VU Research Portal

Agent Systems and Applications

Aylett, R.; Brazier, F.M.; Jennings, N.R.; Luck, M.; Preist, [Unknown]; Nwana, H.S.

published in

Knowledge Engineering Review
1998

DOI (link to publisher)

[10.1017/S0269888998003038](https://doi.org/10.1017/S0269888998003038)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Aylett, R., Brazier, F. M., Jennings, N. R., Luck, M., Preist, U., & Nwana, H. S. (1998). Agent Systems and Applications. *Knowledge Engineering Review*, 13(3), 303-308. <https://doi.org/10.1017/S0269888998003038>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Agent Systems and Applications*

RUTH AYLETT¹, FRANCES BRAZIER², NICK JENNINGS³,
MICHAEL LUCK⁴, HYACINTH NWANA⁵ and CHRIS PREIST⁶

¹Centre for Virtual Environments, University of Salford, UK

²Department of Computer Science, Vrije Universiteit Amsterdam, The Netherlands

³Department of Electronic Engineering, Queen Mary & Westfield College, University of London, UK

⁴Department of Computer Science, University of Warwick, UK

⁵Advanced Research and Technologies Department, BT Laboratories, UK

⁶Agent Technology Group, HP Laboratories, UK

Abstract

As the number of deployed multi-agent applications increases, further and better experience with the technology is gained, enabling a strong evaluation of the field from a more practical perspective. In particular, questions relating to how the theory of multi-agent systems impacts on practice, and how the practical development itself compares with other technologies, can be answered in the light of a heightened level of maturity. Given the tensions between theoreticians and practitioners in computing in general, let alone their spats in AI or multi-agent systems in particular, the discussion on agent systems and applications was both vigorous and enthusiastic.

1 Characteristics of multi-agent applications

What are the characteristics of applications which make them suitable for a multi-agent system solution? Are there applications which are unsuitable for multi-agent systems?

The obvious response to the question of characteristics of suitable applications is that they are worth considering when the application requires both distribution and intelligence. Thus, a multi-agent approach would be sensible for problems that are inherently (physically or geographically) distributed where independent processes can be clearly distinguished. Such problems include, for example, distributed sensor networks, decision support systems, air traffic control, or other networked or distributed control systems. A distributed approach is not in itself enough, however, and there should also be requirements for intelligence or adaptivity in the sub-processes that involve explicit reasoning about behaviour, for example. If the problem can be solved by means of a look-up table at each node of a network, a multi-agent system would be excessive. Physical distribution may not be the only reason for a distributed approach. Minsky's *Society of Mind* paradigm (Minsky, 1986) suggests the use of a multi-agent system where there is a wide range of reasonably self-contained pieces of functionality that require the use of AI, especially if they run asynchronously, or are distributed or independent in the sense of timing. In this way, a multi-agent approach might be applied to a single robot manipulator taking each joint as an agent, or to a single static system with sensors and effectors (an *immobot*) such as a smart building or a spacecraft.

A range of further application areas that qualify for multi-agent solutions can also be enumerated. These include those requiring the interconnection and inter-operation of multiple autonomous, self-interested existing legacy systems, expert systems, and decision systems, or those requiring solutions that draw from distributed autonomous and selfish information sources, such as the Personal Travel Assistance demo from BT; those where the solutions draw from different

* This report is the result of a panel discussion at the *Second UK Workshop on Foundations of Multi-Agent Systems (FoMAS'97)*. All members of the panel are authors, listed alphabetically.

distributed experts, such as health care provisioning, in which some central agent cannot possibly perform the task without help from other experts; problems that naturally cross organizational boundaries for which an understanding of the interactions among societies and organizations is needed; and problems where no single agent has a total view, but several agents have local views. The notion of *ownership* of information and strategies in the application is important here, and in particular when it is distributed over different organizational entities so that no single entity can (or does) have access to all the information. Final particular examples include traders in a marketplace (Chavez et al., 1997), and different entities in a business working on other's tasks (Jennings et al., 1996). In these situations, the problems to be tackled do not have one overall goal, but rather consist of balancing the (possibly conflicting) goals of different entities.

