Artificial sociality

2018 05 22, Web version

Gert Jan Hofstede

Mission

To develop 'artificial sociality': foundational conceptual models of human sociality based on social science, for use in agent-based models of complex systems in the life sciences. Artificial sociality involves two main components: designing individual minds, and modelling self-organization of behaviour at system level. A special focus rests on safety and resilience in socio-'something' systems in which human components are replaced by technical ones.

Contents

Problem issue

Modelling and simulation are powerful ways of comprehending the complexity of our world. I'll give two examples from my research. They involve 'agent-based modelling', a technique in which a system is represented by a two-dimensional world populated by 'agents' that could represent people, institutions, or technologies.

Resilience

How do institutions affect the resilience of pig husbandry in the province of Noord-Brabant against environmental concerns? This question is the focus of resilience brigade REFORIM at WUR SSG.

One of the means of study is to do a historical analysis of the province's development, and another to create an agent-based model for studying the dynamics over the years. The model does not only contain the environmental side (e.g. pig farms produce smell in a certain radius around them), but also the social (e.g. there are farmers, industry workers, and countryside dwellers), economic (income) and the institutional (e.g. there are political parties and lobby groups). During a model run, flows of finance, materials and information happen, in feedback with the institutions.

Safety

How can high-risk industries be made safer? This is the issue tackled by external INF PhD Fred Goede. Advanced technologies did not prevent major incidents such as the nuclear incident in Fukushima, Japan and the oil spill incident in the Gulf of Mexico, USA. Despite engineering efforts to minimise the human factor through automation, behaviour and culture are most commonly blamed for these incidents. Empirical data suggest that reduced incidence of small incidents, and resultant limited learning opportunities for operators in complex systems, contribute to major safety incidents. Fred uses socio-technical agent-based models to investigate the dynamics of accidents. A solution could be to develop sociotechnical system simulations for training employees in safety performance.

In these two examples, agent-based simulation helps understand the dynamics of complex systems that include social and other elements. In what follows I'll create a context for this type of modelling. I start with a wide lens, then zoom in on my contribution.

Gaia and design

For millions of years, our ancestors have been part of nature. Then by and by, human nature gave rise to practices that we recognize as 'cultural' in a broad sense, such as purposeful burning, and later 'agri-cultural' practices (figure 1). Thus, culture started to affect nature. In recent evolutionary times, cultural practices have exploded into development of technology. Technologies of course have also started to affect nature in their own right. Now that technology pervades every part of our society, it has also started affecting culture.

The complex system of systems consisting of nature, culture and technology is what Latour calls 'Gaia' (Latour, 2017). He means to say that these elements are mutually dependent, and we cannot isolate them anymore. According to Latour, climate change is not just a problem of nature, but also a cultural and a technological issue.

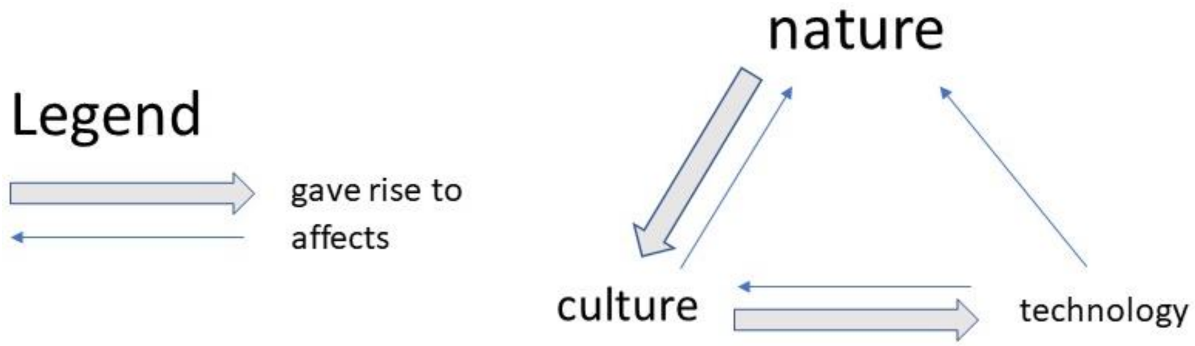

Figure 1: Gaia: the intertwining of nature, culture and technology. Size of letters indicates evolutionary chronology.

In figure 1, ‘culture’ has the very general sense of ‘everything human that is not nature’. For the study of these domains we have the social sciences (sociology, psychology, economics, political science, …) and humanities (history, arts…). In contemporary academic language, the word culture has a narrower meaning: ‘that which distinguishes one group of people from another’, which indicates values and practices, or ‘the unwritten rules of the social game’, indicating values specifically (G. Hofstede, Hofstede, & Minkov, 2010). Culture in this sense is the group-level result of the relational nature of human beings. Unless otherwise indicated, this is how I’ll use the word in this document.

Socio-something systems

This vision deals with subsystems of Gaia that I shall disrespectfully nickname ‘socio-something systems’ to indicate that culture as well as a selection of other aspects fall within the scope. This is not to say I equate nature with technology; unlike nature, technology is explicitly designed and built by people. However, as soon as a technology is invented, it starts to mingle with culture. For instance, we invented currency a few millennia ago, and now, it has developed into a world of its own, part culture, part technology. Hence ‘socio-economical’ is now a meaningful term, alongside ‘socio-ecological’ and ‘socio-technical’.

In the nineteen sixties, Herbert Simon had identified the ‘Sciences of the Artificial’ (Simon, 1996). This seminal book marked the idea that studying engineered parts of reality merited its own methodology. Marrying that thought to the idea of Gaia, we realize that not just the artificial, but Gaia as a whole merits its own methodology. We have to study our world, including technology, in a systemic way, as something parts of which we have designed, or could design. We should not do so purely as engineers though, because for the most part, nature and culture are given to us. We are powerless to change them, but we can take them into account in our models and in our designs.

Socio-technical systems

All ‘socio-something’ systems are alike in that their social subsystem and their other subsystems mutually influence one another in path dependent ways: they are complex adaptive systems. One difference between the natural and the technical element is time scale, and rate of change. Technical subsystems come and go during one life; today’s hype becomes tomorrow’s fossil. Thus in socio-technical systems, the social element is more stable than the technical.

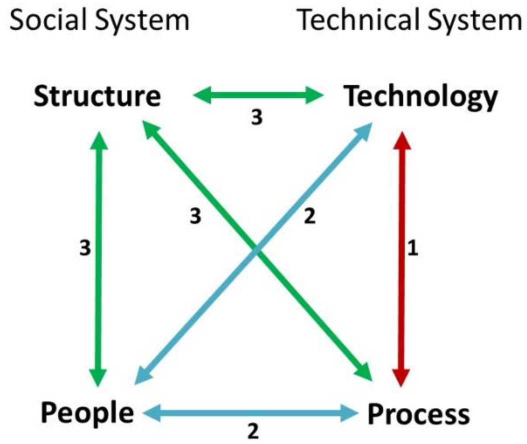

Figure 2 shows four components of a socio-technical system: Technology and process on the technical side, and people and structure on the social side. 'structure' means not only formal organizational structure but also institutions and culture. The figure also shows three orders of change. First-order: New technologies leads to changes in processes. Second-order: Changes in technology and process require changes in roles and skills for people. Third-order: these changes then interact with existing cultural and institutional structure, with unanticipated consequences.

Figure 2: a socio-technical system (Piccoli, 2012). The arrows represent relationships. The colours and number represent orders of change: first (1, red), second (2, blue) and third-order (3, green) change.

