

# Intelligent Agents for Interactive Simulation Environments

Milind Tambe<sup>\*</sup>, W. Lewis Johnson<sup>\*</sup>, Randolph M. Jones<sup>\*\*</sup>, Frank Koss<sup>\*\*</sup>  
John E. Laird<sup>\*\*</sup>, Paul S. Rosenbloom<sup>\*</sup>, Karl Schwamb<sup>\*</sup>

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

<sup>\*\*</sup> Artificial Intelligence Laboratory  
University of Michigan  
1101 Beal Ave  
Ann Arbor, MI 48109  
email: {rjones, laird, koss}@eecs.umich.edu

**Contact Author:** Milind Tambe

**Office address:**

Information Sciences Institute  
University of Southern California  
4676 Admiralty Way  
Marina del Rey, CA 90292  
Office phone number: 310-822-1511  
email: tambe@isi.edu  
Office Fax number: 310-823-6714

**Home address**

8501 South Sepulveda Blvd  
Apartment #108  
Los Angeles, CA 90045  
Home phone number: 310-568-0194

**Alternate Contact Author:** Lewis Johnson

Information Sciences Institute  
University of Southern California  
4676 Admiralty Way  
Marina del Rey, CA 90292  
email: johnson@isi.edu  
Office phone number: 310-822-1511  
Office Fax number: 310-823-6714

## Intelligent Agents for Interactive Simulation Environments

Milind Tambe<sup>\*</sup>, W. Lewis Johnson<sup>\*</sup>, Randolph M. Jones<sup>\*\*</sup>, Frank Koss<sup>\*\*</sup>  
John E. Laird<sup>\*\*</sup>, Paul S. Rosenbloom<sup>\*</sup>, Karl Schwamb<sup>\*</sup>

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

<sup>\*\*</sup> Artificial Intelligence Laboratory  
University of Michigan  
1101 Beal Ave  
Ann Arbor, MI 48109  
email: {rjones, laird, koss}@eecs.umich.edu

### Abstract

Interactive simulation environments constitute one of today's promising emerging technologies, with applications in areas such as education, manufacturing, entertainment and training. These environments are also rich domains for building and investigating intelligent automated agents, with requirements for the integration of a variety of agent capabilities, but without the costs and demands of low-level perceptual processing or robotic control.

Our project is aimed at developing human-like, intelligent agents that can interact with each other, as well as with humans in such virtual environments. Our current target is intelligent automated pilots for battlefield simulation environments. These are dynamic, interactive, multi-agent environments that pose interesting challenges for research on specialized agent capabilities as well as on the integration of these capabilities in the development of "complete" pilot agents. We are addressing these challenges through development of a pilot agent, called TacAir-Soar, within the Soar architecture.

The purpose of this article is to provide an overview of this domain and project by analyzing the challenges that automated pilots face in battlefield simulations, describing how TacAir-Soar is successfully able to address many of them — TacAir-Soar pilots have already successfully participated in constrained air-combat simulations against expert human pilots — and discussing the issues involved in resolving the remaining research challenges.

**Keywords:** Intelligent agents, distributed interactive simulations, real-time systems, multi-agent systems, integrated architectures, Soar.

## 1. Introduction

Interactive simulation environments are becoming ubiquitous, finding applications in education [Moravec 90], training [Jones et al. 93, Webber and Badler 93], entertainment [Bates et al. 92, Maes et al. 94], and manufacturing [Loper et al. 94]. These virtual worlds can be populated not only with humans, but also with intelligent automated agents — AI systems that interact with humans, each other, and their virtual environment. These synthetic environments provide a new laboratory in which intelligent artifacts can be studied.

While the concept of a synthetic environment is closely related to those of software [Etzioni 93], robotic [Brooks 91], and testbed [Hanks et al. 93] environments, it differs from each in significant ways. With respect to software environments, the most significant difference is that synthetic environments usually require real-time behavior in dynamic, limited-information worlds. A strong dependence on traditional planning techniques is thus ruled out in synthetic environments. With respect to robotic environments, the most significant difference is that synthetic environments do not have to meet the challenge of low-level perception and motor control. They thus enable the study of higher-level aspects of intelligence, and their integration, unfettered from the demands and costs of low-level perceptual processing and robotic control. With respect to testbed environments, the most significant difference is that synthetic environments, at least of the type focused on here, are real domains that are developed commercially for government and/or commercial clients who need them for their own use. Developers of agents for synthetic environments thus do not generally have the same freedom to prestructure the environment, to choose which aspects of behavior will matter, nor to instrument the domain for experimental purposes. One additional aspect of synthetic environments that sets them apart from these other three types of environments is the common requirement that their agents act as human surrogates, thus necessitating an important constraint of human-like behavior.

Our research effort is aimed at this task of producing human-like, intelligent automated agents for large scale, realistic simulation environments. Such agents promise several benefits that could enhance the cost-effectiveness and functionality of these virtual worlds. One potential benefit is that intelligent agents could substitute for humans when large numbers of entities are required to populate a virtual world. Thus, in traffic simulators intended for training and experimentation, intelligent agents could be used to create realistic scenarios involving slow or speeding vehicles, pedestrians, traffic jams and other complex traffic patterns [Cremer et al. 94]. Similarly, in large-scale military battlefield simulations, intelligent agents could substitute for troops who are needed to fill out an engagement, but who are not available in person. Another potential benefit is that artificial agents can simplify and speed up experimentation by providing more control of behavior, repeatability of scenarios, and increased rate of simulation (i.e., faster than real-time simulation). This can be critical for testing new products that require repeated experimentation across a wide range of situations, such as collision warning devices in traffic simulators.

