Computational Social Sciences

Frank Dignum *Editor*

Social Simulation for a Crisis

Results and Lessons from Simulating the COVID-19 Crisis

Springer

Computational Social Sciences

Computational Social Sciences

A series of authored and edited monographs that utilize quantitative and computational methods to model, analyze and interpret large-scale social phenomena. Titles within the series contain methods and practices that test and develop theories of complex social processes through bottom-up modeling of social interactions. Of particular interest is the study of the co-evolution of modern communication technology and social behavior and norms, in connection with emerging issues such as trust, risk, security and privacy in novel socio-technical environments. Computational Social Sciences is explicitly transdisciplinary: quantitative methods from fields such as dynamical systems, artificial intelligence, network theory, agent-based modeling, and statistical mechanics are invoked and combined with state-of-the-art mining and analysis of large data sets to help us understand social agents, their interactions on and offline, and the effect of these interactions at the macro level. Topics include, but are not limited to social networks and media, dynamics of opinions, cultures and conflicts, socio-technical co-evolution and social psychology. Computational Social Sciences will also publish monographs and selected edited contributions from specialized conferences and workshops specifically aimed at communicating new findings to a large transdisciplinary audience. A fundamental goal of the series is to provide a single forum within which commonalities and differences in the workings of this field may be discerned, hence leading to deeper insight and understanding.

More information about this series at http://www.springer.com/series/11784

Frank Dignum
Editor

Social Simulation for a Crisis

Results and Lessons from Simulating the COVID-19 Crisis

Springer

Editor
Frank Dignum
Computing Sciences
Umeå University
Umeå, Sweden

ISSN 2509-9574 ISSN 2509-9582 (electronic)
Computational Social Sciences
ISBN 978-3-030-76396-1 ISBN 978-3-030-76397-8 (eBook)
https://doi.org/10.1007/978-3-030-76397-8

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

This book is dedicated to all the people that suffer from the Covid-19 crisis. Both the patients and people that passed away, but also the health care workers working endless shifts, the many SME's that saw their business become unviable, the people that had to work at home while home schooling their children, the youth that has been deprived for a very long time of social contacts in a period of their lives where this is of utmost importance, and all those other groups somehow affected by this crisis.

Preface

This book is the result from an effort of fourteen people that have joined together to work voluntarily (no funding!) on the Agent-based Social Simulation for the Covid-19 Crisis (ASSOCC) framework in order to make a positive difference in this crisis using social simulation. The primary goal of this project was to have real world impact and support decision makers during the crisis. However, doing this project has generated so much valuable experiences for the social simulation community at large and especially for using social simulations for crisis situations that we decided to write this book.

The book is, just like the ASSOCC project, an exceptional case. It seemed not to be possible to have a book with fourteen authors. So, in the end we opted for the solution of an edited book, where I, as initiator of this project, ended up as editor. The separate chapters do have different subsets as official authors, but the whole team has contributed in many ways to all the chapters. To emphasize that the book is actually a joint effort the whole team is co-author of the introduction and conclusion chapter of the book.

Just like the ASSOCC project results have been remarkable, so is this book. If one takes into consideration that half the team consists of (young) Ph.D. students, it is amazing how they have been able to accomplish so much in such a short time. This certainly would not have been possible without Loïs Vanhée who was the chief architect of the implementation and managed to keep a very big and diverse group of code contributors in line. We have all worked countless hours on this project, but he has been always there for anyone at any time of the day to support, encourage and help.

For me, as initiator of the ASSOCC project, it has also been a rewarding experience. All members have been very committed and supportive of each other. They were willing to put up with all my demands and directives. I have learned a lot from all of the team members. I feel guilty sometimes, because most of the media attention has come my way rather than the whole team. However, I have also learned that having regular contact with the media can really increase the impact of our work. With this book I hope that we can give a good foundation for having social simulations being accepted as a valuable and even necessary contribution for crisis management. Both before and during the crisis.

Umeå, Sweden
April 2021

Frank Dignum

Acknowledgements

The simulations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at Umeå partially funded by the Swedish Research Council through grant agreement no. 2018-05973. This research was conducted using the resources of High Performance Computing Center North (HPC2N).

The research presented in this book was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) and WASP—Humanities and Society (WASP-HS) funded by the Knut and Alice Wallenberg Foundation.

We also want to acknowledge here the contributions of Annet Onnes who worked on the transport model. Moreover we like to thank Virginia Dignum who set up and maintained the ASSOCC website, which was essential to distribute our results quick and to a large audience.