Figure 2 helps us to think about how such a system could be modelled. To do this, let me distinguish 'instrumental' from 'relational' logic. Instrumental means bringing something about in the physical world, whereas relational means bringing something about in the relational world. Our actions are usually grounded in both logics. While the technical sub-systems in figure 2 are controlled by the laws of physics (e.g. if we heat the liquid until a certain point, it will boil) and by instrumental logic (e.g. how many calories a day does a person need), the social sub-systems are controlled by a complex mix of logics: some of it instrumental, but some of it based on the relational significance of actions that are only seemingly, or at most partly, instrumental. People do things for a mix of reasons, some instrumental (e.g. feeding oneself to get energy), some relational (e.g. eating out with others, waiting to eat until everyone is seated, loving some food for sentimental reasons, avoiding other food due to social norms or laws). These relational reasons are part of 'culture' in the sense of Latour. In models one can take them into account using concepts such as values, norms, beliefs and rituals. In ordinary life, as well as in quite a lot of sciences, many of them go unnoticed. We tend to underestimate the degree to which our behaviour is relational and socially contagious, and pretend that our actions are rational or instrumental. As a consequence, we underestimate the influence of informal social structure on the functioning of socio-something systems.

Replacing people by computers

Most organizations rely on slight deviance from the formal rules for their smooth operation, and these deviations are mostly due to human decisions with relational logics. Usually they are not documented anywhere. When social subsystems are partially or entirely replaced by technical ones (often dubbed 'computers'), there is a risk that the relational elements are lost. A model system with relational logic can help understand such possible unintended changes in system behaviour. I shall call this relational logic in agent-based models 'artificial sociality', a term coined by (Gilbert & Conte, 1995). In artificial sociality, both formal social structure and cultural 'unwritten rules' matter. Cross-cultural issues will turn out to be important in many socio-technical systems. When people differ in their values, norms, beliefs and rituals, but have to collaborate in a single technical environment, or using a single formal organizational structure, misunderstandings that lead to malfunctioning systems are very likely. Real life abounds with examples about the unanticipated side effects of changes in technology or in policy. Examples in socio-ecological systems include unintended pollution or environmental degradation. In socio-technical systems, they include altered opinion dynamics in networks dominated by social media versus face-to-face interaction. For instance, social media are changing opinion dynamics, while culture also shapes social media usage (Choi, Im, & Hofstede, 2016).

The social subsystem: culture

In the social subsystem we can distinguish various levels of aggregation. For brevity I shall limit the discussion to the individual, the group, the society, and all of humanity. Distinguishing these levels is important for two reasons. First, different social sciences study them, each contributing knowledge to inform models of the social subsystem. Second, lower levels self-organize to yield emergent patterns at higher levels, and this relationship can be modelled with agent-based models. The words 'self-organize' and 'emergent' are two sides of the same coin: in a complex system, self-organization is the process of which emergent patterns are the result.

Let me now sketch the dual nature of our world as both instrumental and relational, using the four levels of aggregation I mentioned.

- At individual level, people do instrumental things in the physical world such as eating, moving, picking up objects; or they do things in the symbolic world such as voting for a candidate, or earning money. They do them because they are motivated to do so, by desires, by obligations, or by norms. They draw these motivations from a deeply relational world. People give relational meaning to their actions.
- People in a group develop a shared understanding of their common activities and their common relational world. This gives rise to observable shared practices, driven by shared culture-specific norms and values. The depth of this culture varies from ephemeral to life-long. People join groups, they play roles in them, they build social identities in which these groups are important. For instance, competition for bravado among groups of young males could lead to ritual competition for drinking, or to destroying bus shelters. The fact that there is such competition would be a cultural attribute of this group.
- A society is not one group, in the sense that not all of its members know one another. Yet a society has a certain structure that is rooted in its shared institutions, as well as in shared values, that is, in a shared understanding of the relational world. These allow its members to get along in their myriad daily interactions without running into misunderstandings. These values are what is called 'culture' in contemporary science. Value systems at societal level are slow to change across generations, despite huge changes in the technological environment and in practices. This is because we acquire our values in infancy.
- Laws of physics are the same for all people, and there is one instrumental world. Only, the same instrumental action can mean something else in one culture or group than in another. As a result, the social world is not necessarily the same to all people. A generic model for humans needs concepts with which to give every agent its own social world in the model.

With a big sweep, I conclude that in the Anthropocene, nature, culture and technology are rapidly merging into one system, which is the focus of Wageningen University for the Life Sciences. As we saw, Latour calls this system Gaia. Gaia is a system of systems consisting of innumerable linked complex adaptive systems. Technologies that are hyped today are part of Gaia: Social media, Internet of Things, Big Data. New hypes will follow, such as subcutaneous devices or device-brain connections. Nature and culture are much slower to change than technology. While we have excellent knowledge about nature, we struggle to understand culture in its relation to complex systems in Gaia. Obtaining better knowledge of culture that can be used in models of complex adaptive systems is thus an urgent mission.

Computational models

Artificial sociality requires that a combination of data and theory-based conceptual models be converted into computational models. This is an art that requires people who are sufficiently at ease with empirics, theory and computation in their field of study. To complicate matters, various types of computational models exist that differ radically in their use of data and theory (figure 3).

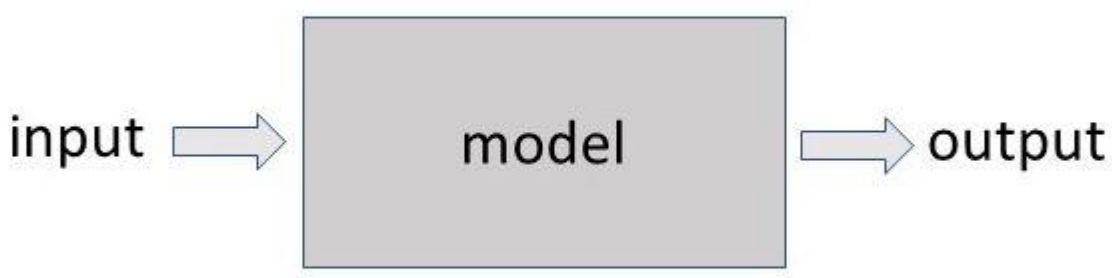

Figure 3: the basic structure of a computational model.

The point of figure 3 is that depending on the case, the researcher may know more about the input, the model, or the output. This has consequences for the type of model that is suitable. I'll discuss two types of models using the figure. In the first, the model is a black box, while in the second, it is a researcher-constructed miniature world. The aim is not to find a 'best' method; each modelling method has its own merits and fields of application, and the methods can also be combined (Lamperti, Roventini, & Sani, 2017; Rand, 2006). The aim is to elucidate the requirements for computational agent-based models.

Machine learning

At one end of the spectrum we have machine learning. With this technique, both inputs and outputs are precisely known. Inputs can be 'big' data, and outputs can be desirable system behaviour. The model is a neural network that is 'trained' to produce the output from the input. Once trained, it can be used to check whether new inputs can also produce the output. This technique is incredibly powerful in situations where the input and output fit within the same conceptual 'world', such as 2-d pictures to recognize faces, candidate-produced texts to match job candidates with personality profiles, or the game of go. A major limitation of neural networks is that they are not transparent: they train themselves, while the researchers do not know how they operate. Another property is that they produce a 'best' result, a match.

