

A multidisciplinary perspective on multi-agent systems

EDMUND CHATTOE¹, KERSTIN DAUTENHAHN²,
IAN DICKINSON³, JIM DORAN⁴ and NIR VULKAN⁵

¹*Department of Sociology, University of Oxford, Oxford, UK, email: edmund.chattoe@sociology.oxford.ac.uk*

²*Department of Cybernetics, University of Reading, Reading, UK, email: kd@cyber.reading.ac.uk*

³*(Panel chair), Hewlett-Packard Laboratories, Bristol, UK, email: Ian_J_Dickinson@hpl.hp.com*

⁴*Department of Computer Science, University of Essex, Colchester, UK, email: doraj@essex.ac.uk*

⁵*Department of Economics, Bristol University, Bristol, UK, email: N.Vulkan@bristol.ac.uk*

Abstract

The theory, principles and practice of multi-agent systems is typically characterised as a computational and engineering discipline, since it is through the medium of computational systems that artificial agent systems are most commonly expressed. However, most definitions of agency draw directly on non-computational disciplines for inspiration. During the 1999 UK workshop on multi-agent systems, UKMAS'99, we invited four speakers to address the conceptualisation of multi-agent systems from their perspective as non-computer scientists. This paper presents their arguments and summarises some of the key points of discussion during the panel.

1 Introduction

Is the study of multi-agent systems (MAS) a branch of computer science? Alternatively put, is it predominantly a computational discipline? To judge from the proceedings of agent systems conferences,¹ in which many and various expressions of computational multi-agent architectures and designs are presented, the answer is assuredly “yes”. Yet it is also typical of MASs that they model, or draw inspiration from, non-computational systems. For example, we talk of *delegating* to agents, or agents *collaborating* with each other in agent *societies*. Definitions of agents often draw on non-computational terminology. For example, Wooldridge and Jennings (1995) define an agent in terms of its properties of *autonomy*, *social ability*, *reactivity* and *proactiveness*.

Indeed, many MAS researchers from outside computer science per se make many valuable contributions to the understanding of the field precisely because different disciplines approach the same questions from very different perspectives (Sengers, 1998). To explore some of the potential contributions and critiques from other academic disciplines, a panel session during the 1999 UK Workshop on Multi-Agent Systems (UKMAS'99) invited four speakers to give their perspectives on a multi-disciplinary approach to MAS. In particular, each speaker was asked to address the following questions:

- Is the study of multi-agent systems predominantly a computational discipline?
- What limitations do we impose by viewing agent systems as a strictly, or predominantly, computational discipline?
- What particular contributions to theory, practice or evaluation of agent disciplines does your discipline or field of study provide?

¹ See, for example, <http://www.atal.org>

- What (or what more) should the agent research community be doing to encourage and integrate contributions from other scientific and humanist disciplines? What should the agent community be giving back in return?

The following sections of the paper present the positions taken by the speakers, followed by a summary of key points in the discussion.

2 Contributions from speakers

2.1 Edmund Chattoe

2.1.1 *Is the study of multi-agent systems computer science?*

Yes. But the important question is whether that is all it has to be or needs to be. Consider a problem, such as the scheduling of activities on a production line. One solution is to abstract from the problem – if (a very big if) it is well defined – and then use “pure” computer science techniques to generate an optimal schedule for the abstracted problem. Ideally the technique chosen will generate not only a good schedule, but one that makes effective use of data, computational resources and so on. However, this approach is *instrumental* in that the measure of success is the solution to the abstracted problem. There is no implication that the technique chosen will be anything like the way that the humans on the production line are currently dealing with scheduling problems. Solutions which attempt to produce some congruence between real and simulated systems (models, abstractions) are *descriptive* and the descriptive understanding of interacting individuals in groups is the province of sociology. (Though, of course, it is also the province of psychology, and also of economics to a lesser extent.)

