INTELLIGENT AGENTS FOR BUSINESS AND PLEASURE[1]

Frederick H. Lochovsky

Department of Computer Science, University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong
Email: fred@cs.ust.hk

1 Introduction

With the ubiquitousness of computing technology, computers are directly touching more and more lives, even in the remotest and least developed places in the world. At the same time, the amount of information available electronically, especially via the World Wide Web (WWW), is also rising dramatically. Thus, on the one hand we see computing technology being made available to an ever wider, but technologically less sophisticated, group of users, while on the other hand, the complexity of using this technology is constantly increasing as more and more information and applications become available and easily accessible over networks like the WWW. These two trends give rise to a major problem faced not only by individuals, but also by businesses: How to make effective use of computing technology. More specifically, how to use it to find useful and relevant information to meet personal and business needs.

With regard to the latter problem, as already stated, the amount of information available via the WWW is increasing daily. While the WWW provides an effective means for delivering this information to individuals and businesses, it does not assist them either in easily locating the information they need or in using it effectively. They are, by and large, left to fend for themselves, having to learn all of the intricacies involved in accessing and searching the WWW. Granted that there are search engines available for locating information on the WWW [19]. However, their coverage is not uniform and many of them are by no means easy to use. WWW searches can be very complex to specify and this assumes that one already has a fairly good idea of what one is looking for and can specify it precisely enough for the search engines to return useful results. As the WWW grows in size and diversity of information available, this problem will only become more acute.

For those individuals and businesses that can afford it, one solution to this problem, of course, is to hire specialists who can write programs to make effective use of computing technology and the WWW. However, this only means that the gulf between those that are "information rich" and those that are "information poor" will widen and that accordingly, many, if not most, individuals and businesses will be unable to participate in and benefit from the new information age quickly coming upon us.

One technology that shows great promise in being able to address this problem is intelligent agents. Intelligent agents can help reduce the workload on people and businesses by hiding the complexity inherent in using computing technology and accessing the WWW. They may also aid in eliminating human errors and biases [12].

In this paper, we will explore the nature of intelligent agents and show how they could be used to assist both individuals and businesses in more effectively using computing technology and the WWW. In Section 2, we discuss the nature and properties of intelligent agents. Section 3 surveys some emerging applications of intelligent agents. Section 4 examines the challenges that intelligent agent technology poses for businesses, individuals and society. Finally, Section 5 summaries our discussion.

[1] This paper was prepared while the author was on sabbatical leave at Boeing Information and Support Services, Bellevue, WA, U.S.A.

2 Intelligent Agents

Many different notions of intelligent agent have appeared in the literature [10, 18, 19]. For our purposes here, we will take a fairly broad view and define an intelligent agent (or agent for short) as a software program that is capable of autonomous (or semi-autonomous) behaviour. Such an agent acts as a representative of a user (individual or business) and interacts directly with either other agents or with users. The behaviour may be dependent on the particular context in which an agent finds itself (i.e., an agent may be able to monitor and react to its environment). Usually, there is also some requirement for an agent to be able to gather and perhaps integrate information from several different information sources. In general, agents abstract away details of their underlying technology and the resources they access. They may also provide distraction for the user in that an agent's character or persona helps to distract a person from what might be a tedious task [5].

generalized
learns
problem-solving based
uses common sense
multiple problem-solving paradigm

specialized
dumb
procedural
no common sense
single problem-solving paradigm

Figure 1. Intelligent agent spectrum.

The "intelligence" of an agent can range from those who monitor user actions and make suggestions on how to proceed to those who completely isolate the user from the intricacies of interacting with computer software and with a few simple instructions can carry out complex and partially unspecified tasks for the user [5, 16]. Figure 1 graphically shows this "intelligence spectrum" of agents. At one end of the spectrum, agents may be merely highly specialized and sophisticated programs that outwardly exhibit intelligent behaviour, but are really quite rigid and inflexible in how they accomplish a task. At the other end of the spectrum, they can react and adapt to their environment and exhibit multiple approaches to solving a problem. For this latter type of behaviour, an agent requires some on-board intelligence, such as planning, reasoning and learning algorithms, to help it achieve intelligent behaviour. It must also be able to accommodate conflicting goals, choose among conflicting goals, and reason about how best to accomplish the goal chosen [11].