Agent-based modelling

At the other end we have agent-based modelling. Here, we typically have some knowledge of all three elements: input, model, output. The input consists of real data, possibly 'big', or might be fictional scenario data. The agent-based model includes explicit knowledge about agents, environment and processes between agents in the environment. This knowledge is based on a mix of data and theory, depending on the case. The agents might all differ from one another, and the researcher controls all the parameters. The output consists of patterns, or indicators, at system level. These patterns are usually of interest to policy makers or decision makers. The decision makers can change both inputs and model parameters to find out which combinations generate interesting outputs. There is usually no 'best' output. What the models allows the user to do is to play around with model hypotheses in order to understand how the dynamics of the model produce the outputs of interest. The model can also be used in interactive mode together with stakeholders, giving them insights into their role in the larger system.

So the spectrum from machine learning to agent-based modelling is one of exact versus approximate knowledge. Policy making typically faces the latter situation. Selecting input data, but even more so adapting theory for use in agent-based models, is an art. Agents need to be taught everything from scratch. They need skills and drives. They need perception, interpretation and action repertoire, like real living beings. Many theories could be used, especially motivational theories from the social sciences, but almost none of these is readily usable. Agent-based modellers need to simplify, combine and disambiguate theory. Conversely, applying theory in agent-based models constitutes a kind of truth test for a theory. If no agent-based model can be created that generates the patterns that a theory purports to explain, then perhaps the theory needs to be re-examined.

State of the art

Individual level: artificial mind

The field of Artificial Intelligence (AI) has given rise to various ontologies for intelligent agents. For instance, there is a wealth of work on logic, on knowledge, and on planning by agents. So far, these streams of research have mainly dealt with 'isolated' artificially intelligent agents, perhaps interacting with a person (Gratch, Wang, Gerten, Fast, & Duffy, 2007; Johnson et al., 2014), but not with one another. Relational motivations to do with group membership and hierarchy that would be needed for artificial sociality have not been as much studied as instrumental ones (G. J. Hofstede, Ambrosius, Bokkers, & Boumans, 2014).

One model with a good track record for describing single decisions used quite a bit in ABM is the Theory of Planned Behaviour and its sequel, the Reasoned Action Approach, by Fishbein and Ajzen (Fishbein & Ajzen, 2010). The BDI framework (Beliefs, Desires, Intentions) has been something of an industry standard for social agents (Georgeff, Pell, Pollack, Tambe, & Wooldridge, 2003). It links the instrumental world with the motivational one, but falls short on the relational side (Dignum, Hofstede, & Prada, 2014; Prada & Paiva, 2005). An enhancement of BDI is provided by the OCC model (Ortony, Clore, & Collins, 1998) dealing with emotions and social valuation. BDI - OCC has been operationalized in the Fatima framework (Dias, Mascarenhas, & Paiva, 2016). It has also been enhanced with considerations of social status, operationalized in Social Importance Dynamics (Mascarenhas, Prada, Paiva, & Hofstede, 2013). Recently, for socio-ecological systems, the MoHub meta-model was proposed, that conceptually captures a number of actor models: economically rational, Reasoned Action Approach, social norms (Schlüter et al., 2017). MoHub includes social motives from the point of a focal actor. Despite these developments, the state of the art in AI is still that knowledge and logic have prevailed over social motivations, with the consequence of obtaining 'autistic' agents (Dignum et al., 2014; Kaminka, 2013).

Across levels: emergence

Methods derived from AI are good at looking into the minds of individuals. They tend to be 'methodologically individualistic' in that they take the individual mind as the carrier of agent agency. Often, they have no ontology for the fact that our world consists of complex adaptive systems that self-organize (figure 4).

Figure 4: emergence, or self-organization, as discussed in the text. If this was modelled in an ABM, the agents would walk around littering, and the resulting system pattern would be a spatio-temporal distribution of litter.

Self-organization here means showing patterns of output that nobody intended, but that 'just happen'. In other words, these unintended patterns emerge from the collective actions of the individuals. Methods that do not take self-organization on board have trouble modelling the emergence of group attributes such as culture, norms and institutions (Castelfranchi, 1998; Gilbert & Conte, 1995). Multi-level methodology is required to capture emergent complexity (Conte et al., 2012). Such work is being carried out in many fields, with different reference theories and emphasis. The following gives a few highlights.

- Game theory: this is a cohesive subfield in which economically rational agents 'play' Prisoner's Dilemma-like scenarios against one another, yielding outcomes in terms of agent strategies and their economic payoffs. These agents have been pitched in repeated games where they encounter one another many times and remember past interactions. Here, it turns out that variants of 'tit for tat' strategies have long-term survival value. In other words, in this simple world, it pays to be nice (Pinker, 2011).
- Sociology: computational sociology with highly abstract agents has been conducted using agent-based models (Deffuant, Amblard, Weisbuch, & Faure, 2002; Grow, Flache, & Wittek, 2015). This strand of work is strong on emergence. It constitutes a new methodological start; 20th century sociological theory with group-level constructs tends to not be used. There are exceptions, e.g. the work by Heise who used a sociological theory on small group interaction in an ABM explaining speaking time in US juries (Heise, 2013).
- Computational social simulation as a new inter-discipline, taking input from a diversity of fields. This has yielded some re-usable meta-models, e.g. the Consumat framework (Jager & Janssen, 2012). Agent-based studies have also been conducted for studying emergent pattern with simplified social agents that can exchange attributes. The results support the notion of a social world that emerges rather than being designed top-down (Axelrod, 1997). It has also proved feasible to model culture in agents that perform single narrowly defined tasks such as purchasing (Roozmand et al., 2011) or negotiating (G. J. Hofstede, Jonker, & Verwaart, 2012).

Over the last decades, spectacular advances have been made in many subfields. Unfortunately, these developments have so far had little impact on mainstream social science (Squazzoni, Jager, & Edmonds, 2013).

Models create a single world

Agent-based modellers need to create comprehensive, 'one world' models of the systems that they are intended for. Therefore, fragmentation between disciplines is a problem for modellers of socio-technical systems if these disciplines are needed together in a model. Comprehensive social scientific models across sub-disciplines, many of which were developed in the past century (Brown, 2000; Coleman, 1990; G. Hofstede et al., 2010; Kemper, 2011; McClelland, 1987; McCrae & Costa, 2003) are usable in modelling complex socio-technical systems, but have tended to go out of fashion in their source disciplines. As a result, comprehensive knowledge with a proven track record of explaining real-world phenomena is nevertheless little used in the design and study of complex systems. We have good models of human motivation, we understand social identity phenomena, we know how cultures systematically differ, but we struggle to use this knowledge in our models to support policy making.

Agent-Based Modelling as a method

Comprehensive models from the social sciences have the promise of integrating disparate concepts into one system, as far as the expressive capacities of the computational modelling framework allow. I'll concentrate on agent-based modelling (ABM), which is the most versatile framework for modelling the social world (Epstein & Axtell, 1996; Gilbert, 2008). Over the last decade, ABM has matured as a methodology that can facilitate the kind of interdisciplinary, comprehensive, dynamic, socially rich models needed for policy making. Typically these models do not pretend to predict the future, but they help policy makers understand the dynamics of the systems they are responsible for (Edmonds & Meyer, 2017). There is still a lot of methodological work to do, but the basics of the method are in place.

An agent-based model contains the actions and interactions, and the resulting simulation generates emergent patterns. Agent-based modelling thus facilitates the consilience between contributing disciplines:

- It allows to integrate 'vertically' between levels of aggregation such as individuals, teams, organisations, countries.
- It allows to integrate 'horizontally' between disciplines that contain knowledge to include.

