AI Magazine Volume 18 Number 3 (1997) (© AAAI)

Intelligent Adaptive Agents

A Highlight of the Field and the AAAI-96 Workshop

Ibrahim F. Imam and Yves Kodratoff

■ There is a great dispute among researchers about the roles, characteristics, and specifications of what are called agents, intelligent agents, and adaptive agents. Most research in the field focuses on methodologies for solving specific problems (for example, communications, cooperation, architectures), and little work has been accomplished to highlight and distinguish the field of intelligent agents. As a result, more and more research is cataloged as research on intelligent agents. Therefore, it was necessary to bring together researchers working in the field to define initial boundaries, criteria, and acceptable characteristics of the field. The Workshop on Intelligent Adaptive Agents, presented as part of the Thirteenth National Conference on Artificial Intelligence, addressed these issues as well as many others that are presented in this article.

If we were to ask 10 researchers from different organizations or institutions what their personal definition of an intelligent agent is, we would most likely get 8 to 10 different answers. Moreover, if a researcher or any curious person wanted to learn about intelligent agents, he/she might get confused after reading even a few papers of the hundreds that were recently published on agency-related subjects. Reasons for this confusion include the following: First, there is no standard definition of what an intelligent agent is. (Today, almost anything can be called an agent and, typically, an intelligent agent). Second, no clear goals or objectives for the agent (for example, the functions of the agent vary from implicit to explicit, systematic-mechanic to environmental, system requirement to user requirement, simple to complex). Third, the agent-user relationship is either missing or vague.

The Workshop on Intelligent Adaptive Agents, part of the Thirteenth National Conference on Artificial Intelligence (AAAI-96), presented state-of-the-art approaches, ideas, and methodologies for research and development of intelligent adaptive agents. The workshop consisted of two invited talks, presented by Brian Gaines and Barbara Hayes-Roth; four discussion sessions, organized and chaired by John Laird, Sandip Sen, Costas Tsatsoulis, and Kerstin Voigt; two commentary evaluations, presented by Yves Kodratoff and Brad Whitehall; and 10 papers, presented by Keith Decker, Karen Haigh, Ibrahim Imam, John Laird, Ramiro Liscano, Daniela Rus, Sandip Sen, Rahul Sukthankar, Kerstin Voight, and Grace Yee.

Intelligence, adaptation, and *agency* are three terms with no standard definitions accepted by researchers in the AI community. Defining the scope of the AAAI-96 workshop and understanding these terms are two associated issues. In general, the workshop focused on research involving the three issues together or in different combinations. For example, the scope of the workshop covered research and development on intelligent adaptive methodologies for agents and intelligent agents that behave adaptively. The definition of these terms were discussed in some presentations as well as some discussion sessions. A summary of these discussions is presented in the next section. The research presented at the workshop can be classified according to varying criteria. These criteria and a classification of the papers are presented in the following section. A brief description of the talks presented at the workshop is described in a later section. The last section introduces a classification of the research presented at the workshop according to conceptual and systematic criteria.

What Is an Agent?

Because the workshop presented diverse definitions of what an adaptive agent is, a discussion on the definition of an agent, a society of agents, and an intelligent agent was also an important part of the workshop. Some issues that are typically discussed in defining any agent were found less important than previously believed, especially when clearly differentiating between an agent and a procedure.

Some of the main reasons for the multiplicity of definitions of an agent include the focus on the source of input to the agent (or the way the agent interacts with the outside world), the focus on the functions of the agent (which in many cases are not dynamic), and the role of the agent as a part of a multiagent society (no clear boundaries or characteristics distinguish a complex agent from a society of agents). We present here abstract definitions of the terms *agent, intelligent agent,* and *society of agents*. Based on the assumption that any agent is a black box, a *multiagent society* is a group of agents that operate independently in a cooperative or a competitive environment. The term *independently* is used here to stress that although agents in the society might serve each other, their objective should not be limited to the service of another agent. For example, a structure of processes where each process depends on other processes to accomplish its task should be considered as a single complex agent with subprocesses rather than as a society of agents.