Our effort has begun with a focus on intelligent automated agents for battlefield simulation environments. These environments are based on Distributed Interactive Simulation (DIS) technology [Loper et al. 94], in which large-scale interactive simulations are built from a set of independent simulators linked together via a network. Since 1983, the Advanced Research Projects Agency (ARPA) has been aggressively pursuing DIS technology as a basis for large-scale battlefield simulations. The SIMNET effort [Thorpe et al. 89], which ARPA undertook in co-operation with the US Army, was an important milestone in this regard. SIMNET provides a successful demonstration of the essential technologies supporting a distributed battlefield simulation environment in which up to a thousand independent combat entities, mainly US Army systems such as tanks and armored personnel carriers, can participate in training exercises in a virtual battlefield. Within SIMNET, humans and

computer generated forces — referred to as *semi-automated forces* (SAFORs) — may engage in unscripted, *free play* simulated combat. SAFORs are able to perform many simple activities autonomously (such as driving point-to-point in formation or engaging an enemy) up to some (relatively low) echelon of command. Higher echelons of command are vested in humans, who are also responsible for coordinating the activities of multiple SAFORs as well as intervening at lower echelons when the automated responses are inadequate (thus justifying the term "semi-automated").

The success of SIMNET has led ARPA to pursue the extension of DIS technology across all services (Air Force, Army, Marines, Navy). The goal is to provide a cost-effective and realistic environment for training and rehearsal, as well as for testing new doctrine, tactics and weapon system concepts [Loper et al. 94]. The extension involves expanding the types of vehicles (to include essentially all active ground, surface and air vehicles), as well as greatly increasing the number of participating entities to a level of approximately fifty thousand (thus enabling effective training at higher levels of command). Despite this large desired growth in the number of entities, cost factors dictate an upper bound of approximately one thousand on the number of humans that can directly participate. Thus the remaining forty-nine thousand entities required to fill out the battlefield must be computer generated forces.

SAFORs can provide many of these automated forces. However, they have sufficient weaknesses to have led ARPA to consider alternatives. One major weakness is that the autonomy of SAFORs is extremely limited. Although they can perform simple maneuvers, SAFORs still require significant human management when it comes to the fine points of executing a complex coordinated activity, such as moving into a defensive position on a hill. This requires knowledge that SAFORs do not have about the potential avenues of attack, how to overlap fire for mutual protection, and how to hide behind features of the landscape. The resulting need to micro-manage SAFOR entities can easily overload a human controller during the height of an engagement. Moreover, the human must attempt to simulate the combined reasoning and interactions among many independent entities in accordance with the appropriate military tactics and doctrine. Finally, the human controller must be highly skilled — it takes approximately six months to train a new controller. Such considerations led ARPA to initiate the intelligent forces (IFOR) program, with the goal to develop fully automated intelligent forces that can participate in the simulated combat environment. The ultimate idealization is to have "plug compatibility" between humans in physical simulators and computer generated forces. Either could be plugged in to play a particular role in an engagement.

The key research challenge in building IFORs is generating intelligent, human-like behavior in a complex, dynamic and uncertain battlefield simulation environment. The ideal is automated behavior that is indistinguishable from human behavior, at least with respect to the modalities through which agents can be observed (such as visual and radar sensing of the vehicle) [Tambe *et al* 94]. However, pragmatically, all that is really necessary is the production of behavior that is close enough to that of humans to force the other entities in the simulation, whether human or IFOR, to interact with the IFORs in the same way they would interact with humans. This is still an ambitious goal, but because this is a synthetic environment, the IFORs at least do not need to deal with issues of low-level perception and robotic control. Furthermore, the forms of interactions between entities on a battlefield are greatly restricted by the specific goals, doctrine, and missions of the participants. For instance, there are essentially no verbal interactions between opposing entities, and cooperating entities restrict their communication to the details of current missions — they rarely if ever discuss poetry.

The goals of this article are to explore the challenge of creating intelligent agents for distributed battlefield simulations, to outline our progress to date in meeting this challenge, and to discuss the remaining research challenges. It is difficult in general to say what challenges an environment presents

without understanding what tasks are to be performed in that environment. For this effort, these tasks are primarily defined by the demands of a large-scale simulated tactical exercise to take place in 1997 — called Simulated Theater of War '97 (STOW-97). STOW-97 is to involve approximately fifty thousand entities from all branches of the military. Our challenge is to provide synthetic pilots for all of the missions in STOW-97 that involve military aircraft, including fighters, bombers, troop transports, reconnaissance aircraft, and helicopters.

The body of this article begins, in Section 2, with an examination of the simulated air-combat environment, and in particular an in-depth exploration of a fragment of a typical air-to-air mission, as might occur during STOW-97. This exploration brings out a set of requirements for automated pilots that includes a variety of important capabilities: goal-driven and knowledge-intensive behavior, reactivity, real-time performance, conformance to human reaction times and limitations, performance of multiple overlapping tasks, coordination with other agents, communication, agent modeling (plan recognition), temporal reasoning, planning, maintaining episodic memory, and explanation.

Section 3 describes in detail the design and status of our automated pilots, along with the work that remains to enable them to participate in STOW-97. The creation of effective, intelligent, automated pilots is facilitated by an underlying architecture that can support the above listed requirements in an integrated fashion. To develop our pilot agents, we have focused on the Soar integrated architecture [Laird, Newell, and Rosenbloom 87, Rosenbloom, et al. 91]. The system we have developed within Soar, TacAir-Soar [Jones 94, Johnson et al. 94] represents a generic automated pilot. Specific automated pilot agents are created by specializing it with specific parameters and domain knowledge. These pilot agents then participate in battlefield simulations via simulated aircraft provided by ModSAF [Calder et al. 93], a distributed simulator that has been commercially developed for the military. Versions of these automated pilot agents have already been demonstrated on several occasions to Navy personnel, ranging from current and former Navy pilots up to a handful of admirals, at Naval Air Station Oceana (in Virginia Beach, VA). Two of the earlier demonstrations focused on small-scale simulated air-combat engagements against expert human pilots. While these engagements were simplified both in terms of the numbers of aircraft involved in the combat and the types of maneuvers the human pilots could use, the automated pilots' performance was impressive enough for the Navy personnel to provide very enthusiastic feedback. Capt. McLane, Commander of Navy Fighter Weapons School (TOPGUN), was present at one of the demonstrations and commented as follows on the set of air combat simulation capabilities he had witnessed (of which TacAir-Soar provided the only intelligent automated pilots):

I am totally enthusiastic. I have been looking for six months for this type of capability. I want one now.