Contents

Contributors

Paul Davidsson Department of Computer Science and Media Technology, Internet of Things and People Research Center, Malmö University, Malmö, Sweden

Bart de Bruin Faculty of Technology, Policy and Management, TU Delft, Delft, The Netherlands

Frank Dignum Department of Computing Science, Umeå University, Umeå, Sweden

Amineh Ghorbani Faculty of Technology, Policy and Management, TU Delft, Delft, The Netherlands

Maarten Jensen Department of Computing Science, Umeå University, Umeå, Sweden

Christian Kammler Department of Computing Science, Umeå University, Umeå, Sweden

Kurt Kreulen Faculty of Technology, Policy and Management, TU Delft, Delft, The Netherlands

Fabian Lorig Department of Computer Science and Media Technology, Internet of Things and People Research Center, Malmö University, Malmö, Sweden

Alexander Melchior Department of Information and Computer Science, Utrecht University, The Netherlands Ministry of Economic Affairs and Climate Policy and Ministry of Agriculture, Nature and Food Quality, Utrecht, The Netherlands

René Mellema Department of Computing Science, Umeå University, Umeå, Sweden

Cezara Păstrăv Department of Computing Science, Umeå University, Umeå, Sweden

Mijke van den Hurk Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

Loïs Vanhée Umeå University, Umeå, Sweden

Harko Verhagen Department of Computer and Systems Sciences, Stockholm University, Kista, Sweden

Part I
ASSOCC Theory and Platform

In part I, we lay the foundations of the ASSOCC platform for the simulations that we have run with ASSOCC. We describe the theories that are used for the agent deliberation processes and we describe extensively how these theories are implemented in a practical and efficient way. We also show how we provide a proper user interface to the simulations that provide decision-makers with possibilities to follow the runs and also analyze them in several ways.

Chapter 1
Introduction

Frank Dignum, Loïs Vanhée, Maarten Jensen, Christian Kammler, René Mellema, Fabian Lorig, Cezara Păstrăv, Mijke van den Hurk, Alexander Melchior, Amineh Ghorbani, Bart de Bruin, Kurt Kreulen, Harko Verhagen, and Paul Davidsson

Abstract The introduction of this book sets the stage of performing social simulations in a crisis. The contents of the book are based on the experience of creating a large scale and complex social simulation for the COVID-19 crisis. However, the contents are reaching much further than just this experience. We will show the general contribution that social simulations based on fundamental social-psychological principles can have in times of crises. In times of big societal changes due to a pandemic or other disaster, these simulations can give handles to support decision makers in their difficult task to act in a very short time with many uncertainties. Besides giving our results, we also will indicate why the results are trustworthy and interesting. Finally we also look what challenges should be picked up to convert the successful project into a sustainable research area.

F. Dignum (✉) · L. Vanhée · M. Jensen · C. Kammler · R. Mellema · C. Păstrăv
Department of Computing Science, Umeå University, SE-901 87 Umeå, Sweden
e-mail: dignum@cs.umu.se

L. Vanhée
e-mail: lois.vanhee@umu.se

M. Jensen
e-mail: maartenj@cs.umu.se

C. Kammler
e-mail: ckammler@cs.umu.se

R. Mellema
e-mail: renem@cs.umu.se

A. Ghorbani · B. de Bruin · K. Kreulen
Faculty of Technology, Policy and Management, TU Delft, Jaffalaan 5,
2628 BX Delft, The Netherlands
e-mail: a.ghorbani@tudelft.nl

B. de Bruin
e-mail: bdb785@gmail.com

K. Kreulen
e-mail: kurtkreulen@gmail.com

F. Dignum (ed.), *Social Simulation for a Crisis*, Computational Social Sciences,
https://doi.org/10.1007/978-3-030-76397-8_1

1.1 Crisis

In March 2020, the gravity of the pandemic caused by the corona virus slowly became apparent. While most people (including us) thought that the consequences would be limited to Wuhan it became clear that the virus had already spread throughout Europe as well. The reason why COVID-19 could wreak such a havoc is not because it is very virulent and kills its host in a short time. The reason why it can spread so easily and is so persistent is exactly that not everyone is affected at the same level and that it might take quite some days before symptoms become clear, if at all. This means that people can carry the virus and spread it without being aware of their infection for a considerable amount of time. The COVID-19 virus is, thus, placed between Ebola in one side and flu at the other side. Ebola is very virulent and because of that isolation of people being infected can be done quite effective. Hence, although the virus kills most people that it infects, it can usually be contained pretty well. The flu viruses are usually not well contained because they often have an incubation time of several days and symptoms only appear after some time as well. Thus, the virus can be spread quite easily during the incubation time and by patients that have relatively mild symptoms and keep going to work and other places where they meet other people. However, the flu viruses are not very lethal and, thus, the disruption of society is relatively small.

The characteristics of the COVID-19 virus made it difficult to contain. The standard procedure of the health care authorities in cases of a pandemic outbreak is to try to track and trace all contacts of an infected person and isolate these persons as quickly

F. Lorig · P. Davidsson
Internet of Things and People Research Center, Department of Computer Science and Media Technology, Malmö University, 205 06 Malmö, Sweden
e-mail: fabian.lorig@mau.se

P. Davidsson
e-mail: paul.davidsson@mau.se

M. van den Hurk · A. Melchior
Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands
e-mail: m.vandenhurk@uu.nl

A. Melchior
e-mail: a.t.melchior@uu.nl

A. Melchior
Ministry of Economic Affairs and Climate Policy and Ministry of Agriculture, Nature and Food Quality, The Netherlands, Bezuidenhoutseweg 73, 2594 AC Den Haag, The Netherlands

L. Vanhée
GREYC, Université de Caen, 14000 Caen, France

H. Verhagen
Department of Computer and Systems Sciences, Stockholm University, PO Box 7003, 16407 Kista, Sweden
e-mail: verhagen@dsv.su.se

as possible. However, someone could have infected other persons in the previous six days while having been to a pub at Friday night, having gone shopping in a shopping mall on Saturday and visited a soccer match with 40.000 other people it becomes very difficult to trace all possible contacts that might have been infected. Although the traditional track and tracing is still valuable it is not enough in this situation and other measures are needed.