A *single agent* is a system or a machine (entity description) that provides assistance to the user (objective description). Describing the entity and the objective constitutes a sufficiently abstract definition of an agent. An autonomous vehicle can be considered an agent because it can provide transportation to the user. An information navigator system can be considered an agent because it searches and retrieves information for the user. A robot agent can observe the external world and inform the user about its observation. Such agents can retrieve their information from a variety of sources, but their common objective is to serve the user.

Defining an intelligent agent requires defining the methodology used by the agent. An *intelligent agent* is a system or a machine that utilizes inferential or complex computational methodologies to perform the set of tasks of interest to the user. These methodologies should enhance the agent's capabilities to learn. This definition distinguishes between intelligent agents in general and nonintelligent adaptive agents. An example of an adaptive agent that is not intelligent is the irrigation agent. The *irrigation agent* is a robot that irrigates a green house based on a simple equation involving the humidity and temperature inside the green house. Such an agent can irrigate the green house different number of times at different hours during each day. The external actions of the agent surely reflect adaptive behavior; however, the agent cannot be considered intelligent (some researchers, however, consider it to be intelligent).

Another way to describe the difference between intelligent and nonintelligent agents is by the degree of freedom the agent has when accomplishing the given task. In other words, the agent should have more flexibility in controlling the way it accomplishes a task independent of any changes in the environment. This characteristic can be observed in agents using different problem-solving methodologies, control parameters, knowledge bases, different objects, and so on. The ability to infer or select among different alternatives is one of the main characteristics of an intelligent agent. Intelligent agents act as superior or control systems. Nonintelligent agents act as a function or a procedure. These functions receive a set of parameter values (representing the environment) and apply the same set of steps or operators to accomplish the task. It is true that most research on agent modeling does not distinguish between intelligent and nonintelligent agents. Moreover, distinguishing between a procedure that uses a set of parameters and an algorithm that selects the best problem-solving scenario is considered of no importance when modeling an intelligent agent. Also, it is difficult to judge whether an agent is intelligent by observing its external actions (behavior). An interesting approach for distinguishing intelligent agents can be based on the relationship between the external actions (behavior) and the internal adaptation recognized in some workshop presentations.

The workshop audience also raised an interesting set of questions about the definition of an agent: What isn't an agent? Do we consider the search for a word or the word-count routines in a word processor program as agents? Why don't we consider animals that assist handicapped or other people as agents? Following Sherlock Holmes's strategy, we started eliminating the improbable to reach the accurate. The first issue is that as long as we are limited by the boundaries of computer science, agents are either computer systems or machines using computerized systems. The second and most confusing issue is associated with the term *assistance* in most agent definitions. We can always claim that any computer program, system, or machine provides assistance to the user in one way or another. At the workshop, we agreed that a word-count routine is an agent (but not an intelligent agent). Also, to call a system or a machine an agent, it should be designed to serve the user.

To clarify further, we associate what is referred to as agent assistance with the user task, then categorize the different user tasks and use the categories to clearly identify the difference between agents and nonagent entities. We categorize the *user tasks* (jobs that can be requested by the user from the agent) into three groups:

First is the *master task,* which specifies a final requirement; either, it has no outcome, or its outcome is an atomic object that cannot independently be processed inferentially or computationally. Examples of master tasks include counting the words in a documents, purchasing a plane ticket, transportation from point A to point B, checking if the conference room is empty, picking up an object, and identifying a face.

Second is the *generic task,* which specifies a final requirement but whose outcome can be processed or used to accomplish another task. Examples of such tasks include retrieving information or knowledge about a subject of interest to the user and allocating an appropriate web page relevant to a given query.

Third is the *intermediate task,* which specifies a general requirement needed to accomplish other tasks. Examples of such tasks include learning or discovering knowledge from data, planning a course of action, and determining a list of web pages to be searched for relevant information.

