

Editorial

Intelligent agent technology is at an interesting point in its development. Commercial strength agent applications are increasingly being developed [1]. Furthermore commercially supported development environments are available, e.g. [2,3] and design methodologies [4,5], reference architectures [6] and standards [7] are beginning to appear. These are all strong indicators of a mature technology. However, the uptake of the technology is not as rapid or as pervasive as its advocates have expected. Is intelligent agent technology simply in need of the killer application, or are there more fundamental reasons as to why a technology that promises so much has not seen more widespread acceptance?

An intelligent agent, as defined in [8], is a computer program that is capable of acting autonomously according to the input from the environment in order to achieve its design objectives. From a software engineering perspective, one would expect to gain major benefits from the deployment of such technology in complex distributed applications such as virtual enterprise management and the management of sensor networks. Based on our experience in developing industry-strength agent applications in the defence and manufacturing sectors over the past decade, we would suggest that the following reasons are responsible for the limited uptake of agent technology:

1. The software engineering reason: the agent-based software engineering process is not sufficiently mature to be applied to the large, mission-critical projects that would most benefit from agent technology.
2. The research reason: the existing agent theories and reasoning models are inadequate and that further basic research into these areas is required before the agent-oriented paradigm will be widely accepted
3. The autonomy reason: acquisition of the behaviours required for an agent to exhibit autonomous operation is too difficult

Each of the papers in this special issue presents research results pertinent to one or more of the reasons listed above. The two papers by Jarvis et al. and the papers by both Lam and Lee and Huang et al. indicate that industrial strength agent frameworks are readily available and can be used for the development of complex agent-based applications. However, more applications that explore the boundaries of existing frameworks and theories need to be developed. Not only are such applications needed to extend both frameworks and theories, as exemplified by Jarvis et al's paper on BDI plan step execution, they are also needed to both ground and direct research into methodological issues. This latter point underpins the papers by both Lam and Lee and Huang et al. In terms of theories and reasoning models, the BDI model [9] and Newell's unified theory of cognition [10] have provided the theoretical underpinnings for the frameworks that have been preferred for the development of knowledge intensive agents [11]. However, as Jarvis et al. demonstrate, these models, while having a stable core, will continue to evolve to address issues such as coordinated team behaviour and resource management. Mobile agents, as exemplified in the paper by Lovrek and Sinkovic, present another area that promises to extend the boundaries of existing agent theories.

However, perhaps the biggest challenge facing intelligent agent technology is the issue of autonomous behaviour. The current practice in fielded applications is to explicitly model the reasoning processes involved in agent behaviour. This invariably involves a significant behaviour acquisition or construction phase and results in agents that can only function within well defined environments. Providing agents with the ability to learn behaviour presents an attractive mechanism to overcome this limitation and to adapt to new situations. The papers by Fyfe et al., Leng et al. on the one hand, and Lovrek and Sinkovic on the other, indicate contrasting approaches towards this goal. The first two papers are concerned with the development and subsequent application of generalised algorithms for the learning of agent behaviour. In contrast, Lovrek and Sinkovic develop a mobility management model for personal agents which will then be used as a basis for the learning of regularities in movement patterns. Both approaches are needed to advance the field.

At the beginning of this introduction, we posed the question as to why intelligent agent technology has not achieved the success that has been expected by its advocates. We would argue that the answer can be found in the papers presented in this issue. While there are major software engineering issues that need to be addressed and while agent theories will continue to be extended and new theories developed, the key issue is the provision of autonomy. The successes of existing intelligent agent frameworks have typically been achieved in domains that are characterised by large bodies of predetermined behaviour, such as manufacturing or the military. These applications are characterised by explicit reasoning about particular situations, a process that is both time consuming and inherently limiting. Agent learning provides one approach for overcoming these difficulties, but as the papers in this issue indicate, much work needs to be done before it can be deployed in mission critical systems

Ajith Abraham, Dennis Jarvis, Jacquie Jarvis and Lakhmi Jain (Guest Editors)

References

1. Evertsz, R., Fletcher, M., Jones, R., Jarvis, J., Brusey, J. and Dance, S.: Implementing Industrial Multi-agent Systems using JACKTM. In: Dastani, M., Dix, J. and El Fallah Seghrouchni, A. (eds): Programming Multi-agent Systems, First International Workshop, PROMAS 2003. Lecture Notes in Artificial Intelligence, Vol. 3067. Springer, (2004) 18-48
2. <http://www.agent-software.com>
3. <http://www.soartech.com>
4. Bussmann, S., Jennings, N. and Wooldridge, M.: Multiagent Systems for Manufacturing Control. A Design Methodology. Springer, (2004)
5. Padgham, L. and Winikoff, M.: Developing Intelligent Agent Systems: A Practical Guide. Wiley, (2004)
6. Van Brussel, H., Wyns, J., Valckenaers, P., Bongaerts, L. and Peeters, P.: Reference Architecture for Holonic Manufacturing Systems: PROSA. Computers in Industry. 37 (1998) 255-274
7. <http://www.fipa.org>
8. Wooldridge, M.: An Introduction to MultiAgent Systems. Wiley, (2002)
9. Rao, A. and Georgeff, M.: BDI Agents: from theory to practice. In: Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS'95), San Francisco, California, USA.(1995) 312-319.
10. Newell, A.: Unified Theories of Cognition. Cambridge, MA: Cambridge University Press, (1987)
11. Jones, R.M. and Wray, R.E.: Comparative Analysis of Frameworks for Knowledge-Intensive Intelligent Agents. AI Magazine, 27, No. 2, (2006) 57-70.