The desirable/necessary properties that intelligent agents should have are numerous and diverse, and not everyone will agree with any one given list [5, 20]. Here, we present a somewhat composite list. An intelligent agent should be:

- *active*: An agent should be able to decide when to act based on its environment. For example, it should be able to anticipate a user's need for information and to initiate problem solving activities and integrate information from different sources to satisfy those needs.
- *adaptive*: An agent should be able to accommodate changing user needs and task environments and to customizes itself to preferences of its user based on previous experience.
- *autonomous*: An agent should be able to accept high-level requests (from either humans or other agents) and exercise non-trivial control over its actions. In particular, an agent should not always be under direct human control, though the degree of autonomy may be controlled by a human. This is, arguably, the most important and distinguishing property of agents versus other computer programs. Without autonomy (the ability to initiate independent actions), an agent is really no different than a traditional computer program albeit a sophisticated one.

- *collaborative*: An agent needs to be able to interact and cooperate with humans and other agents. For example, it may need to engage in complex communication with other agents (including people) to obtain information or enlist help.
- *goal-oriented*: To use an agent, a user request should only need to specify what is wanted. The agent decides how, where and when to satisfy the request.
- *persistent*: An agent should be a continuously running process and not just a one-time computation. This allows an agent to possibly learn new behaviours over time and to be available when needed by its users.
- *trustworthy*: An agent needs to be able to reliably serve its users' needs so that users will develop trust in the agent's performance. Without such trust, users will be reluctant to delegate tasks to agents.

In the rest of this section, we elaborate on some additional important properties that will help us better understand the nature and requirements of intelligent agents.

2.1 Types of Agents

Having one "giant" agent to handle all aspects of agent communication and processing is not a good idea because such an agent would need an enormous amount of knowledge to deal effectively with user requests. It would also be a processing and communication bottleneck and a possible single point of failure. In very complex information processing environments, the required processing would probably overwhelm an agent because of the complexity of the information finding and filtering task. Thus, different types of agents, that specialize in performing only certain functions are needed. One way to divide agent functionality is to have at least three types of agents—interface agents, task agents, and information (extraction) agents (see Figure 2) [20].

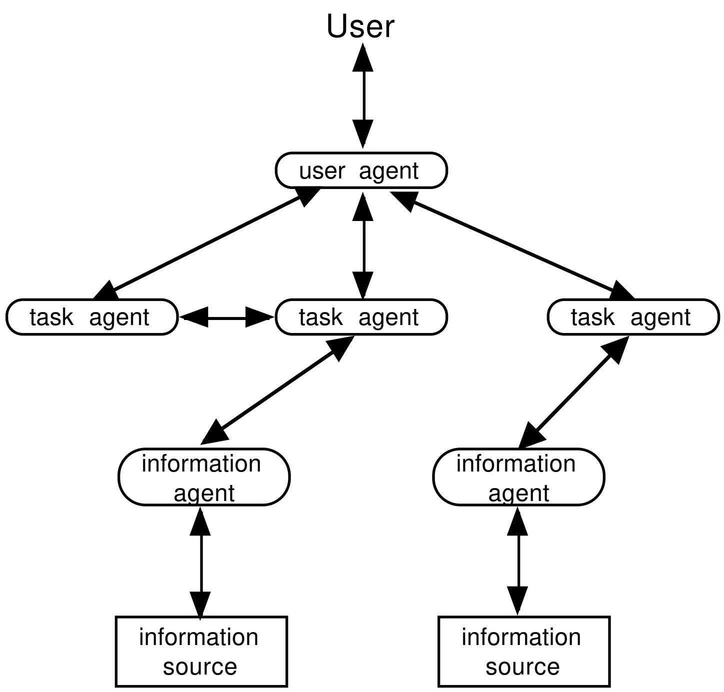

Figure 2. Types of agents.

An *interface agent* is an agent that interacts with a human. Such agents acquire, model and utilize information about user preferences to guide them in coordinating the user's tasks. To do this, they need to be able to:

- collect relevant information from the user to initiate a task;
- present relevant information, including results and explanations to the user;
- ask the user for more information during problem-solving;
- ask for confirmation of actions when necessary.

User interface agents hide the underlying distributed nature of information-gathering and the problem-solving complexity by accepting high-level requests from users and breaking such requests down into smaller tasks that are assigned to task agents.

