
Intelligent Agents \equiv Learning Agents

Peter Edwards

Department of Computing Science

King's College

University of Aberdeen

Aberdeen, AB9 2UE

pedwards@csd.abdn.ac.uk

1 Introduction: Why Agents Must Learn

The title of this workshop (*Intelligent Agents*) will almost certainly mean different things to different people; some will interpret it as being a forum for discussion of issues in multi-agent systems, while others will consider it the appropriate place to present the latest results of intelligent user-interface research or mobile agent technology. The term *agent* is troublesome, simply because it is used by so many different communities within Computer Science. Each community has its own unique definition of the term; however, one characteristic which seems to be common is the concept of agents as *autonomous* entities.

In essence, my thesis is that to be truly autonomous, agents must employ some form of learning capability, i.e. that intelligent agents are learning agents. By *learning agent*, I mean any agent which learns about the environment in which it operates, the user(s) it interacts with, or other agents. In the next section I present a number of scenarios to illustrate this point.

2 Some Not So Intelligent Agents ...

1. The *Phil* system is an information filtering agent capable of assisting users with management of Internet-based information, such as electronic mail and USENET news. The system is programmed using a scripting language. Scripts contain rules which describe the characteristics of mail messages/news articles and the actions which should be performed on them, such as deleting, forwarding, etc. One day, *Phil* encounters a mail message, executes the appropriate script and deletes the message. Unfortunately the message was from the user's superior and should not have been deleted. The next day, another message is received from the same person and this too is deleted. The only option available to the user is to hand edit the rules.
2. A mobile network agent is despatched to find and book a skiing holiday for its user. The agent traverses the network to the HoliSun booking agency. Unfortunately, this host only deals in Summer holidays, and as a result, the agent has to move on to another to find an appropriate holiday. Visiting an inappropriate host obviously cost the agent (and therefore the user) valuable resources. To avoid repetition of this failure in future, the user must explicitly re-program the agent.
3. A group of agents are provided with knowledge about a domain and one another's capabilities. This

community is presented with a problem and the agents cooperate in an attempt to solve it. Due to the incorrect knowledge of one agent, the group fails to achieve its goal. The knowledge of the faulty agent must be repaired, but how ?

It should be noted that the agents described above are purely fictitious; any relationship to real agents, either living or dead, is purely coincidental.

Although each of these agents demonstrates some degree of autonomy, I would argue that their behaviour falls short of the level required for a truly intelligent agent. In each case, a learning component would have allowed the agents to adapt to their environment and/or user. Work within the *Novel Agent Systems & Applications* group at Aberdeen focuses on the role of learning within agent technology. The next section provides an overview of our activities.

3 Research at Aberdeen into Agents that Learn

Our work can be divided into two broad areas:

- **Intelligent Information Management/Retrieval** We have developed a number of systems which employ learning techniques to observe users performing tasks such as processing electronic mail [8, 9, 11], reading USENET news articles [6, 7, 11] or browsing World-Wide Web pages [1, 7]. All of these systems are capable of learning a user's preferences over time so that tasks may be automated; in addition, they are able to adapt to changes in user behaviour. Work on these systems led directly to a new learning algorithm [10], designed to be used in agents of this type.

Agent-Based Knowledge Discovery [4, 5] provides a new technique for performing data-mining over distributed data. Specialized learning agents are employed which communicate using KQML performatives.

Other current activities include the development of a personalized electronic newspaper agent which learns a profile of a user's interests, and then assembles a newspaper by retrieving stories and pictures from several Web-based newspapers; a personalized author/book recommendation system; an agent-based system to support distance learning.

- **Adaptive Multi-Agent Systems** As stated earlier, a community of agents which interact to solve a problem are only as good as the knowledge they possess about the domain and each other. The DRAMA system [2, 3] addresses this problem by providing a mechanism for knowledge refinement within agents. Learning techniques are used to identify the key characteristics of successful and unsuccessful community behaviour; these are then used to guide the refinement process.

4 Summary

The agents briefly described above all employ learning to overcome weaknesses of the kind outlined in Section 2. To close, I believe that developers of agents (of whatever hue) must come to terms with the fact that their systems should learn, if they are to be worthy of the title *intelligent agents*.

Acknowledgments

I am indebted to the following, all of whom have contributed to the ideas expressed here by developing learning agent technology at Aberdeen: David Bayer, Ciara Byrne, Winton Davies, Claire Green, Ron Mackenzie, Terry Payne, Andrew Waterson. Our work has been supported by the UK Engineering & Physical Sciences Research Council, the University of Aberdeen Research Committee, and British Telecom.

References

- [1] D. Bayer. A Learning Agent for Resource Discovery on the World Wide Web. MSc Thesis, Department of Computing Science, University of Aberdeen, Scotland, 1995.
- [2] C. Byrne and P. Edwards. Collaborating to Refine Knowledge. In *ML95 Workshop on Agents that Learn from Other Agents, Tahoe City, California*, 1995.
- [3] C. Byrne and P. Edwards. Refinement in Agent Groups. In *Proceedings of IJCAI95 Workshop on Adaptation & Learning in Multi-Agent Systems*, 1995. To be published by Springer-Verlag as part of the Lecture Notes in Artificial Intelligence series.
- [4] W.H.E. Davies and P. Edwards. Agent-Based Knowledge Discovery. In *Proceedings of AAAI Spring Symposium on Information Gathering from Distributed, Heterogeneous Environments*, 1995.
- [5] W.H.E. Davies and P. Edwards. Distributed Learning: An Agent-Based Approach to Data-Mining. In *ML95 Workshop on Agents that Learn from Other Agents, Tahoe City, California*, 1995.
- [6] C.L. Green. USENET News Agent. BSc Final Year Project Report, Department of Computing Science, University of Aberdeen, Scotland, 1995.
- [7] C.L. Green, D. Bayer, and P. Edwards. Towards Practical Interface Agents which Manage Internet-Based Information. Submitted to Intelligent Agents Workshop, BCS Specialist Group on Expert Systems and Representation & Reasoning Special Interest Group, 1995.
- [8] T.R. Payne. Learning Email Filtering Rules with Magi, A Mail Agent Interface. MSc Thesis, Department of Computing Science, University of Aberdeen, Scotland, 1994.
- [9] T.R. Payne and P. Edwards. Interface Agents that Learn: An Investigation of Learning Issues in a Mail Agent Interface. Submitted to Applied Artificial Intelligence, 1995.
- [10] T.R. Payne and P. Edwards. Learning Mechanisms for Information Filtering Agents. Submitted to Intelligent Agents Workshop, BCS Specialist Group on Expert Systems and Representation & Reasoning Special Interest Group, 1995.
- [11] T.R. Payne, P. Edwards, and C.L. Green. Experience with Rule Induction and k -Nearest Neighbour Methods for Interface Agents that Learn. In *ML95 Workshop on Agents that Learn from Other Agents, Tahoe City, California*, 1995.