The interventions and measures taken have differed widely between countries all over the world. There have been debates about what are the "best" measures and countries have been blamed and praised (and sometimes both at different times) for the measures they took. Unfortunately, this book will not give an answer on which is the best measure to take. Basically, because we do not believe there is one best measure. The measures that can and should be taken depend on the country/region, its infrastructure, its culture and many other aspects of its society.

However, the realisation that the spread of the COVID-19 virus and also the success of measures in a country or region depends crucially on human behaviour led us to the conviction that social simulations could have a huge added value in this type of pandemic. Thus, on March 16, 2020 Frank Dignum wrote e-mails to Ph.D. students and other colleagues to see who would want to collaborate to build a simulation for the COVID-19 crisis. So, that is how the Agent-based Social Simulation of the Coronavirus Crisis (ASSOCC) project started. In all respects, it is an extra-ordinary project.

First of all, it is not funded! All members of the team participate on voluntary basis and do a lot of work in their spare time. Fortunately, enough work could be combined with "normal" research work in order to keep the project moving. However, we have all the time been carefully balancing between being enthusiastic and spending many extra hours on the project and preventing people from burn-outs due to an unrelenting schedule driven by the events during the COVID-19 crisis.

Secondly, the project did not have a project plan, not even a start and end date or predefined milestones. However, we agreed on the cognitive models that would serve as foundation for the simulations and we knew that we wanted to make not just one simulation on one aspect but rather a sandbox in which many scenarios could be developed and run. We also knew that interfacing for non-specialists would be important and, thus, we set up a separate module to provide an adequate interface to the simulation and its results. The deadlines for the project were set by the events of the crisis. After a quick set up of the basic components of the system we wanted to be ready and show results in time to inform the national discussions on the major measurements that were considered. Should schools be closed, people work at home, etc.

Thirdly, the members of the team did not apply for a position, but were all asked if they were willing to spare some time for the project. Each member contributes as far as possible next to a normal job. It means the members of the team are highly motivated and all believe that the ASSOCC approach is not just of some academic interest, but can be of added value in the real world. This commitment, not to a job, but to a common goal and ideal has made a huge difference in the outcome of the project! Without the dedication and countless hours spend on the project we could not have achieved any of the results in such a short time. Thus we see that the disadvantage of working in times of a crisis can also be an advantage as it focuses efforts and also shows very concrete the impact one can make with one's research. It maybe should be mentioned at this place that we did not have an epidemiologist on the team. The social simulations of the ASSOCC framework were about the COVID-19 crisis as a whole and not specifically the epidemiological part of it. So, we have regularly consulted with epidemiologists and used their models as part of our framework rather than incorporating the discipline itself in the team.

Despite the unusual circumstances in which the ASSOCC project has been conducted it has been very successful in a number of respects. First of all, we have achieved a number of interesting results from our simulations that proved to be a real contribution to the debates on measurements in diverse countries. That in itself is a good result for any social simulation project.

However, a more interesting result is that all the scenarios that were run on very different aspects of the crisis have been using the same implementation model! Thus, we have shown that one can base a simulation framework like ASSOCC on a fundamental model that connects different aspects of life in a coherent way and allows to make all kinds of combinations of factors to create new scenarios. An ultimate example of this is the curfew scenario which is not a separate chapter in this book, because it was run at request of some party in The Netherlands during the debate leading up to the curfew in February 2021. We were able to set up, run and analyse this scenario within two days (and come up with believable results that seem to be corroborated since by the real world situation)! It provides a powerful argument for the use of abstract models based on sound social-psychological principles in this type of simulations.

Maybe more important than the specific results that we got from our simulations were the lessons that we learned from running these simulations during a crisis. These lessons were the direct reason for writing this book as they seem to be valuable for the whole social simulation community. It would have been very nice and helpful if there would have been tools and methodologies available at the start of the ASSOCC project specifically for social simulations for crises. So, a large part of this book is dedicated to lessons learned from the ASSOCC project and especially discussing the biggest challenges when trying to create social simulations for crisis situations. In the rest of this chapter, we will already briefly position this type of agile social simulations in the field of simulations in general and give an overview of the contents of the book, describing the role of each chapter in the main message of book:

Agent-based Social Simulation can make a valuable contribution, not only to science, but also to society in times of crisis!