Our view of what isn't an agent also includes systems or machines that perform intermediate tasks. In real life, agents that perform generic tasks are less popular than agents that perform master tasks. Now to specify our abstract definition of what an agent is, we propose a more concrete definition: A *single agent* is a system or a machine that can accomplish master or generic tasks that are of interest to the user.

A Summary of Issues Discussed at the Workshop

Considering a clear understanding of what an agent is, the workshop focused on the adaptive process and its characteristics in intelligent adaptive agents. Intelligent adaptation in agents was characterized by many themes, including the goal, the cause, the systematic architecture, and the methodologies of adaptation.

The goal of adaptation: Optimization (in different forms) was the main goal in most papers. Papers that dealt with the optimization using a single agent differed widely (they mainly learn to adapt) from those that dealt with multiple agents (they mainly react among themselves to adapt). In the multiagent society, all agents are based on architectures that ensure accurate and optimized interactions among agents. In single-agent applications, agents optimize either their knowledge or their problem-solving methodology.

The cause of intelligent adaptation: Changes in the environment and inconsistency are the two main causes for starting any adaptation process with intelligent agents. The workshop speakers argued about the relationship between both causes. We attempt to present initial definitions to illustrate the difference between the two causes. Environmental changes can be viewed in many ways: new object or active agents in the environment, unusual readings from a sensor, actions taken by other agents, unexpected or new requests from a user, availability of some resources, or other observations. Self-improvement is the main motive for an agent to detect and correct inconsistencies internally and externally. Self-improvement is a somewhat more complex problem, and it might not reflect any changes in the behavior (external actions) of the agent. Agents can optimize the way they function or adapt their knowledge. Different views of the problem were introduced by many papers at the workshop.

The systematic architecture of adaptation: Two architectures for performing adaptation were presented in most of the papers. First is a fixed architecture that is finely tuned to the task that the global system must perform. In the other category, adaptation is performed inside the system's architecture itself (that is, the agent's architecture), which evolves over time. Another remarkable issue in multiagent societies is that they offer different types of "rewards" to the individual agents to make them evolve and improve.

It is to be noted that the difference between the two approaches (the fixed system architecture and the evolving system architecture) was also illustrated by the two invited speakers. It is also worth mentioning that the two approaches were opposed also by the kind of reward their agents receive. In Barbara Hayes-Roth's case, it is a hedonist "mutual satisfaction," but in Brian Gaines's case, agents compete for Spartan fulfillment of the most difficult task. One should not, however, generalize such a relationship among architectures and inducement techniques. An interesting issue for future research is to show that the inducement technique (for the agents to adapt) is good enough to lead to an efficient architecture found by the system itself.

The methodology used for adaptation: Most of the presented papers demonstrated agent systems that use inferential or complex computational algorithms. Because machine learning offers the most successful inferential algorithms, two approaches for using learning systems were presented for building intelligent adaptive agents. The first approach is to build a system (agent) that acts as an intermediary between the learning system and the environment. The agent serves as an intelligent user interface. The second approach is to modify the learning system to include timer changes, sensory input, constraint evaluation, and so on, in the learning function.

Other issues discussed at the workshop included (1) competitive and cooperative adaptive agents in a multiagent society, (2) common experiences or policies among agents in a multiagent society, (3) adaptation at different levels of knowledge, (4) adaptive control versus learning control (resolving the problem for different constraints versus solving different or similar problems), and (5) user expectation of adaptive agents.

A Summary of the Presented Talks

Barbara Hayes-Roth (Stanford University) presented the first invited talk at the workshop, entitled "Directing Improvisational Actors." The talk introduced agents as improvisational actors that can freely or jointly interact to cope with the changing scenario of a story. The agents are highly interactive, and they can answer to unusual constraints such as the "mood" (chosen by the operator) of the other agents. Such agents follow a structure of functions that separate the story line they are engaging in from the role they are ordered to perform and the performance they are expected to achieve. The talks introduced a set of interesting issues, including the handling of the generality of task requirements by agents, the interfacing between behavioral models and a real personality, and the viewing of agents (for example, actor, slave) differently.

