Towards Realism for Policy Testing

Maarten Jensen[1], Loïs Vanhée[1], and Frank Dignum[1]

Umeå University, Umeå, Sweden

Abstract. Essential in a social simulation model that has the purpose of *"understanding the effectiveness of a policy on society"*, is that the agents can not only follow the policy, but also have ways of breaking the policy. Especially modelling the motivations for breaking a policy requires additional aspects of life in the simulation. However, adding all these aspects can have a toll on the scalability of the model. This paper is dedicated to explaining why these aspects together are relevant to the realism of modelling for policy testing. The paper ends with an illustration that state-of-the-art pandemic simulations achieve some but not all requirements. And concludes as a way forward to consider both conceptual and engineering improvements.

Keywords: Policy testing, Realism, Scalability, Behavioural model

1 Introducing the requirements

Policies imposed by the government can have effect on many different aspects of a society. This has become especially clear during the COVID-19 pandemic, where policies aimed at preventing the spread of the disease also affected people in their daily lives. For example, children could not go to school, bars and restaurants were closed, and people were sometimes obliged to stay home. However, it was difficult to determine whether the policy would have the desired effect. How many people would deviate from the policy? And in what way would people deviate? To understand the effectiveness of a policy on society, as many relevant aspects as possible related to the phenomena should be taken into account. However, this completeness is not trivial and depends on a number of aspects.

Essential in a social simulation model that has the purpose of *"understanding the effectiveness of a policy on society"*, is that agents adapt their behaviour to a new policy. This can be both policy following or policy breaking; see various models of the pandemic [4, 5, 12]. For modelling the policy following aspect it is usually enough to model the phenomena itself and the normative aspects (the policy). In a COVID-19 model, one could think of a contagiousness system (the virus and locations where people can meet) and a policy such as social distancing (making people less infectious) or quarantine (people will stay at home).

However, for policy breaking, the side effects of the policy, other aspects of life also need to be taken into account. During the pandemic it became clear that the imposed policies were often not followed [14]. To name a few examples of breaking polices: breaking quarantine, not wearing mouth masks, not social distancing, not vaccinating. Behind each of these policy violations, there could be numerous possible motivations to break the policy. For example, in many countries people were expected to stay at home (social norm policy); however, people broke this policy for various reasons, e.g. to get some fresh air, to see friends, to buy food, to work to earn money, to do physical activity. To be as complete as possible, as many of these aspects as possible need to be taken into account. And these aspects should be taken into account by the behavioural model of the agents. Agents should be capable of breaking the policy based on an explicit motive.

Including many aspects of life and having a behavioural system that can consider all these aspects of life, can have a toll on the computational complexity of the model. This in turn decreases execution time and can lead to scalability issues, making the model too slow to efficiently work with or get output from in a reasonable time [15]. Thus, when modelling to understand the effectiveness of a policy on society, scalability becomes an important requirement as well.

Thus, for a model to be complete, it requires multiple aspects of life, an interdependent behavioural model, and scalability. Together, these requirements are what we call the realism of the model within the purpose of *"understanding the effectiveness of a policy on society"* (see Figure 1). The rest of the paper explains these three requirements in more detail, explains the state-of-the-art models that achieve these aspects to some extent, and gives future recommendations in the form of Context-Sensitive Deliberation and scalability techniques.

2 The requirements

2.1 R1: Multiple aspects of life

As indicated before, policies can be both followed and broken. Modelling the following of policies requires at least a normative model (to represent the policies). Modelling the breaking of policies is what requires more aspects of life, such as economical, shopping, social aspects [8]. Each of these aspects are not just represented as one parameter, but actually are a system or sub model. This is how interesting behaviour can appear. For example, an agent could be forced to stay at home based on the normative system as there is a global lockdown policy. Meanwhile, all the friends of the agent decide to break quarantine, and due to

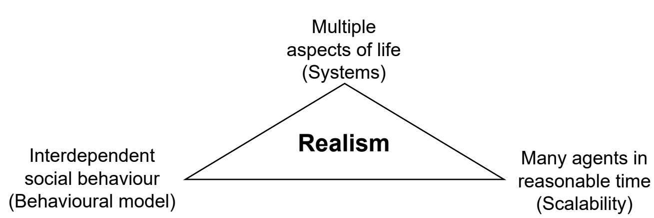

Fig. 1. Three requirements for realism for models for policy testing

this influence of the social system, the agent itself also decides to break quarantine. However, having all systems explicitly in a model requires a behavioural model that can take this into account, which is explained in the next section.