1.2 Simulations for Crisis Situations

Before getting into the details of the ASSOCC project and the rest of the book it is important to first place simulations for crisis situations in the broad spectrum of simulations being performed. One of the main determinants of a simulation is the purpose for which it is built. Reference [1] describes seven core purposes for simulations:

- prediction: anticipate well-defined aspects of data that are not already known
- explanation: establishing a possible causal chain from a set-up to its consequences in terms of the mechanisms in a simulation.
- description: an attempt to partially represent what is important of a specific observed case (or small set of closely related cases)
- theoretical exploration: establishing and characterising (or assessing) hypotheses about the general behaviour of a set of mechanisms (using a simulation).
- illustration: communicate or make clear an idea, theory or explanation
- analogy: use a simulation to describe another process that is hard to access
- social learning: encapsulating a shared understanding (or set of understandings) of a group of people.

So, what is the main purpose of a simulation for a crisis situation? Right away it becomes clear that in a crisis several of the above purposes are important if the simulation is to support the decision makers during the crisis. Decision makers want to have at least some form of predictions in order to shape their preferences between different courses of action (restrictions or policies). But the simulation should also be able to explain what is happening. In a fast moving world during a crisis the decision makers need to have some sort of understanding how their decisions affect the world. To a lesser extent, one would like the simulation to highlight which are the determining factors that will define the effects of decisions. Due to the high interdependency of many factors in a crisis it is often difficult to distinguish determinant variables from confounding factors. Simulations can be used to get a grip on this. An example of this is the question whether closing basic schools will effectively help contain the spreading the virus? Which are the determining factors and how will they be affected by the closure of schools? Finally, we have also actively used ASSOCC for social learning. Using the simulation we could show people why track and tracing apps might be handy for the health care organisations, but will have a very limited effect on the spread of the COVID-19 virus.

Given the above, very brief, description showing that simulations for crisis situations have inherently multiple purposes it is easy to understand that these simulations are also inherently complex. One could argue that separate simulations should be built for each purpose. However, it is very difficult to keep these simulations consistent and also how to combine results of the different simulations. Indeed, we see that the ASSOCC framework and system is inherently quite complex, but can indeed be used for several purposes due to the principled architecture and wide coverage of the model.

So, how does the complex ASSOCC framework fit in the classical taxonomy of simulations described in [2] as a prototypical simulation for crisis situations? We will briefly describe each dimension.

Abstract versus Descriptive: The Abstract versus Descriptive axis from [2] denotes two modelling purposes: simulating for the sake of reproducing a general phenomena, generally using on abstracted mechanisms (Abstract) or for the sake of reproducing a very specific situation, often including a wide array of detailed elements that are specific to the situation (Descriptive).

ASSOCC is in the middle of these two extremes. It should be abstract to model many possible situations in a quickly changing world in crisis. E.g. people will violate lockdown rules due to unfulfilled needs. But it also contains enough details to make the results relevant for decision makers at the time of the crisis. E.g. will track and tracing apps be useful, as studied in Chap. 7. This bipolar orientation is a central aspect of the design methodology for building simulations for crisis, as described in Chap. 14.

Artificial versus Realistic: The Artificial versus Realistic axis from [2] denotes the goal for building simulations for either observing the behaviour of possible societies (Artificial) or for replicating the behaviour of existing societies (Realistic).

ASSOCC is again in the middle of these poles. It is meant to simulate potential effects of policies during the crisis. In such it is meant to simulate possible societies and alternatives. But these societies should be clearly anchored in the current society. However, we do not try to just explain phenomena of the current situation and are thus not completely realistic.

Positive versus Normative: The Positive versus Normative axis from [2] denotes the goal of building simulations for either studying a phenomenon, with a generative social-science mindset (Positive) or to be used for guiding decision-makers (Normative).

In this dimension, ASSOCC is purely based on the normative pole. It is clear that simulations for crisis situations are meant for supporting the decision makers during the crisis.

Spatial versus Network: The Spatial versus Network axis from [2] distinguishes two modelling method concerns: whether the simulation is laid in a space such as a 2D grid or a map (Spatial) or whether distances are abstracted away (Network).

The ASSOCC framework is strongly based on the Network pole. However, this was a choice purely based on pragmatic arguments. Although a spatial map would be good to have, it would also make the simulation far more complex and inefficient. Thus, we chose to leave the spatial component out only for efficiency reasons and not for any conceptual reason.

Complex versus Simple Agents: The last dimension distinguishes whether agents rely on advanced cognitive models (Complex) or simplified if-then kind of statements (Simple).

ASSOCC is squarely positioned on the complex agent pole. We will argue in the next chapter why this is necessary for any simulation for a crisis situation.

If we take the position of a framework like ASSOCC with respect to all dimensions and compare it with other simulations, we see that it has a quite a unique position.

It has complex agents, based on an abstract model. However, the complexity is not primarily caused by trying to fit as closely as possible to all the details of a specific situation, but rather by the combination of many aspects of reality. Thus, ASSOCC does not require loads of data. We have used data mainly to calibrate certain aspects of the simulation rather than the simulation as a whole. ASSOCC is, in it present form, also not meant to give very detailed predictions. The scale of the simulations is too small to be able to do that. However, ASSOCC simulations can indicate some timelines and general trends. E.g. a curfew will reduce the number of newly infected people, but not enough to prevent a new wave after the curfew is lifted. So, other measures are needed in combination with a curfew. The positive thing is that we can show with the ASSOCC simulations that having a principled, abstract agent decision making model facilitates creating reasonable realistic simulations in a crisis situation. This property is especially important in these situations where data about the situation is scarce and normal behaviour is no longer normal. In these situations, having a model that is not very dependent on lots of empirical data is very useful! Thus, it seems that with the ASSOCC framework we have shown that social simulations for crisis situations do take a unique place in the field of simulations. And, moreover, this place requires some type of characteristics of the simulation that are not well supported by the common simulation tools yet, while crucial for working in crisis situations. We will use the rest of this book to argue why this is the case.