The descriptive approach has both pragmatic and theoretical advantages. First, if the abstraction of the problem is inept or too extreme—a fact that cannot easily be established without investigating how real humans are performing the task—the solution will be no use, even if it is optimal relative to the abstracted problem. (Of course the Procrustean solution to this failure is then to argue that humans should be replaced by robots.) Second, both selection mechanisms and human deliberation mean that it is very unlikely that nothing useful can be learned from the way that humans actually do the scheduling task. (Another way of putting this is that if the problem is poorly defined, what humans do may be the only obvious place to start in narrowing the set of possible solutions.) One important thing that needs to be established in all problem-solving is what is part of the problem and what is part of the solution. Doing this is not trivial. For example, computer science can produce agents with extremely large memories, seeing forgetting as a “weakness” of humans. In fact, forgetting may be an effective way of culling unused information that does not require any global oversight. The third advantage of a descriptive solution is that it may provide understanding, generalisations and new ideas, rather than simply “cookbook recipes”. This is because descriptive systems are meaningful. Learning systems based on neural networks are extremely effective, but it is very hard to tell whether they have over-generalised. Classifier-learning may not be as effective, but it is much easier to understand what the system has learnt. Furthermore, being able to tell what has been learnt may provide insight into what is lacking from the system. For example, large numbers of very similar classifiers for dealing with visual input may draw attention to ineffective generalisation or a representation language that does not permit it.

2.1.2 *What limitations do we impose by viewing agent systems as a computational discipline?*

Strictly speaking, it is impossible to do this. There is no such thing as a “purely” computational model of an agent. Even if the approach taken is strictly instrumental, the capacities of agents will reflect assumptions about what is useful and functional in biology, human behaviour and so on. It is always better that these assumptions be explicit and defensible. Failure to recognise this fact makes it extremely hard for social science to contribute.

What tends to happen is that the background to an approach gets “left behind” when it is absorbed into computer science. It thus appears to be a formal technique, when actually its

effectiveness depends to a considerable extent on descriptive plausibility for reasons given above. Rational choice theory, originating in economics, has been adopted in multi-agent systems and proved very successful in solving *some* problems. However, it has also been criticised within social science as being empirically unfounded, foundationally inconsistent and even potentially unscientific: if agents always choose what they prefer then no behaviour is ever inconsistent with the basic theory. Sociology has made an important contribution to these criticisms and the debate continues. (Effective critique is bound to be a multidisciplinary enterprise. Multi-agent systems based on evolutionary metaphors will be more appropriately criticised by biologists.) If computer scientists are not aware of the presuppositions of the formalisms they adopt, they will neither recognise their limitations nor have any insight into when they should be applied effectively.

On a more positive note, distinctively sociological theories of behaviour, like social construction (the idea that meaning and facts are not given objectively but arise in the process of negotiation and interaction), can contribute to the (undoubtedly synthetic) task of building “effective” agents. Historically, sociologists have been uncomfortable specifying theories rigorous enough to be built directly into agents, but this is changing with the development of social simulation (Conte & Castelfranchi, 1995; Epstein & Axtell, 1996; Gilbert & Conte, 1995; Gilbert & Doran, 1994).

2.1.3 *What particular contributions to MAS does sociology provide?*

Many of the problems for which multi-agent systems are used have interesting (and sometimes diverse) social analogues. In fact, one could see societies and social practices as evolved solutions to messy problems. If these social analogues cannot tell us anything about how we can find solutions, it would be interesting to know why this is. If they can, we would be foolish not to learn from them. Sociology is already heavily involved in the study of communication, small-group interaction, organisations and so on.

Probably the most important thing that sociology can contribute generally is a range of data collection tasks for understanding social behaviour. In particular, sociology has pioneered techniques in the elicitation and organisation of *qualitative* data: interviews (Mason, 1996), focus groups (Morgan, 1997), conversation analysis (Hutchby & Wooffitt, 1998), participant observation and ethnography (Hammersley & Atkinson, 1995). Since multi-agent systems aspire to link cognitive content to behaviour to find effective solutions, the relatively “behaviourist” disciplines like experimental psychology and economics are less well equipped to help with this task.

