Building Believable Agents for Simulation Environments: Extended Abstract

Milind Tambe[*], Randolph M. Jones[**], John E. Laird[**], Paul S. Rosenbloom[*], Karl Schwamb[*]

[*] Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292
email: {tambe, rosenbloom, schwamb}@isi.edu

[**] Artificial Intelligence Laboratory
University of Michigan
120 ATL Building
1101 Beal Ave
Ann Arbor, MI 48109
email: {rjones,laird}@eecs.umich.edu

1. Introduction

The goal of our research effort is to develop generic technology for intelligent automated agents in simulation environments. These agents are to behave believably like humans in these environments. In this context, believability refers to the indistinguishability of these agents from humans, given the task being performed, its scope, and the allowable mode(s) of interaction during task performance. For instance, for a given simulation task, one allowable mode of interaction with an agent may be typewritten questions and answers on a limited subject matter. Alternatively, a different allowable mode of interaction for the same (or different) task may be speech rather than typewritten words. In all these cases, believability implies that the agent must be indistinguishable from a human, given the particular mode of interaction. Such an agent technology can potentially provide *virtual humans* for the multitude of *virtual reality* environments under construction. Its applications can be found in many fields, including entertainment [1], education [5, chapter 3], and training [2].

To begin this effort, we have focused on creating specific automated agents for simulated tactical air combat. The automated agents act as the *virtual pilots* for simulated aircraft, and will participate in exercises with real Navy pilots. These exercises will aid in training Navy pilots, development of tactics, and evaluation of proposed hardware. This is a non-trivial task, with many real-world complexities, and as such it offers several advantages. It pushes research based on real-world needs on topics such as reactivity, real-time reasoning, planning, episodic memory, agent modeling, temporal reasoning, explanation, and natural language understanding/generation. Furthermore, it forces the integration of all of these component AI technologies, because it requires a single automated agent to perform all of the functions performed by a pilot in air combat. Simultaneously, however, as a simulation task, it delimits the component technologies to be integrated. For example, it does not force the integration of vision or locomotion components. Finally, the task also imposes external metrics for success.

The task also poses an important constraint: the automated agents must believably act and react like trained human pilots. These agents are to take part in exercises with other human pilots. If human trainees identify our agents as automated pilots, they may take advantage of specific known characteristics of their behavior. Training in such a situation could actually be harmful. For instance, if the automated agents do not react as quickly as other human pilots (or react too quickly), trainees may learn to act too aggressively (or not aggressively enough) in a real aerial combat. Additionally, if the agents behave unrealistically, observers and tacticians at "ground control" (who can watch the simulated combat from different perspectives), may not be able to develop realistic tactics and strategies.

Thus, this task requires the development of believable automated pilots. For this fixed task, believability refers to the indistinguishability of the automated pilot from a human pilot, given the scope of the task, and the allowable modes of interaction. The scope of the task depends on (at least) the number of aircraft involved on each side, e.g., whether it is a one "friendly" aircraft versus one "enemy" aircraft (1v1) air-combat situation, or a 2v1, or 2vN situation. The allowable modes of interaction depend on whether it is a Beyond Visual Range (BVR) combat situation, where pilots only get radar information about the enemy aircraft, or Within Visual Range (WVR) combat situation, where the pilots can also directly see the enemy aircraft. In 2v1 (or 2vN) combat situations, additional modes of interaction are possible: the pilots of two or more "friendly" aircraft may communicate via radios, electronic data links, or even by executing simple maneuvers. A human observer at "ground control" adds even more modes of interaction. He/she can observe the combat in progress on a TV monitor, zoom in and out on it, focus on the maneuvers of a particular aircraft, and so on. A passive observer can only observe the combat in progress, while an active observer can supply the pilots new information or commands over the radio.

The specific scope of the task, together with the choice of certain modes of interaction, dictates the capabilities an agent must possess for believability. These capabilities define a certain *level of believability*. If the agent possesses these capabilities, then we refer to it as having (or being at) this level of believability. For instance, consider a 1v1 BVR air-combat situation, with no observers, and with a single human pilot engaged in combat with a single automated agent. The only mode of interaction is what the human pilot can view of the automated agent's actions on its radar. The capabilities required for believability are that these actions must appear like those of a trained human pilot. An agent with these capabilities has a certain (moderate) level of believability. Suppose we add a passive observer to this situation. Since the observer can watch the automated agent's actions much more closely, the agent must have a higher level of believability. As we add more aircraft, an active observer, and switch to WVR, the agent must have even higher levels of believability, with requirements for capabilities such as natural language (and speech) understanding/generation to support different types of radio communication.