2.2 R2: The behavioural model

The behaviour of the agents should have interdependence among the different aspects of life. There are quite some pandemic simulations that do not model interdependency as they use, for example, a fixed probability to motivate changes in agents behaviour, e.g. in Hinch [5], for example breaking lockdown. These models are less realistic in modelling the side effects of long-term lockdown, as there is no difference between an agent being one day in lockdown and one who has been in lockdown for 30 days [11]. An additional parameter could be added that tracks the lockdown days over time and adjusts the probability. However, this will not work when the agents also have other norms to consider. For example, introducing quarantine when people have already been following other norms such as social distancing and wearing mouth mask, might make people fed up with following rules in general. To properly model this interdependency, which we also call *the ripple effect*, the behavioural system should contain interdependency among the aspects of life modelled in the simulation. An example of such an architecture is described in Dignum [3]. An example framework where agents can follow and violate norms is described in [9].

2.3 R3: Scalability

Including many aspects of life and having a behavioural system that can consider all these aspects of life, can have a toll on the computational complexity of the model and thus the scalability of the model. The ASSOCC simulation [4], which models the COVID-19 pandemic in Netlogo, had problems simulating much more than 1.000 agents in a reasonable time [15]. It would require about 3-4 minutes to simulate 1.000 agents; however, simulating 10.000 agents would easily take 2 hours since ASSOCC has not been optimised for these large agent numbers. This lengthy execution time would hinder the inspection of the effects of changes in code and parameters on the model and hinder the rapid model development. This is particularly relevant when modelling for crisis situations as results are required on a daily basis [4]. Due to this 1.000 agents limitation, it was not possible to expand the ASSOCC model further, as that would make the model more computationally complex and increase the execution time even further. Higher agent numbers are, however, required if additional subgroups needed to be made, for example when one wants to study the effects of specific hospital personal getting sick. Those subgroups would have too few agents to realistically represent them. With a population of about 1000 agents, there are about 20 agents working in the hospital. If these were divided into nurses and doctors, there would probably be only 2 to 4 doctors. This number would be too small and would be susceptible to having a situation where literally all doctors in the society are sick at the same time.

3 State-of-the-art pandemic models

Quite some advanced models have been made to understand the effects of policies in the COVID-19 pandemic [10]. PanSim [13, 12] is a state-of-the-art model that simulates the spread of COVID-19 in the state of Virginia. It successfully simulates millions of behaviourally complex agents and thus satisfies the scalability requirement (R3). It includes multiple aspects of life (R2), i.e. variety of activities for the agents to do (work, shop, study, religious and other), contagiousness and disease progression, and up to eleven norms that can be enforced. The agents can also reason about these norms and adjust their behaviour based on a BDI behavioural model [1]. If we consider the requirement of interdependent social behaviour (R2), the model seems less realistic, as the agents are, for example, not more likely to break quarantine if they have not seen their friends for a long time. Breaking a norm is modelled using a probability instead, which is dependent on trust in the government and the behaviour of agents around. But since it is a probability, it lacks transparency as to why the agent breaks quarantine. Despite of this, it is an advanced model due to its scalability capabilities.

The ASSOCC model [4] implemented in Netlogo incorporates many aspects of life (R1): working, studying, leisure, shopping, a contagiousness and disease model, economic model, social model, etc. [8]. The behavioural model is a need-based deliberation model that spans and connects all aspects of life [3] (R2). As discussed above, this combined complexity of the aspects of life and the behavioural model comes at the cost of scalability (R3). It cannot be extended anymore without adjusting the internal architecture or changing to a different modelling framework, e.g., Repast high performance [2]. The model is realistic for the intended scenarios, but not for studying other scenarios since it cannot be further extended without a change of architecture. A promising approach could be Context-Sensitive Deliberation [6, 7]. This approach conceptually changes the behavioural model to decrease computational complexity, and thus increase scalability, while potentially retaining realism.

