1999

Distributed artificial intelligence in a virtual reality setting: a case study

Rebecca Lohr
University of Northern Iowa

Distributed Artificial Intelligence in a Virtual Reality Setting:

A Case Study

by Rebecca Lohr

Senior Thesis Project

March 7, 1999

Advisor: Eugene Wallingford, Associate Professor of Computer Science

Introduction

Artificial intelligence, or AI, is a fascinating area of research. AI refers to the attempt to create a computer program, known as an intelligent agent, which can "think" and operate in a complex, changing environment. Many problems have been discovered in attempting to develop these intelligent agents, including the simulation of learning, planning, and natural language understanding. While researching these issues, a new branch of AI research has developed. Researchers discovered that there are other problems associated with how intelligent agents work together to solve problems. This field of research has become known as distributed artificial intelligence, or DAI.

Discussion of Basic DAI Concepts

Distributed artificial intelligence is essential to the development of a truly intelligent being. Little of what we do in our lives is done alone. As humans, we have developed sophisticated techniques for negotiating, compromising, synchronizing, communicating, and sharing knowledge. To develop an intelligent agent that can perform in a truly complex environment, it is necessary that the agent be able to work with other intelligent agents to solve a problem.

In their book, <u>Readings in Distributed Artificial Intelligence</u>, editors Alan Bond and Les Gasser divided the field of DAI research into two primary arenas: distributed problem solving, or DPS, and multi-agent systems, or MA. The editors distinguished between the two by saying that DPS research concentrates on how a problem can be divided up and its parts solved by different agents. This allows each agent to specialize in a certain area and studies how intelligent agents depend on one another to solve complex problems. MA research, on the other hand, is concerned with studying how multiple intelligent agents work together to solve a common problem.

This area studies how agents coordinate their own ideas and beliefs with those of others, and includes such areas of research as negotiating, knowledge sharing, and communication (Bond 3).

Potential Uses of DAI Systems

The potential uses of DAI systems are widespread. One problem encountered when developing these systems for interaction with humans is the potential need for these sophisticated systems to be incorporated into robotics. This may be necessary so that the agents can operate in these complex environments using senses humans can relate to, such as vision, sound, and touch. Robotics would allow intelligent agents to perform tasks with humans; however, developing these robots presents a whole new set of problems.

An alternative to using robotics to allow for interaction between intelligent agents and humans is to develop a virtual environment for the agents to operate in. Through the use of virtual reality technology (VR), researchers and developers can avoid such complex problems as vision and touch and can concentrate their efforts on problems related strictly to intelligent behavior. Virtual reality would allow individuals to work in a three-dimensional simulation of any given environment in which no distinction would be made between humans and intelligent agents.

Case Scenario

For the purposes of studying how this combination of DAI and VR can be accomplished, I have selected a case scenario that could potentially have a strong impact on the field of medical training. One of the major parts of a surgical student's training is experience in an operating room. Unfortunately, the only way for students to gain surgical training is for them to be in an operating room (OR) during surgery. This takes place only after several years of schooling and residency. However, this kind of surgical training is dependent on several factors. First, students can only get experience on operations that take place at the hospital in which they do their residency. Second, students can only get training when someone's life is hanging in the balance. It would be better if there was another way for surgical students to practice, rather than on actual patients. If a mistake were going to happen, it would be better if it could happen to an inanimate object.

A virtual operating room would offer students another option. A computer program could simulate the operating room environment and allow a student to practice a variety of surgical techniques on simulated patients. That way, if something did go wrong, no one would be hurt or could die.

In a traditional operating room, there are many doctors and nurses working together to solve problems. A basic virtual OR would not solve the problem of the lack of time that these professionals have to assist the students, but by adding intelligent agents and DAI technology, this problem could be solved. By simulating the other members of the surgical team, the students could practice their techniques without depending on someone else's schedule. Students could also have the opportunity to attempt surgeries that they otherwise would not have exposure to.

Also, because there are multiple roles to be played, any number of students could theoretically work together with any number of artificially intelligent agents. In other words, a surgical student and a training nurse could each play their respective roles in the same simulation. They could practice working not only with artificially intelligent agents, but also with each other. Ideally, a student would not be able to distinguish between another student and a simulated worker.