Another important area where sociology can contribute is in drawing attention to the implications of different kinds of mental content for the effective functioning of agents. As well as the strictly “functional” (and often implicitly individualistic) beliefs and desires, what are the roles of norms, emotions, collective representations, ideologies and so on (Conte & Castelfranchi, 1995)?

Finally, sociology has been instrumental in challenging naive positivistic approaches to the scientific method (Latour & Woolgar, 1986). In consequence, it has long been preoccupied with the theoretical and practical implications of reflexivity, self-awareness, inter-subjectivity and so on (Crossley, 1996). All of these phenomena become potentially important in dealing with autonomous agents. It is fair to say that some of these contributions have been hampered by the lack of a “computational mindset” but there are nonetheless useful insights to be found.

2.1.4 *What more should the agent research community do to encourage and integrate contributions from the social sciences?*

At a general level, there needs to be greater awareness that multi-agent architectures are based on behavioural *assumptions* and that they are not logical formalisms. More specifically, it matters which behavioural assumptions are built into multi-agent architectures and there is already an extensive literature in the social sciences about the strengths and weaknesses of some common assumptions. These debates are not mere footnotes to the engineering choice of architectures, but in some cases may invalidate whole approaches—except as purely instrumental techniques. Without this kind of awareness, computer scientists will have no motivation to listen to what social scientists have to say. At the same time, cooperation would be facilitated by specifying challenging multi-

agent problems in terms that were less tied to the computational aspects. This would allow social scientists to identify appropriate social analogues, data collection techniques and existing literature that could provide insights. Conversely, social scientists need to make their insights more rigorous and suitable for computation. Unfortunately, the only way that the two sides will ultimately get together is if each is humble enough to realise that there is always more to be learned; this frame of mind is not one that can be inculcated by argument.

2.2 Kerstin Dautenhahn

2.2.1 *Is the study of multi-agent systems computer science?*

The answer to this question depends on what we mean by “computer science” and “multi-agent systems”. Obviously, multi-agent simulations and systems widely employ computer technology. But the field of computer science seems to be very weakly defined, referring to tools (computers) and techniques (programming) rather than to theory, research goals or subjects of investigation. For example, Webopedia² gives the following definition of computer science:

The study of computers, including both hardware and software design. Computer science is composed of many broad disciplines, including artificial intelligence and software engineering.

This definition substantially differs from definitions of, for example, biology, a natural science that can be defined as follows: “The branch of science dealing with properties and interactions of physico-chemical systems of sufficient complexity for the term ‘living’ (or ‘dead’) to be applied” (*Penguin Dictionary of Biology*, 1994). Thus biology is the study of life; it is defined by its research subjects (animals and plants) which share a particular property (“life”, or more precisely a list of properties of living systems). In the same way as a complete list of properties of “life” does not exist in biology, a complete list of properties of an agent or a multi-agent system does not exist, either. Interestingly, even formal approaches to agents and MAS are often inspired by behavioural or cognitive skills of “real-life” agents, such as notions of mobility, goals, beliefs, intentions and many more. This is particularly obvious in areas where researchers are interested in complete agent architectures, e.g. the software pets and their environment developed in Creatures.³ The field of multi-agent research is relatively new, so it is hoped that in future the notions of “agent” and “multi-agent system” are becoming more precise. In my view research into multi-agent systems could benefit from aiming at a similar development as we have observed in natural sciences like physics or biology, which are based on a (socially constructed) framework of theories, methods and methodologies and a strong grounding in experimentation.