4 Conclusion

This paper argued that to create a simulation model that is realistic with regard to the purpose of "understanding the effectiveness of a policy on a society", the following three requirements are necessary. R1: As many relevant aspects of daily life as possible should be included. R2: The model needs agents that can respond to policies in various ways, not only following policies but also breaking policies. R3: Despite the previous two requirements the simulation should be scalable, i.e. run many agents in reasonable time. As illustrated, these requirements are partially fulfilled by existing simulations; however, these simulations can still be made more realistic. As a possible solution to push forward towards more realistic simulation for policy testing, it is necessary to not solely consider conceptual improvements (e.g. Dynamic Context-Sensitive Deliberation [6, 7]) or solely engineering improvements (e.g. PanSim's parallelisation approach [1]) but to bring them together in a unified approach.

References

1. Bhattacharya, P., de Mooij, A.J., Dell'Anna, D., Dastani, M., Logan, B., Swarup, S.: Pansim+ sim-2apl: A framework for large-scale distributed simulation with complex agents. In: International Workshop on Engineering Multi-Agent Systems. pp. 1–21. Springer (2021)
2. Collier, N., North, M.: Parallel agent-based simulation with repast for high performance computing. Simulation 89(10), 1215–1235 (2013)
3. Dignum, F.: Foundations of social simulations for crisis situations. In: Social simulation for a crisis: Results and lessons from simulating the COVID-19 crisis, pp. 15–37. Springer (2021)
4. Dignum, F.: Social Simulation for a Crisis. Springer (2021)
5. Hinch, R., Probert, W., Nurtay, A., Kendall, M., Wymant, C., Hall, M., Lythgoe, K., Cruz, A.B., Zhao, L., Stewart, A., et al.: Effective configurations of a digital contact tracing app: a report to nhsx. Retrieved July 23, 2020 (2020)
6. Jensen, M., Vanhée, L., Dignum, F.: Dynamic context-sensitive deliberation. In: Multi-Agent-Based Simulation, MABS 2023 (2023)
7. Jensen, M., Vanhée, L., Dignum, F.: Dynamic context-sensitive deliberation for scalability in realistic social simulations. In: Advances in Social Simulation. Springer (2023)
8. Jensen, M., Vanhée, L., Kammler, C.: Social simulations for crises: From theories to implementation. In: Social Simulation for a Crisis. pp. 39–84. Springer (2021)
9. Kammler, C., Dignum, F., Wijermans, N.: Utilizing the full potential of norms for the agent's decision process. In: Conference of the European Social Simulation Association. pp. 193–205. Springer (2022)
10. Lorig, F., Johansson, E., Davidsson, P.: Agent-based social simulation of the covid-19 pandemic: A systematic review. JASSS: Journal of Artificial Societies and Social Simulation 24 (2021)
11. Lorig, F., Jensen, M., Kammler, C., Davidsson, P., Verhagen, H.: Comparative validation of simulation models for the covid-19 crisis. In: Social Simulation for a Crisis: Results and Lessons from Simulating the COVID-19 Crisis, pp. 331–352. Springer (2021)
12. de Mooij, J., Bhattacharya, P., Dell'Anna, D., Dastani, M., Logan, B., Swarup, S.: A framework for modeling human behavior in large-scale agent-based epidemic simulations. Simulation 99(12), 1183–1211 (2023)
13. de Mooij, J., Dell'Anna, D., Bhattacharya, P., Dastani, M., Logan, B., Swarup, S.: Quantifying the effects of norms on covid-19 cases using an agent-based simulation. In: International Workshop on Multi-Agent Systems and Agent-Based Simulation. pp. 99–112. Springer (2021)
14. Nivette, A., Ribeaud, D., Murray, A., Steinhoff, A., Bechtiger, L., Hepp, U., Shanahan, L., Eisner, M.: Non-compliance with covid-19-related public health measures among young adults in switzerland: Insights from a longitudinal cohort study. Social science & medicine 268, 113370 (2021)
15. Vanhée, L.: Engineering social simulations for crises. In: Social Simulation for a Crisis: Results and Lessons from Simulating the COVID-19 Crisis, pp. 353–378. Springer (2021)