A *task agent* supports the achievement of a user's goals by formulating problem-solving plans and carrying them out through querying and exchanging information with other agents. Task agents have specialized knowledge of a task domain. They know what other agents to call to carry out even more specialized tasks. Task agents perform most of the autonomous problem-solving in an agent system.

An *information agent* is essentially a type of task agent. However, its role in an agent system is very important as it provides intelligent access to heterogeneous information sources by interacting with and extracting information from those information sources. To do so, it needs to have a model of the associated information source as well as strategies for information source selection, information access, conflict resolution and information fusion. It may also monitor an information source for pre-specified conditions and notify the user or another agent when a pre-specified condition arises.

The preceding discussion illustrates the need for several types of agents in an agent system to distribute and share the workload. Other specialized types of agents, other than those discussed above, may also be required depending on the application domain. The fact that more than one agent may be involved in accomplishing a given user task leads to the need for agent cooperation and communication which we discuss further later in this section.

2.2 Instructing Agents

Just like a newly hired human personal assistant, agents may come with some generic information gathering and problem-solving knowledge, but they will not have specific knowledge of the task domain in which they will be used. Yet an agent's user will want their agent to become competent in the specific task domain [15]. Therefore, users of agents will have to have some means of instructing their agents in the specific knowledge required to become competent in the task domain. In addition to competence, users also want their agents to be trustworthy. That is, a user should feel comfortable delegating tasks to his agent. These two issues go somewhat hand in hand as the more a user instructs an agent the more likely he is to trust that agent.

To allow an agent to acquire competence in a task domain, agent systems have to provide a mechanism for allowing users to instruct agents. There are two main existing approaches to allowing users to instruct agents. In the first, the user "programs" his agent(s). Often this programming is in the form of specifying rules for agents to follow. For example, in the Oval system a user specifies an email processing agent by creating a number of rules that process incoming messages [13]. To be able to program agents in this way, a user needs to have competence to program an agent to do the right thing, at the right time and in the right way. An agent's competence is therefore just as good as the user's ability to specify the correct rules. Provided that the user has faith in his programming abilities, such agents should be highly trusted as the user has explicitly instructed the agent in all details of task processing.

In the second approach, agents are endowed with extensive domain-specific background knowledge about the task domain and the user. Knowledge engineers are required to provide agents with this knowledge. Agents use this knowledge to recognize user goals and plans and to find opportunities to contribute to them. The UCEgo agent, which helps users solve problems related to using the Unix operating system, is an example of this type of agent [1]. Usually in this approach, agent knowledge is fixed, is not reusable in another domain, and cannot be customized to individual users. Thus, while agent competence may be high, it can decline over time as the environment changes. Also, the degree of trust that the users have in such agents may be lower than in the previous approach as agents are a "black box" to users who may not be sure of exactly what an agent does and does not know.

A third, hybrid, approach is being explored in several research projects [2, 3, 15]. In this approach, an agent is initially endowed with some minimum background knowledge as in the second approach above. However, unlike the first or second approach, the agent then learns additional task knowledge by observing its environment and interacting with its user(s). Thus, like a personal assistant, over time the agent acquires more and more knowledge about the task domain as well as customizes itself to the particular individual with which it is interacting. Such an agent learns in four ways [15]:

- by observing the actions of a user and noticing repetitive patterns of behaviour which it can offer to automate;
- by receiving direct and indirect feedback from the user;
- from examples given explicitly by the user to the agent;
- from other agents that perform similar tasks for other users.

This third approach better addresses the problems of agent competence and trust. Agents built using this approach become more and more competent over time as they learn from their environment and user. Such agents can also be more readily trusted by their users since their knowledge is acquired from their users and demonstrated to be trustworthy over time.

2.3 Agent Cooperation

While an agent working on its own can be very useful, some people claim that real competence emerges from a large number of agents working together [9]. To allow agents to work together cooperatively, they need mechanisms to exchange information and services so that they can solve problems that cannot be solved by each alone. Several issues need to be considered to allow agents to cooperate. Among these are inter-agent communication capabilities, agent competence finding/matching capabilities, and developing a common view of the world to facilitate inter-agent communication and task sharing.