The second invited talk, "Adaptive Interactions in Societies of Agents," was presented by Brian Gaines (University of Calgary). The continuous homogeneity of societies of agents highly depends on the adaptive interactions among the agent members. Modeling adaptive interactions among agents allows agents to account for their capabilities and task allocation. Knowledge is viewed as the state variable in these models. This talk presented a simple training strategy of allocating tasks with increasing difficulties (as the agent adapts to optimize the rate of learning or attempts to linearize the sigmoidal learning curve) that keeps the agent's performance constant.

The talk also addressed some issues that were shared by other presentations, including (1) how a program is manipulated to filter or restructure the information sources of another agent, (2) how agents communicate with each other to change the state of knowledge for one of them, (3) how to treat agents as systems at the knowledge level by assigning knowledge and goals to them, (4) how to decide for each agent what task to carry, and (5) what to learn about the relationships between tasks (an agent accomplishing certain tasks might be able to accomplish other tasks). In such models, adaptive agents can be characterized by the task they perform. Also, failing to accomplish a task can be cause to assign a simpler task.

The first presented paper, "Adaptive Intelligent Vehicle Modules for Tactical Driving" by Rahul Sukthankar, Shumeet Baluja, John Hancock, Dean Pomerleau, and Charles Thorpe (all of Carnegie Mellon University), presented one of the difficult projects for autonomous agents, where adaptation is crucial and necessary. The project is concerned with building intelligent vehicles that can drive on real highways in mixed traffic environments. One can view the vehicle as a global intelligent agent that controls different groups of agents. Each group of agents is responsible for performing different functions, including driving tasks (for example, recognizing cars ahead or upcoming exits) and self-state tasks (for example, managing car velocity). Each agent utilizes low-level sensors and a large number of internal and external parameters to accomplish its task. Agent parameters are automatically selected by a novel population-based incremental learning algorithm. Agents work independently but cooperatively. Adaptation occurs both internally and externally. The vehicle is evaluated by an evaluation metric that studies the number of collisions, lane-keeping consistency, speed versus desired speed, and so on.

The second presented paper, "Adaptation Using Cases in Cooperative Groups" by Thomas Haynes and Sandip Sen (both of University of Tulsa), introduced adaptation as a key component of any group of intelligent agents. Each agent learns from problem-solving cases how to adapt its model of other agents of the same group. Adapting the agent's model of other agents usually changes the course of actions that the agent can follow in different situations. The paper demonstrated a distributed AI problem called PREDATOR-PREY. The problem concerns four agents (predators) attempting to capture another agent (prey). Each predator adapts its moves based on the potential moves of other predators to avoid conflicts. The paper also proposed a solution to avoid deadlock situations that result from overlearning.

The third presented paper was "Knowledge-Directed Adaptation in Multilevel Agents" by John Laird and Douglas Pearson (both of University of Michigan) and Scott Huffman (Price Waterhouse Technology Center). Because adaptation is desired whenever errors or unexpected changes in the environment occur, it is important to detect this error or change, determine the cause (if possible), determine the correct course of modification, and adapt the agent functions to resolve such situations. The paper introduced a nice approach to model adaptation in agents, which characterizes the adaptation process as three levels of knowledge and control: (1) the reflex level for reactive response, (2) deliberate level for goal-driven behavior, and (3) a reflective layer for plan deliberation and problem decomposition. The paper demonstrated adaptation at both the reflex and the deliberate levels. At the reflex level, the domain theory is modified and extended to determine needed actions for similar situations. At the deliberate level, the agent uses the reflective knowledge to update its course of action. The paper included an example for explaining the approach: An agent (robot) attempts to perform a task with (and without) external instructions.

