

MAS and Social Simulation: A Suitable Commitment

Rosaria Conte¹, Nigel Gilbert², and Jaime Simão Sichman³

(1) Div. of AI, Cognitive and Interaction Modelling, Inst. of Psychology/CNR, Rome, Italy
rosaria@pssc2.irmkant.rm.cnr.it

(2) Department of Sociology, University of Surrey, Guildford, UK
N.Gilbert@surrey.ac.uk

(3) Department of Computer Engineering, University of São Paulo, São Paulo, Brazil
jaime@pcs.usp.br

Abstract. The goal of this introduction is to point out several similarities and differences between the research fields of multi-agent systems and social simulation. We show that these fields are complementary in several aspects, thus each one can benefit from results that emerge from the other. We finish the introduction by presenting and classifying the contributions in this volume.

1 Multi-Agent Systems and Social Simulation: Objective Affinities

The research fields of multi-agent systems and social simulation have some interesting points in common. We characterize each of the fields next, stressing their mutual influences in the last years.

1.1 Multi-Agent Systems

The field of Multi Agent Systems (MAS) is a well-established research and applied branch of AI, which has taken its impetus from the problems encountered in the implementation of tasks on distributed computational units interacting with one another and with the external environment (Distributed AI). A report on the results achieved within DAI, and a synthesis of the reasons underlying the development of the MAS field, is beyond the scope of this introduction (for a quite comprehensive picture, see O'Hare and Jennings 1996). Suffice it to say that distributed AI systems soon revealed a need for autonomy. The more autonomous the local units of the system from a central one, the more efficient the task distribution and execution, and the lower the computational load of the overall system. This discovery stimulated AI researchers and designers to turn their attention to intriguing and apparently philosophical issues, such as how to conceive of an autonomous system and how to design it. In turn, the development of autonomous systems brought about another perhaps even trickier question, i.e. how to obtain coordination and cooperation among autonomous systems executing a common task?

Application-oriented solutions to these questions have often been attempted (e.g. blackboard architectures, master-slave and benevolence assumptions; see Huhns, 1987). Nevertheless, in the last decade, the conceptual question of autonomy has increasingly become a focus of AI scientists' attention. This is shown by several scientific events: from mainly European gatherings such as the early Modelling Autonomous Agents in a Multi-Agent World (MAAMAW) events which characterized the MAS field in its early days (see, Demazeau and Mueller, 1990, Demazeau and Werner, 1991), a larger community has grown (e.g. the International Conference on Multi-Agent Systems, ICMAS; for the last one, see Demazeau, 1998).

The MAS field is increasingly characterized by the study, design and implementation of societies of artificial agents. Fruitful contributions are made by other AI sub-fields. Among these, one which deserves particular attention for its recent developments and increasing popularity is the Agent field with its highly reputed scientific events (the ATAL workshops, the Autonomous Agents Conference, etc.) and journals (the Journal of Intelligent Systems).

If the AI, logic-based and cognitive science approaches have contributed considerably to developments of MAS, the social sciences have exerted relatively less influence. An exception to this rule is offered by economics and game theory, which have rapidly invaded the MAS field (for a critical review, see Castelfranchi and Conte, 1998). The hegemony of these fairly specific areas of the social sciences on MAS is essentially due to the attention paid by economists and game-theorists to the study of the evolution of cooperation from local interactions among self-interested agents, also the quintessential problem of MAS scientists.

The role played by economics has prevented the MAS field itself from taking advantage of the whole range of theories, models, and conceptual instruments that abound in the social sciences and that have received a great impulse thanks to the spread of computer simulation.