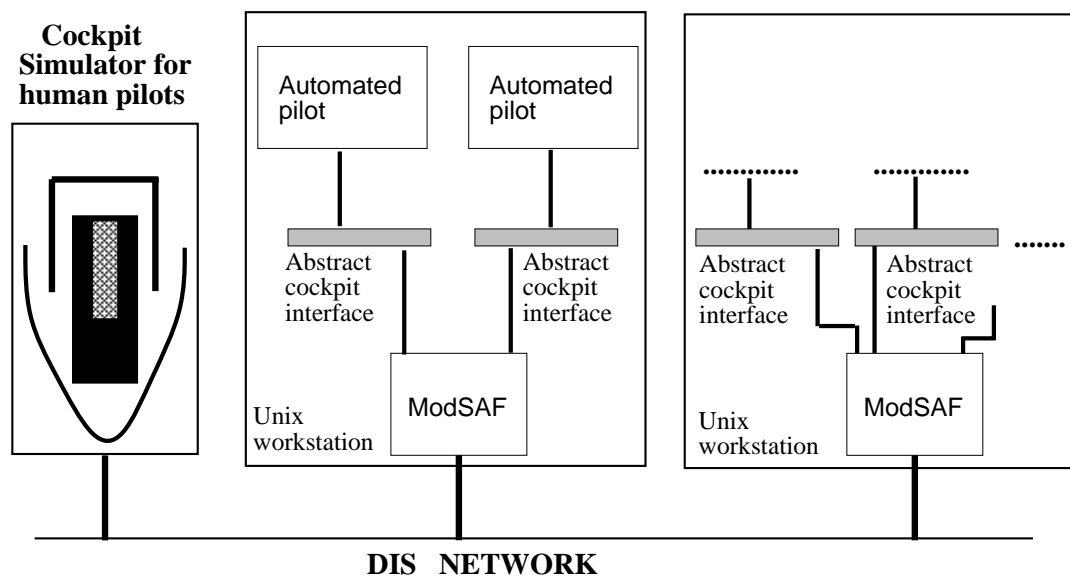
The most recent demonstration was part of a much larger engagement called Simulated Theater of War, Europe (or STOW-E), which was held November 4-7, 1994. STOW-E was an operational (and international) military exercise — in fact the largest distributed simulation exercise to date — involving approximately 2000 entities originating from nineteen sites across the United States and Europe. These entities represented everything from aircraft, ships, and ground vehicles to individual soldiers and missiles. Some entities corresponded to actual manned vehicles (e.g., actual F18 fighter jets) that were broadcasting their status over the DIS network, or to manned simulators. However, most of the entities corresponded to computer generated forces. Our role in STOW-E was to provide IFORs for a range of air-to-air, air-to-ground and air-to-surface missions. This was a real learning experience for us in terms of understanding what happens in such complex situations, both in terms of the stress it puts on the low-level technology (such as the number of cycles required to constantly maintain and display the status of so many entities on a workstation), as well as the many surprises it provides for the agents (such as their being shot down by surface-to-air missile batteries that we had no idea existed, or the unexpected forms

of behavior exhibited by unconstrained human pilots). Nonetheless we were able to successfully participate in a range of missions across all four days of the exercise, and on more than one occasion shot down human pilots flying in simulators. Our understanding is that this makes TacAir-Soar the first AI system to have participated directly in an operational military exercise.

Finally, Section 4 provides a broader perspective on TacAir-Soar's relationship with other AI systems, and Section 5 concludes.

## 2. The Air-Combat Simulation Domain

In DIS, large-scale, interactive simulations are developed from a set of independent simulator nodes linked together via a network. These nodes communicate via the DIS standard network protocol [IEEE 93]. Each node simulates the activities of one or more entities in the environment and communicates relevant information to the other nodes in real time. To control the cost of the simulation, a central hypothesis underlying DIS is that there is no need to strive for completely realistic (i.e., *high fidelity*) simulations of the battlefield. Instead, the emphasis is on simulating the environment to a sufficient level (i.e., *selective fidelity*) so as to trigger tactically relevant responses from humans [Crooks et al. 92].



**Figure 1:** Human and automated pilots interact with the DIS environment via distributed simulators.

Figure 1 illustrates a portion of the DIS network. It shows three interacting simulator nodes: one of them a physical cockpit simulator for human pilots, and the remaining two UNIX workstations, each running a copy of the ModSAF simulator [Calder et al. 93]. Each copy of ModSAF can simultaneously simulate a number of different fighter aircraft.<sup>1</sup> These aircraft can engage in simulated air combat with and against each other, either on a single workstation or across the network on multiple workstations, in real time. The total number of aircraft simulated on a single workstation depends on the complexity of the automated pilots flying the aircraft, and obviously the type of workstation. We are presently able to run up to six TacAir-Soar piloted aircraft, plus ModSAF, on a single R4400-based Silicon Graphics

<sup>1</sup>ModSAF also provides other types of air, ground and sea vehicles, not discussed here, plus facilities for creating and controlling semi-automated forces.