The fourth presented paper was "Adaptive Methodologies for Intelligent Agents" by Ibrahim Imam (George Mason University). In this paper, intelligent adaptive agents are considered as systems or machines that utilize inferential or complex computational algorithms to modify or change control parameters, knowledge bases, problem-solving methodologies, a course of action, or other objects to accomplish a set of tasks required by the user. Considering that environmental changes are the main cause of adaptation, adaptation for intelligent agents is classified into three categories based on the relationship between the internal actions and the external actions (behavior) of the agent. The first category is *internal adaptation*, where changes in the external environment are matched by internal changes to provide the same solution to the given task. The second category is *external adaptation*, where changes in the external environment directly reflect changes in the external actions of the agent. The third category is *complete adaptation*, where changes occurred in both the internal and external actions of the agent, and adaptation is not necessarily caused by changes in the environment. The paper illustrated these categories with the applications of an intelligent travel agent and an identification agent.

The fifth presentation was "Autonomous and Adaptive Agents That Gather Information" by Daniela Rus, Robert Gray, and David Kotz (all of Dartmouth College). In a virtual reality simulation, agents can independently transfer through a network of computers to accomplish a task. This paper introduced adaptive agents as systems that can completely terminate their existence at a given location, transform to a new location, and resume the task they are accomplishing. Such agents have the ability to (1) sense the state of the network (for example, to check if the local host is connected, find out if a site is reachable, or estimate the load of the network), (2) monitor conditions of software resource (for example, monitor activities of a site or another agent that expected to receive or obtain information relevant to the current task), and (3) interact with other agents (for example, an agent might need to locate other agent locations in the network, gather information about the tasks they can perform, or request a task from another agent).

Two types of adaptation can be categorized in this work: (1) external adaptation observed by transforming the agent from one location to a better one over the network to achieve the given task and (2) internal adaptation viewed by modifying the problem-solving strategy based on the information obtained from monitoring and sensing different resources and agents in the network.

The sixth presentation was "Intelligent Adaptive Information Agents" by Keith Decker, Katia Sycara, and Mike Williamson (all of Carnegie Mellon University). In a multiagent society, predicting environmental changes is another approach to planning for intelligent adaptation. To predict and adapt to environmental changes in an information-based multiagent society, the paper presented an approach where a matchmaker information agent gathers organizational information about all agents' functions, each agent plans the control flow of its actions or decisions using information about the relationships between all current tasks, all agents use a flexible scheduling mechanism, and each agent can control its active execution load.

Adaptation occurs at all levels. At the organizational level, a brokering matchmaker agent can optimize the distribution of tasks on different agents to balance the load of each agent. At the planning level, adaptation is needed in certain situations, for example, when an agent becomes unavailable or goes offline, the number of actions needed to accomplish a task is reduced, or the reduction of some tasks depends on the completion of other tasks by other agents. At the scheduling level, agents can adjust the scheduling of new tasks whenever related tasks are about to miss their deadlines. At the execution level, agents can control their availability whenever they are overloaded with tasks. When an agent is overloaded, it can inactivate itself with respect to new tasks in the current agent group and create a clone of itself to accomplish these tasks.

The seventh presentation was "Sacrificing versus Salvaging Coherence: An Issue for Adaptive Agents in Information Navigation" by Kerstin Voigt (California State University at San Bernardino). Navigation for highly relevant information at a lower cost is the goal of any information-navigation system. The paper introduced an approach for reducing the access cost of retrieving information relevant to the given query. The paper presented a utility function for measuring the access time and the coherence of information. Agents can use the access scores from measuring the cost of retrieving specific information and restructure the information to minimize this score. Information coherence is measured from knowledge about the constraints among different information items. These constraints can include the degree of generality, where general information is presented first. The paper also introduced a penalty function for any violation of precedence constraints among information items and a multiobjective optimization of hierarchical information structures to recover coherence.