1.2 Social Simulation

The computer simulation of social phenomena is a promising field of research at the intersection between the social, mathematical and computer sciences. The use of computer simulation in the social sciences ranges from sociology to economics, from social psychology to organization theory and political science, and from demography to anthropology and archaeology. The use of computers in some social scientific areas can be traced back to the fifties (Halpin, 1998). In its early days, and up to the seventies, computer simulation was essentially used as a powerful implementation of mathematical modelling (Troitzsch, 1997). More recently, computer simulation is more often used in its own right, "as a means of manipulating the symbols of programming languages" (Troitzsch, 1997: 41). Nowadays, the computer simulation of social phenomena and processes can be considered a well established field of research, as is witnessed by a large numbers of publications and scientific events and its own journal, the Journal of Artificial Societies and Social Simulation (for a review, see Gilbert and Troitzsch, 1999). In particular, in the last two decades, the field of

computer simulation has been able to benefit from a number of increasingly accessible facilities such as the development of high-level languages; the appearance of learning algorithms and systems; etc. DAI and the MAS have provided architectures and platforms for the implementation of relatively autonomous agents. This greatly contributed to the establishing of the agent-based computer simulation, an approach which has produced a vast body of simulation research, including rebuilding the Cellular Automata tradition, thanks to new technical and theoretical instruments (for a good example of simulation studies based on Cellular Automata modelling, see Hegselmann, 1996).

The agent-based approach enhanced the potentialities of computer simulation as a tool for theorizing about social scientific issues. In particular, the notion of an extended (multiple) computational agent, implementing cognitive capabilities (cf. Doran 1998), is giving encouragement to the construction and exploration of artificial societies (Gilbert and Conte, 1995; Epstein and Axtell, 1996), since it facilitates the modelling of artificial societies of autonomous intelligent agents.

If the MAS field can be characterized as the study of societies of artificial autonomous agents, agent-based social simulation can be defined as the study of artificial societies of autonomous agents. One could argue that the operation result should not be affected by the operators' order. However, the two fields are far from self-sufficient, as the following discussion will try to show. In particular, we shall argue that:

1. despite their evident affinities, the two fields in question have suffered and still suffer from an inadequate interface;
2. their cross-fertilisation would encourage research in both fields and at the same time stimulate innovative research arising at the intersection between them.

2 MAS and Social Simulation: An Unwarranted Gap

MAS and social simulation differ in terms of the formalism used (logic- and AI-based in the MAS domain, and mathematically based in the social simulation domain). But they also differ in other, more substantial ways.

2.1 Background Theory

Although decision and game theory have had a significant influence on both, theoretical differences between the two fields abound. MAS has inherited a large share of the AI and cognitive science conceptual and theoretical endowment, which entailed (a) long experience with the design and implementation of integrated architectures, rather than elementary automata; (b) a strong emphasis on the whole agent, rather than solely on its actions; (c) careful attention paid to the process of plan-construction, not just decision-making and choice; (d) familiarity with the normalization and implementation of agents mental, as well as their behavioral states; (e) a tendency to provide the social agent with specific capacities for actions answering social requests

and tasks (e.g., obligations, commitment and responsibility, etc.), rather than modelling social processes as mere emerging properties of agents' interaction.

The area of social simulation, benefited from the social sciences to a far greater extent than MAS. Among others, the following factors contributed to the field's progress: (a) a tendency to use computer simulation to test theoretical hypotheses, rather than the computational system's efficiency; (b) more familiarity with the interpretation of real-life social phenomena; this in turn implied (c) the production of vast bodies of data relative to artificial large-scale populations. All these features converged to consolidate the scientific methodological reputation of computer simulation, and lessen the toy-world character of its applications. Arising at the intersection of several social sciences, the field of social simulation could profit from their most recent and significant advances, such as (d) the development of the paradigm of complexity, facilitated by a close interaction with the sciences of physical and biological systems; and, in particular, (e) the development of theories, models and techniques for implementing and exploring social dynamics and evolution: social learning (cf. Macy and Flache, 1995), evolutionary game theory (cf. Weibull, 1996), cultural evolution (cf., for one example, Reynolds, 1994) and memetics (see the Journal of Memetics), etc.

