Delft University of Technology

Social agents: Bridging simulation and engineering
Seeking better integration of two research communities

Dignum, Virginia

DOI
10.1145/3148265
Publication date
2017
Document Version
Accepted author manuscript
Published in
Communications of the ACM

Citation (APA)
Dignum, V. (2017). Social agents: Bridging simulation and engineering: Seeking better integration of two research communities. *Communications of the ACM*, *60*(11), 32-34. https://doi.org/10.1145/3148265

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

Social Agents: Bridging Simulation and Engineering

Virginia Dignum

December 23, 2016

The use of the agent paradigm to understand and design complex systems occupies an important and growing role in different areas of social and natural sciences and technology. Application areas where the agent paradigm delivers appropriate solutions include online trading [16], disaster management [10], and policy making [11]. However, the two main agent approaches, Multi-Agent Systems (MAS) and Agent-Based Modeling (ABM) differ considerably in methodology, applications and aims. MAS focus on solving specific complex problems using autonomous heterogeneous agents, while ABM is used to capture the dynamics of a (social or technical) system for analytical purposes. ABM is a form of computational modeling whereby a population of individual agents are given simple rules to govern their behavior such that global properties of the whole can be analyzed [9]. The terminology of ABM tends to be used more often in the social sciences, whereas MAS is more used in engineering and technology.

Although there is considerable overlap between the two approaches, historically the differences between ABM and MAS are often more salient than their similarities. For example, it is often remarked that a main difference between ABM and MAS is that ABM models are *descriptive* aiming at explanatory insight into the collective behavior of agents at the macro level, whereas MAS are *operational* systems, acting and affecting its (physical) environment, with a focus on solving specific practical or engineering problems, and emphasizing agent architectures with sophisticated reasoning and decision processes. This has lead to the development of two research communities proceeding on nearly independent tracks.

However, this division is not as black and white as it may seem. In fact, much ABM work goes beyond descriptive simulations of a situation, and, as input for decision-making and policy-setting, indirectly affects the environment. And, the design of MAS is often geared to analytic insights and simulations towards the understanding of how configurations of agents behave in different circumstances. Currently, applications of MAS are broader than pure distribution problems, including interactive virtual characters, where the focus is on the cognitive, affective and emotional characteristics of the system, and game-theoretic models, focusing on the design of incentive mechanisms that guarantee a given strategic behavior.

Social abilities are central both in ABM, where agents represent humans and their interactions, and in MAS, that enable game-theoretic analyses of decision strategies, or provide interactive virtual agents in varied situations. It is precisely in this area that the need for integration of ABM and MAS is undoubt-

edly the most necessary. In social simulation, the benefits of combining MAS and ABM have been advocated for many years, and are the focus of the long-lasting workshop series on Multi-Agent Based Simulation (MABS) [2]. ABM has increasingly and successfully been used for social simulations [3], but it is in the MAS area that fundamental research on agent architectures implementing psychological traits and social concepts such as norms, commitments, emotions, identity, and social order, has been most prominent [5, 4]. Bridging these somewhat parallel tracks, requires a new grounding for agent architectures.

1 Questioning Rationality

Traditionally, one of the most salient aspects shared by both ABM and MAS approaches is the premise of *rationality*. This is derived from the traditional definition of agents as autonomous, proactive and interactive entities where (a) each agent has bounded (incomplete) resources to solve a given problem, (b) there is no global system control, (c) data is decentralized, and (d) computation is asynchronous [21]. Agent rationality can be summarized as follows:

- Agents hold consistent beliefs;
- Agents have preferences, or priorities, on outcomes of actions;
- Agents optimize actions based on those preferences and beliefs.

This view on rationally entails that agents are expected, and designed, to act rationally in the sense that they choose the best means available to achieve a given end, and maintain consistency between what is wanted and what is chosen [14]. Even though multiple alternatives have been proposed, in both the ABM and MAS approaches, individual agents are still typically characterized as bounded rational, acting towards their own perceived interests. The main difference is that agent behaviors in ABM are used to capture the dynamics of a system for analytical purposes, grounded whenever possible on existing data about system outcomes, whereas MAS focuses on solving specific problems using independent agents, through the formalization of the complex goal-oriented processes, such as the Beliefs-Desires and Intentions (BDI) model proposed by Bratman (cf. [20]) or game-theoretic approaches.