The eighth presentation was "Learning Reliability Model of Other Agents in a Multiagent Society" by Costas Tsatsoulis (University of Kansas) and Grace Yee (Lockheed Martin Missiles and Space). In a multiagent society, learning to anticipate actions of other agents improves the quality and accuracy of agent performance. The paper presented an approach for learning reliability models of other agents in a multiagent society. Each agent learns a belief model to avoid erroneous data and optimizes the global problem-solving process. A system called the error adaptation communication control system (EACCS) was introduced for solving such problems. EACCS consists of three components: (1) dependency trace for keeping track of the path of a received message, (2) blame assignment for assigning blame to the paths supplying inconsistent data and allocating reliability values to the agent communication data, and (3) contradiction resolution for resolving contradictory information. Adaptation occurs whenever the agent encounters a new contradiction or inconsistency. The agent traces the source of inconsistency from all agents involved in providing the information. The agent models of the error-generating agent are then updated, and the system isolates the agent. Adaptation can also occur during the conflict-resolution phase when, for example, the sensory data are accurate, and the agent knowledge base is not.

The ninth presentation was "Using Perception Information for Robot Planning and Execution" by Karen Haigh and Manuela Veloso (both of Carnegie Mellon University). If the mind controls the body, the body is the main information tributary to the mind. This paper presented an approach for adapting the domain models of a planner based on a robot's direct observations of the environment. The approach introduced the agent ROGUE, which uses the planning and learning system PRODIGY to support the robot XAVIER in performing physical tasks. After performing each task, the robot XAVIER provides the agent ROGUE with its observations about the environment. ROGUE responds to the observations by dynamically updating the domain model. This update can affect the set of tasks that the robot needs to perform. The robot then starts performing tasks of the modified plan.

The paper illustrated the approach with an example where information is transferred between the two systems. Goals are classified according to importance, and the system attempts to opportunistically achieve less important tasks while it accomplishes the important tasks. The approach demonstrates the system's ability to respond to the normal dynamics of a real-world office environment while it performs common office tasks. Adaptation occurred on the planning level and the plan-execution level (by the robot). On the planning level, plans are updated online and in accordance with new observations. On the execution level, the robot gets new feedback that guides its future actions.

The tenth presentation was "Cooperative Agents That Adapt for Seamless Messaging in Heterogeneous Communication Networks" by Suhayya Abu-Hakima, Ramiro Liscano, and Roger Impey (all of the National Research Council of Canada). When information is mixed with voices, multimedia, and faxes in one carrier, real-time actions are crucial to the success of communication systems. This paper presented an approach for utilizing groups of agents to solve the problem of seamless messaging in heterogeneous communications networks. The paper introduced a framework for solving the problem using different groups of agents. Two types of agent were used for solving different aspects of the problem: (1) fine grained and (2) coarse grained. The *fine-grained agents* are responsible for simplifying the communication load between the user (or the device agents) and the agent responsible for message distribution. *Coarse-grained* (user) *agents* manage the user environment by observing actions and learning models of behaviors. *Surrogate agents* transfer across the network as messengers among user agents. *Message-transfer agents* mediate in delivering messages. Other coarse-grained agents with different functions are also introduced in this paper.

A Synthetic Summary of the Presentations

As described in the opening, the research presented at the workshop can be categorized according to different criteria. Other views of the presented papers are possible and might well bring some light on the topic.

Single Agents

In most papers concerned with intelligent adaptive single agents, it seems that the presented research can be described as choosing a preferred machine-learning system and modifying it to include either time changes or adaptation under constraints, yielding an (often highly) intelligent agent. Under this categorization fall the following papers:

Haigh and Veloso: Their machine-learning system is PRODIGY, which modifies its model of the world under failure. It is worth noting that their application field is robotics, advanced enough to compete at the American Association for Artificial Intelligence competition.

Laird et al.: Their machine-learning system SOAR represents three kinds of knowledge. Adaptation is done by modifying the *deliberate knowledge*, which is a kind of model of the world.

Imam: His machine-learning system is AQDT, which adapts by modifying its knowledge base, again a sort of model of the world.