2.2 Objectives

In the field of social simulation, objectives vary as a function of the discipline of interest. Applied objectives prevail in political, economical and management science, and especially in the science of organizations, which aims to optimize certain effects (for instance, international cooperation, resource allocation and distribution, or organizational performance). Other disciplines such as archaeology and anthropology are more clearly aimed at increasing scientific knowledge by formulating and testing interpretative models of existing phenomena through computational reconstruction. Overall, therefore, objectives vary between the purely normative in application-oriented disciplines like economics and the descriptive in interpretation-oriented disciplines like sociology, anthropology and archaeology.

In MAS, scientific objectives are unfortunately increasingly subordinate to producing software for various applications such as (micro-) robotics for manufacture and surgery, air traffic control and military defence. Nevertheless, this field has already contributed to improve our understanding of social intelligence and agent modelling by means of its architectural approach. By this is meant the integrated design of the modules, or "specialists", responsible for the different competencies involved in action. In particular, an intelligent autonomous agent architecture should respond in an intelligent adaptive way to a complex environment such as an agent society.

While it can hardly be doubted that social structures are created by interaction among agents, it is also true that the latter are in turn shaped by the social demands that they are supposed to meet. Good MAS theories (see, for the best-known example, the Belief-Desire-Intention (BDI) architecture proposed by Rao and Georgeff 1991)

have attempted to model the traits and competencies that enable autonomous agents to cooperate and coordinate with one another for a common task (citations abound; see any volume of proceedings of the ICMAS conference); elaborate and execute multi-agent plans; communicate with and influence one another; form and pursue collective intentions; commit themselves to a given action in a flexible intelligent way; assume and eventually abandon given social responsibilities; and so on.

2.3 Outstanding Issues

Despite the results achieved separately within the two fields, the potential of the computational study of social phenomena has not yet been fully exploited. Several questions are still unanswered and several points are still missing. Generally speaking, these fall into three related areas:

1. How should one combine a more sophisticated agent model and design with the simulation of qualitatively and quantitatively significant social phenomena? While multi agent systems are aggregates of small number of computational units, social simulators obtain data about large-scale populations. Is it possible to obtain a significant volume of data on, say, a BDI platform?
2. What is the role of the agent in the science of complexity? Should we accept the "trendy" view shared by physicists, mathematicians and some social scientists that complex systems can be described with the vocabulary and models of physics and biology only? Or should other vocabularies be developed to describe independently the various levels of complexity that are displayed in any reality of interest (be it natural or artificial)? And if the latter option is preferred, what are the specific vocabularies of social and mental complexity, and how should they be related to one another? In other words, how are the paradigm of emergence and the study of (social) intelligence to be related?
3. How is the two-way micro-macro link, from behaviors to social structures to be explained? The emergent paradigm is insufficient: it accounts for one direction only, namely from behaviors to social structures. What about the other direction? Cultural evolution and memetics try to give an account of the diffusion of so-called second-order emergent properties (mental representations of social constructs). Somehow this paradigm tries to model the way back from social structures to mental representations. However, how can this process be reconciled with agents' autonomy? Why, when and how do agents decide to form, accept and disseminate these representations? The agent's mind seems to play a mediator role here. How is the mediator role of cognition to be explained?

The above questions are intimately intertwined as is shown when one considers some more specific questions:

1. Agents make rational choices; they decide among alternatives for action in a way that is consistent with their internal criteria (rationality, utility maximisation, goal satisfaction, or any other). How do they form such alternatives? We need to integrate models of social decision-making with planning and problem solving.