As with any other piece of technology, there are plenty of applications that do not require a multi-agent approach. Multi-agent systems are not required merely to produce modularity (though they reduce complexity), extra speed (though this may be an effect of their inherent parallelism), reliability (though they provide redundancy), flexibility or re-usability at Newell's knowledge level (Newell, 1982). In the same way, they are not required simply because a problem is too large for a centralized single agent due to resource limitations, nor because of the sheer risk of a centralized system, nor merely for reasons of efficiency, heterogeneous reasoning, etc.

For example, a payroll system might benefit in a software engineering sense from an object-oriented approach (providing modularisation and re-use), but such standard data-processing problems do not really need the communications overhead or functionality of a multi-agent approach. Such applications usually require neither distribution nor intelligence. In the same way, a small free-standing expert system used in a single location requires the intelligence but not the distribution, and the human interface is insufficiently complex to be worth thinking of it as an agent.

Finally, it is worth noting that a multi-agent system approach may be useful, though not necessary, when tackling problems that are easiest *visualized* in a way that appears to have the above characteristics, such as combat simulation (Rao and Selvestrel, 1992).

2 Applications development

What does a multi-agent approach to applications development buy you over more standard approaches such as object-oriented, expert systems, or distributed computing approaches.

In general, Object-Oriented (OO) systems, expert systems and distributed computing techniques do not offer solutions to the kind of problems for which multi-agent systems are used, for a range of reasons (Wooldridge, 1997).

OO techniques are good in general, but are rather low-level for intelligent applications. They can be used, for instance, to implement knowledge representations, but they do not themselves provide a knowledge representation. OO development methodologies can, however, be seen as a low-level underpinning for a multi-agent methodology.

The same might be said of distributed computing methodologies and indeed, many multi-agent systems (e.g. ADEPT) are built on top of distributed platforms such as CORBA. However, if it can be argued that if OO approaches are still relatively new, these are even newer and less generally accepted. Again, however, the level is wrong—for example, communications protocols do not operate at the high level of Speech Acts as one might wish for a multi-agent system. Yet distributed computing approaches could deal with some of the lower levels as well as providing some of the basic techniques (such as protocol definition and validation). An OSI model with an elaborated applications level might be ideal or, alternatively, current approaches to multi-agent applications development might form such a layer.

Expert systems are even more problematic in terms of development methodology. There is a reasonable consensus on the life-cycle, but while KADS has made some impression, many avoid it. Developments such as KRL and other movements for standardization are very much in their infancy, while ontologies are helpful but not widely used. The most successful parts of knowledge

engineering relate to knowledge acquisition (elicitation and structuring) and are not incompatible with a multi-agent approach in the case where such a system has to embody expert knowledge. This is not always the case, however.

More importantly, multi-agent applications require a *cooperation knowledge level* (Jennings and Campos, 1997), while expert systems (and others) typically operate at the symbol and knowledge levels (at most) (Newell, 1982). However, if using a *cooperation knowledge level* buys nothing for the application, it should be avoided. For example, Huberman's thermal economy multi-agent system has been criticised in that it provides an inferior solution to a single 'agent' control solution. Arguably, this is a clear *distributed problem solving* issue, but if it did not involve just *one* owner, then an inferior multi-agent solution might be justifiable.

The benefits of a specifically multi-agent methodology would be a reduction of the semantic gap between analysis on the one hand, and design and implementation on the other, leading to a reduction in the time to design and implement, with the usual trade-off between better expandability and losses in execution efficiency and design specificity. The Internet revolution is resulting in increased communication between distinct entities with different goals and distinct boundaries which must be secure. The new class of applications which service this need will inevitably use some form of multi-agent approach. Current methodologies emphasise top-down design, but multi-agent systems adopt a different approach—top-down within the agent, and bottom-up in the agent community.

If the utility of using a multi-agent solution (however it is measured, and it will depend on every new situation) is greater than the utility gained from a single "agent" application, then such an approach should be used. In summary, multi-agent systems research can be regarded as developing a way of looking at problems, (a nascent design methodology) rather than a technology¹. Hence multi-agent systems can, and do, use object-oriented, expert systems and distributed computing technologies to implement applications and toolkits that embody this approach. Multi-agent systems will not replace these technologies, but provide a different way of using them to tackle new kinds of problems. (See the panel report on methodological foundations from FoMAS'96 for a more detailed discussion of some of these issues (Fisher et al., 1997).)