Voigt: An apparent exception to this classification is the paper by Voigt because it focuses on a problem instead of a system. The paper dealt with the so-called *unsupervised learning paradigm*, that is, learning to change the structure of knowledge under comprehensibility constraints. The proposed solution to the problem is far less sophisticated than machine-learning systems such as COBWEB; hence, the learning is quite primitive. Inversely, the approach contains a sophisticated system to handle the efficiency of changes.

Multiagent Societies

In most papers concerned with intelligent adaptive multiagent societies, two approaches can be recognized: (1) an approach with a fixed system architecture and (2) an approach with an evolving system architecture. With regard to the multiagent society with a fixed architecture, the following papers were presented:

Hayes-Roth presented highly interactive agents because they can answer to nonclassical constraints such as the mood (chosen by the operator) of the other agents.

Abu-Hakima et al. showed an architecture for messaging in heterogeneous networks and defined the different types of agent needed for this task.

Decker et al. described a system with a large number of different agents, each of them enabled with some simple adaptability, such that the whole of them is highly adaptable.

Rus et al. addressed a problem quite similar to Abu-Hakima et al. Their solution relies on a definition of *sensing the network*.

With regard to a multiagent society with an evolving architecture, the following papers were presented:

Gaines's system architecture will be less precisely defined, and the game is now to prove that the inducement technique (for the agents to adapt) is good enough to lead to an efficient architecture found by the system itself.

Tastsoulis and Lee defined agents that learn and evolve following the reliability of their competitors.

Haynes and Sen used a case-based–reasoning technique to learn and induce changes following the agent's capacity to evolve.

Sukthankar et al. is yet another apparent exception to this classification. It considers the system of agents as black boxes, of which the system sees only a set of parameters. Evolving here involves optimizing a set of parameters. The optimization technique is hill climbing, which allows faster convergence than most sophisticated approaches. Genetic algorithms or neural networks might yield more precise results, but they require more computation time.

Conclusion

The workshop was successful in addressing interesting and difficult problems associated with research and development of intelligent adaptive agents. One of the most recognizable achievements is the general framework drawn by all the attendees to define the field of intelligent adaptive agents. Thoughts about this framework include that the general goal of intelligent adaptive agents should be oriented toward serving the user, and intelligent adaptation should generally be motivated to improve either the agent's services or performance. In multiagent societies, more problems and research issues should be considered; however, the whole society, as well as any single agent, should be evaluated by its objectives and services.

Acknowledgments

The authors would like to thank Keith Decker, Karen Haigh, Ken Kaufman, Jim Mitchell, and Katia Sycara for their comments, suggestions, and a careful review of an earlier draft of this article.

Ibrahim F. Imam is a research associate at George Mason University. He received his B.Sc. in mathematical statistics from Cairo University in 1986 and his M.S. in computer science and his Ph.D. in information technology and engineering from George Mason University in 1992 and 1995, respectively. Imam's research interests are in intelligent agents, machine learning, adaptive systems, knowledge discovery in databases, and knowledge-based systems. Since 1993, he has authored or coauthored over 30 publications in refereed journals, books, and conference and workshop proceedings. His e-mail address is iimam@site.gmu.edu.

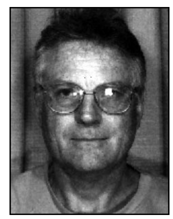

Yves Kodratoff is a director of research at the French Centre National de la Recherche Scientifique at the University of Paris-Sud. His interests cover all kinds of inductive inference system, including those that merge symbolic and numeric learning techniques, especially with respect to applications in data mining and knowledge discovery in databases. He is also working on the C3I to gather the field knowledge that allows an intelligence officer to guess what the enemy is planning. Others interests include inductive logic programming, the application of machine learning to vision, the use of explanation patterns in plan recognition, the inclusion of knowledge in neural networks, and epistemological aspects of machine learning. He helped to create three different spin-offs: (1) Intellisoft in 1988, (2) ISoft in 1990, and (3) EPCAD in 1995. He has published over 100 books, papers, and so on, on machine learning. His e-mail address is yves.Kodratoff@lri.lri.fr.