Knowledge Management Agents

Leon Sterling
University of Melbourne, Australia

APPLYING AGENTS WITHIN KNOWLEDGE MANAGEMENT

The agent has existed as a concept for thousands of years. In the human context, an agent is a person that performs some task on your behalf, for example, a travel agent planning flights and accommodation for your holiday, or a real-estate agent helping you buy or sell a house, or someone arranging marriages. Some Biblical laws specifically refer to agents.

In the much more recent software context, an agent is loosely a program that performs a task on your behalf. Agents have grown in popularity since the introduction of the PC (personal computer) as the target environment for application software has increased in complexity. Software systems must now operate robustly in a networked, global environment comprised of diverse, distributed technologies. Furthermore, the environment is dynamic, and frequent change is inevitable. Having automated help is almost a necessity.

Despite many attempts, there is no universally agreed technical definition of agents. An oft-cited reference by Franklin and Graeser (1996) gives almost a dozen different definitions. Let us consider a textbook definition given by Wooldridge (2002, p. 15). An agent is "an encapsulated computer system, situated in some environment, and capable of flexible autonomous action in that environment in order to meet its design objectives."

Essential characteristics of the agent paradigm that can be elicited from this definition are:

- The autonomy of individual agents, or their ability to act for themselves and to achieve goals
- The reactivity of individual agents in response to changes in the environment
- The modularity of individual agents and classes to allow the easy development of complex systems
- The ability of agents to communicate effectively and interact with legacy systems

Optional characteristics of the agent paradigm, which emerge from broader considerations of agents than the above definition, include mobility in moving around a network and the ability to reason.

This article rests on the metaphoric view of agents as entities performing tasks on one's behalf. Agents are presumed useful for building software to interact with complex environments such as the Internet or within complex organizations such as universities and multinational corporations. Expected of a program being viewed as an agent is an ability to sense and be aware of the environment in which it is situated, an ability to communicate with other agents, and an ability to take action in its situated environment. According to these three expectations, sophisticated e-mail programs such as Microsoft's Outlook and Qualcomm's Eudora can be viewed as agents. They are situated on the Internet and sense various aspects of the Internet, including when Internet connections are live and when new mail arrives. They communicate with other e-mail clients by sending and receiving messages. They take actions such as raising alerts when mail has arrived, sending mail that has been queued once an Internet connection is restored, or filtering messages according to rules.

We now connect with knowledge. Organizations operating in today's software environment need to represent, interact with, and above all, maintain a large collection of knowledge, including, for example, business practices, trade secrets, intellectual property, organizational hierarchies, promotional organizational descriptions, and knowledge of both its own policies and policies of relevant, external regulatory bodies. There is out of necessity great diversity in the form, content, and context of the knowledge. Most of this knowledge is in unstructured or semistructured form. The problem of the representation and maintenance of such knowledge within an organization can be loosely called the knowledge management problem.

For the purposes of this article, there is no need to define the knowledge management problem or knowledge management, for that matter, more precisely. However, we note that the term knowledge management subsumes the term content management. Referring to knowledge rather than content suggests some concern with formalizing knowledge explicitly.

How might agents be applicable to the knowledge management problem? As a running, concrete example, consider knowledge management issues related to the responsibilities of a university lecturer in charge of a subject[1]. She or he must prepare, deliver, and maintain content in a variety of forms, possibly including lecture notes, papers, and media presentations. Let us particu-

larly focus on one component of the task, namely, maintaining a Web site for the subject.

Several possibilities exist for enlisting the help of agents. An obvious first task for agents is to help with the acquisition of knowledge, which is obtaining content and placing it on the Web site.

What type of software agent might be useful for the acquisition of knowledge? It is natural to envisage a custom Web crawler (http://en.wikipedia.org/wiki/Web_crawler), Programs that trawled specified Web sites looking for content were early applications built to exploit the World Wide Web. Building a Web-crawling agent immediately raises important considerations. The agent should be aware of important regulatory issues such as the fact that downloading mp3 files is illegal in some countries without the authorised permission of the copyright holder. The agent should also be aware of conventions such as the robots.txt protocol (http://en.wikipedia.org/wiki/Robots.txt_protocol) in which guidelines are given about parts of a file hierarchy that should be ignored by well-behaved agents. There are many similar policies of which a knowledge-acquiring agent would need to be aware. These policies demonstrate some of the complexities that need to be taken into account in building agents.

Search engines are based on exhaustive trawls and efficient indexing of files using techniques from information retrieval. Agents can also be constructed using techniques derived from experience in building knowledge-based expert systems. Consider the task of tracking down a particular paper by a particular author. One may have been referred to the paper by word of mouth or by the need to cite a final version of the paper for which you only had a preliminary version. A prototype citation-finding agent, CiFi (Loke, Davison, & Sterling, 1996), was built for this task. CiFi used the following three strategies for finding papers. First, CiFi tried to find a link from the author's home page[2] using heuristics about possible keywords such as research and publications. Then CiFi looked for a link from a page of publications or technical reports linked from the author's department. Finally, CiFi sent an e-mail to the author asking for the file or citation.

A challenge in building CiFi was making it work on a variety of Web sites. Ideally, a single agent is desirable that can operate successfully over a range of Web sites. CiFi was not particularly intelligent or effective. It clearly reflected a bias to papers written by researchers within universities. It failed, for example, to find white papers written by companies. It would have had difficulty adapting to current spam filters if its e-mail message was blocked. Having agents adapt to changing circumstances is a desirable property. However, CiFi is indicative of an agent that might be applied to a knowledge management task.

Another task that might be assigned to an agent is to look for new articles by particular authors. Suppose you respect the work of a particular researcher and want to be notified of any of his or her new publications. It is possible for an agent to look for changes on a Web site and alert you that a new publication may be present[3]. In general, reporting changes or the presence or absence of documents is a task that the reader should have no difficulty in identifying as being potentially useful within his or her own organization. Providing new information or reminding participants that the next step in a work flow needs to happen can be helpful. Such an agent can be viewed as being a facilitator. Facilitation was espoused by Winograd and Flores (1987) as an alternative model for agents rather than artificial intelligence.

Let us return to specific tasks within Web-site management. Content on a Web page may be made available through links to other resources. However, it is frustrating when browsing to find outdated links on Web pages. An agent could check periodically whether links are still live. It would need to sense the result of its search and update the links on the page.

Here is another task for an agent. Some of the knowledge on a Web site can be usefully cross-linked. For example, an online quiz would be enhanced for self-study by having links from questions to material where the correct answer can be found. These links may be provided once the student has attempted the quiz. An agent could construct these links automatically. Of course, any changes to content would mean that the cross-links would need to be checked. A prototype agent called QuizHelper that can perform this task has been described in Chan and Sterling (2003).

Several of the above suggestions for agents address the performance of maintenance activities. Maintenance is key for knowledge management. The reader can doubtless imagine maintenance activities in his or her own environment that might be performed by an agent. Some maintenance activities are already happening automatically, for example, through alerts about software updates or the downloading of security patches.

A different type of task that an agent can perform is monitoring the use of a program.

For a program developed to help students learn material, an educator may want an agent to assess if the program is being used properly by the students. The assessment may be used to give feedback to the software developers or to try to ascertain whether the student is meeting learning objectives. The conceiving of agents to monitor student interaction with a program suggests good design questions. How is the agent going to sense