Agent-based models are themselves complex systems. This means that modellers have to study the path-dependent behaviour of the agent-based models they created, to find out why they generate certain patterns.

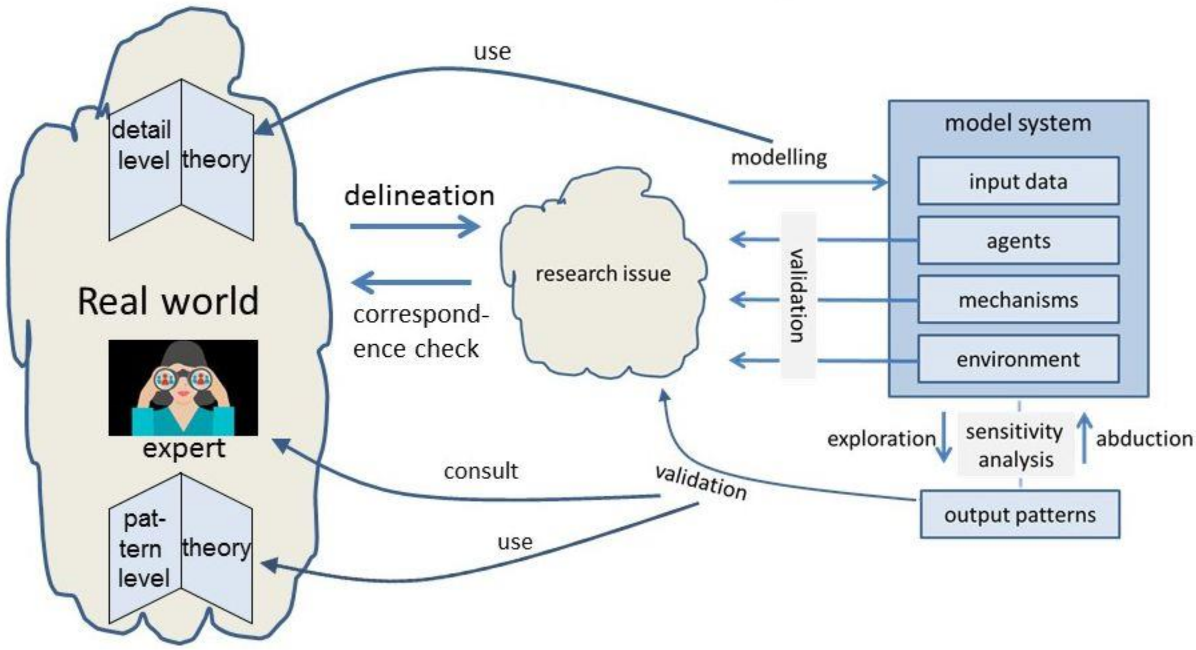

Figure 5: the place of agent-based modelling in a project (see text)

Consider figure 5, that I developed based on methodological experiences in six PhD projects in the Complex Adaptive Systems (CAS) strategic programme. Figure 5 shows the 'real world' on the left, where 'experts' have experience with the research issue. Knowledge is available at two levels: the 'detail', that will appear as agents and environment in the models, and the 'pattern', which is a system-level observable that will become an output variable of the models.

In the middle, a 'research issue' is taken from the real world by 'delineation'. This delineation is essential; how to select what to take into account is determined by the aim of the modelling exercise. Agent-based models are usually unable to *predict* future system behaviour, but they are very useful for acquiring *understanding* into possible futures of the system under study. They can show second-order effects, such as trends, tipping points or other nonlinearities.

Now, on the right-hand side, the actual modelling starts, using detail-level theory, leading to a 'Model system' (the ABM), containing the elements 'input data', 'agents', and 'environment'; and whose dynamics depend on 'mechanisms' for the agents in their environment using the input data. Both the agents, the mechanisms and the environment must be validated against the research issue, to argue that they make sense according to the acronym KIDS: 'Keep it Descriptive, Stupid!' (Moss & Edmonds, 2005).

Bottom right we see how 'output patterns' caused by the model system are explored through model runs, and can lead to their explanation ('abduction') through the model system's elements. An important aspect of the exploration / abduction cycle is 'sensitivity analysis', analysing which model variables contribute to the output patterns to what extent. Specific sensitivity analysis techniques are still being developed for agent-based models (ten Broeke, van Voorn, & Ligtenberg, 2016).

To close the cycle, the output patterns must be validated against the research issue, consulting experts and using pattern-level theory.

Gap: artificial sociality

So, what is missing? For the current gap I would like to use the term 'artificial sociality', as a complement to 'artificial intelligence' (AI). Artificial sociality as I see it needs to encompass

- An individual-level operable, dynamic model of the relational side of human social behaviour that is grounded in solid social science and will be usable across disciplines for strategic use in policy making.
- Multi-level theory of self-organization. It should link individual behaviour with resulting emergent patterns of behaviour at system level.

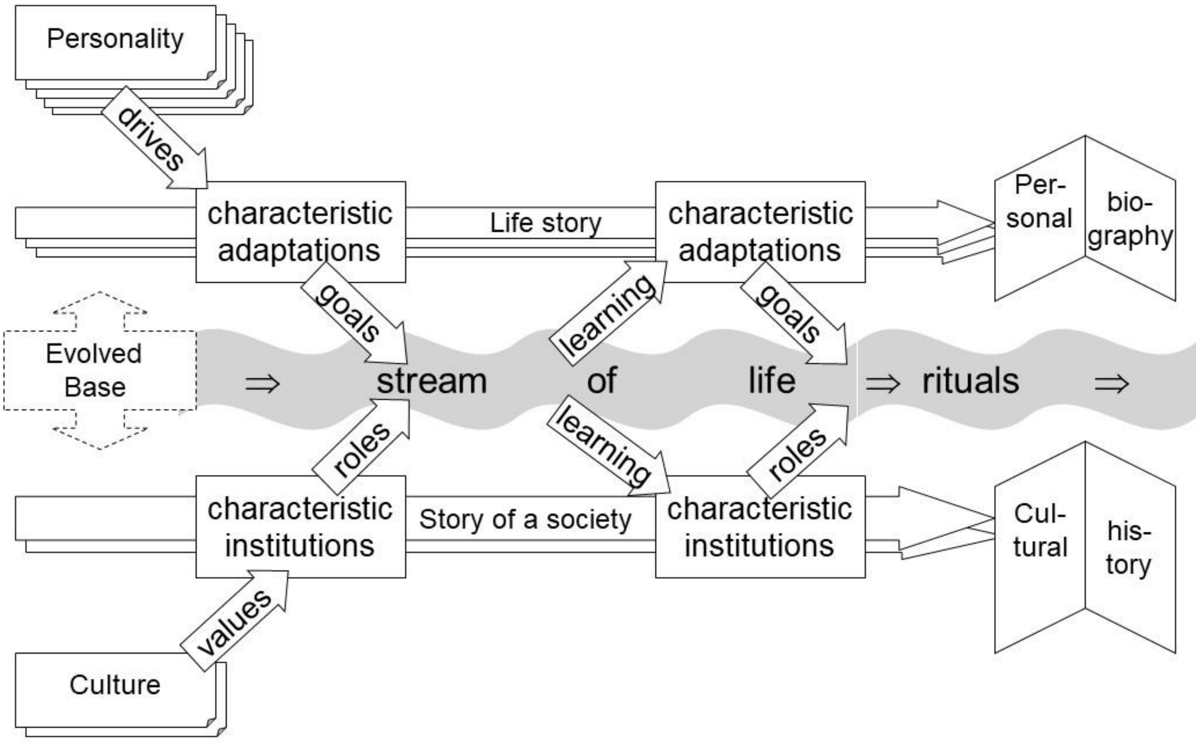

Figure 6: rituals in the stream of life, influenced at individual and societal level (G. Hofstede et al., 2010) (Chapter 12)