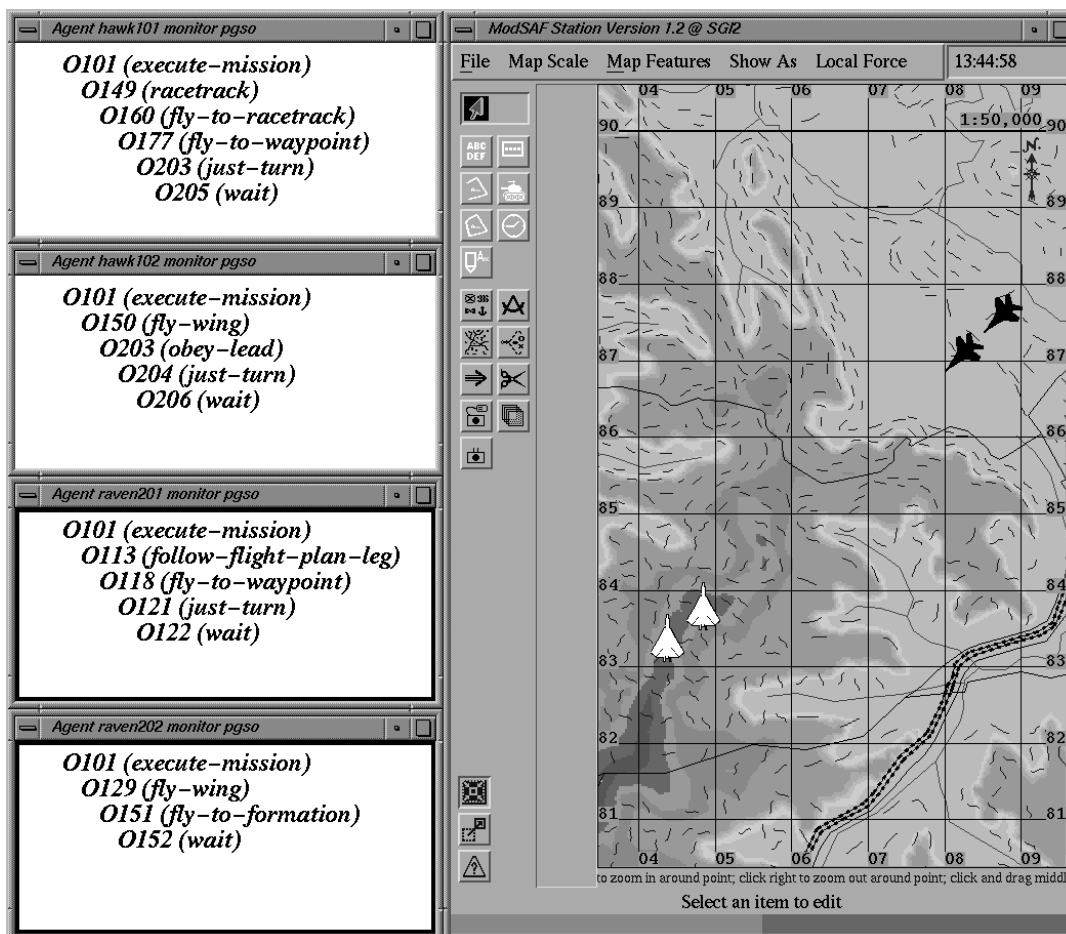
Indigo workstation. ModSAF models the relevant details of the aircraft, including their motion dynamics, sensors, and weapons capabilities, for selective fidelity simulation. In running the simulation, ModSAF's basic cycle starts by sequentially invoking each agent. It then updates its simulation model based on the actions of these agents, data received over the network about other agents, and the predicted movement of other agents (based on elapsed time). Finally, it broadcasts changes in its agents' motions over the network. The basic cycle time is flexible, adjusting from between two and fifteen times a second based on available computational resources.

As shown in Figure 1, an automated pilot makes use of an abstract cockpit interface [Schwamb et al. 94] to interact with the aircraft provided by ModSAF. This interface is intended to emulate the environment that a human pilot encounters — the aircraft cockpit. The interface provides an automated pilot with data corresponding to sensory information it would obtain from an aircraft cockpit, e.g., radar displays, radio messages, vehicle status instruments, and visual sightings out of the cockpit canopy. However, to avoid expensive low-level signal processing, the interface provides this data not in the form of actual images, or digitized audio, but rather in the form of numeric or symbolic encodings. For instance, the interface supplies numeric data corresponding to the aircraft's altitude, say 10000 feet, rather than an image of its altimeter. The sensors do not provide unlimited access to the environment, but instead attempt to model the perceptual limits actual pilots would have. Currently, there are a total of 31 static inputs and 19 dynamic inputs available to an agent, which describe its own state, plus 3 static inputs and up to 10 dynamic inputs for each radar or visual contact.

In order to allow the simulated aircraft to be controlled by an automated pilot, the interface also accepts commands from the pilot, again in numeric or symbolic form. These commands are then sent to ModSAF in order to control the aircraft's motion (desired speed, heading, altitude, rate of turn), adjust the radar system (orientation, mode), select and fire missiles, or communicate via radio. Currently, there are a total of twenty-six types of output commands that an agent may send to the interface. The underlying model of the vehicle constrains the agent to perform only those actions that would be possible in a real plane. This has also been extended to include some constraints on pilots. For example, the automated pilot can never cause the plane to perform maneuvers that would exceed the G force sustainable by a human pilot. One verification of the sufficiency of the cockpit abstraction is that we have been able to build software on top of it that allows humans to fly aircraft effectively in ModSAF via a simple, graphical, mouse-driven interface [van Lent and Wray 94].

Figure 2 shows a snapshot taken from ModSAF's plan-view display of an air-combat simulation. The background shading indicates the terrain. The grid is used for measurement; in this case, each grid cell indicates an area of 10x10 kilometers. The four aircraft icons in the figure have been artificially magnified for visibility. The light-shaded aircraft are simulated F-14D aircraft and the dark-shaded ones are simulated MiG-29 aircraft. The smaller windows along the side of the figure display information about the active goal(/operator) hierarchy of the TacAir-Soar-based pilots controlling these aircraft. This is a small portion of the entire forest of goals that the pilots can possibly have (Figure 3).