2. This integration leads to another more crucial confrontation, between activation and goal-directed action, on one hand, and rationality and utility-maximisation, on the other hand. The social scientific view of the agent is strongly, if not uniquely, influenced by the economic interpretation of rationality. But this view has little to say about how agents should be constructed or how they concretely act. The mechanisms that activate agents are not yet clearly connected with the principles which govern action and decision-making.
3. The emergent paradigm aims to explain social convergence on given behavioral regularities (spatial segregation) on the grounds of a rather simplified characterization of the agents' motivations (for example, the attitude to be next to and imitate in-groups). The same characterization has been applied to explain the emergence and diffusion of conventions (Lewis, 1969). However, if social learning elucidates some mechanisms responsible for the spread of conventions, it does not seem to provide a sufficient account of innovation: how can new rules and conventions break the customary ones? How should one reconcile stability and innovation? Genetic algorithms bridge this hiatus by introducing mutation (for a discussion on different mechanisms for obtaining innovation, see DeJong and Spears, 1995) also the co-learning algorithm introduced by Shoham and Tennenholtz, 1994). However, the view of innovation as accidental mutation does not do justice to the agents' active role in the establishment of conventions. Agents' representations and interpretations seem to have a fundamental part in such a phenomenon.
4. This raises the more general problem of the connection between learning and reasoning. Social dynamics are not only an effect of learning and evolution. Social agents modify their behavior and that of their fellows by means of reasoning, planning and influencing. How should one combine the flexibility and social responsiveness of the agent with its autonomy and intelligence?
5. The emergent paradigms explain conventions in terms of the diffusion of behavioral regularities. Apparently, there is no relationship between conventions on one hand, and obligations, prescriptions, moral rules, etc. on the other. Therefore, either the existence of the latter is denied, or these two phenomena are claimed to be unrelated. However, the former is counter-intuitive. The latter is anti-economical: why should one ignore the relationships between phenomena so evidently intertwined? And, moreover, how is the "mandatory" character of conventions to be accounted for without a model of the agent where that character is somehow represented?

3 A Fruitful Cooperation

Cross-fertilisation among the two fields is needed. This would strengthen the agent-based approach in social simulation, add dynamics into MAS and reinforce simulation as a means of testing MAS systems. More specifically, a number of positive consequences can be expected from a closer interaction between the fields in question. MAS are likely to:

1. profit from the more refined and well-established theories, concepts and models of social organizations and institutions developed within the social sciences;
2. adopt the more dynamic approach shared by the social scientists using computer simulation;
3. acknowledge the importance of theory-driven computational studies even in addressing applied objectives;
4. import an approach to computer simulation from the social sciences which sees it as a tool for making and testing theories, rather than applications.

Social scientists interested in computer simulation, in their turn, are likely to:

1. give up both the static view of the agent as proposed by some rationality theories, and the behavioral view as proposed by theories of social learning,
2. refine their view of the agent and start to conceive of it as a computable although complex entity,
3. discover the role of the mind as a necessary intermediate between social structures and social behaviors;
4. familiarise themselves with more sophisticated agent architectures.

4 The Contributions in this Volume

Fifteen papers were presented at the first workshop on Multi-Agent systems and Agent-Based Simulation, which was organized by the editors of this volume and held as part of the Agents' World conference in Paris, July 4-6, 1998. Most of the social sciences were represented, with contributions touching on sociology, management science, economics, psychology, environmental science, ecology, and linguistics. There were a total of 69 registered participants. The workshop was organized in association with SimSoc, an informal group of social scientists who have arranged an irregular series of influential workshops on using simulation in the social sciences beginning in 1992.

Over 50 abstracts were received in response to the call for participation, but time constraints meant that only 16 could be accepted. The program included:

1. papers considering the value of agent-based approaches to social simulation compared with other techniques better known in the social sciences such as system dynamics and the implications for methodology (Parunak; Moss);
2. modelling the dynamics of markets (Terna) and the interactions between markets and natural resources such as tropical forests (Antona et al.);
3. explaining the emergent behavior of macro-level entities (e.g. groups and societies) from the actions of individuals (Sumpter, whose paper was on honey bee colonies; Hashimoto on the emergence of language; and Servat et al on water runoff processes);
4. group formation based on cultural evolution (Hales);
5. explorations of social dependency, agent interaction and organisational problem-solving (Conte and Pedone; Hannoun et al; Verhagen; Garrido et al);
6. and a study of how best to represent time in a multi-agent simulation (Fianyo et al).

While the papers were quite heterogeneous in substantive domain and in their disciplinary origins, there were several themes which recurred during the workshop. One of these was considered in more depth in a round table discussion led by Jim Doran at the end of the workshop on 'Representing cognition for social simulation', which addressed the issue of whether and how cognition should be modeled. Quite divergent views were expressed, with some participants denying that individual cognition needed to be modeled at all, and others arguing that cognition must be at the center of social simulation.