1.3 Guide to the Book

The rest of this book is split up in three parts. In the first part, we describe the background and foundations of the ASSOCC framework. In Chap. 2, we give a detailed overview of the theories that we have used to base the agent decision models on and also the arguments why we used exactly these theories. The main claim that we make is that an abstract model is needed for the decision models of the agents and we give some arguments why the theories and models that we have chosen are particularly well suited for simulations for crisis situations.

In Chap. 3, we give an extensive overview of the way the foundations have been implemented. We more or less follow the ODD protocol in describing the elements of the implementation, but adjust this to better explain the very extensive submodels of the ASSOCC framework. People that are mostly interested in the actual results of the simulations might want to skip this chapter. However, this chapter shows the actual complexity of the simulation and especially the agent decision making model. Anyone who wants to use the ASSOCC framework for their own purposes can find all the details necessary of all parts of the model to adjust them, discard them or extend them. It also important to re-iterate that this implementation is used for all the results shown in Part II of the book. Thus, it can also be used to analyse all kinds of details of these results.

In the last chapter of part I (Chap. 4), we describe the user interface module of ASSOCC. A unique feature where we create an interface for stakeholders from which

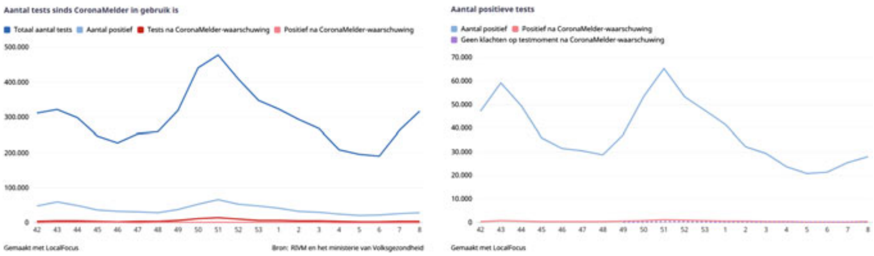

Fig. 1.1 Effect of track and tracing app in The Netherlands

they can see all scenarios, adjust some of the parameters in a controlled way and can explore the results of the simulations in various ways. A user interface like this is a necessity for any simulation as complex as the ones from the ASSOCC framework. We describe the architecture of the whole system in a way that others can use a similar set up if they want to provide a high level user interface for stakeholders of a simulation.

In Part II of the book, we collected six scenarios and their results that were run on the ASSOCC framework. Although many more scenarios could be run and included, we chose for these six scenarios as they are representative for certain types of applications and were in several cases directly used in national debates on the measures simulated in those scenarios.

Chapter 5 gives some insights in the effectiveness of closing (basic) schools. This was particularly relevant in the beginning of the COVID-19 crisis. Countries have chosen different strategies and the effectiveness of them is still not completely clear. In this chapter, we show which aspects play a role here and how their interdependence leads to some counter-intuitive results that still appear to be corroborated by reality.

The next chapter (Chap. 6), discusses some testing scenarios. This was done at the request of a regional government that wanted to know whether testing a large enough group randomly would work as well as giving priority to testing certain risk groups (like health care workers) regularly. Unfortunately, the results from the simulation were not very promising and this policy was never followed up.

The simulation results of the track and tracing apps have probably had the biggest societal impact of ASSOCC. We ran these simulations in April 2020 while the public debate in The Netherlands was questioning the benefits and fearing the consequences for privacy. Our results of the effectiveness of the apps deviated substantially from the most used epidemiological models. We predicted that the app would have a very limited effect on the spread of the virus. The following Fig. 1.1 which denotes the effect of the app in The Netherlands in the end of 2020 and beginning of 2021 shows that we were basically right with our prediction.

In the left figure the dark blue line indicates the number of tests being taken, the light blue line indicates the number of positive tests. At the bottom we see the red line denoting the number of tests taken after a warning from the app and the hardly visible pink line underneath indicates the number of those test that were positive. In the right figure the top blue line indicates the number of positive tests. The pink line shows the number of positive tests taken after a warning from the app and the purple line shows how many of those had no symptoms yet. Of course, one should also look at how many people actually used the app, which other measures were in place, etc. So, we will not make a scientific claim of having made the right prediction. But it gives a good indication and some of the analysis behind this figures is completely in line with the analysis that we will provide in Chap. 7.