It is important to realize that while a tool of this nature could potentially be very useful in the training of medical students, it should not replace the valuable experience that can only be attained by working with other experts in the field. While this kind of simulation would require the development of a sort of expert system for each participant, it is important to recognize the differences in personal style and to acknowledge the importance of observation in the learning process. This tool should not be used as a primary source of experience, but rather as a method for attaining more experience of a different nature.

Relating Case Scenario to DAI

This scenario lends itself to both divisions of DAI. While each participant has his or her own responsibilities, two doctors performing the surgery may have different ideas of what techniques to use and when to use them. Both DPS and MA research is necessary to developing a virtual OR.

Each area of DAI research concerns itself with different types of problems. DPS is often concerned with how problems are broken down and how agents are assigned different roles. In the medical case scenario, this problem is solved. The case scenario defines a role for each participant to play and each role has several responsibilities attached to it. The recombining of the sub-problems, another major concern of DPS, is also defined by the example. This allows research to be concentrated on other important issues.

MA research presents a whole new set of problems. While DPS is concerned with the breakdown of problems, MA is concerned with the negotiating of different goals and plans. Compromising is a major area of study for MA researchers. Studying how agents coordinate their goals and ideas with those of others is essential to allowing intelligent agents to solve a common problem.

Some problems are common to both subdivisions. One of those problems is the idea of knowledge sharing. Studying how agents communicate is essential to understanding how problem solving takes place between multiple intelligent agents. Another common problem is dealing with uncertainty. Much of what we do in life involves some level of uncertainty. We do not always know the results of our actions ahead of time. Sometimes, we must choose a course of action based on what information we have access to at the time. Often, we wish we had more information about a subject, but we realize that we have neither the time nor the resources to know everything about a problem.

The specialization of roles helps to solve the problems of uncertainty. One participant may depend on the expertise of another to solve a common sub-problem. Though this does decrease the need for complete knowledge of a situation, one must consider the role of emergency situations in an operating room environment. An inexperienced individual may attempt a technique and make a mistake. In that situation, the virtual patient may react in a manner that the participants do not expect. Agents playing the roles must be able to deal with spontaneous situations. Things can go wrong, and decisions often have to be made in a split second. Experience may lead agents to better react to situations they have previously encountered, but agents must also have the flexibility to develop new techniques for solving new problems.

Exploration of Research in Selected Areas of DAI

Much of the research I encountered studied ideas related to planning and synchronization of plans and actions of multiple agents. In his paper "Communication and Interaction in Multi-Agent Planning," Michael Georgeff studies how the actions of one agent can affect the actions of another. He compares multi-agent plans with single-agent plans and proposes a method for transforming single-agent plans to resemble a format that can be useful in multi-agent plans. The idea is to identify potential problems in interaction between agents and to work to solve the problems associated with synchronization of activities to allow for the greatest possible amount of simultaneous activity (200).

Stephanie Cammarata, David McArthur, and Randall Steeb also study the problem of synchronization of activity. In their paper, "Strategies of Cooperation in Distributed Problem Solving," they suggest that "a main challenge to DPS is that the solutions which a distributed agent produces must not only be locally acceptable, achieving the assigned task, but also they must be interfaced correctly with the actions of other agents solving dependent tasks" (102). The authors suggest the importance of global consistency through local computation.

Another major area of research concerned itself with communication between agents. In studying the problems of synchronization, Cammarata, McArthur, and Steeb also studied how and when communication should occur between cooperating agents. They studied such ideas as broadcast versus selective communication, unsolicited versus on-demand communication, acknowledged or unacknowledged communication, and single-transmission versus repeated-transmission of a message (103).

Broadcast versus selective communication deals with the decision of who should receive the transmission of the message, whether it is a single other agent or all other agents working on the problem. Unsolicited versus on-demand communication deals with whether an agent should wait for the information to be requested or transmit the information when it becomes available or the agent views the information as necessary. Acknowledged versus unacknowledged communication simply poses the question of whether or not it is necessary for the agent to make it clear that they received the message. Finally, single versus repeated transmission deals with whether a piece of information needs to only be communicated once or needs to be repeated.