Another theme which was repeatedly mentioned in the presentations was the idea of 'emergence': that features observable at the level of a group or society should emerge from individual behavior (although Servat et al argued that it was appropriate to model macro-level features directly and then study ways in which individual agents could discover and adapt to those macro-level features).

Multi-Agent Based Simulation may be a good occasion for social scientists, on one hand, and AI and cognitive scientists, on the other, to meet on a level ground. Let us hope this volume makes the interested reader willing to take advantage of such an occasion.

References

- Castelfranchi, C. and Conte, R. (1988) Limits of economic rationality for agents and MA systems. *Robotics and Autonomous Systems*, Special issue on Multi-Agent Rationality, Elsevier Editor, (in press).
- DeJong, K. and Spears, W. (1995) On the state of evolutionary computation, *Proceedings of the Sixth Conference on Genetic Algorithms*, 618-23.
- Demazeau, Y. (ed.) (1998) *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS 98)*. IEEE Press.
- Demazeau, Y. and Mueller, J.P. (eds.) (1990) *Decentralized AI*, Elsevier, North-Holland.
- Demazeau, Y. and Werner, E. (eds.) (1991) *Decentralized AI - 2*, Elsevier, North-Holland.
- Doran, J. (1998) Simulating Collective Misbelief, *Journal of Artificial Societies and Social Simulation*, vol. 1, no. 1, <<http://www.soc.surrey.ac.uk/JASSS/1/2/4.html>>.
- Epstein, J. M. and Axtell, R. L. (1996) *Growing Artificial Societies: Social Science from the Bottom Up*. The Brookings Institution Press: Washington D.C. The MIT Press: Cambridge, Mass.
- Gilbert, N. and Troitzsch, K. G. (1999) *Simulation for the Social Scientist*. Milton Keynes: Open University Press, (in press).
- Gilbert, N. and Conte, R. (eds.) (1995) *Artificial Societies: the Computer Simulation of Social Life*. UCL Press: London.
- Halpin, B. (1998) Simulation in Sociology: A review of the literature. Paper read at the Workshop on "Potential of the computer simulation for the social sciences", Centre for Research on Simulation in the Social Sciences (CRESS), University of Surrey, 14-15 January.
- Hegselmann, R. (1996) Modelling social dynamics by cellular automata. In K.G. Troitzsch, U. Mueller, N. Gilbert, and J. Doran (eds.) *Social Science Microsimulation*. Heidelberg: Springer.

- Huhns, M.N. (ed.) (1987) *Distributed Artificial Intelligence*. Pitman, Morgan Kaufmann, San Mateo, CA.
- Lewis, D. (1969) *Convention*. Cambridge, MA: Harvard University Press.
- Macy, M. and Flache, A. (1995) Beyond rationality in models of choice. *Annual Review of Sociology*, 21, 73-91.
- O'Hare, G. and Jennings, N. (eds.) (1996) *Foundations of Distributed AI*, Wiley Inter-Science.
- Rao, A. S. and Georgeff, M. P. (1991). Modelling rational agents within a BDI architecture. In J. Allen, R. Fikes, and E. Sandewall (eds.), *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, San Mateo, CA: Kaufmann, 473-485.
- Reynolds, R.G. (1994) Learning to cooperate using cultural algorithms. In N. Gilbert, and J. Doran (eds.) *Simulating Societies: the computer simulation of social life*. London: UCL Press.
- Shoham, Y. and Tennenholtz, M. (1994) Co-Learning and the Evolution of Social Activity, CS-TR-94-1511, Stanford University.
- Troitzsch, K.G. (1997) Social science simulation - Origin, prospects, purposes. In R. Conte, R. Hegselmann, P. Terna (eds.) *Simulating social phenomena*, Heidelberg: Springer.
- Weibull, J.W. (1996) *Evolutionary Game Theory*. Cambridge, MA: The MIT Press.