One of the first requirements for agents to cooperate is that they speak the same language. Such a common language allows agents to easily exchange information and make requests of each other. Many of the agent research efforts are using the Agent Communication Language (ACL) developed by the ARPA Knowledge Sharing Effort [6, 8]. This language consists of three parts:

- a vocabulary which is a large open-ended dictionary of words described by both an English description for use by humans and a formal notation (Knowledge Interchange Format) for use by agents;
- an inner language called Knowledge Interchange Format (KIF) which is an extended form of first order predicate calculus and is used as the basis for the Knowledge Query and Manipulation Language;
- an outer language called Knowledge Query and Manipulation Language (KQML) which uses the vocabulary to provide a linguistic layer on top of KIF in which task context is taken into account.

For agents that do not speak KQML, translation agents can be built that translate between KQML and the language understood by the agent (see Section 3.1).

If agents are to cooperate, then they need to be aware of the capabilities possessed by other agents. An agent could be told about other agents by its user or by other agents, but this may not be a very good way of finding out about new agents and their capabilities and places a large burden on the user. One approach to dealing with this problem more automatically is to make use of agents to make "connections" with other agents [7]. Agents can advertise their capabilities to specialized agents (called facilitators or match-making agents) which keep track of agent capabilities. Alternately, or in addition, the match-making agent can also query other agents about their capabilities. In this way, the

match-making agent can be used to help agents locate other agents with the requisite capabilities. Such match-making agents can further be organized hierarchically (forming a network of match-making agents) to distribute the processing load and to reduce the potential for the failure of one agent bringing the system to a halt.
Finally, for agents to pool their knowledge and work together, they need to share enough common foundational knowledge in terms of which their specialized knowledge can be communicated [16]. This situation is somewhat analogous to humans with different specialized skills trying to work together (e.g., engineers and doctors). While each may be competent in their particular field, neither understands much of the specialized knowledge of the other. However, both have a wealth of common knowledge that they share and to which they can relate specialized knowledge in their respective specialities. For humans, this common foundational knowledge is acquired over many years and is related to shared culture and shared viewpoints of the world, among other things. However for agents, who do not experience the real world, such common foundational knowledge is lacking. The Cyc project is attempting to build such a common foundational knowledge base which agents can then use to more closely mimic human behaviour [9].

3 Applications of Intelligent Agents

Having discussed some of the technical issues related to the construction and use of agents and thus, hopefully, having a better understanding of their capabilities and potential as well as their limitations, we now briefly examine some of the applications for which agent technology is being used. The application of agent technology both in a business context and a personal context are numerous. Some current applications that are being explored for intelligent agents include [11]:

- management of air traffic and other transport services;
- aircraft flight control and mission analysis;
- control of telecommunications, power and computer networks;
- transaction management in banks and insurance companies;
- monitoring and control of industrial processes;
- on-line fault diagnosis and malfunction handling;
- provision and monitoring of medical care;
- control of industrial robots.

In the following subsections we will examine in more details some of the applications of agent technology.

3.1 Accessing Legacy Computer Systems

In Section 2 we saw that agents that use the Agent Communication Language (ACL) are able to exchange requests and information. However, there are many existing systems (often referred to as legacy systems) that agents may need to interact with that do not understand ACL. How are agents to communicate with such systems? One approach to addressing this problem is to simply reimplement legacy systems using agent technology [8]. While this is a rather drastic solution to the problem, in the long run it may be the best solution for updating and revitalizing business applications.

In the short term, however, such an approach may not be acceptable to many organizations. Therefore, two other approaches using agent technology are possible (see figure 3) [8]. In the first approach, called the transducer approach, a mediator agent acts as an intermediary between a legacy system and other agents. This approach only requires that the mediator agent know the communication behaviour of the legacy system. The second approach, called the wrapper approach, injects agent technology into a legacy system so that it can communicate in ACL. This approach requires that some of the code of the legacy system be changed, but it has the advantage of greater communication efficiency between the legacy system and other agents. Note that both transducers and wrappers can also provide additional intelligent functionality not available in the legacy system.

