Realistic Intelligent Agents for Training Simulators

Randolph M. Jones
Artificial Intelligence Laboratory
University of Michigan
1101 Beal Avenue
Ann Arbor, MI 48109-2110
rjones@eecs.umich.edu
(313)764-0683

For many skills, an effective teaching method is to have students learn and practice the skill in a simulated environment. Simulation is advantageous because it can provide effective training in a more controlled environment than the real world. In addition, training by simulation is generally cheaper and often less dangerous than the real situations it models. This research focuses on skills that require interaction with other humans. Example domains include air-traffic control, public transportation, cooperative activities, economics, and military tactics. Appropriate behavior in these domains relies heavily on the interaction between multiple intelligent agents, each with their own goals, skills, and equipment. In a training simulation, some of these agents may be other humans taking part in the simulation, whether they be other trainees or teachers. However, for large simulations it becomes necessary for most of the agents to generate their behavior automatically, without requiring attention from a human operator.

In order for automatic simulated agents to provide an effective training situation, they must behave as realistically as possible. This implies that the agents must behave like human agents would, within the limits of the domain. This paper discusses some of the issues involved in creating realistic intelligent, automated agents for simulation and training. In addition, it presents our efforts at constructing such an agent for the domain of tactical flight.

Requirements for a realistic agent model

The overall goal of a realistic agent model is to demonstrate "human-like" behavior appropriate to the simulation domain. In order to accomplish this, the model must incorporate a number of characteristics of general intelligent behavior. First, the agent must be able to reason flexibly, adapting gracefully to new situations, and developing creative solutions to accomplish its goals when necessary. The agent must also be able to reason effectively about its current situation and the goals and actions of the other intelligent agents in the simulation environment. Finally, the agent should learn with experience and adapt its behavior appropriately.

The particular domain our research focuses on is tactical flight (Jones, Tambe, Laird, & Rosenbloom, 1993). Our goal is to use automatic agents in flight simulators to help train navy pilots. This domain highlights a number of additional characteristics that an ideal intelligent agent would include. For example, flight is inherently a dynamic, fast-paced activity, and any realistic simulation must take place in real time. This implies that intelligent simulation agents must reason and execute actions in real time as well. In addition, the agent must plan when it has time to do so, but it must be able to react quickly without planning when the dynamics of the situation require immediate action.

Selecting an agent platform

Given the above constraints, we have chosen to use SOAR (Rosenbloom, Laird, Newell, & McCarl, 1991) as the architectural platform for intelligent agents. Flexible and adaptive behavior is one of SOAR's primary strengths. In addition, SOAR's automatic learning mechanism provides it with the capability of improving its performance with experience. The current version of SOAR incorporates state-of-the-art matching techniques and runs quickly enough to perform multi-level control in complex, real-time domains (Pearson, Huffman, Willis, Laird, & Jones, 1993). The architecture also allows the smooth integration of planning and reaction in decision making. Finally, SOAR serves as the foundation for the development of a proposed unified theory of human cognition (Newell, 1990), so it addresses a number of the cognitive issues of interest.

Research issues

Having chosen an agent architecture, there remain a number of important research issues that must be resolved in constructing an intelligent agent. The issues we are currently addressing include the acquisition and representation of knowledge, integration of real-time planning and execution, situation interpretation and agent modeling (or plan recognition), and reasoning about multiple interacting goals.

Our initial knowledge acquisition for the tactical flight domain involved analyzing navy training scenarios and tactics, and carrying out extensive interviews with former pilots and radar intercept officers. However, we did not simply encode their high-level decision-making into a set of rules. We feel that a successful agent must ground its decisions in the variety of knowledge sources available to the pilots, such as knowledge about geometry, flight, airplane and missile characteristics, different types of missions, and navy doctrine. In order for the agent to generate flexible behavior, it must understand the behavior it generates. Thus, we have chosen to implement the base knowledge in each of these areas, such that high-level tactical actions emerge from the decision-making process. This is in contrast to a system that might have a set of scripts for appropriate actions to take in different situations. Much of our initial research effort has involved gleaning these "first principles" and implementing them into various SOAR problem spaces and data structures.

We have also implemented a planning system for the agent. When the system has a choice of actions to take, it sets up an internal state representation of the world. Then it uses internal simulation knowledge to predict the consequences of the actions. Evaluation of the possible actions depends on the predicted outcomes and the agent's mission goals. After planning, the system learns chunks: production rules that compile the results of the planning process. These chunks make planning unnecessary in future similar situations.

In general, an automated agent taking part in a simulation will be "pre-trained," so it does not have to plan every decision it makes (similar to experienced pilots). However, the planning capability allows the system to perform appropriately in unanticipated situations that were not pre-trained. We also expect planning to aid in the knowledge-acquisition process, so we do not have to hand-code all the chunks that might eventually be learned.