2.2.2 *What limitations do we impose by viewing agent systems as a computational discipline?*

In my view, some of the most exciting and scientifically challenging agent systems are those with a multidisciplinary approach (see recent proceedings of the Autonomous Agents conferences, or Huhns & Singh, 1998). Linguistics, psychology, the arts and other fields play an increasingly important role in particular in agent systems that are supposed to interact with people. Here, issues of “believability” can be as important for user acceptance (and the commercial success) of a product as computational issues, e.g. speed and realism of agent interfaces. The “human-in-the-loop” as observer and user (and designer) of an agent system constrains but also enhances the development of interactive systems which show their full potential only through the interaction dynamics with a human, or a group of humans (in a multi-user context). For an extensive discussion of these issues of the “human-in-the-loop” see Dautenhahn & Nehaniv (1999) and Dautenhahn (1998).

²See <http://webopedia.internet.com>

³See <http://www.creaturelabs.com/>

2.2.3 What particular contributions to MAS do biology and cybernetics provide?

My fields of former and current study are biology and robotics. A major contribution of robotics to agent research is the importance of “embodiment” and real world dynamics. The issue of agent and embodiment is discussed in more detail in Aylett et al. (2000).

Biology is the only “real” science of agents, since it studies *agents as we know them*. Biology can serve as a model for the development of a science of *agents as they could be*. The strong version of this argument can be phrased as follows: The only way to build complete agents (and agent architectures, accordingly) is to develop and evolve agents in the same way as animals evolved and adapted to an ecological niche, being able to cope with a widely unpredictable and dynamic environment. A similar argument is given in Grand (1999). A *niche* for an agent is defined through the application area, e.g. the Internet for mobile and e-commerce agents, or the “social space” for user-interface agents that are interacting with humans. Also, a niche is not an empty space which has to be filled, it co-evolves along with its inhabitants. And there is no optimal inhabitant of a niche, in the same way as there is no general-purpose animal. A diversity of solutions (animal designs) are equally successful and coexist. Similarly, I do not expect that narrowing agent research by strictly defining a single agent architecture and agent language across different application areas and problem domains is a promising way to go. Standards are very useful as far as they support investigation and experimentation, but they should not prevent researchers from exploring the design space of agents.⁴

2.2.4 What more should the agent research community do to encourage and integrate contributions from other disciplines?

I see the following main points:

- There is a general problem in funding and evaluating interdisciplinary work. This problem also applies to agent research.
- Another general problem of interdisciplinary work is a communication problem, i.e. the difficulty of researchers in finding a common language for scientific discussions. For example, the term “social” is often used very vaguely in agent research and in a way which is not compatible with how e.g. linguists or ethologists use the term. I do not believe that a common language can be developed among agent researchers in different fields, at least not in the near future. The best short-term solution is therefore to be very explicit when using a term that is defined differently in other fields, e.g. the terms “social” (e.g. used in social sciences and ethology) or “communication” (e.g. used in linguistic and signal theory).
- Agent research can provide frameworks and computational testbeds that could be used by people from other fields, e.g. social scientists can study phenomena of migration and culture (e.g. Epstein & Axtell, 1996).

2.3 Jim Doran

2.3.1 Is the study of MAS part of computer science?

Yes, MAS studies are part of computer science in the same way that artificial intelligence is. But clearly it is also possible rigorously to explore the properties of multi-agent systems other than by using computers, for example by using classical mathematics and formal logic. Taking human beings to be agents, we might be tempted to go further and suggest that *all* of social science, to the extent that it is about the properties of human groups and organisations (both in general and in relation to specific instances), may be viewed as a special case of MAS studies. But that seems a step too far. The key characteristic of current MAS studies is the *precise* specification of multi-agent systems and hence the *precise* derivation and study of their properties.

⁴Due to space limitations the relationship between biology and agents cannot be developed in full in this paper.

2.3.2 What are the limitations and benefits of viewing MAS in CS terms?

Benefits

- Viewing MAS in computer science terms enables detailed working out of the implications of precise assumptions about MAS—so it's a “waffle-killer”.

Limitations/Dangers

- There is a temptation to drift into “hard-core” computing for its own sake, getting lost in the “nitty-gritty” of graphical interface design.
- There is a temptation merely to “play” with the computer (network), leaving scientific investigation all but forgotten.
- It sometimes seems that computation is not quite the right modelling medium. Put otherwise, it is not clear that the (mathematical) symbol systems manipulated by computers have the ideal “texture” to act as social models. But what is the alternative?