Figure 6 can serve as a reference when we want to simulate a 'socio-something' system. It can be seen as an elaboration of the left side of figure 2. In the middle we see time pass as the 'stream of life' that can be rendered as an agent-based model. Human action in it will be in the form of partly instrumental, partly ritual practices. The possible technical, environmental, and other instrumental sides of that system are not shown, but will have to be modelled, probably as environment characteristics. The agents that act in the system will have their individual personalities and drives, will learn through experience, and adapt their goals and behaviour. A very similar patterns holds at societal level: institutions, in our case perhaps formal organizational hierarchies, or safety legislation, are informed by culture, applied in practices, leading to potential learning and adaptation of these institutions. So, the agents in our simulation will need to hold cultural values that affect the rules of the institutions that they operate in. Figure 6 also allows for self-organization of what happens in the 'stream of life', and of which patterns result. For instance, in a hierarchical culture, workers are unlikely to communicate of their own initiative about work processes to superiors, regardless of formal rules about reporting in the company; this kind of phenomenon could be elegantly modelled in such a framework. Figure 6 does not offer any specific modelling guidance; it merely sketches a worldview that allows creating models with artificial sociality.

In recent years researchers have begun publishing on artificial mind and artificial sociality. We shall mention some highlights per level of aggregation.

Artificial mind

Humans are cultural beings in essence: our brains are plastic and we are programmed for life during childhood. As a result, far from being uniform, even very generic psychological attributes of people are culture-dependent (Smith, Bond, & Kagitcibasi, 2006). Across cultures, we do not perceive the world in the same way. This is reflected in norms and values, such as the relative importance of obedience versus initiative that is vital for employee behaviour in high-tech environments. Another example is the social contract, or lack of it, between citizens and government. This kind of empirical fact is indispensable in making sense of cross-cultural experience. Yet has as yet been very little used for designing artificial minds. How culture affects mental activity in individuals is largely *terra incognita*, relevant for agents with cognitive abilities (Dignum et al., 2014; G. J. Hofstede, 2018).

Artificial sociality

At group and society level, artificial sociality includes attempts to situate agents in their social world (figure 7).

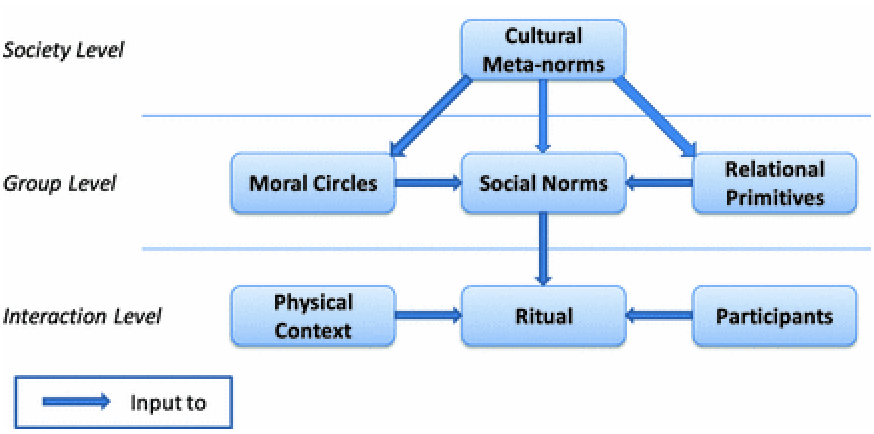

Figure 7: conceptual model for socio-cultural agents (Degens et al., 2014)

Figure 7 has been used in the development of TRAVELLER (Mascarenhas et al., 2015), a virtual agent world for cross-cultural training. It reads as follows: The place where agents interact is the ritual, where participants (agents, perhaps also human users of the system) meet in a physical context that differs from case to case. The minds of the agents share group level primitives: they have social norms per group. A group with norms is called a 'moral circle', indicating shared rights and obligations (G. J. Hofstede, 2009). These norms will be about 'relational primitives' such as respect. All these factors are themselves informed by the cultural values of society (Mascarenhas et al., 2015; Mascarenhas et al., 2013; McBreen, DiTosto, Dignum, & Hofstede, 2011).

At group level, cultural values can mediate emergent patterns of group formation (G. J. Hofstede, 2013; G. J. Hofstede, Dignum, Prada, Student, & Vanhée, 2015; G. J. Hofstede, Student, & Kramer, 2018) and ritual (G. J. Hofstede, Mascarenhas, & Paiva, 2011). Enabling this in ABM is pioneering, quarter making work. These ideas need to be given sample applications, and to be used in real-world cases. A recent and promising result is a meta-model for the relational side of human social behaviour called GRASP (for Groups, Rituals, Affiliation, Status, Power); (G. J. Hofstede, 2017). This framework allows dealing with social identity issues between groups. Such issues are often hidden causes of unexpected behaviour of socio-something systems. In a rapidly globalizing world this is an urgent issue.

Research priorities

Some groundwork for artificial sociality has been laid, as we saw. Three aspects have to be improved on, in order to create operable knowledge for supporting policy making on complex socio-something systems. These are: individual agent minds, emergence in social aggregates, and case-based work.

Artificial mind

A meta-model is needed for human social behaviour, based on solid social scientific theory about the relational world, linked to the instrumental world, usable in ABM. In fact, different applications will call for a family of meta-models, some of which lay more focus on certain elements, depending on the world they model and the model's aim. For instance, cross-cultural situations may require their own extra model constructs. Artificial sociality should have usable individual-level concepts of mind and volition, and it should have usable group-level concepts of social contagion that can give rise to self-organisation. At system level this yields emergent patterns.

Artificial sociality

Methodology is needed for modelling the dynamics of socio-something systems, that satisfies three conditions:

- Allowing for heterogeneity between elements. ABM makes this easy. Agents can have different personalities, relationships, social status, and roles. The environment can also be heterogeneous.
- Allowing for complexity. An agent-based model facilitates this in that it includes at least two levels of analysis: the agents and their actions, and the system-level output patterns. The analytical framework (figure 5) distinguishes these two levels. Models should allow relational motivation at individual level to yield culturally specific patterns at group level.
- Allowing not just different individuals, but also several different groups, so that issues of social identity and group affiliation can be tackled.

The GRASP meta-model mentioned above will serve as a point of departure for work on all three levels. GRASP has the potential of consolidating artificial sociality, but it still has to prove itself in applications. The link between the instrumental and relational world in particular, that is mediated by rituals, requires work that is application-specific. Figure 8 shows a GRASP world without any instrumental side.

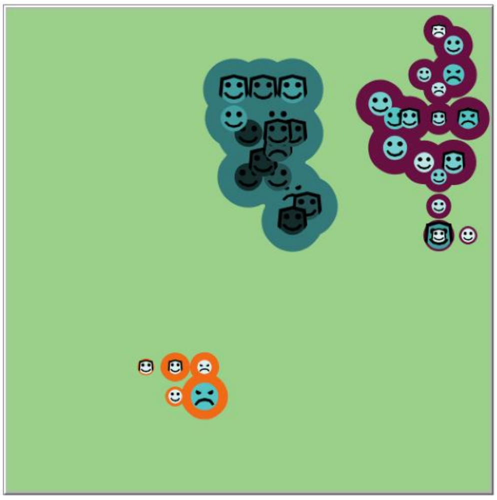

Figure 8: impression of 'Norm-Value Arena', a work-in-progress GRASP agent-based model, at the end of a session. Agents meet in groups and stay as long as they are happy. There are three groups of agents, each with its aura colour. Agents show current happiness through their mouth. Size connotes social status. Shade stands for norms; they have self-organized along lines of similar norms. The size of the aura indicates how long an individual has been in that group.