Figure 4 shows this same snapshot, but with the terrain and side windows removed for clarity, and an airfield added (shown by the crossed lines in the bottom left corner). This is the first snapshot in a scenario that will be used to illustrate the key requirements for building intelligent automated pilots in this environment. The scenario involves four fighter aircraft: the two light-shaded aircraft are defending the airfield, and the two dark-shaded aircraft are attempting to clear out the light-shaded aircraft so that bombers can later make it to the airfield uncontested (these bombers are not shown in the figure). As the scenario progresses (Figures 5 through 9), the two light-shaded aircraft fire missiles at the dark-shaded aircraft; however, these missiles fail to hit their targets. The aircraft then re-position themselves for a new



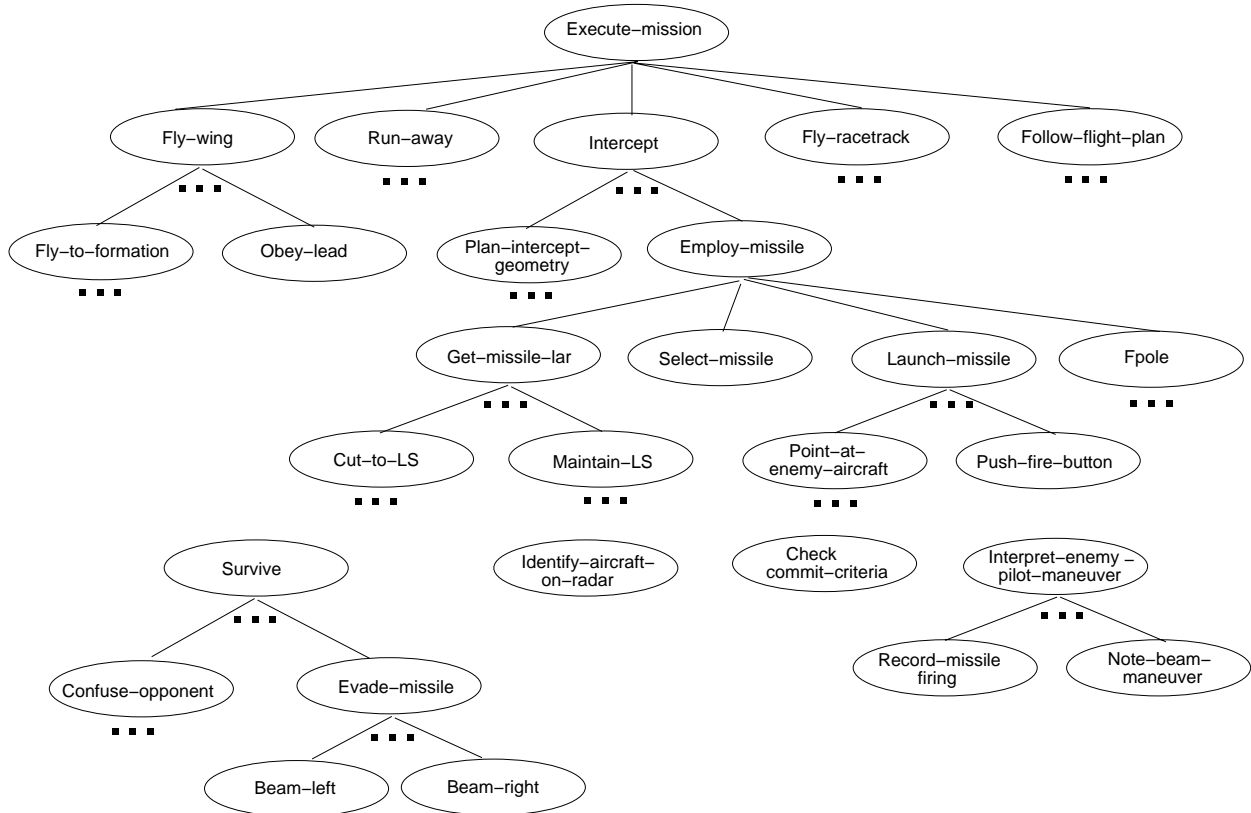
**Figure 2:** A snapshot of ModSAF's simulation of an air-combat situation.

attack. This scenario is representative of simulation runs involving TacAir-Soar-based automated pilots, although there are some minor modifications for illustrative purposes.

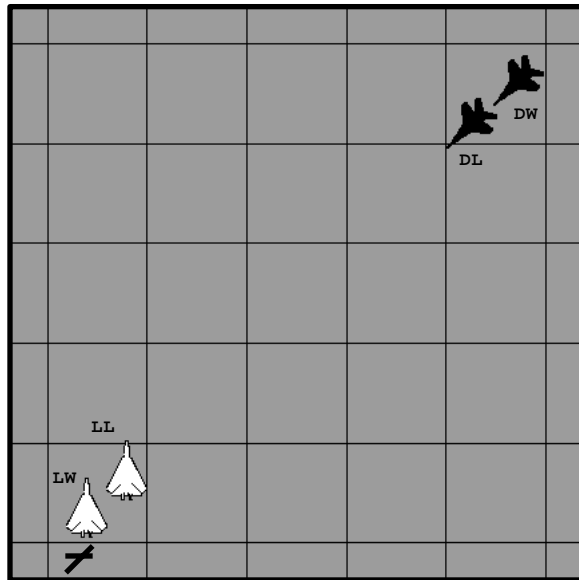
At the beginning, in Figure 4, the two aircraft groups are separated by a large distance, and cannot sense each other, even on long-range radars. To defend their airfield, the light-shaded aircraft are flying a Combat Air Patrol; that is, they are patrolling in a repetitive pattern between the airfield and the anticipated direction of attack. In service of this, they are flying as a closely coordinated pair, called a *section*. As specified by the mission parameters, pilot LL (i.e., light lead) is the leader or *lead* of this section. Pilot LW (i.e., light wingman) is the *wingman* of this section — it is maintaining formation and closely following any commands issued by LL. LL and LW are flying in a formation called the *fighting wing*, with LW slightly behind and offset from LL. The two dark-shaded aircraft — DW (dark wingman) and DL (dark lead) — are flying in a *trail* formation, with DW directly behind DL. DL and DW are attempting to clear out the light-shaded aircraft by performing a *Sweep*; that is, they will fly in and engage LL and LW in an attempt to either destroy them or drive them away.