2.3.3 What contributions can archaeology make to MAS studies?

At first sight, none! But archaeology is all about the recovery of the history of societies of (human) agents from trace evidence. Just possibly, it can help us understand what dynamic societies of artificial agents (e.g. on the Internet) have been doing (*e.g. Good grief! What did those guys DO last night?!*).⁵

More specifically, prehistoric archaeology (a) deploys a detailed methodology for the recovery of evidence of past (unrecorded) human activity, (b) interprets recovered evidence at micro and macro levels of social activity, and (c) studies the long-term dynamics of “simple” agent societies (e.g. slow centralisation and hierarchy formation over a thousand years; revolution; sudden socio-cultural collapse). This expertise is well developed and tested, and its projection into the domain of artificial societies is thought-provoking at the very least. Notice, for example, that one year on the Internet, with agent generations 20 minutes long, would correspond to about 500,000 years of human history. A lot has happened in that time!

A specific example of potential cross-fertilisation is that the long-term dynamics of human societies draw attention to the impact of *collective (mis)belief systems* (ideologies, systems of religious beliefs) (Doran, 1998) and of *collective emotional states* (e.g. the confident and aggressive “let's go for it” society, or the society that is defeated, subjugated and demoralised) (Doran, 2000b), topics that have so far been rather little explored in artificial societies work.

2.3.4 What contributions can MAS studies make to archaeology?

Agent-based modelling promises to be a key tool for exploring social processes, including those of interest to archaeologists. Primarily, this is because it supports explicit modelling of individual cognition, which is surely central to effective models and simulations of human societies. Indeed, over the last 25 years there have been a number of attempts to use agent-based modelling/simulation in archaeology (Hodder, 1978; Doran, 1990), not always using the word “agent”. Disappointingly, although there have been isolated successes there has been little real impact on the field (Doran, 2000a). Why has agent-based modelling and computer simulation not been more successful in archaeology? It seems to be because archaeologists

- cannot validate agent-based models in any detail and therefore do not trust them,
- do not take seriously the idea of social science *theory building* (Epstein & Axtell, 1996) using computer simulations (agent-based or otherwise) which arguably does not need exact validation and
- do not anyway have the resources (mainly of human expertise), to do much in the way of computer experimentation.

⁵A detailed examination of this suggestion will be presented to the AISB'2000 Symposium “Starting from Society”.

2.3.5 What more should the agent research community do to encourage and integrate contributions from the social sciences, such as archaeology?

This question is difficult to answer. But perhaps we have reached the point where there can usefully be a detailed and systematic match of fragments of social science theory (perhaps this implies first a choice of a social science “school of thought”) with corresponding fragments of MAS theory, aiming at unified theory endorsed by both sides. Candidate topics are Durfee (1999) on the design of efficient hierarchical organisations, and the wealth of current research on markets and auctions.

2.4 Nir Vulkan

First, my understanding of multi-agent systems is that in this framework different components represent different entities, like sellers and buyers. The system controls the rules by which these agents interact, but does not control the behaviour of the agents. In almost all MASs I have seen or been involved with, agents have potentially conflicting goals (e.g. the selling agent prefers a high price, the exact opposite preferences are held by the buying agent). These conflicts are normally the result of agents having to share limited resources.

Given this view of MASs, it should be clear why I see economics and game theory as extremely useful disciplines for the study and design of MASs. Economics is the study of the allocation of limited resources. Game theory is a formal theory that studies the interactions between rational, self-interested agents. By rationality I mean that agents are time-consistent and utility-maximising. These two assumptions are often doubtful when applied to humans, but seems extremely likely for automated agents: these computer codes are programmed to maximise a given function and, once running, are not capable of changing their minds (compare that to a person who decides to quit smoking tomorrow, and changes his mind when “tomorrow” arrives). In fact, game theory is much more suitable for automated agents than it is for humans.