3 Killer applications

Is there a killer application for multi-agent systems and, if so, what is it?

The search for a *killer application* for multi-agent systems brings quite different responses. Nwana and Brazier suggest applications such as interactive collaborative design, air-traffic control, network management, command and control and distributed data mining.

By contrast, Aylett and Preist take a different stance. While recognizing the killer application as the "holy grail" of the agent world, Preist believes that it does not yet exist because the problems where multi-agent systems are necessary are only just appearing. Current inter-organization applications do not require multi-agent systems, for two reasons. First, they *support* human communication to carry out tasks, rather than *automating* tasks completely (e.g. FastParts component trading). Second, they are proprietary systems that link a small number of companies in a predetermined way (e.g. the WalMart Quick Response system (Kalakota and Whinston, 1996)). Nevertheless, Preist does suggest applications that are both automated and open such as on-line automated trading as possibilities.

Aylett questions the need for *killer* applications as such, but sees computer pets as catching the popular imagination, while in the corporate world, information retrieval agents ride on the back of Internet hype. In these cases, the *agenthood* is obvious, but for other multi-agent applications only the results are obvious, and only a few people will ever care about how they were achieved.

¹Note that there is some multi-agent systems technology, in particular in dealing with the social aspects of problem-solving (e.g. cooperation, coordination, etc.).

4 Barriers to the uptake of technology

Are there barriers to be crossed before a broader uptake of multi-agent systems technology will be seen? For example, do we need to have a stock of specialized tools or development methods (such as languages, etc.) for multi-agent systems for them to be a widely-accepted class of system technology?

Barriers to the uptake of multi-agent systems technology can be divided into two classes, which we can call social and technical barriers. Social barriers cover issues relating to public relations, training and experience, while technical barriers cover tools and techniques. These classes of barrier are inter-related but are considered in turn below.

For a technology to be used “in anger”, there must be a clear notion of its purpose and usefulness. Typically, this is obtained via small industrial trial projects that train those in the field in the concepts and methods and allow some basic experience to be laid down. There is no substitute for this process. Only if some clear gains are realized during the course of it will any further progress be made. Slick and expensive commercial packages are not required at this stage, but it must be possible to use cheap or public domain software reasonably easily, even if this has limited functionality. Government initiatives and industrial clubs are a good way of getting this process off the ground. Large high-tech organisations (such as BT, Siemens, Daimler-Benz and Logica) are already successfully carrying out this process but it is not yet active in the wider industrial context.

The existence of trained people helps significantly. Industry now contains many people who studied expert systems at university and a growing number who also studied neural nets and fuzzy logic. But how many are there who studied agents? This barrier should not be underestimated.

One outcome of this process—if it is successful—is a demand for better tools and methods. At the same time, particular tools can have the effect of increasing experience, and thus demand, as the experience with expert system shells has shown. Limited as they were, expert system shells popularized the technology widely, and laid the basis for a second generation of more sophisticated tools.

On the technical side, there are a number of tools and techniques that might push back the barriers, and in particular concerning the interaction between agents. Because inter-organizational applications are one of the major drivers for the uptake of multi-agent systems, some form of inter-agent communication standard is necessary, either pre-designed, or emerging *de facto* from a community. The standards are likely to be problem-specific initially with, for example, the possibility of a standard language of trade negotiation emerging. If a pre-designed standard is to be adopted, however, it is important that it is easy to use and is designed with the upcoming needs of the business community in mind. Some efforts in this direction are already underway with, for example, FIPA, Agent Society, MAF, W3 consortium, OPS, etc. Note the related impact CORBA is having with distributed objects.

Other technical areas of concern include issues of security so that agents can communicate securely and be trusted with important business tasks; the need for clear, non-*ad hoc* development methodologies addressing *real* agent issues to support the design and development of multi-agent systems; and tool-kits to facilitate their construction with good debugging environments, hence ZEUS (Nwana et al., 1998) and dMARS.