The specific mission and mission-related parameters for each of the pilots are specified to them in a *briefing* session prior to the mission, i.e., before the aircraft are airborne. (The actual selection process for particular missions or mission-related parameters will not be of concern in this scenario.) The mission parameters typically specify a pilot's role as the lead or wingman in the mission, its initial flight formation, and the specific numbers and types of missiles that it is to carry on its aircraft. In addition, the





**Figure 3:** A portion of an automated pilot’s forest of goal hierarchies.



**Figure 4:** Snapshot 1: The initial positions of the aircraft in the air-combat simulation scenario.

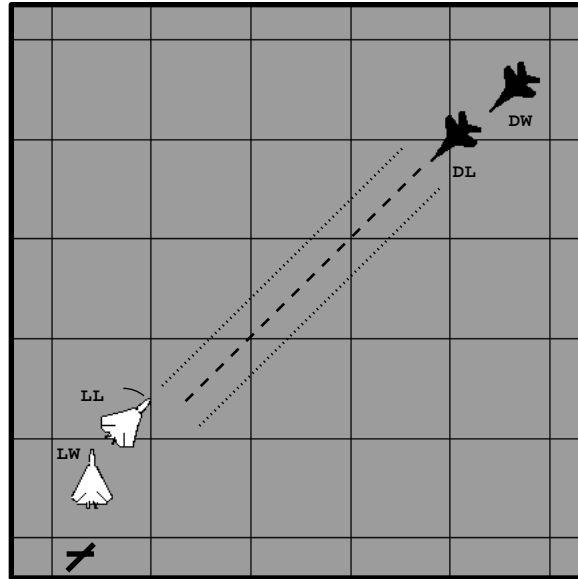
mission parameters also specify the *rules of engagement* and the *commit criteria*, both important parameters. Rules of engagement are a set of directives regarding the general conduct of a mission. For

fighter pilots, these directives may specify that they must positively identify an aircraft to be enemy military before firing. This identification can be done using a variety of methods, such as the use of special electronic beacons. Commit criteria are used by a pilot to decide whether to intercept, i.e., engage in combat with, an enemy aircraft. For instance, the commit criteria may specify that a pilot must intercept any enemy aircraft approaching its airfield, but not intercept other enemy aircraft (since that may leave the airfield unprotected).

The important point to note in the discussion of Figure 4 is that the mission and mission-related parameters help to specify a pilot's goals. Figure 3 illustrated a portion of the forest of goal hierarchies that are embodied in the pilots. Not all of these goals will be active at any one time — for example, the pilot will usually only be executing one mission at a time, so only one child of Execute-mission will be active at a time. However, goals from multiple hierarchies, such as survive and execute mission, will often be simultaneously active. The pilot's behavior needs to be driven by these active goals. For the pilot to be able to generate this behavior, it needs to possess a large amount of knowledge relevant to the various missions and mission-related parameters. Figures 3 and 4 thus illustrate a requirement for knowledge-intensive and goal-driven (top-down) behavior. In addition, the previous discussion also illustrates a requirement for reactive (bottom-up) behavior. In particular, a wingman (either LW or DW) constantly needs to monitor its position with respect to its lead (either LL or DL), and continually modify its heading, speed and altitude, so as to maintain the formation. This level of reactivity on the part of the wingman is important — due to the realism in this simulation, continuous monitoring and correction is necessary even while attempting to fly straight towards a pre-determined destination.

Now consider the situation shown in Figure 5. Here, LL and LW, by virtue of their longer range radar, spot DL and DW before they are themselves detected. As the lead of its section, LL checks the commit criteria and, as per its rules of engagement, attempts to make a positive identification of the enemy aircraft. In this case, we assume that it succeeds in doing this. In addition, since these aircraft are approaching in the direction of LL's airfield, its commit criteria are also satisfied. Hence it decides to intercept them. LL's first step is to ensure that, while prosecuting this intercept, it will maintain advantageous conditions for itself and its partner (LW). For instance, it is advantageous for LL to maintain roughly ten kilometers of *lateral separation* during the intercept. Lateral separation is the perpendicular distance between LL's position and its opponent's (in this case DL's) projected straight line of flight. In this case, LL gains lateral separation with respect to DL rather than DW because DL is closer and hence considered the primary threat. DL's projected line of flight is indicated by the dashed line extending from its nose in the figure. Flying along either of the two dotted lines running parallel to this line can provide the desired ten kilometers of lateral separation. This lateral separation is advantageous because, among other things: it makes it easier to perform a visual identification (should one be necessary); it isolates the opponents on one side of the aircraft (if done to the correct side); and it makes it easier to turn back to reengage if the opponents are passed and not destroyed. In addition to achieving lateral separation, it is advantageous to maintain high speed and altitude during the intercept (for improved maneuverability), and to not cross an enemy aircraft's projected flight path.

Based on the above considerations, LL could conceivably generate a detailed plan prescribing the precise sequence of actions for an optimal intercept of DL and DW. However, such a plan is almost guaranteed to fail because of the uncertainty and dynamics of this environment. For example, any turn by the opponent will invalidate the details of the plan. However, the obvious alternative of just reacting to the situation is also inadequate because the implications of the situation may extend considerably beyond the immediate choice of action. As a compromise, TacAir-Soar commits to abstract plans, such as to prosecute the intercept from DL's right side (which keeps the opponents on one side without crossing their path). Such plans constrain future actions [Bratman, et al. 88], but can themselves be discarded if



**Figure 5:** Snapshot 2: LL attempts to achieve lateral separation.

the situation changes enough to warrant it (such as if an attack were now required from the other side).