The main advantages of such rationality assumptions are their parsimony and applicability to a very broad range of situations and environments, and their ability to generate falsifiable, and sometimes empirically confirmed, hypotheses about actions in these environments. This gives conventional rational choice approaches a combination of generality and predictive power not found in other approaches. In fact, rationality approaches are the basis of most theoretical models in the social sciences, including economics, political science, or social choice theories.

Unfortunately, from a modeling perspective, real human behavior is neither simple nor rational, but derives from a complex mix of mental, physical, emotional and social aspects. Realistic applications must consider situations in which not all alternatives, consequences, and event probabilities can be foreseen. This renders rational choice approaches unable to accurately model and predict a wide range of human behaviors.

2 Towards Social Agents

Human sociability refers to the nature, quantity and quality of interactions with others, including both pro-social, or cooperative, behaviours, and conflict, competitive, or dominating behaviours. Sociability is also the ability to influence others, by changing their behaviours, goals and beliefs, the emotional reaction to others and to the environment, and how actions are affected by emotions, and the ability to create, structure and 'rationalize' the environment to fit ones expectations and abilities (leading e.g. to the design of organizations, institutions, and norms)

Following an increasing number of researchers in both ABM and MAS that in recent years have come to similar conclusions [13, 7, 18, 19], we claim that new models of preference and belief formation are needed that show how behavior derives from identities, emotions, motivation, values and practices [6]. The endeavor required to construct such agent models that are socially realistic requires the effort and the capabilities of both the MAS and ABM communities, bringing together formalization and computational efficiency, and planning techniques as in MAS, with the ABM expertise on empirical validation and on adapting and integrating social sciences theories into a unified set of assumptions [1], furthering the fundamental understanding of social deliberation processes, and developing techniques to make these accessible for simulation platforms. This article is therefore an appeal to join the strengths of both communities towards *sociality-based agents*.

Without claiming a readily available solution, we propose the concept of *sociality* as the leading principle of agency, as an alternative for *rationality*. Following the above description of rational behaviour, the main characteristics of sociality-based reasoning are:

- Ability to hold and deal with inconsistent beliefs for the sake of coherence with identity and cultural background. That is, beliefs originate from other sources than observation, including ideology or culture.
- Ability to fulfill several roles, and pursue seemingly incompatible goals concurrently, e.g. simultaneously aiming for comfort and environmental-friendliness, or for riches and philanthropy.
- Preferences are not only a cause for action but also a result of action. Moreover, preferences change significantly over time and their ordering is influenced by the different roles being fulfilled simultaneously, which requires the need to deal with misalignment and incompatible orderings.
- Action decisions are not only geared to the optimization of own wealth, but often motivated by altruism, fairness, justice, or by an attempt to prevent regret at a later stage.
- Understand when there is no need to further maximize utility beyond some reasonably achievable threshold.
- Understand how identity, culture and values influence action, and use this knowledge to decide about reputation and trust about who and how to interact.

The first step towards sociality-based agents is a through understanding of these principles, and open discussion across disciplines on the grounds and requirements for sociality from different perspectives. This discussion will be fundamental to the development of formal models and agent architectures that make sociality-based behavior possible and verifiable.

Moreover, it is necessary to identify and formalize which mechanisms, other than imitation, can describe how agents can adapt to pressures in the environment to behave in a socially acceptable, resource-sustainable fashion. Resulting models support the understanding or predicting human behaviour, including rich models of emotions, identities, culture, values, norms, and many other socio-cognitive characteristics. Such models of social reality are also needed to study the complex influences on behavior of different socio-cognitive characteristics and their relationships. The integration of psychological models of motivation and cognition, sociological theories of value and identity formation, and philosophical theories of coherence and higher-order rationality, together with different formal methods, quickly yields intractable models. However, it is important to identify what is the model being developed for. In fact, richer models are not always the most appropriate ones.

Once these characteristics are well understood, then simplified models can be developed to suit different needs. That is, implementing sociality-based agents, will require other techniques than those currently used in either MAS or ABM [8], including the use of simpler, context-specific decision rules, mimicking how people themselves are able to deal with complex decision-making, e.g. using social practices as a kind of short-cuts for deliberation [15, 17]. Where it concerns utility, *satisficing* can be more suitable approach than maximizing [12]. This also allows to integrate agents of varied richness levels, e.g using rich cognitive models to *zoom-in* the behavior of salient agents in a simulation, whereas other agents just follow simple rules. This approach can counter the obvious criticism that sociality-based agents will become too complex for use in computational simulations.