In economic theory, distinction is made between models in which we analyse the optimal behaviour of individuals or firms *given* the underlying mechanism (or rules of the game), and models in which we study optimal mechanism design, *given* that agents behave optimally. In the current early stages of multi-agent design these two approaches are being developed simultaneously. One cannot compare two different protocols (mechanisms) without specifying the behaviour of the interacting agents. Similarly, one cannot design optimising agents without some information about the protocols governing their interaction. However, we think it clarifies the underlying philosophy to maintain a clear distinction between the design of protocols and the design of the agents who operate within the rules specified by the protocols.

In current MAS bargaining, automated agents are programmed with rules of thumb distilled from intuitions about good behavioural practice in human negotiations. The danger is that the programmer may not be fully aware of the circumstances to which human behavioural practice is adapted, and hence use behavioural rules that are capable of being badly exploited by new agents that have been programmed to take advantage of the weaknesses of the agents currently in plan. When protocols have been deliberately constructed to take advantages that are available within the artificial environment of a computing system, the risks of creating the opportunity for such destabilising invasions by new agents are particularly large.

Economists believe that their approach provides an escape route from these difficulties. In principle, an agent should be designed to *optimise* on behalf of the decision-maker whose role it usurps. The revelation principle of mechanism design applies also to agent design, and so the designer of a properly engineered agent can tell his client that his programming takes care of all the *strategic* problems involved in bargaining optimally. This leaves the client to report *truthfully* on his preferences and his information. This may not always be easy for the client to understand. For example, evidence from the recent Guttman and Maes electronic-agent marketplace experiment at MIT (Maes et al., 1999) shows that users consistently *lied* to agents that they had designed for themselves, because they thought they could get more by giving the impression of being tougher than they are. Similar problems have occurred with other negotiation-based MASs. With a properly

designed agent, it would always be a mistake to tell such lies. If it is optimal to pretend to be tough, the agent will do all the pretending necessary.

It is an interesting fact that artificial intelligence (AI) and economics have had many overlapping interests over the years. John von Neuman's pioneering work laid the foundations for modern AI as well as modern game theory. Along the same lines, Herbert Simon's work on rationality and bounded rationality greatly influenced researchers in both fields. It is therefore not surprising that we find ourselves these days in a situation where researchers from both fields work together in pursuit of what may become one of the more important technological changes of modern life.

In the short term, computer science is likely to benefit more from this cooperation, because economic wisdom on the efficiency of systems consisting of self-interested agents can be almost directly applied to multi-agent systems. In addition, the economic methodology that stresses the importance of looking for the underlying incentives of participants can provide important insights into automated negotiations and electronic trade. But in the longer run, economists have also much to gain from this joint adventure. By focusing on what is essentially an application of economic theory and mechanism design to automated environments, we can learn about the usefulness of our theories and intuitions.

3 Discussion

Some of the key points from the open discussion that followed the panel's presentations are summarised below.

Are you proposing that there should be a common theory between these various disciplines?

The panel's view was that a common theory was neither feasible nor especially useful. However, it was felt that there were many shared points of view between different disciplines, and that such commonalities could be exploited to illuminate problems in one field with models, stances or solutions from another. It may also be that points of view that appear radically different have more in common viewed from a higher level of abstraction.

Neural networks could be viewed as either models of brain function or practical tools for solving certain kinds of problem. Which is the right analogy for MAS?

In practice, it may not be possible (or at least useful) to separate abstract problems from the context of their real application. In addition, if we consider that the terms borrowed by MAS researchers from social disciplines are often, in essence, *analogies*, the proper understanding of the social science may be crucial to the correct interpretation of the analogy. Examples range from MASs which fail due to incorrect understanding of the social context of use (for example, the organisation in which the system is deployed), through to the FIPA standard⁶ auction protocol being open to unfair exploitation that proper game-theoretic mechanism design could have prevented.