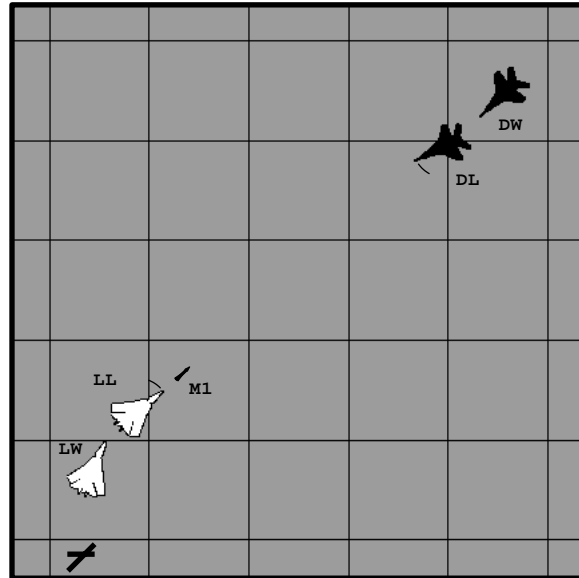
Once the abstract plan is established, LL attempts to position itself to fire a missile at its opponent. LL is flying an F-14D, which in this case carries two long-range radar-guided missiles, as well as several medium and short-range missiles. For LL to fire its long-range missiles, it needs to be within a particular range (distance) from the enemy aircraft. Furthermore, as just discussed, it should also have about 10 kilometers of lateral separation. As shown in the figure, LL achieves these goals by specializing its abstract plan so as to fly to DL's right side, along the ten kilometer lateral separation line. (The small arc along LL's nose in the figure indicates its turning in service of this plan.)

If at this point, DL suddenly turns and increases or decreases lateral separation, LL needs to turn as well, to attempt to regain the desired lateral separation. While LL is maneuvering in this fashion, LW continues to maintain its position relative to LL based on its perception of LW's motion. In general, LL will not directly communicate its maneuvers to LW, but will instead expect the wingman to observe and match the turns on its own. This reduces the cognitive load on LL and minimizes radio communication. The few exceptions to this rule are for particularly large turns, for which the wingman's understanding of the turn is more important than the communication costs.

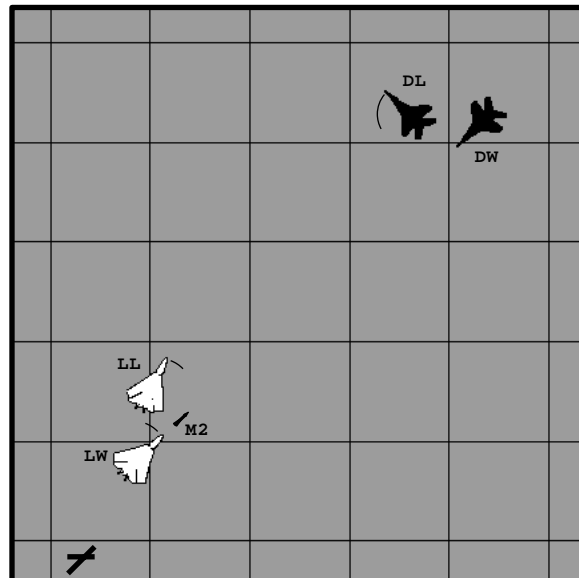
There are two important points that come out of the description of Figure 5. The first is again the pilots' mix of goal-driven and reactive behavior, as seen in LL's decision to intercept its opponent, and its continued maneuvers to maintain lateral separation. The second is that fragments of planning are also seen, particularly in LL's determination of the direction for the intercept.

As shown in Figure 6, LL eventually reaches its long-range missile firing position, and fires a missile (M1) at DL. In the process it had to decide whether it should target DL, or go after DW instead. This decision is made by first selecting a *sorting* criterion, which specifies how the opponents should be divided up between LW and itself. Based on the relative positions of the four aircraft, LL decides here to sort by *range*, and communicates this criterion to LW. LW, as the wingman, then targets the furthest opponent (DW), while LL, as the lead, targets the closest (DL). LL therefore targets DL, by getting a radar lock and turning to *attack heading*, i.e., the collision course for its missile. It then fires missile M1, which obtains its guidance information from LL's radar lock. To continue to provide radar guidance to

M1, while reducing the rate at which the two aircraft are closing, LL then turns its nose away from DL — a maneuver known as an *Fpole* — as shown in Figure 7. As a result of the *Fpole* and the decreased rate of closure between the two planes, DL will not achieve its own missile launch criteria until after M1 has reached it (MiG-29's do not have missiles with as long a range as an F-14D).



**Figure 6:** Snapshot 3: LL fires a missile at DL.



**Figure 7:** Snapshot 4: DL and DW begin a Post-hole maneuver.

Meanwhile, back in Figure 6, the aircraft come sufficiently close to allow DL and DW to acquire radar contact with LL and LW. They then maneuver themselves to reach an advantageous position in their intercept of LL and LW. While maneuvering, they observe LL's own maneuvers on their radars. These particular maneuvers are closely associated with missile firings. Therefore, DL and DW infer a possible missile firing, although they have no means to physically detect the missile (the radar cross-section of a missile is too small to be detected) and thus they cannot be absolutely certain about it. Nonetheless, they assume the worst and act as though a missile has actually been fired. They therefore suspend their

intercept maneuvers and instead attempt to survive the missile. Of several available alternatives, in this case they execute a *Post-hole* maneuver. This maneuver aims not only to defeat the missile, but also to confuse LL and LW. The maneuver involves DL's turning full circle, so that it ends up behind DW. It also involves DL's switching roles as the lead with DW. The intended result is for DL and DW to again end up in trail formation (their previous formation), but with DW as the lead, and with LL and LW confused about how many aircraft there are out there, and where they are.