Sociality-based agents are also fundamental to the new generations of intelligent devices, and interactive characties in smart environments. These artifacts not only must build (partial) social models about the humans they interact with, but also need to take social roles in a mixed human/digital reality. To me a interesting challenge would be to use the same technologies in real-time mixed human/artificial interactions, and criticisms could also be on the feasibility to use these architectures (or controlled reductions/simplifications) in real-time or near-real time.

3 Moving forward

The aim of this article is to appeal for a collaborative research effort towards fundamental formal theories and models that increase our understanding of the *principles* behind human deliberation (such as the ones listed above), before deciding on which modeling techniques we need to implement them. Even though, several approaches to model social aspects in agent behavior are available, there is not suffient consensus on which characteristics are needed for what, nor on how to specify and integrate them. The paper has identified an initial set of characteristics for sociability, proposed a research path linking theory, model,

and implementation, and suggested possible theories and techniques to develop sociality-based agents. These incorporate expertise from both ABM and MAS and require integration of both areas in order to succeed. We welcome the discussion of these issues towards a novel area of research on Social Agents, which take sociability as the basis for agent deliberation and enable interaction.

References

[1] S. Chai. *Choosing an Identity: A General Model of Preference and Belief Formation*. University of Michigan Press, 2001.

[2] R. Conte, N. Gilbert, and J. Sichman. MAS and social simulation: A suitable commitment. In J. Sichman, R. Conte, and N. Gilbert, editors, *Multi-Agent Systems and Agent-Based Simulation*, volume 1534 of *Lecture Notes in Computer Science*, pages 1–9. Springer Berlin Heidelberg, 1998.

[3] P. Davidsson. Agent based social simulation: A computer science view. *Journal of artificial societies and social simulation*, 5(1), 2002.

[4] J. Dias, S. Mascarenhas, and A. Paiva. Fatima modular: Towards an agent architecture with a generic appraisal framework. In *Proceedings of the International Workshop on Standards for Emotion Modeling*, 2011.

[5] F. Dignum, V. Dignum, and C.M. Jonker. Towards agents for policy making. In *MABS IX*, pages 141–153. Springer, 2009.

[6] F. Dignum, V. Dignum, R. Prada, and C.M. Jonker. A conceptual architecture for social deliberation in multi-agent organizations. *Multiagent and Grid Systems*, 11(3):147–166, 2015.

[7] F. Dignum, R. Prada, and G.J. Hofstede. From autistic to social agents. In *AAMAS 2014*, May 2014.

[8] V. Dignum. Mind as a service: Building socially intelligent agents. In V. Dignum, P. Noriega, M. Sensoy, and J. Sichman, editors, *COIN XI: Revised Selected Papers*, pages 119–133. Springer International Publishing, 2016.

[9] J.M. Epstein and R. Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. The Brookings Institution, Washington, DC, USA, 1996.

[10] F. Fiedrich and P. Burghardt. Agent-based systems for disaster management. *Commun. ACM*, 50(3):41–42, March 2007.

[11] A. Ghorbani, V. Dignum, F. Dechesne, and C. Jonker. Enhancing abm into an inevitable tool for policy analysis. *Policy and Complex Systems*, 1(1), 2014.

[12] G. Gigerenzer. Moral satisficing: Rethinking moral behavior as bounded rationality. *Topics in cognitive science*, 2(3):528–554, 2010.

[13] G. Kaminka. Curing robot autism: A challenge. In *AAMAS 2013*, pages 801–804, May 2013.

[14] S. Lindenberg. Social rationality versus rational egoism. In *Handbook of sociological theory*, pages 635–668. Springer, 2001.

[15] A. Reckwitz. Toward a theory of social practices. *European Journal of Social Theory*, 5(2):243–263, 2002.

[16] A. Rogers, E. David, N. Jennings, and J. Schiff. The effects of proxy bidding and minimum bid increments within ebay auctions. *ACM Trans. Web*, 1(2), August 2007.

[17] E. Shove, M. Pantzar, and M. Watson. *The Dynamics of Social Practice*. Sage, 2012.

[18] B. Silverman, D. Pietrocola, B. Nye, N. Weyer, O. Osin, D. Johnson, and R. Weaver. Rich socio-cognitive agents for immersive training environments: case of nonkin village. *Journal of Autonomous Agents and Multi-Agent Systems*, 24(2):312–343, March 2012.

[19] L. Vercouter, S. Casare, J. Sichman, and A. Brandão. An experience on reputation models interoperability based on a functional ontology. In *IJCAI'07*, pages 617–622, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

[20] M. Wooldridge. *Reasoning about rational agents*. MIT Press, 2000.

[21] M. Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2009.