All discussions about which country was taking the right measures at which time led us to investigate what could lead to the differences in effectiveness of measures between countries. Of course, countries differ in many aspects such as geography, population density, infrastructure, institutions and culture. Because taking up everything at the same time would take a multi-year project, we decided to pick one aspect that we already had previous experience in our simulations with: culture. So, in Chap. 8 we investigate the influence of culture on the effectiveness of the diverse measurements taken during the COVID-19 crisis.

At the beginning of the COVID-19 crisis, it was clear that this crisis was seen as a health crisis. Economic aspects were playing a minor role. Governments have given massive subsidies to industry to keep the economy from going bankrupt. However, after some time discussions started about which form of subsidy would be effective and how much and how long this financial support should be given. In Chap. 9, we investigate some economic aspects and effects of measurements of the government. Here, we also see that a macro perspective on the economy might lead to different measures than a social perspective that tries to support all individuals. More details on this are given in Chap. 9.

The last chapter of part II of the book appropriately investigates the consequences of different exit strategies. Which restrictions should be lifted first? In what order and when can restrictions be lifted without getting into a new wave of the pandemic. In Chap. 10, we look at several groups of exit strategies as they were applied around the world. Some exit strategies focus on getting the economic activity started again. Others mainly look at public life and how that can be restored safely. In this chapter, we see that some exit strategies have surprisingly similar consequences even though they are based on quite different principles. We analyse why this might be so and also give some heuristics that could be used to choose a good exit strategy.

After all the chapters of part II that gave an overview of the diverse set of scenarios that were run on the ASSOCC platform, in part III we turn to the analysis of the project as a whole. What did we learn from this experience and how can this help us for the future? In the first chapter of part III, Chap. 11, we discuss the actual impact we have had with the ASSOCC simulations. Not surprisingly, we were not part of governmental advisory committees. That could not be expected as newcomer in the field and in a time of crisis. However, our simulations have played a major role in the public debates in several countries in Europe and have indirectly also steered decisions that way. In this chapter, we discuss more in depth what we learned from the process, and what should be done for the future to get a place on the table for a next crisis.

When we started publicising our results and the media started picking up on that, (legitimate) questions were raised on the validity of our results. Especially the results of our simulations of the track and tracing app gave results that were at first sight counter intuitive. Because we did take these questions on validity serious we have done an extensive investigation into the validity of our simulations by comparing them to a state of the art epidemiological model [3]. In Chap. 12, we report on this comparison and how this can be used to show the validity of our simulations. It has been a long and difficult journey to get to all the details of both simulations and see exactly how they can be compared. But it is also worthwhile, because by itself also gave a better insight in the ASSOCC simulations. We actually would promote these kind of comparisons to be done more often.

Already quite early in the project we realised that scalability of the simulations would be an issue. Using NetLogo together with a complex cognitive agent model means that one can run maximally around 2000 agents in a run. But besides this, obvious limitation there were many issues to deal with while creating one of the most complex NetLogo simulations. In Chap. 13, we describe the software engineering aspects of running this big and complex project that also had to deal with external deadlines and an ever shifting focus on new aspects that became important during the crisis. The main reason we could manage this was that we had a very solid foundation to start with on which we could easily add and change all other components. Keeping very good software engineering principles in managing the code and coders was also of prime importance.

Many times, we have thought during the project how nice it would be if we already would have had some tools prepared beforehand. Although we did manage to build and adapt most support tools that we needed for the ASSOCC project it is clear that a better starting point would have helped in many ways to achieve even more, get quicker analysis, better communication, etc. In Chap. 14, we describe which are the main areas that have to be developed and what is needed for that in order to be ready for a next crisis. There are some fundamental conceptual and design aspects that can support a flexible and scalable simulation platform.

In Chap. 15, we recapitulate the challenges that were found during the project and indicate the most important research directions. These are not challenges for the ASSOCC project, but more fundamental issues for social simulations for crisis situations. They are about creating a flexible decision making mechanism for the agents that is also scalable. About which software engineering techniques can be used to support the scalability issues of these agile social simulations. In short, this chapter describes a first step towards a research agenda for the community that wants to give social simulations real impact on crisis situations.

The book is closed by Chap. 16, where we draw some general conclusions and give a vision of future work for social simulations for crisis situations based on the experiences of the ASSOCC project for the COVID-19 crisis.

References

1. B. Edmonds, Different modelling purposes. in *Simulating Social Complexity* (Springer, 2017), pp. 39–58
2. Nigel Gilbert, Agent-based social simulation: dealing with complexity. Complex Syst. Netw. Excell. **9**(25), 1–14 (2004)
3. R. Hinch et al., OpenABM-Covid19-an agent-based model for nonpharmaceutical interventions against COVID-19 including contact tracing. medRxiv (2020)

Chapter 2
Foundations of Social Simulations for Crisis Situations

Frank Dignum

Abstract Simulating human behaviour in times of crisis requires models of human decision that are include aspects beyond directly visible actions. In crisis times the behaviour of people will change based on the changing environment and needs. Without an underlying model that can represent how and when people will change their behaviour it becomes difficult to incorporate these behavioural changes in the simulation. In this chapter we will introduce the foundations of the model that we used to model the human behaviour for the COVID-19 crisis. We argue that these foundations are not only useful for this application but are broadly applicable for simulations that need to capture behavioural change due to crises or other external influences.