While all these aspects might contribute to a greater uptake of agent technology, one thing that must not be standardized or preordained is the internal design of agents. As agent communities will cross organisational boundaries, it is important that different organisations are free to adopt whatever internal approach they choose.

Nwana also sounds a note of warning, suggesting that a new field is only defined by its *problems*, not its *methods*. The field of multi-agent systems is falling into the trap that has befallen artificial intelligence, of deluding itself that the methods (e.g. cooperation, achieving rationality, competition, believable agents, conceptual and theoretical foundations, multi-agent planning, negotiation) are the really important issues. It is the problems of that are of foremost importance.

5 Parallels with other technologies

Are there parallels with other applications and development technologies that can illuminate the future of multi-agent systems, such as OO technology, or are multi-agent systems something fundamentally different?

All the panellists agreed that multi-agent systems are not fundamentally different. Indeed, many solutions to multi-agent systems problems already exist in several parts, and researchers need to do *creative synthesizing* of some of these already invented wheels. Where necessary, “new wheels” may be invented to link up the old ones. This position is vindicated by the ZEUS project which synthesizes in one system many old wheels (Nwana et al., 1998). Other successful projects such as dMARS (d’Inverno et al., 1998) and ADEPT (Jennings et al., 1996) also support this position. The new wheels, such as coordination and visualization, are simply there to link up old concepts. As a result, the whole then becomes greater than the sum of its different parts, and hence novel, but this does not mean it is fundamentally different.

Other technologies such as Distributed Objects, the Internet, OO techniques, etc., have interesting histories from which to learn. For example, without HTML, HTTP, CGI standards, the Internet would not be the phenomenon it is today. Similarly, CORBA has made Distributed Objects a worthy investment, and a potential route to solving the legacy problem. These technologies are successful because at their core, they are quite simple, and all the complexity resides at their edges. At a more detailed level, however, while OO technology may provide a means to implement agent systems, it is not aimed at providing a conceptual representation of knowledge, reasoning, intelligent behaviour and interaction.

6 Objective measures of benefits and drawbacks

How can the benefits and drawbacks of a multi-agent system solution be assessed in an objective manner?

In trying to find ways of assessing the contribution of multi-agent systems, one view argues that this is only really with experience. Strictly speaking, we need to know how good or bad the multi-agent approach is in comparison with the best alternative solutions. Nearly always, the answer is unknown because resources do not allow multiple approaches to be tried. Moreover, in the real world there are usually any number of different ways of solving a particular problem and it is impossible to compare them in any objective fashion. In practice, this is where the “suck it and see” applications mentioned above come in. From experience of producing small systems in a variety of different organizations to solve a variety of different problems, it becomes possible to generalize about the benefits and drawbacks and to see what the evidence is for what was claimed to be good and bad.

Again, the comparison with expert systems or neural nets is illuminating. In the former case, the benefits were expected to be in the area of efficiency (because of fewer experts needed) but turned out to be in the area of effectiveness (with better use of existing experts and more widespread distribution of expertise). In the case of neural nets, some comparisons have been made with equivalent mathematical and statistical techniques, and the benefits are less in performance and more in the time and skill required for development and extension. Without the experience in both of these cases, we would not have been at all sure what to start measuring.

The alternative view is that though rigorous evaluation is the academic answer to this, the market will eventually be the most objective evaluator. Companies are buying into Distributed Object technologies because there are clear needs and benefits for them, and the same is true of multi-agent systems, which will not be taken up unless they demonstrate their need and potential. In the short term, we need to define criteria for judging multi-agent solutions. Do they offer a *value-added* solution to other conventional applications such as control, expert systems, distributed problem solving approaches, and blackboard approaches? This simply requires the definition of *value-added* which could range from excellent (where no solution was probable without a multi-agent approach),

through minimal (where a multi-agent solution is marginally of value) to poor (where conventional approaches offer better solutions).