If it needs to be repeated, then a determination of frequency is required (Cammarata 103).

A final major area of research deals with the idea of cooperation between agents. Reid Smith and Randall Davis study this idea in their paper "Frameworks for Cooperation in Distributed Problem Solving." In their paper, the authors compare DPS to distributed processing systems. One of the major distinctions that can be made between these two similar systems is the focus of the system. Distributed processing systems focus on "an optimal static distribution of tasks, methods for interconnecting processor nodes, resource allocation, and prevention of deadlock" (61). Distributed processing systems often concern themselves with multiple machines working on several different yet related applications, each node specializing in one arena, and then recombining the results into a final answer. Issues of access control and data protection are major considerations in this field.

DPS, on the other hand, is concerned primarily with the breakdown of a single problem or application. DPS attempts to create a system that is open to cooperation of multiple agents. In other words, DPS works to create a team environment in which all agents do what they can to help one another solve the common problem. (Smith 61-62)

Another way to look at the contrast between DPS and distributed processing systems is to consider the fact that DPS tries to create an environment which allows and encourages cooperation between willing entities, whereas distributed processing systems enforce cooperation through a centralized control unit and data protection. (Smith 62)

Problems Encountered During Research

One problem I ran across while researching these topics is that most of the research and examples are concerned with systems that only interact with other artificially intelligent agents. I found no examples of how these techniques would be different if human interaction was necessary. One of the major advantages of ignoring human interaction is that it is not necessary to consider the massive problems associated with natural language understanding, a major part of the communications research. However, if DAI systems are only developed to work within themselves, we close ourselves off to an entire world of possible applications. Due to the limits of our current understandings, it is easily understood why these problems are avoided. However, it is beneficial to at least consider the possible implications involved with adding human interaction to the equation.

Suggestions for Solving the Problems Associated with DAI Research

In our case scenario, communication would operate much differently if agents did not have to worry about communicating with humans. Unfortunately, our problem requires communication between agents and humans and does not allow for the communication to take place through a monitor and a keyboard. This problem creates the need for the development of some sort of natural language interpreter in order to allow for communication between agents. Each agent must use this interpreter in every interaction with other agents to allow for the flexibility of not distinguishing between artificially intelligent agents and student participants. Also, it is not necessarily known what information is needed by all agents and what information must be communicated to only one other agent. In other words, one agent may want to say something to another agent, but the information may also be useful to a third. If an attempt were made to reduce usage of the interpreter, vital information may not be communicated to the proper individuals. While this interpreter would add a great deal of complexity to the virtual OR, it would also add a great deal of flexibility, and could potentially be very useful in a variety of other applications.

In solving problems related to synchronization and negotiation, we must first be able to better understand the different ways in which humans negotiate. Psychological and sociological research can be very useful in understanding how and why humans interact as they do. This understanding of negotiation and compromise is necessary to develop intelligent agents that can interact with humans. We must also keep in mind that, while the participants, both human and artificial, are working together to solve a problem, they also must remain as autonomous as possible in order to best replicate a surgical OR. Doctors must develop their own ways of handling difficult situations so that once they finish their training, they can operate on their own.

Conclusion

Human interactive technology is something that has yet to be widely studied in the field of distributed artificial intelligence. It encompasses several problems from several different fields, including distributed artificial intelligence, artificial intelligence, robotics, and virtual reality. However, the development of such advanced technology could prove to be extremely useful in the long run.

Works Cited

Bond, Alan H. and Les Gasser, eds. Readings in Distributed Artifical Intelligence. San Mateo: Morgan Kaufmann, 1988. 3-35.

Cammarata, Stephanie, David McArthur, and Randall Steeb. "Strategies of Cooperation in Distributed Problem Solving." Bond 102-105.

Davis, Randall and Reid G. Smith. "Frameworks for Cooperation in Distributed Problem Solving." Bond 61-70

Georgeff, Michael. "Communication and Interaction in Multi-Agent Planning." Bond 200 204.