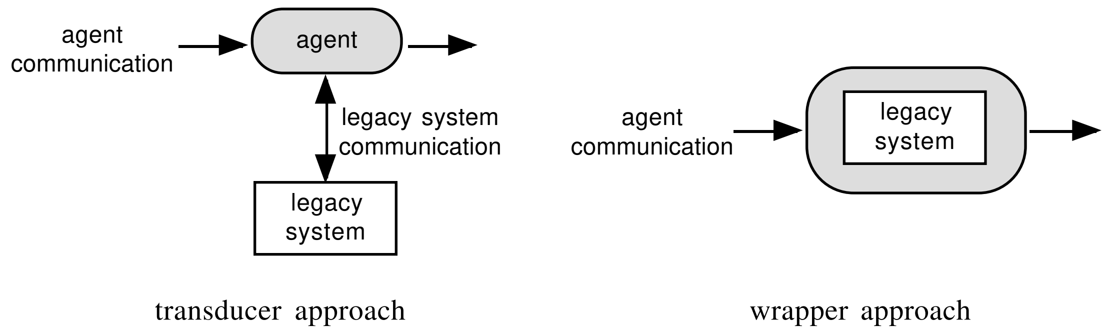

Figure 3. Accessing legacy systems.

3.2 Searching the WWW

As we have alluded to earlier in this paper, finding information on the WWW is emerging as perhaps the most important current application of agents. WWW information is increasing at an exponential rate. Information sources on the WWW are diverse which makes searching them by individual users somewhat problematic. While many search services exist for finding information on the WWW, using even these can pose difficulties. Many of the search services each only cover a portion of the WWW. Querying multiple services is quite tedious as each has a different user interface that needs to be learned. In addition, the results returned may be irrelevant, outdated or unavailable.

Several research projects are addressing this problem and some commercial products are starting to appear. As an example of a research project, the MetaCrawler agent developed at the University of Washington [http://metacrawler.cs.washington.edu:8080/] provides a service that allows users to pose a single query to a MetaCrawler agent [19]. The agent then submits this query to every search service that it knows about in parallel. A list of references to WWW pages is returned as the result of each of these queries. The MetaCrawler agent merges the returned results listing a source only once. The query syntax and associated query options help insure that a high proportion of the returned results are relevant to the query.

WebTamer is a commercial WWW search agent from AgentSoft [http://www.agentsoft.com] that is similar in functionality to MetaCrawler [10]. Users can create personalized intelligent browsing agents. The agent uses multiple Internet search engines simultaneously combining their results into one report which is filtered by a users's needs. A simple search agent can be programmed by a user and does the same thing every time. A sophisticated search agent interacts with several sites to collect information required by the user.

3.3 Monitoring and Predicting Stock Markets

A rapidly emerging application of agent technology is to monitor and predict stock and currency exchange markets. Like WWW searching, the information sources for this application are numerous and diverse. Gathering and summarizing relevant information is quite tedious. In the financial markets, various specialists find, filter and evaluate relevant information for portfolio management. Based on these evaluations and on predictions of the economic future, the specialists make suggestions about buying or selling various financial instruments. Several attempts are being made to automate this process using agents [20].

As an example, in the Warren system, developed at Carnegie Mellon University, various agents are used to perform these functions [20]. The *portfolio manager agent* is an interface agent that interacts with the user to obtain information about the user's goals. The *fundamental analysis agent* acquires and interprets information about a stock. The *technical analysis agent* uses numerical techniques to predict the near future in the stock market. The *breaking news agent* tracks and filters news stories and determines whether they are important enough for the user to know about them immediately. Finally, the *analyst tracking agent* tries to collect intelligence about what human analysts are thinking about a company. These agents gather data through requests to information agents that monitor things like stock prices, financial news, and SEC filings from on-line information sources via the WWW.

3.4 Customized Information Presentation

More and more news organizations are making their news stories available on-line via the WWW. However, this information is not customized to individual users of these services. Instead, users have to sift through the huge volumes of information provided to find articles of interest to them. Agent technology can aid in customizing the presentation of this information to the tastes and preferences of individuals.

WiseWire from Empirical Media [http://www.empirical.com] provides agents that can personalize a newspaper for its user. The agent asks for preferences as well as feedback from the user who rates his interest in articles that are assembled. The agent learns from a user's reading interests and uses machine learning to find new WWW pages that might be of interest to the user. The newspaper content is obtained from various Internet sources such as usenet news, the WWW, news wires, etc. The agent can analyze the content of articles. It learns as users rate articles that it sends them. It customizes the news articles so that over time they are closer and closer to a user's preferences.