Case base

A variety of applications in different contexts can be useful, provided there is enough cross-fertilization between the artificial sociality concepts used in those cases. The case base will also help to create 'Good Modelling Practice', that is, a methodological knowledge base.

Cases are also the best vehicles for translating these abstract ideas into forms that non-specialists will understand. I want these theories to prove themselves in practice. This is why cases used for policy making are likely the most appropriate: they do not require building perfect predictive models, merely models that will allow improvement in policy over current practice.

Not all cases will use GRASP for the social part. Other, simpler meta-models can also be used, depending on modelling aim and available data. This can be a matter of expediency, for instance because a meta-model has a track record in an area; and it can be useful for triangulation purposes.

One important challenge is marrying the substantive and instrumental aspects of cases with the relational meta-model. Here the work of Searle on construction of social reality can serve, in particular his ‘counts-as' concept (Searle, 1995). The modeller can use this to convert case-specific elements into their relational impact.

Action plan

I will normally hold my current job until 2024. I intend to use these six years for building critical research mass, education scope, and capacity.

Aim for 2024

- Artificial sociality is a normal ingredient of ABM for modelling socio-technical, socio-economical or socio-ecological systems
- The GRASP meta-model for agents is available as a plug-in meta-model in one or more showcase agent-based models. It is used in research at many places in the world, a.o. for models that support policy making.
- ABM is one of the available methods all around WUR, and the INF group can, through Silico Centre, give methodological support about ABM along with Biometris.
- The Socio-technical systems track flourishes at the INF group, and has found successful integration with the software engineering and knowledge engineering tracks.
- ABM has become better integrated in optional parts of curricula at MsC level.
- Others can carry on where I leave off at retirement.

In order to reach these aims I will become more visible and more active in ‘pulling the cart'. New hires, Silico activities, and an education portfolio that matches my research agenda, are now creating proper conditions. Also my work is creating more traction, and given its urgency I expect this to expand in the coming years.

Research priorities

- Improve modelling ontology for human social behaviour in socio-something systems.
    - Heart of the social side is GRASP and whatever improvements I, and others, can make to it, in discussion with similar meta-models such as MoHub (Schlüter et al., 2017).
    - Create foundational model for socio-technical systems with multiple people in more than one group, and multi-unit technology. This links to the idea of ‘systems of systems' that is one of the current trends in IT.
    - Use this model to study unintended side effects of changes in technology, given certain social structures in the system.
- Methodologically refine ABM method
    - Create a typology of patterns that ABM of socio-something systems can yield.
    - Work on matching modelling aims and validation methods. How to validate depends on the model's aim: illustrate, understand, describe, predict?
- Develop case base
    - Work on cases in industrial safety in my role of extraordinary professor at Northwestern University, Johannesburg
    - Work on cases in resilience with various WASS groups and with WECR
    - Use the case-based tracks to practice-test the modelling ontology

- Make use of available data sources, ‘big’ or other, to empirically ground the models
- Seek ontological integration across cases.

Themes

Currently I am concentrating on the following themes.

- Safety
  Safety, whether food safety or industrial safety, is of obvious societal importance. It is also invariably hyped, and linked to social identity, due to incidents and their emotional aftermath that includes inter-group dynamics and politics in culture-dependent ways. In academic terms, it has strong emergent dynamics, due to its association with contagious social motives such as fear and anxiety. Agent-based models can lengthen the time and expand the scope of investigation, both in prevention and in curative measures.
- Resilience
  Resilience is concerned with the capacity of a system to keep providing certain output despite pressure. It is essential for human survival, and it is an important strategic theme at Wageningen University. It also has strong emergent aspects, due to its common good characteristics; there are usually conflicts between individual and system interest, and between now and later, when resilience is the aim.

Around both themes, rapid automation, virtualization, and robotization are happening, for instance in ‘smart’ cities. Citizens may be involved in generating data. Models are often useful in discussions with stakeholders.

I certainly do not exclude that other themes will pop up in the coming years, and will apply my methodological expertise where relevant.

Sources of funding

- WASS: I will try to get PhD projects funded on modelling and simulation of complex socio-technical or socio-something systems. This will often take the form of joint projects between INF and another WASS group. New links are possible with groups that have empirical data, and disciplinary knowledge, but could have more grip on systemic dynamics of their research field.
- WASS: I am involved in preparations for acquisition and ABM development in several Resilience Brigades initiated under WU’s Strategic Plan’s Resilience theme: REFORIM (with a.o. BEC: Miranda Meuwissen, PAP: Katrien Termeer, WECR: Krijn Poppe), ABM Meets Tourism (with a.o. ENP: Machiel Lamers, ESA: Bas Amelung), TRECA (with a.o. CPT: Peter Feindt, LAG: Kai Purnhage).
- WU: I currently collaborate with all knowledge units at WU: Animal (APS: Imke de Boer), Plant (Biometris: Jaap Molenaar, Hilde Tobi, George van Voorn), CSA: Paul Struik, FSE: Jeroen Groot, Louis Bolk Institute: Edith Lammerts van Bueren) Environment (Biodiverse Environment Programme: Lawrence Jones-Walters, GIS: Arend Ligtenberg, SPL), Technology (FQD: Vincenzo Fogliano), Social (BEC, ORL, MST, PAP, SCG, UEC, WEcR). Silico Centre has a University-wide scope, and so has the biannual Summer School on Agent-Based Modelling. The new course ‘ABM for busy people’ is drawing people from both within WUR and elsewhere. I expect to derive projects for building up my case base from these relationships.
- Netherlands: I am now active in TKI-Dinalog project DaVinc3i (PostDoc: Giulia Salvini), and will remain active in acquisition with N W O and STW. The sectors of logistics and horticulture are the most likely application context of these acquisitions. Methodological

work can be carried out in 4TU, where I am involved in a proposal around ABM for resilience.
- Europe: I am regularly involved in EU acquisition (In 2017 the H2020 project MIMOSA on ABM in migration issues just failed to be funded).
- World: as an extraordinary professor at Northwestern University, Johannesburg, I expect to be able to find funding for work in South Africa related to industrial safety.
- Industry: as systemic thinking and ABM start to receive interest from industry, I am increasing my contacts with potential funding sources in industry. Currently I am involved in a PhD with the Food Technology & Design group funded by industry. I am also using artificial sociality thinking in my role in the scientific council of Seedlinktech.

In 2018

I will use my sabbatical (1 September – 1 December 2018) to

- Develop a generic GRASP ABM
- Visit Shanghai to work with Seedlinktech on links between personality profiles and longitudinal team performance
- Visit INESC-ID, Portugal, for work on social agents and GRASP
- Visit Johannesburg in the context of my extraordinary professorship for work on socio-technical systems in industrial safety, applying GRASP and acquiring case research contacts.

Why me?