Another important ability for intelligent agents is interpreting the actions of other participants in the simulation, as well as determining the goals they are trying to achieve. We have only begun addressing this issue, but our basic approach is to use the agent's own knowledge to model the other participants. For example, when an enemy plane turns, a fighter must identify the turn and determine why he turned to that particular heading. Our agent will put itself in the enemy's situation and determine which goals would cause it to turn to such a heading. It could be a random turn to keep from being predictable, or it could be an attempt to break a radar lock. If the enemy is pointing straight at us, it is possible that he is shooting a missile. In general for this domain, the agent will identify the different possible reasons for observed actions and assume the worst (in order to be as safe as possible).

An additional necessary reasoning process involves generating behaviors that meet multiple interacting goals. In tactical flight, pilots generally have mission goals, individual goals for intercepting enemies, and the overall goal of surviving. Different combinations of these goals lead to different tactical behaviors. For example, after firing a missile, our agent has a range of possible maneuvers to choose from, including flying straight toward the target, turning slightly to decrease closing velocity while maintaining a radar lock, or turning away to defeat the enemy's radar lock. The automated agent considers each possible action in light of the goals that the action helps achieve and attempts to choose the best compromise.

Future research issues

Developing realistic intelligent agents for simulations requires addressing a large number of research issues. The issues we described above pertain mostly to developing complete, robust intelligent agents for simulation. In the near future, we plan to address additional ways in which the agent model can aid instruction.

One such topic is the ability to generate explanations of the agent behavior. This is important for instructors to understand how their simulation tools are working. Explanation could also be used in an assessment model, in which the simulated agent generates possible explanations for a human trainee's behavior, to aid in future training.

The agent model will also be useful for developing and testing new training curricula, by examining the effects of new training scenarios on the agent's learning. For example, the model could be used to determine which types of instruction would be easier for a student to understand and learn from (Huffman & Laird, in press). In addition, simulation with automated agents can be used to develop new techniques for solving problems. In our domain, for example, the simulated agents could be used to generate and test new tactical behaviors without using the valuable time of expert pilots or risking danger to the pilots and expensive equipment.

Finally, a potential use for simulated agents involves carrying out "what if" simulations. After a pilot has completed a training scenario, his recorded behavior is played back in a debriefing session. The simulated intelligent agent could take over the human's position to answer questions such as, "What if you had dived at this point instead of turning?" The agents should have the knowledge and reasoning powers necessary to carry out each hypothetical branch in the training scenario. This will allow trainees to examine the consequences of possible actions, leading to more effective instruction.

Summary

The basis of this research is that intelligent, automated agents can provide realistic simulation models for skill training. Constructing such intelligent agents requires us to focus on core research issues in artificial intelligence and machine learning. Solutions to these issues, as well as large amounts of domain knowledge, must be implemented in a fully integrated system, such that it generates human-like behavior for the particular simulation domain. We have constructed an initial agent within the SOAR architecture for the domain of tactical flight. This agent includes knowledge of the first principles of the domain, from which high-level tactical actions emerge. In addition, the current agent reasons in real time and has limited capabilities to plan, model agent behavior, and reason about multiple interacting goals. Our future research will involve investigating additional ways in which the simulated agent model can aid instruction, and developing those abilities.

Acknowledgements

This research is a group project involving the cooperative efforts of Milind Tambe, John Laird, Paul Rosenbloom, Lewis Johnson, Jill Fain Lehman, Robert Rubinoff, Karl Schwamb, and Frank Koss. It is funded by a contract from DARPA/ASTO.

References

Huffman, S., & Laird, J. E. (in press). Learning procedures from interactive natural language instructions. In P. E. Utgoff (Ed.), Machine Learning: Proceedings of the Tenth International Conference. Los Altos, CA: Morgan Kaufmann.

Jones, R., Tambe, M., Laird, J. E., & Rosenbloom, P. S. (1993). Intelligent automated agents for flight training simulators. In Proceedings of the Third Conference on Computer Generated Forces and Behavioral Representation (pp. 33-42). Orlando, FL.

Newell, A. (1990). Unified theories of cognition. Cambridge, MA: Harvard University Press.

Pearson, D. J., Huffman, S. B., Willis, M. B., Laird, J. E., & Jones, R. M. (1993). Intelligent multi-level control in a highly reactive domain. In F. C. A. Groen, S. Hirose, & C. E. Thorpe (Eds.), Proceedings of the Third International Conference on Intelligent Autonomous Systems. IOS Press.

Rosenbloom, P. S., Laird, J. E., Newell, A., & McCarl, R. (1991). A preliminary analysis of the Soar architecture as a basis for general intelligence. Artificial Intelligence, 47, 289-325.