Again, Nwana points out the dangers, suggesting a critical appraisal of some current efforts whose benefits are questionable. He cites Stuart Russell who recently argued against what he calls *premature mathematization*:

“There is always a danger, in this sort of claim, that its acceptance can lead to ‘premature mathematization’, a condition characterized by increasingly technical results that have increasingly little to do with the original problem” (Russell, 1997).

According to Nwana, multi-agent systems research is in danger of falling into the same trap that Russell asserts of AI.

Despite this warning, the prospects look good. Indeed many researchers in the field have acknowledged the difficulties that arise when formal theoretical work is not used to inform practice, and *vice versa*, and there is an emerging trend of research that aims to bring these two strands together, e.g. (Jennings, 1995; Luck et al., 1997; Rao, 1996). In a sense, difficult experiences with AI in general over the last thirty years or so have provided the impetus for a strong effort to avoid similar problems in multi-agent systems. The development of real systems and applications forms a major part of the field and, as has been suggested, the experiences gained can and should be fed back into the basic research. Relative youth here can be seen to be an advantage in this way. Ultimately, the enthusiasm for agent systems and applications, coupled with cautious and perhaps suspicious voices such as Nwana’s, suggests a productive future.

References

- Chavez, A, Dreilinger, D, Guttman, R and Maes, P, 1997. “A real-life experiment in creating an agent marketplace”. In: *Proc. Second Conference on Practical Applications of Agents and Multi-Agent Systems* 159–178.
- d’Inverno, M, Kinny, D, Luck, M and Wooldridge, M, 1998. “A formal specification of dMARS”. In: *Intelligent Agents IV: Proceedings of the Fourth International Workshop on Agent Theories, Architectures and Languages*, M Singh, A Rao and M Wooldridge (eds), *Lecture Notes in Artificial Intelligence*, 1365, 155–176, Springer-Verlag.
- Fisher, M, Müller, J, Schroeder, M, Staniford, G and Wagner, G, 1997. “Methodological foundations for agent-based systems” *The Knowledge Engineering Review*, 12(3) 323–329.
- Jennings, NR, 1995. “Controlling cooperative problem solving in industrial multi-agent systems using joint intentions” *Artificial Intelligence* 75(2) 195–240.
- Jennings NR and Campos, JR, 1997. “Towards a social level characterisation of socially responsible agents” *IEEE Proceedings on Software Engineering*, 144(1) 11–25.
- Jennings, NR, Faratin, P, Johnson, MJ, Norman, TJ, O’Brien, P and Wiegand, ME, 1996. “Agent-based business process management” *International Journal of Cooperative Information Systems*, 5(2&3) 105–130.
- Kalakota R and Whinston, AB, 1996. *Frontiers of Electronic Commerce* Addison Wesley, pp. 36–37.
- Luck, M, Griffiths, N and d’Inverno M, 1997. “From agent theory to agent construction: A case study”. *Intelligent Agents III: Proceedings of the Third International Workshop on Agent Theories, Architectures and Languages; Lecture Notes in Artificial Intelligence* 1193 49–63, Springer-Verlag.
- Minsky, M, 1986. *The Society of Mind* Simon and Schuster.
- Newell, A, 1982. “The knowledge level” *Artificial Intelligence*, 18(1) 87–127.
- Nwana, HS, Ndumu, DT and Lee, LC, 1998. “ZEUS: An advanced tool-kit for engineering distributed multi-agent systems”. In: *Proceedings of the 3rd International Conference on Practical Applications of Intelligent Agents and Multi-Agent Technology* 377–391.
- Rao AS, 1996. “Agentspeak(l): BDI agents speak out in a logical computable language” In: W Van de Velde and JW Perram (eds), *Agents Breaking Away: Proceedings of the Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Lecture Notes in Artificial Intelligence* 1038, pp. 42–55. Springer-Verlag.
- Rao, A and Selvestrel, M, 1992. “Agent-oriented architecture for air combat”. In: *Proc. Australian Joint Conference on Artificial Intelligence*.
- Russell, SJ, 1997. “Rationality and Intelligence” *Artificial Intelligence*, 94(1) 57–77.
- Wooldridge, M, 1997. “Agent-based software engineering” *IEEE Proceedings on Software Engineering*, 144(1) 26–37.