3.5 WWW Shopping

Another important emerging application for agents on the WWW, and one that connects with electronic commerce, is shopping. However, the idea is not just merely to buy things over the WWW, but to comparison shop and find the best buy. Using agents, a user would describe to the agent the item he wishes to purchase. The agent would then find available sources on the WWW for the item and compare between the sources to find the best buy for the user. The comparison need not be only in terms of price, but might also be in terms of quality of the product, warranties offered, etc. The user would set the comparison criteria for the agent and the agent would do the comparison shopping for the user.

ShopBot [http://www.cs.washington.edu/research/shopbot/] developed at the University of Washington, is a domain-independent comparison-shopping agent [4]. Given the home pages of several on-line stores, ShopBot autonomously learns how to shop at those stores. After its learning is complete, ShopBot can visit other on-line stores, extract product information, such as availability and price, and summarize the results for the user. Studies have shown that ShopBot enables users to find both better prices as well as reduce WWW shopping time.

A slightly different approach to WWW shopping makes extensive use of agent collaboration [10]. Firefly agents from Firefly Networks Inc. [http://www.firefly.net] automate the word-of-mouth process of consumer purchasing. Agents act like "friends" that are knowledgeable in some area, such as music, movies or books, and make recommendations to users. To come up with their recommendations, Firefly agents, besides using preferences specified by the user, also consult other user agents with similar interests. Besides being able to purchase items on the WWW, Firefly agents can also connect their users with other people on the WWW with similar interests.

4 Our Future with Intelligent Agents

There is no doubt that agent technology, applied to the right application domains, can be a boon to businesses and individuals in more effectively applying computing technology and taming the WWW. In this section we will examine the possible effects of agent technology on businesses, individuals and society as a whole.

How will agents change the way that business is conducted? As agent technology becomes more prevalent and pervasive, businesses will have to be on the WWW or have to forego a large (and ever increasing) customer pool. Geographic location of a business will become less and less important as just being on the WWW will ensure a world-wide business presence. As customers are able to more easily comparison shop on the WWW, price competition should become more fierce. As a consequence, regional price variations are likely to become less and less pronounced as anyone can shop anywhere at any time. Businesses are likely to never close as they become 24 hour operations on the WWW. At the same time, the way that businesses deliver information about their products to customers will change as more and more customers use their agents to shop on the WWW. Businesses will have to worry not only about how products are advertised to people, but also how they are advertised to agents who are immune from appeals to emotions and self-esteem.

For individuals, agent technology will reduce the complexity of some tasks, such as comparison shopping, and make others easier to perform, such as news source selection. However, this may be replaced by the complexity and difficulty of instructing and interacting with one's agent(s). More of our time is likely to be devoted to this task, at least until agents become "smart" enough to instruct themselves and we develop trust in their abilities. Agent technology is likely to put some people out of work. For example, for mail order houses, as more and more of their business moves to the WWW, there will be less need for sales-people and order clerks as agents take over these tasks.

Currently, instructing agents is not for the faint-hearted or technically unsophisticated. Thus, agent technology is not yet for the masses. However, as we learn more about how to allow agents to learn and interact with each other, it should become easier for people to interact with and instruct agents. We also have much to learn still about how agents interact with and learn from each other. How will agents know what other agents to trust, not only for obtaining knowledge but also for giving their knowledge to other agents? Will agents become deceptive as businesses and individuals try to find out ever more information about potential customers and associates? This concern raises issues of security and privacy which few people have yet considered in depth in the context of agent technology [17].

As a society, should we be concerned about the increasingly widespread use of agents? As we delegate more and more tasks to our agents, how do we know that they are doing the "right" thing and are looking after the interests not only of their users, but also society as a whole? This concern has to do with the issue of agent safety [21]. Also, how much autonomy do we give to our agents? Do we allow them to override human actions? Who will be the final arbiter of who should be in control? As an example to illustrate this concern, the crash of an Airbus A-320 at Nagoya airport in Japan in 1994 was caused because the airplane's automatic flight system, which had not been turned off, was counteracting the pilot's actions [12].

Today's agents are still quite brittle. While they work well in situations that have been anticipated by the designer, they usually "break" when they encounter unexpected situations. As we learn more and more about how to build and use agents, this may become less and less of an issue. Nevertheless, no matter how intelligent agents become, the possibility of human intervention needs to be allowed in the event of a crisis. Accordingly, responsibility for and control of an agent should always rest with the person, or persons, who instructed it.