One cannot build an agent-based model without mastering all its elements to some degree. As a population biologist, experienced software designer, social scientist specialized in culture, and designer of simulation games and agent-based models, I fit the bill. Specific competencies I have are

- I have a comprehensive grasp of human social behaviour.
    - The world over, researchers and other people ask me for assistance on culture. I am a co-author to the book *Cultures and Organizations* that has appeared in twenty-odd languages, the first author of *Exploring Culture* that appeared in five*,* and I created a large number of simulation games on culture in organizational settings. Later I worked on trust and transparency across cultures. This has resulted in me being sought out for counsel, supervision, lecturing and keynotes on all continents.
    - I can connect to a wide variety of audiences in business, non-profit or academia
    - I have a unique take on dynamics and evolution of culture (G. Hofstede et al., 2010)
- I am experienced in suitable modelling techniques:
    - Simulation gaming. I introduced the concept of 'Synthetic Cultures' for cross-cultural simulation games (G. J. Hofstede & Pedersen, 1999; G. J. Hofstede, Pedersen, & Hofstede, 2002) and investigated what makes simulation games work in practice (G. J. Hofstede, Caluwé, & Peters, 2010).
    - Computational social simulation, ABM. Besides the content I developed, discussed in this vision, I made a case for using this method in cross-cultural and strategic management (G. J. Hofstede, 2015).
- I laid the groundwork for an operable theory on human social behaviour for use in ABM during NIAS fellowship.
- I have a wide international interdisciplinary network in academia and industry.
- I am well connected through PhD and other projects in all knowledge units of WU.

Culture in artificial social agents

My special expertise in culture, that I have harnessed for modelling, may be my strongest unique selling point. In the context of artificial sociality, culture translates into norms and values that differ per cultural group. In terms of the GRASP meta-model, one could call culture the 'small print' of GRASP rules. Let me give a few examples. Culture is about e.g. the unit of status worthiness. If one person misbehaves, is that occasion for individual guilt, or for shame of the entire community? Culture is also about the relative status worthiness of conferring status (being nice) versus showing power (being strong): is there sympathy for a strong-handed leader, or for the victims of that leader? As a third example, culture is also about the elaborateness of ritual needed to achieve changes in groups. Does the establishment of a new business relation require a phone call, a meal, or perhaps a family relationship? Using this knowledge about cross-cultural differences, in which theory on culture is linked to stylized casuistry, culture in agents can be built on top of a GRASP-like basis. In socio-technical systems this allows to study how the 'same' system, with the same technology and formal organizational structure, can differ in behaviour when populated by agents with different culture.

Artificial sociality at the INF group

Artificial sociality in agent-based models, like other branches of information technology, requires conceptual ontologies and computational models. Its natural home in Wageningen is the INF group. The fact that INF group is part of the social sciences makes it even more suitable, since connections with social scientific input are easy to make.

Currently, ABM is mainly in the hands of Mark Kramer, Sjoukje Osinga and myself. Other INF members could step in as needed, e.g. Ioannis Athanasiadis and Vahid Garousi, both of whom have worked with ABM. PhD students, of which I currently have seven, are usually shared with an application area-oriented group, while INF provides methodological and modelling guidance.

The VSNU recently indicated the urgency of an expanded mission for IT. I quote: "To ensure effective connections between technology and society, new knowledge of technology and natural sciences must be properly integrated with knowledge of, for instance, economic, social science-related, psychological, cultural and political factors." (VSNU 2016). If the INF group is to keep up with the wave of AI sweeping through science and society, then artificial sociality should be a vital component of its activities.

My impact at Wageningen

Wageningen university is engaged with all of Gaia. It has specialized into five knowledge units in accordance with figure 1: nature is represented by plant, animal, and spatial sciences; culture is represented by social sciences; technology is represented by technological sciences. It remains one faculty, and the current credo 'OneWageningen' expresses a desire not to fragment. We are thus well placed for tackling any subsystem of Gaia using agent-based models that integrate all the disciplines that are relevant to the system under study.

It is my conviction that if Wageningen is to successfully address today's worldwide challenges, more research is needed that uses a system perspective embracing not only nature and technology, but also culture. Culture pushes both nature and technology into directions that are unanticipated and often undesirable. We need to understand its impact better in models of subsystems of Gaia. Dealing with Gaia subsystems as wholes is a challenge that I can help Wageningen to meet.

Acknowledgements

Over the course of development I received comments from many colleagues: Ioannis Athanasiadis, Jacqueline Bloemhof, Ayalew Kassahun, Eveline van Leeuwen, Miranda Meuwissen, Sjoukje Osinga, Bedir Tekinerdogan, Emely de Vet.

References

Axelrod, R. (1997). The Dissemination of Culture; A model with local convergence and global polarization. *Journal of Conflict Resolution, 41*(2), 203-226. doi:10.1177/0022002797041002001

Brown, R. (2000). *Group Processes: Dynamics within and between groups* (2nd ed.). Malden, MA: Blackwell.

Castelfranchi, C. (1998). Modelling social action for AI agents. *Artificial Intelligence*(103), 157-182.

Choi, K. S., Im, I., & Hofstede, G. J. (2016). A cross-cultural comparative analysis of small group collaboration using mobile twitter. *Computers in Human Behavior, 65*, 308-318. doi:https://doi.org/10.1016/j.chb.2016.08.043

Coleman, J. S. (1990). *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.

Conte, R., Gilbert, N., Bonelli, G., Cioffi-Revilla, C., Deffuant, G., Kertesz, J., . . . Helbing, D. (2012). Manifesto of computational social science. *The European Physical Journal Special Topics, 214*(1), 325-346. doi:10.1140/epjst/e2012-01697-8

Deffuant, G., Amblard, F., Weisbuch, G., & Faure, T. (2002). How can extremism prevail? A study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation, 5*(4).

Degens, D. M., Hofstede, G. J., McBreen, J., Beulens, A. J. M., Mascarenhas, S., Ferreira, N., . . . Dignum, F. (2014). Creating a World for Socio-Cultural Agents. In T. Bosse, J. Broekens, J. Dias, & J. van der Zwaan (Eds.), *Emotion Modeling: Towards Pragmatic Computational Models of Affective Processes* (Vol. LNAI 8750, pp. 27-43). Zürich: Springer.

Dias, J., Mascarenhas, S., & Paiva, A. (2016). FAtiMA Modular: Towards an Agent Architecture with a Generic Appraisal Framework. In *Emotion Modeling* (Vol. LNCS 8750, pp. 44-56): Springer.

Dignum, F., Hofstede, G. J., & Prada, R. (2014). *From autistic to social agents.* Paper presented at the AAMAS 2014, Paris.

Edmonds, B., & Meyer, R. (2017). *Simulating social complexity: a handbook*: Springer.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: social science from the bottom up*. Washington D.C.: The Brookings Institution.

Fishbein, M., & Ajzen, I. (2010). *Predicting and Changing Behavior: the Reasoned Action Approach*: Psychology Press, Taylor & Francis.

Georgeff, M., Pell, B., Pollack, M., Tambe, M., & Wooldridge, M. (2003). The Belief-Desire-Intention Model of Agency. In (Vol. LNCS 1555, pp. 1-10): Springer.

Gilbert, N. (2008). *Agent-based models* (Vol. 153): Sage.

Gilbert, N., & Conte, R. (1995). *Artificial Societies: The Computer Simulation of Social Life* University College London Press.

Gratch, J., Wang, N., Gerten, J., Fast, E., & Duffy, R. (2007). *Creating Rapport with Virtual Agents.* Paper presented at the Intelligent Virtual Agents, Paris.

Grow, A., Flache, A., & Wittek, R. (2015). An Agent-Based Model of Status Construction in Task Focused Groups. *Journal of Artificial Societies and Social Simulation, 18*(2), 4.