2.1 Introduction

During the COVID crisis it has become very apparent that the spread of the corona virus heavily depends on (changing) human behaviour. Where for other epidemics of less lethal viruses the human behaviour could be approximated using statistical models of normal behaviour, this was no longer sufficient for the corona virus. Due to a combination of a long incubation time where people are contagious but have no symptoms yet, the fact that many people do not show any easily recognisable symptoms at all, the fact that older people are much more likely to suffer severe consequences of being infected and the lethality of the virus meant that very strict restrictions were considered necessary to prevent the virus to spread to the most vulnerable groups and cause huge amount of deaths. Another important factor that made human behaviour and behaviour change important is that the pandemic and the various restrictions stretched over several months and thus impacted every aspect of life. Thus models of human behaviour during the crisis would also need to include different aspects of life, like social effects of long term isolation, economic consequences of closures of shops, public places, leisure places, etc.

F. Dignum (✉)
Department of Computing Science, Umeå University, 901 87 Umeå, Sweden
e-mail: dignum@cs.umu.se

F. Dignum (ed.), *Social Simulation for a Crisis*, Computational Social Sciences,
https://doi.org/10.1007/978-3-030-76397-8_2

Although the issues described above are special for the COVID crisis they are by no means exclusive for this crisis. In many crisis situations the above issues play a major role in the way a crisis evolves. There are several places in the world where regular natural disasters like war, draught, flooding or earth quakes create a crisis situation. In these situations the evacuation of people from the affected area, providing "temporary" shelters and recuperation of a "normal" life are important. Whereas modelling these crisis situations might at first focus on the evacuation process and the creation of the refugee camps, it should also include the social relations and status when allocating places in the camps. Moreover, the crisis is not finished with people having moved out of a disaster area. It is over when those people have some way of existence in another place or back in the original area after it has been restored. This longer term perspective might have a huge impact on how the short term aspects are or should be handled. E.g. where to place a refugee camp. In an example of flooding situations in Indonesia it is known that people are reluctant to evacuate from their homes out of fear of plundering and fear that they are not able to return to their often illegal dwellings. This hampers many long term solutions for this crisis. In all of these situations, social, economical and psychological aspects play a role and are not easily disentangled.

In this chapter we will first investigate what are the consequences of the above observations for the type of models that the social simulations for crisis situations should be based on. Next we will describe the foundations of a model that fulfils these requirements and is used in the ASSOCC project for the COVID-19 crisis. We will show later in this book that this model can be used to get insightful results in the COVID-19 crisis on many different aspects.

2.2 Crisis Situations Require Abstract Models

The points described in the previous section give strong arguments to create an agent model that is based on some fundamental abstract notions that can be used to link all of the different aspects mentioned above. This differs from social simulations that use (and often only require) statistics of real-world behaviours to model behaviours. E.g. if "23% of the people decide on A" this is modelled by having agents randomly 23% of the times decide on A. However, this behaviour cannot be explained afterwards, but more important we loose a possible consistency of behaviour. It might that in general 23% of the population has some property that makes it decide for A. (e.g. living in an area, having a certain profession, being of a certain age, etc.). This dependency of A on that property is now lost and the results of the simulation might differ substantially because of it. We will see some of this in the simulations about the effectiveness of the track and tracing apps in Chap. 7.

In Fig. 2.1 we very schematically compare the two approaches. We are aware that shows a very black and white picture and is grossly oversimplified. However, it shows the crux of the differences and the choice to be made.

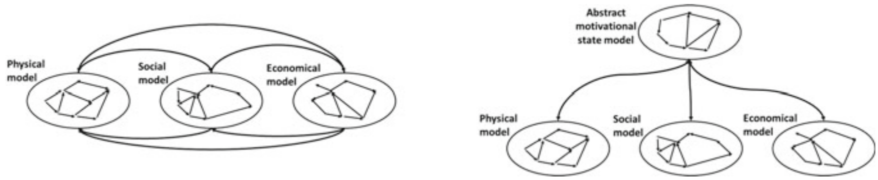

a:connected concrete models b:abstract model

Fig. 2.1 Concrete models versus abstract model

In Fig. 2.1a one can see a few concrete models that each focus on a different aspect of reality. Each of these models can be validated on information from the real world. However, there also all kinds of dependencies between elements in the different models. These are denoted by the arrows connecting the models. Usually the existence of these dependencies is discovered through correlations in data and unexpected phenomena where a condition in one model will lead to a different action in another model than might be expected. E.g. when there is no money to buy petrol for transport (logistics model), evacuation by bus and car will not work even though it might be the most preferred option in terms of mobility and flexibility (preference model). The main disadvantage is that the dependencies between the models are generally not covered or through a collection of sometimes contradicting theories. This can lead to inconsistencies, incompleteness and ad-hoc solutions that are difficult to explain and justify.

