The Knowledge Engineering Review, Vol. 14:2, 1999, 151–152
Printed in the United Kingdom

Diversity and agent technology

MICHAEL WOOLDRIDGE

Department of Electronic Engineering, Queen Mary & Westfield College, University of London, London E1 4NS, UK (email: M.J.Wooldridge@qmw.ac.uk)

When I began studying for my PhD, I knew exactly what field I wanted to work in. It was a subfield of Artificial Intelligence (AI) that was known as Distributed AI (DAI). In DAI, we studied how a number of (semi-)autonomous problem solving systems known as 'agents' could be made to cooperate and coordinate their activities to produce effective overall system behaviour (Bond and Gasser, 1988). At the time, not many people were working in distributed AI (only two books had been published on the subject). I found that when I talked about my research, I had to repeatedly explain what 'agents' were all about. It was depressing business. More than once, I wished I had gone into a more glamorous area of research—the expert systems people were getting all the glory.

Today, the situation is very different. When I talk to someone working in AI or computer science about my research, there is now a good chance that they will have at least heard of agents. It is no exaggeration to say that agents are currently one of the most vibrant and active areas of research and development, not just in artificial intelligence, but in computer science generally. The degree of interest, but more importantly, the *breadth* of interest, continues to surprise me. From the KER special issue on *Perspectives on Recent Intelligent Agents Research*, we can see that agents are of interest not just to what might be called 'good old fashioned AI', but also to researchers from user interface design, robotics, networks, electronic commerce, online publishing, and many other disciplines besides. From the KER special issue, we can identify at least four major schools of thought that go under the 'agent' banner.

1. One view of agents comes from the AI/DAI communities, and originates from work on AI planning (Allen et al., 1990) and distributed problem solving (Smith, 1980). See Wooldridge and Jennings (1995) for a detailed overview. Here, the main concern is that many problems may naturally be solved by the interaction of semi-autonomous agents.
2. More recently, a community has emerged that is concerned with using agents in user interfaces (see, e.g., Maes (1994) for a typical view). Here, the main concern is that of moving away from the so-called *direct manipulation* paradigm that has dominated in user interface design, towards applications that can take the initiative and act more like a cooperative assistant to a user working on some task.
3. Another community that has evolved from the object-oriented programming paradigm is that of *mobile agents* (White, 1994). Here, the issue is to do with program code that can execute a 'go' instruction, which causes the program and its state to transmit itself across an electronic network to a remote site, whereupon it recommences execution. Issues such as privacy and security become paramount when considering such agents.
4. Finally, the robotics community has a vested interest in autonomous decision-making in robotic systems, which links the field closely with agents.

O'Brien et al. discuss agents of the first type; Crabtree, Johnson and Hayes-Roth, Guttman, and Rao all discuss agents of the second type; Johnson and Hayes-Roth discuss agents of the third type; and Bekey discusses agents of the fourth type.

While this breadth of interest is in many ways gratifying, I believe it also raises some uncomfortable issues. In particular, it is no longer clear who 'owns' agents. In 1989, when

computational multi-agent research was, by and large, limited to the DAI community that I mentioned above, the computational multi-agent research agenda *was* the DAI research agenda, by default: nobody else was very interested. But the community that originated from DAI research would certainly not lay sole claim to the field today.

The diversity of interest in agent technology can in one sense be seen as a major weakness. In particular, after reading *The Knowledge Engineering Review* (KER) **13**(2), someone unfamiliar with the literature and history of agent technology might be forgiven for leaping to the conclusion that the term 'agent' had become utterly devoid of meaning. There are so many different types of agents on offer in KER **13**(2) that it is hard to see any commonality of interest or purpose.

There is some strength to this argument. Dave Cliff, speaking at the third workshop of the UK Special Interest Group on Agent Technology (UK-MAS), joked that probably the only people currently making money from agent technology were 'snake oil salesmen', cashing in on the general interest in agents by offering content-free agent products to a gullible public. Associated with this phenomenon is what appears to be a trend to 'rebrand' work that, to my mind at least, has little or nothing to do with agents. Presumably, it is believed that a research paper or grant proposal is more likely to be supported if it is somehow agent-related. The diversity of interests in agent technology contributes to the phenomenon—for who is to say what does and does not count as an agent?

For me, this tendency to rebranding is potentially the most dangerous single threat to the long-term future of agents. If it continues, it will surely lead to cynicism on the part of the computer science community at large. And this community represents the ultimate marketplace in which any information technology will succeed or fail.

But this diversity, and the fact that nobody 'owns' agents is, I believe, also a strength. It could be argued, for example, that the key weakness of expert systems technology was that, in the perception of the computing community at large, it was tied to a single technique—rule-based systems. Agent technology cannot be accused of having all its technological eggs in one basket—the field is far too diverse.

The fact that the concept of an agent is common currency to so many widely different fields surely suggests that the concept of an agent is in some sense *right*. To draw an analogy, just as the world is populated by objects that have state and behaviour but which are essentially *passive*, so it is also populated by *active* entities. Object-oriented programming has proved to be such a successful paradigm at least in part because it allows us to model the world in terms of entities—objects—that more closely reflect the way in which we ourselves understand it (Booch, 1994). Agents are another such abstraction: they allow us to understand and model the world in terms of entities—indivduals, organisations, groups, and so on—that we see in the everyday world. For me, this idea is now second nature. I sincerely believe that agents will have value in helping us to understand how to design and implement a range of computer systems, in many disciplines.

References

Allen, JF, Hendler, J, Tate, A, editors, 1990. *Reading in Planning* Morgan Kaufman.

Bond, AH and Gasser, L, editors, 1988. *Readings in Distributed Artificial Intelligence* Morgan Kaufman.

Booch, G, 1994. *Object-Oriented Analysis and Design (second edition)* Addison-Wesley.

Maes, P, 1994. 'Agents that reduce work and information overload' *Comm ACM* **37** 31–40.

Smith, RG, 1980. *A Framework for Distributed Problem Solving* UMI Research Press.

White, JE, 1994. 'Telescript technology: The foundation for the electronic marketplace' White paper. General Magic, Inc., 2465 Latham Street, Mountain View, CA 94040.

Wooldridge, M and Jennings, NR, 1995. 'Intelligent agents: Theory and practice' *The Knowledge Engineering Review* **10**(2) 115–152.