The levels of believability provide us a means of staging an attack on this problem (and correspondingly staging the system development effort). Thus, to begin this effort, we have focused on an agent at a moderate level of believability: an agent for 1v1 BVR air-combat, with a passive observer. Even at this level, the task remains highly knowledge- and capability-intensive. Trained Navy pilots possess vast knowledge about different mission types, tactics and maneuvers, performance characteristics of the aircraft, radar modes, missile types and so on. The challenge for constructing an automated agent is then to integrate this knowledge into a single system, along with the following capabilities:

1. *The agent must be extremely flexible in its behavior:* Situations in air combat can change very rapidly. Unexpected events can occur, e.g., an on-target missile may fail to explode, or an aggressive adversary may engage in some preemptive action disrupting an ongoing maneuver. Accordingly, the agent must respond flexibly to the evolving situation.

2. *The agent must act/react in real-time:* Since a human may be interacting with the agent in real-time, the agent must act/react in real-time as well.

3. *The agent must try to interleave multiple high-level goals*: For this task, the agent must continuously attend to at least three high-level goals: (a) executing maneuvers to destroy the opponent; (b) surviving opponents' weapon firings; and (c) interpreting opponents' actions. Given the need for real-time response, the agent must be capable of rapidly switching among these goals (or achieving them in parallel).

4. *The agent must conform to human reaction times and other human limitations:* As

discussed earlier, the agent must not react to input data faster (or slower) than a human pilot would. The agent must also not maneuver the simulated aircraft like a "superhuman", e.g., it must not make very sharp turns. Finally, the agent must exhibit some unpredictability in its behavior, when appropriate.

5. *Others:* Some other capabilities such as planning, temporal reasoning, are also required for this task in limited proportions.

Note that, because a passive observer can watch an automated agent more closely than what is visible on radar, this additional level of believability requires more accurate modeling of human reaction time and physical limitations.

2. Developing Believable Pilot Agents

The basis of our work on developing automated agents is the Soar integrated architecture [4, 6] (Due to space constraints, we will assume that the reader has some familiarity with the Soar architecture). Some of the characteristics of this task are particularly well-suited for Soar. First, Soar is a single unified architecture for the research, development and integration of various component AI technologies. Second, Soar represents a developing unified theory of cognition, which is advantageous, given the constraint of psychological verisimilitude (e.g., limitation on reaction time) in this task.

The automated pilots for the 1v1 BVR air-combat task are based on TacAir-Soar, a system developed within the Soar architecture, which currently includes about 1100 productions. TacAir-Soar encodes the basic task knowledge for an agent in a set of problem spaces. A particular automated agent is realized by initializing TacAir-Soar with a specific set of parameters, such as its mission, the level of risk it can take for the mission, and the kind of weapons it has available.

The current design of TacAir-Soar is guided by two sets of constraints: the task requirements (as specified by the targeted level of believability), and the Soar architecture itself. Consider the key requirement of flexibility of behavior. This has turned out to be a strong constraint on the design of problem spaces and operators. For instance, any maneuver consisting of a sequence of actions is implemented not as a single monolithic plan, but rather as a sequence of appropriately conditioned operators in a problem-space. This allows TacAir-Soar to respond flexibly to an evolving situation, and not remain rigidly committed to a specific plan. Furthermore, this constraint discourages highly specific, narrowly focused problem spaces. For instance, a problem space devoted solely to employing one type of missile may not allow the system to switch quickly to employing a different type of missile, as the situation rapidly evolves. In contrast, a problem space that combines the operators for employing different types of missiles facilitates such actions.

TacAir-Soar's highly reactive behavior derives at least in part from Soar's ability to react at a number of different levels [3]. Specifically, Soar can respond to new inputs at three levels: (i) in a single production firing, (ii) in a single decision, which involves firing multiple productions, or (iii) in a problem-space, which involves executing multiple decisions. Thus, as the situation changes, Soar can respond very quickly within the time-span of a single production firing. If needed, it may also respond after much deliberation in a problem space. Additionally, Soar's efficient implementation technology plays a large role in allowing it to respond in real time.