The architecture where an abstract model is used that somehow governs the concrete models solves the ad-hoc representations of the inter-dependencies by using the abstract model that should have its own properties. This right away indicates the disadvantage of this approach. We need an abstract model that people can agree upon. Moreover, this abstract model can usually not directly be validated by information from the real world. It needs indirect validation through the other models. The advantage of using more concrete, simple models is that each of these models can be validated against (historical) data. But it should be noted though that in crisis situations the concrete models are often no longer correct.

Concrete models usually make implicit and (probably) unintended assumptions as they are usually based on stable situations where people will react in reasonably predictable ways to the situation. In these situations people will act according to standard social practices, norms, habits, etc. Thus simple models that connect the situation to an action are sufficient. This can be illustrated as in Fig. 2.2a where the red graph shows the actual behaviour of the people and the blue straight line shows the simple approximation. It works well in the left part of the graph, but gets worse results when the red graph changes direction (e.g. due to a crisis situation).

In a crisis the context of decision making changes drastically and thus the data used for previous situations does no longer predict the behaviour in the current context. Thus more abstract models might be necessary for these situations that contain several internal states (represented by additional parameters). It does not mean that these

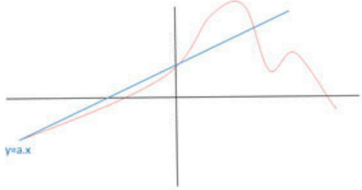

a:simple linear model

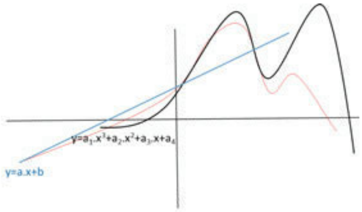

b:wrong abstract higher order model

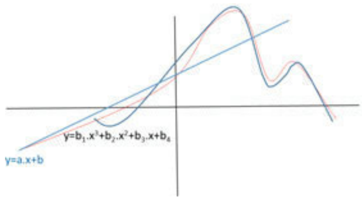

c:better abstract higher order model

Fig. 2.2 Simple model versus complex abstract model

abstract models will by definition give better results! If these abstract models are not well constructed or internal relations between concepts are not well defined they rather confuse than clarify behaviour. This is illustrated in Fig. 2.2b where we use a higher order function to approximate the real behaviour, but we apparently did not get all the parameters a_1 to a_4 right.

In Fig. 2.2c we have the parameters better calibrated and the resulting graph fits pretty well with reality. Of course, it is clear that to validate these more abstract models we need either a lot more data over many more situations or a good theory on how the parameters are related.

So, neither approach is per definition better than the other (unless we have a universally accepted abstract deliberation model). This is in line with [1] who argues that the type of model to be used depends on the situation that is to be simulated. Our argument is that the abstract model approach is more useful under the conditions of crisis situations that we want to model and simulate:

1. There are many dependencies that play a role in the situation and they clearly influence the outcome a lot.
2. changes in the environment (either natural or social) play a big role in the situation and the reaction to these changes depends on several dimensions.
3. the simulated situation spans over a time frame that is long enough to be more dependent on the interactions between the different dimensions.

With respect to 1. take as example the COVID crisis. The main perspective is that of health and thus an epidemiological model seems to be the most appropriate to base a simulation on. However, it soon became apparent that human behaviour is a determining factor in the spread of the virus and this behaviour could not easily be captured through the usual statistical estimates of behaviour. Thus at least a epidemiological model and behavioural model are needed and ways to connect these models.

With respect to 2. in any crisis situation the environment is changed in unexpected or new ways and society will react to those changes. However, the way people react to a crisis or a new policy/restriction in a crisis depends on many factors that are not always part of the same dimension. E.g. keeping isolation when having corona symptoms clearly is advantageous from the health perspective. However, old people that cannot see their (great)grandchildren due to this isolation might prefer to violate the isolation and take the risk of getting corona. So, the need for social contact overrides the health concerns.

Finally, in 3. we emphasise that dependencies between different social dimensions become more apparent and important over longer time periods. E.g. rescuing refugees from a war zone leads to a strong focus on safety and survival. However, in a long term perspective the refugees should also be given a perspective for the rest of their lives and their children's lives. This perspective leads to considerations of other aspects of life than just safety.

We argue that in the situation of the COVID crisis the above conditions are all present and play a big role. Therefore we took an abstract deliberation model as the basis for our agent models. We already stated that there is no universally accepted abstract model for agent deliberations. We also will not argue that the model that we will describe in this chapter and that forms the foundation of the rest of the project is the only possible model or the "best" model. However, we will argue that there are a number of characteristics that we would like to have from such an abstract model when we use it to model agents in crisis situations. We will discuss these characteristics in the next section. After that we describe the way we have filled in the foundations for our simulation model for the ASSOCC project. We do not claim to have the one and only foundation for abstract agent deliberation models. However, we claim that the considerations that we use to compose this foundation are important for any simulation for crisis situations. One might make different decisions on how to fill in the different components due to the importance of some aspects. However, some general properties of the model will be preserved if the considerations that we lay down are followed. That our foundations for the model do work can mainly be seen from the second part of this book in which several scenarios of the COVID crisis and their results are described. All of these scenarios have been made using the same conceptual model! This shows the power of the approach and also that the model gives at least interesting and explainable insights on several aspects of the crisis.