentage of recovered agents in the simulation plotted against the percentage of reported cases ( $\times 30$ ) in the same counties (3b).

hospitalizations or deaths, may give better results. However, while the number of hospitalizations or deaths is generally assumed to be more accurate, like the number of positive tests, they are only a proxy for actual disease spread, and the effect of disease spread on these metrics needs to be known or calibrated as well.

### 3. Experiment

In this section, we briefly show how our simulation framework can be used to help policy makers evaluate interventions through counterfactual simulations. For the interventions implemented in Virginia, we ran 10 experiments  $E_0 \dots, E_9$ . In  $E_0$ , we disabled all interventions, in  $E_1$  we included only the interventions from  $EO_0$  ( $n_1$  and  $n_4$ ); in  $E_2$  we included the interventions from the first two EOs, in  $E_3$  the first three EOs, and so on up to  $E_9$  in which we simulate all interventions reflecting the actual policy implemented in Virginia. Note that, in each experiment, all activated EOs still follow the timeline from Figure 1. For each experiment we ran 10 simulations to account for stochastic variation. Figure 4 shows the average progression of the disease in each of the 10 experiments plotted against the reported (and scaled) number of cases. Compared to no interventions at all ( $E_0$ ), the policy implemented in Virginia ( $E_9$ ) can be considered effective in the sense that it reduced the number of infections in the first four months from  $\sim 50\%$  of the population to only  $\sim 7\%$ . Moreover, the results show that each additional set of interventions both reduced the number of infections and delayed the exponential spread compared to the previous set of interventions. The largest gain appears to be made by going from  $E_3$  (in yellow) to  $E_4$  (max group size of 10, stay home when sick, reduce business capacity to 10), and from  $E_4$  to  $E_5$  (Closure of all non-essential businesses, *everyone* asked to stay home as much as possible, restaurants closed except for take-away).



**Figure 4:** Cumulative simulated (blue lines) and reported ( $\times 30$ , orange lines) in  $E_0 - E_9$

## 4. Conclusion

We showed the feasibility of using data-driven multi-agent simulation to evaluate the effectiveness of behavioral intervention policies. We explained how both the behavior and disease model components can be calibrated with real-world data, and demonstrated the utility of such a model through a simple experiment in which we quantify the effect of the 9 different sets of behavioral interventions that were implemented in Virginia in the early months of the COVID-19 pandemic through counterfactual simulations. Given the availability of suitable data, such as a synthetic population, the model can be calibrated and applied to different regions with different interventions or varying rates of compliance.

Our model is not able to correctly approximate the spread of disease, both temporally and in magnitude. In future work, we plan to both incorporate age-stratified infectivity and calibrate against better estimations of true infections, which may help improve the calibration and modeling of the spread of disease. This may also prove useful for studying the effects of interventions in regions with different age distributions.

We also plan to investigate the use of the *Stringency Index* [11], a measure to express the stringency of a set of measures as a single dimension to allow easy comparison of intervention policies between regions. Because this metric was designed for comparing regions, it does not capture the finer details of any specific region, as categories to which measures are assigned are relatively broad. For this reason, we have not used the stringency index to determine the impact of mitigation measures in this work. Instead, we are investigating how the stringency index can be used to define the cost of a policy, as part of a methodology to efficiently find effective mitigation policies that balance the effect on transmission with societal impact.

## Acknowledgments

We thank Cuebiq; mobility data is provided by Cuebiq, a location intelligence and measurement platform. Through its Data for Good program, Cuebiq provides access to aggregated mobility

data for academic research and humanitarian initiatives. This first-party data is collected from anonymized users who have opted-in to provide access to their location data anonymously, through a GDPR and CCPA compliant framework. To further preserve privacy, portions of the data are aggregated to the US Census block group level.

PB and SS were supported in part by NSF Expeditions in Computing Grant CCF-1918656.

## References

- [1] G. R. Shinde, A. B. Kalamkar, P. N. Mahalle, N. Dey, J. Chaki, A. E. Hassanien, Forecasting models for coronavirus disease (COVID-19): a survey of the state-of-the-art, *SN Computer Science* 1 (2020) 1–15.
- [2] F. Lorig, E. Johansson, P. Davidsson, Agent-based social simulation of the COVID-19 pandemic: A systematic review, *JASSS: Journal of Artificial Societies and Social Simulation* 24 (2021).
- [3] F. Squazzoni, J. G. Polhill, B. Edmonds, P. Ahrweiler, P. Antosz, G. Scholz, É. Chappin, M. Borit, H. Verhagen, F. Giardini, et al., Computational models that matter during a global pandemic outbreak: A call to action (2020).
- [4] S. Milan, Techno-solutionism and the standard human in the making of the COVID-19 pandemic, *Big Data & Society* 7 (2020) 2053951720966781. URL: <https://doi.org/10.1177/2053951720966781>. doi:10.1177/2053951720966781. arXiv:<https://doi.org/10.1177/2053951720966781>.
- [5] A. Adiga, A. Agashe, S. Arifuzzaman, C. L. Barrett, R. J. Beckman, K. R. Bisset, J. Chen, Y. Chungbaek, S. G. Eubank, S. Gupta, M. Khan, C. J. Kuhlman, E. Lofgren, B. L. Lewis, A. Marathe, M. V. Marathe, H. S. Mortveit, E. Nordberg, C. Rivers, P. Stretz, S. Swarup, A. Wilson, D. Xie, Generating a Synthetic Population of the United States, Technical Report NDSSL 15-009, Network Dynamics and Simulation Science Laboratory, 2015.
- [6] J. de Mooij, D. Dell’Anna, P. Bhattacharya, M. Dastani, B. Logan, S. Swarup, Quantifying the effects of norms on COVID-19 cases using an agent-based simulation, in: *Proceedings of the The 22nd International Workshop on Multi-Agent-Based Simulation (MABS)*, 2021.
- [7] P. Bhattacharya, A. J. de Mooij, D. Dell’Anna, M. Dastani, B. Logan, S. Swarup, PanSim+ Sim-2APL: A framework for large-scale distributed simulation with complex agents, in: *International Workshop on Engineering Multi-Agent Systems*, 2021.
- [8] M. Dastani, 2APL: a practical agent programming language, *Autonomous agents and multi-agent systems* 16 (2008) 214–248.
- [9] M. Dastani, B. Testerink, From multi-agent programming to object oriented design patterns, in: *International Workshop on Engineering Multi-Agent Systems*, Springer, 2014, pp. 204–226.
- [10] J. A. Nelder, R. Mead, A simplex method for function minimization, *The Computer Journal* 7 (1965) 308–313. arXiv:<https://academic.oup.com/comjnl/article-pdf/7/4/308/1013182/7-4-308.pdf>.
- [11] T. Hale, N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, et al., A global panel database of pandemic

policies (Oxford COVID-19 government response tracker), *Nature Human Behaviour* 5 (2021) 529–538.