Heise, D. R. (2013). Modeling Interactions in Small Groups. *Social Psychology Quarterly, 76*(1), 52-72. doi:10.1177/0190272512467654

Hofstede, G., Hofstede, G. J., & Minkov, M. (2010). *Cultures and Organizations, Software of the Mind* (3rd ed.). New York: McGraw Hill.

Hofstede, G. J. (2009). The Moral Circle in Intercultural Competence: Trust Across Cultures. In D. K. Deardorff (Ed.), *The SAGE Handbook of Intercultural Competence* (pp. 85-99). Los Angeles, CA: Sage.

Hofstede, G. J. (2013). *Theory in social simulation: Status-Power theory, national culture and emergence of the glass ceiling.* Paper presented at the AISB 2013 Social.Path, Exeter.

Hofstede, G. J. (2015). Culture's causes: the next challenge (Distinguished scholar essay). *Cross Cultural Management: an International Journal, 22*(4), 545-569.

Hofstede, G. J. (2017). GRASP agents: social first, intelligent later. *Ai & Society*, 1-9. doi:https://doi.org/10.1007/s00146-017-0783-7

Hofstede, G. J. (2018). Mental Activity and Culture: The Elusive Real World. In C. Faucher (Ed.), *Advances in Culturally-Aware Intelligent Systems and in Cross-Cultural Psychological Studies* (pp. 143-164). Cham: Springer International Publishing.

Hofstede, G. J., Ambrosius, F. H. W., Bokkers, E., & Boumans, I. (2014). *Who's friends, who's boss? Affiliation and hierarchy in multi-agent systems.* Paper presented at the Workshop Social.Path 2014, Computational Social Science and Social Computer Science: Two Sides of the Same Coin, 23-24 June, University of Surrey, Wates House (Treetops), Guildford, UK, Guildford.

Hofstede, G. J., Dignum, F., Prada, R., Student, J., & Vanhée, L. (2015). Gender differences: the role of nature, nurture, social identity and self-organization. In F. Grimaldo & E. Norling (Eds.), *Multi-Agent-Based Systems Workshop (MABS) 2014, LNAI 9002* (pp. 72-87): Springer.

Hofstede, G. J., Jonker, C. M., & Verwaart, T. (2012). Cultural Differentiation of Negotiating Agents. *Group Decision and Negotiation, 21*(1), 79-98. doi:10.1007/s10726-010-9190-x

Hofstede, G. J., Mascarenhas, S., & Paiva, A. (2011). *Modelling rituals for Homo biologicus*. Paper presented at the ESSA 2011 (7th Conference of the European Social Simulation Association), Montpellier. www.gertjanhofstede.com/research/

Hofstede, G. J., & Pedersen, P. B. (1999). Synthetic Cultures: Intercultural Learning Through Simulation Gaming. *Simulation & Gaming, 30*(4), 415-440.

Hofstede, G. J., Pedersen, P. B., & Hofstede, G. (2002). *Exploring Culture: Exercise, Stories and Synthetic Cultures*. Yarmouth, Maine: Intercultural Press.

Hofstede, G. J., Student, J., & Kramer, M. R. (2018). The status–power arena: a comprehensive agent-based model of social status dynamics and gender in groups of children. *Ai & Society*, 1-21. doi:10.1007/s00146-017-0793-5

Jager, W., & Janssen, M. A. (2012). *An updated conceptual framework for integrated modelling of human decision making: The consumat II*. Paper presented at the ECCS, Brussel.

Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Riemsdijk, M. B. v., & Sierhuis, M. (2014). Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction, 3*(1), 43-69. doi:DOI 10.5898/JHRI.3.1

Kaminka, G. A. (2013). *Curing Robot Autism: A Challenge.* Paper presented at the Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Minnesota, USA.

Kemper, T. D. (2011). *Status, Power and Ritual Interaction; A Relational Reading of Durkheim, Goffman and Collins*. Farnham, UK: Ashgate.

Lamperti, F., Roventini, A., & Sani, A. (2017). Agent-Based Model Calibration using Machine Learning Surrogates. *HAL Archives Ouvertes, hal-01499344*.

Latour, B. (2017). *Facing Gaia: Eight Lectures on the New Climatic Regime*. Cambridge, UK: Polity Press.

Mascarenhas, S., Degens, D. M., Paiva, A., Prada, R., Hofstede, G., Beulens, A., & Aylett, R. (2015). Modeling culture in intelligent virtual agents. *Autonomous Agents and Multi-Agent Systems*, 1-32. doi:10.1007/s10458-015-9312-6

Mascarenhas, S., Prada, R., Paiva, A., & Hofstede, G. J. (2013). Social Importance Dynamics: A Model for Culturally-Adaptive Agents. In R. Aylett, B. Krenn, C. Pelachaud, & H. Shimodaira (Eds.), *Intelligent Virtual Agents* (Vol. 8108, pp. 325-338): Springer.

McBreen, J., DiTosto, G., Dignum, F., & Hofstede, G. J. (2011). *Linking Norms and Culture.* Paper presented at the 2nd International Conference on Culture and Computing, Kyoto.

McClelland, D. C. (1987). *Human Motivation*: Cambridge University Press.
McCrae, R. R., & Costa, P. T. j. (2003). *Personality in Adulthood: A Five-Factor Theory Perspective* (2nd ed.). New York: Guildford.
Moss, S., & Edmonds, B. (2005). Towards good social science. *Journal of Artificial Societies and Social Simulation, 8*(4).
Ortony, A., Clore, G., & Collins, A. (1998). *The Cognitive Structure of Emotions*. Cambridge, UK: Cambridge University Press.
Piccoli, G. (2012). *Information Systems for Managers*: John Wiley & Sons.
Pinker, S. (2011). *The better angels of our nature: why violence has declined*: Viking Books.
Prada, R., & Paiva, A. (2005). *Believable groups of synthetic characters.* Paper presented at the AAMAS, fourth international joint conference on autonomous agents and multi-agent systems.
Rand, W. (2006). Machine learning meets agent-based modelling: when not to go to a bar. Retrieved from https://ccl.northwestern.edu/2006/agent2006rand.pdf
Roozmand, O., Ghasem-Aghaee, N., Hofstede, G. J., Nematbakhsh, M. A., Baraani, A., & Verwaart, T. (2011). Agent-based modeling of consumer decision making process based on power distance and personality. *Knowledge-Based systems, 24*(7), 1075-1095.
Schlüter, M., Baeza, A., Dressler, G., Frank, K., Groeneveld, J., Jager, W., . . . NandaWijermans. (2017). A framework for mapping and comparing behavioural theories in models of social-ecological systems. *Ecological Economics, 131*, 21-35.
Searle, J. R. (1995). *The Construction of Social Reality*. London: Penguin.
Simon, H. (1996). *The sciences of the artificial, third edition*: MIT Press.
Smith, P. B., Bond, M. H., & Kagitcibasi, C. (2006). *Understanding Social Psychology Across Cultures* Sage.
Squazzoni, F., Jager, W., & Edmonds, B. (2013). Social Simulation in the Social Sciences: A Brief Overview. *Social Science Computer Review, 31*, 1-16.
ten Broeke, G., van Voorn, G., & Ligtenberg, A. (2016). Which Sensitivity Analysis Method Should I Use for My Agent-Based Model? *Journal of Artificial Societies & Social Simulation, 19*(1), 1-35.

(VSNU), V. v. S. N. U. (2016). The Digital Society. from http://www.vsnu.nl/files/documents/Publications/VSNU_The_Digital_Society.pdf