5 Summary

As computing technology becomes more pervasive and the amount of information available on-line rapidly increases, users will need help in dealing with this information overload. Intelligent agent technology holds much promise for meeting this need. At the same time, however, it drastically changes the nature of people's interactions with computers. In some sense we have come full circle in terms of our interaction with computers, but with a difference. In the early days of computing we needed to give them detailed instructions. Currently we use higher-level, direct manipulation interfaces. While the style of interaction has become more user-oriented, still little real assistance is provided to users in performing their computing-related tasks. Intelligent agents again require us to instruct the computer but now with a difference—the instructions are at a much higher, more human-oriented level and we delegate a good deal of the details of actually accomplishing the instructions to agents [14]. Thus, we are slowly but inexorably moving toward having agent-based computers as our partners intimately involved in every facet of our lives.

References

[1] D. Chin, "*Intelligent interfaces as agents,*" in Intelligent User Interfaces, J. Sullivan and S. Tyler, eds., ACM Press, N.Y., 1991.

[2] A. Cypher, "*EAGER: Programming repetitive tasks by example,*" Proc. CHI'91, pp. 33-39, 1991.

[3] L. Dent, J. Boticario, J. McDermott, T. Mitchell, and D. Zabowski, "*A personal learning apprentice,*" Proc. Natl. Conf. on Artificial Intelligence, 1992.

[4] R.B. Doorenbos, O. Etzioni, and D.S. Weld, "*A scalable comparison-shopping agent for the World-Wide-Web,*" Tech. Rep. UW-CSE-96-01-03, Dept. of Computer Science and Engineering, Univ. of Washington, Seattle, WA, 1996.

[5] O. Etzioni and D.S. Weld, "*Intelligent agents on the Internet: Fact, fiction and forecast,*" IEEE Expert, pp. 44-49, August 1995.

[6] T. Finn, J. Weber, G. Wiederhold, M. Genesereth, R. Fritzson, D, McKay, J. McGuire, R. Pelavin, S. Shapiro, and C. Beck, "*Draft specification of the KQML, Agent Communication Language,*" 1993.

[7] L.N. Foner, "*Yenta: A multi-agent, referral-based matchmaking system,*" Proc. First Intl. Conf. on Autonomous Agents, ACM Publications, 1997.

[8] M.R. Genesereth and S.P. Ketchpel, "*Software agents,*" Comm. ACM Vol. 37, No. 7, pp. 48-53, 1994.

[9] R.V. Guha and D.B. Lenat, "*Enabling agents to work together,*" Comm. ACM Vol. 37, No. 7, pp. 127-142, 1994.

[10] S. Hedberg, "*Agents for sale: First wave of intelligent agents go commercial,*" IEEE Expert, pp. 16-19, December 1996.

[11] J.A. Hendler, "*Intelligent agents: Where AI meets information technology,*" IEEE Expert, pp. 20-23, December 1996.

[12] H. Kitano, "*Nausicaa and the Sirens: A tale of two intelligent autonomous agents,*" IEEE Expert, pp. 60-61, December 1996.

[13] K. Lai, T. Malone, and K. Yu, "*Object Lens: A 'spreadsheet' for cooperative work,*" ACM Trans. Office Inf. Syst. Vol. 6, No. 4, pp. 332-353, 1988.

[14] P. Maes, "*Intelligent software: Easing the burdens that computers put on people,*" IEEE Expert, pp. 62-63, December 1996.

[15] P. Maes, "*Agents that reduce work and information overload,*" Comm. ACM Vol. 37, No. 7, pp. 31-40, 1994.

[16] M. Minsky, "*A conversation with Marvin Minsky about agents,*" Comm. ACM Vol. 37, No. 7, pp. 23-29, 1994.

[17] D.A. Norman, "*How might people interact with agents,*" Comm. ACM Vol. 37, No. 7, pp. 68-71, 1994.

[18] D. Riecken, "*Intelligent agents,*" Comm. ACM Vol. 37, No. 7, pp. 18-21, 1994.

[19] E. Selberg and O. Etzioni, "*Multi-service search and comparison using the MetaCrawler,*" Proc, World Wide Web Conf., 1995.

[20] K. Sycara, A. Pannu, M. Williamson, D. Zeng, and K. Decker, "*Distributed intelligent agents,*"IEEE Expert, pp. 36-45, December 1996.

[21] D. Weld and O. Etzioni, "*The first law of robotics (a call to arms),*" Proc. AAAI, 1994.