Would a standardisation activity be a good way to help integrate these various different approaches?

No.

4 References

- Aylett, R, Dautenhahn, K, Doran, J, Luck, M, Moss, S and Tennenholtz, M, 2000, "Can models of agents be transferred between different areas?" *Knowledge Engineering Review* **15**(2) 199–203.
- Conte, R and Castelfranchi, C, 1995, *Cognitive and Social Action* UCL Press.
- Crossley, N, 1996, *Intersubjectivity: the Fabric of Social Becoming* Sage Publications.
- Dautenhahn, K, 1998, "The art of designing socially intelligent agents—science fiction and the human in the loop" *Applied Artificial Intelligence Journal* **12** 573–617.
- Dautenhahn, K and Nehaniv, C, 1999, "Living with socially intelligent agents: a cognitive technology view" In K Dautenhahn (ed.) *Human Cognition and Social Agent Technology* John Benjamins Publishing Company.

⁶See <http://www.fipa.org>

- Doran, J, 1990, "Computer-based simulation and formal modelling in archeology" in A Voorrips (ed.) *Mathematics and Information Science in Archeology: A Flexible Framework* HOLOS.
- Doran, J, 1998, "Simulating collective misbelief." *Journal of Artificial Societies and Social Simulation* **1**. Available from: <http://www.soc.surrey.ac.uk/JASSS/1/1/3.html>.
- Doran, J, 2000a, "Prospects for agent-based modelling in archaeology" *Archeologia e Calcolatori* **X**.
- Doran, J, 2000b, "Trajectories to complexity in artificial societies: rationality, belief and emotions" in T Kohler and G Gummerman (eds) *Dynamics in Human and Primate Societies* Oxford University Press.
- Durfee E, 1999, "Exploring Hierarchical Representations for Efficient and Flexible Coordination in Multiple Agent Systems" *Invited Presentation, UKMAS'99*. Hewlett-Packard Laboratories, Bristol.
- Epstein, J and Axtell, R (eds), 1995, *Growing Artificial Societies: Social Science from the Bottom Up* MIT Press.
- Gilbert, N and Conte, R (eds), 1995, *Artificial Societies: The Computer Simulation of Social Life* UCL Press.
- Gilbert, N and Doran, J (eds), 1994, *Simulating Societies: The Computer Simulation of Social Phenomena* UCL Press.
- Grand, S, 1999, "Creating souls from cells" In K Dautenhahn (ed.) *Human Cognition and Social Agent Technology* John Benjamins Publishing Company.
- Hammersley, M and Atkinson, P, 1995, *Ethnography: Principles in Practice* Tavistock.
- Hodder, I (ed.), 1978, *Simulation Studies in Archeology* Cambridge University Press.
- Huhns, MN and Singh, MP (eds), 1998, *Readings in Agents* AAAI Press/MIT Press.
- Hutchby, I and Wooffitt, R, 1998, *Conversation Analysis: Principles, Practices and Applications* Polity Press.
- Latour, B and Woolgar, S, 1986, *Laboratory Life: The Social Construction of Scientific Facts* Princeton University Press.
- Maes, P, Guttman, R and Moukas, A, 1999, "Agents that buy and sell" *Communications of the ACM* **42** 81–7, 90–1.
- Mason, J, 1996, *Qualitative Researching* Sage Publications.
- Morgan, D, 1997, *Focus Groups as Qualitative Research* Sage Publications.
- Penguin Dictionary of Biology*, 1994, Penguin.
- Sengers, P, 1998, "Anti-boxology: agent design in cultural context" Ph.D. Dissertation, Carnegie-Mellon University, Pittsburgh, PA, USA. Also available from: <http://www.cs.cmu.edu/afs/cs.cmu.edu/user/phoebe/mosaic/work/thesis.html>.
- Wooldridge, M and Jennings, N, 1995, "Agent theories, architectures and languages: a survey" In M Wooldridge and N Jennings (eds.) *Intelligent Agents* Springer-Verlag.