In achieving multiple high-level goals, TacAir-Soar faces an interesting issue: as limited by the Soar architecture, it cannot construct multiple goal/problem-space hierarchies (in parallel) in service of the high-level goals. TacAir-Soar can and does construct a goal hierarchy in an attempt to achieve the high-level goal of destroying the opponent. For instance, to achieve the goal of destroying the opponent it creates a subgoal to "destroy-with-missile". To achieve destroy-with-missile, it generates subgoals to get into missile firing range, and so on. However, TacAir-Soar cannot construct goal-hierarchies for its remaining high-level goals — survival and interpretation of opponent actions — in parallel. To address this limitation, TacAir-Soar opportunistically installs operators for these high-level goals into its existing goal hierarchy (without eliminating the hierarchy). This avoids the overhead of rebuilding the goal hierarchy, while allowing it to switch attention among different types of goals rapidly. While this solution has allowed TacAir-Soar to exhibit reasonable performance so far, it does have some disadvantages. First, by not representing the different goal hierarchies explicitly, the solution does hinder TacAir-Soar's ability to reason about the interactions between multiple goals. Second, it is unclear if the scheme will generalize beyond the targeted level of believability. For instance, it is unclear if natural language understanding/generation will fit into this scheme. Alternative solutions are currently under investigation.

TacAir-Soar's ability to adhere to human reaction times is hindered by the artificiality of the interface to the simulation environment. In particular, TacAir-Soar does not spend time physically manipulating different instruments (e.g., turning a knob), or decoding actual instrument displays (e.g., decoding radar displays). As a result, TacAir-Soar tends to react faster than human pilots in some situations. Therefore, deliberate delays have been set up to slow down some of TacAir-Soar's responses. Similarly, TacAir-Soar's turning maneuvers have been constrained so as not to exceed human capability. As for unpredictability, much of it occurs "naturally" in TacAir-Soar. In particular, while two complex situations may appear very similar to a human observer, they may be quite different from TacAir-Soar's perspective, leading TacAir-Soar to two different actions. To add to this unpredictability, TacAir-Soar does random selection among operators that are considered to be equally appropriate in a given situation.

3. Current Status and Future Plans

Currently, even with approximately 1100 productions, the TacAir-Soar system continues to perform well within real-time constraints. Agents based on TacAir-Soar are fairly capable and robust within a narrow range of missions for 1v1 BVR combat. Recently, in a demonstration organized for Navy personnel, these agents were tested against (constrained) human pilots. The demonstration was a success in that the agents were able to function adequately at this targeted level of believability, i.e., they were able to react realistically to the humans pilots.

We are currently extending TacAir-Soar to deal with co-ordinated multi-aircraft air-combat simulations. Essentially, we are extending TacAir-Soar agents to higher levels of believability, and hence need integration of capabilities such as natural language understanding/generation. Thus, so far, for this task, the levels of believability appear to be useful as a means of staging development, as well as for measuring believability. Whether this usefulness will continue in the future, or for other tasks, remains to be seen.

References

1. Bates, J., Loyall, A. B., and Reilly, W. S. Integrating reactivity, goals and emotions in a broad agent. Tech. Rept. CMU-CS-92-142, School of Computer Science, Carnegie Mellon University, May, 1992.

2. Jones, R. M., Tambe, M., Laird, J. E., and Rosenbloom, P. Intelligent automated agents for flight training simulators. Proceedings of the Third Conference on Computer Generated Forces and Behavioral Representation, March, 1993.

3. Laird, J.E. and Rosenbloom, P.S. Integrating execution, planning, and learning in Soar for external environments. Proceedings of the National Conference on Artificial Intelligence, July, 1990.

4. Laird, J. E., Newell, A. and Rosenbloom, P. S. "Soar: An architecture for general intelligence". *Artificial Intelligence 33*, 1 (1987), 1-64.

5. Moravec, H. *Mind Children.* Harvard University Press, Cambridge, Massachusetts, 1990.

6. Rosenbloom, P. S., Laird, J. E., Newell, A., and McCarl, R. "A preliminary analysis of the Soar architecture as a basis for general intelligence". *Artificial Intelligence 47*, 1-3 (1991), 289-325.

Table of Contents

1. Introduction
2. Developing Believable Pilot Agents
3. Current Status and Future Plans
References