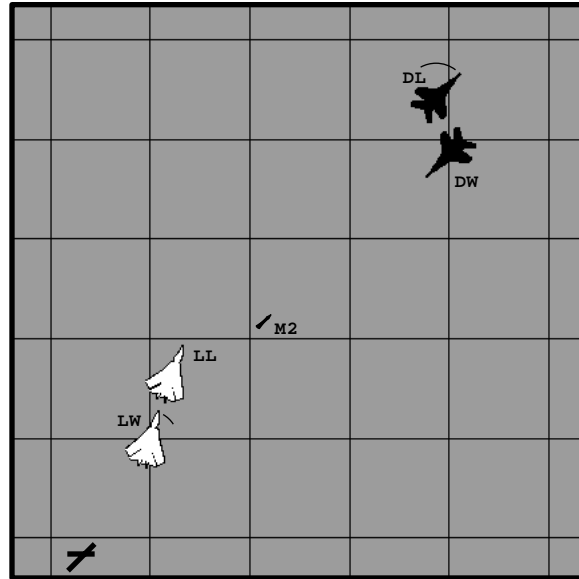
Figure 7 illustrates why the Post-hole maneuver can help DL survive the missile attack. As DL begins executing this maneuver, it turns roughly perpendicular to the path of LL's radar emission. Since LL's radar operates on the "Doppler principle", it cannot detect this perpendicular motion.<sup>2</sup> This causes LL to lose radar contact with DL, thus defeating missile M1 by depriving it of the required radar guidance (hence M1 is not shown in Figure 7). This loss in radar contact can also potentially confuse LL. To minimize this confusion, LL has to infer DL's on-going maneuver based on DL's turn. LL can then anticipate the likely loss in radar contact, recognize that its missile is likely to be defeated, and switch radar modes to maintain contact (which is not to be done lightly because it guarantees that the missile will no longer track DL). During all this, as shown in Figure 7, LW reaches its own missile firing range and fires a missile (M2) at DW.

Figures 6 and 7 again illustrate the knowledge-intensive nature of this task — the pilots need to be well-informed about various maneuvers, such as Fpole and Post hole. More importantly, the description also illustrate some new requirements. The first new requirement is for agent modeling [Anderson, et al. 90, Ward 91], which involves observing other agents' low-level actions — in this case an opponent's turns — to infer its higher-level behaviors. This capability is closely related to plan recognition [Kautz and Allen 86], although the emphasis here is not so much on recognizing other agents' plans, as it is on recognizing their flexible combination of goal-driven and reactive behavior. The second new requirement is simultaneous performance of multiple tasks. This is seen in DL's modeling of its opponents' actions, while simultaneously maneuvering its own aircraft, and possibly monitoring quantities such as lateral separation for its own intercept. A third new requirement is real-time performance. This is seen in DL's need to infer a missile firing, interrupt its on-going maneuvers, and initiate a Post-hole maneuver, all before the missile gets too close. A fourth new requirement is the ability to act in a coordinated fashion. This is seen in their flying and turning in formation, performing coordinated maneuvers, and dynamically changing responsibilities between lead and wingman. The fifth, and final, new requirement is the ability to communicate with other pilots. This is seen in the communication between LL and LW that occurs during the process of sorting the opponents.

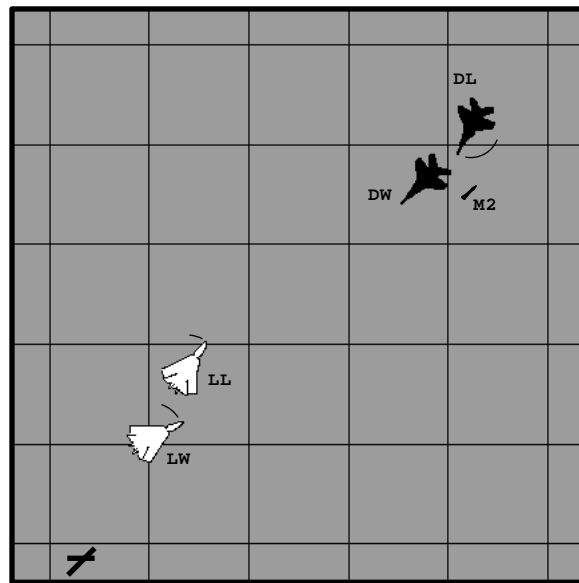
In Figure 8, as DL continues to turn full circle, it loses radar contact with the other aircraft. As DL turns even further in Figure 9, its radar again shows three aircraft. However, DL cannot automatically assume them to be DW, LL and LW. It is quite possible that during the time that DL lost contact, three new enemy aircraft have approached the combat arena. Such a situation would be extremely unfavorable for DL. Thus, DL needs to re-establish the identities of the aircraft on its radar. In military parlance, this is the problem of maintaining *situational awareness*. To this end, DL estimates the total time interval for which the aircraft have been invisible to it. Based on this estimate, and the speed of each aircraft, it projects the most likely position of each aircraft. Then, based on this extrapolation, it concludes that the three aircraft on its radar are most likely to be DW, LL and LW.

---

<sup>2</sup>LL's radar has other modes of operation, not all of them based on the Doppler principle. However, the Doppler mode is necessary for long-range missile firing.



**Figure 8:** Snapshot 5: DL's turn causes it to lose radar contact with all other aircraft.



**Figure 9:** Snapshot 6: All four fighter aircraft have re-positioned themselves for a new attack.

Meanwhile, in Figure 8, DW, which has taken over as the lead since the beginning of the Post-hole maneuver, continues to fly straight towards the opponents. Unfortunately, in this case, DW focuses its radar on LL and fails to observe LW's maneuvers. As a result, it fails to recognize LW's missile firing. Fortunately for DW, missiles sometimes simply fail to reach their targets, and this one passes by harmlessly (Figure 9). Meanwhile, both LL and LW recognize that their missiles have failed to hit their targets — LL recognizes this since it loses radar contact with DL, while LW keeps track of the missile's expected time of impact and notices that this time has been exceeded, without its target disappearing from the radar. Furthermore, since LL and LW are neither fooled nor confused by their opponents' maneuvers, they turn back for a new attack (Figure 9).

The new requirements brought out in Figures 8 and 9 are centered on DL's turn. The plane that DL is flying could theoretically execute a full-circle turn extremely quickly. However, there are physical

limitations on how quickly a human can turn a fighter aircraft without blacking out. Furthermore, after completing its turn, DL must identify the radar blips on its screen. Although information is available within the simulator so that this identification could be done immediately, that would cause DL to react more quickly and more accurately than could humans, thus impacting the realism of the exercise. Therefore, the inputs available to DL are made to correspond to the same inputs that would be available to a real pilot, leading to DL's need to perform a mental projection to maintain situational awareness, just as humans would have to. For this projection, DL must track the time interval during which aircraft were invisible, and reason with that. Thus, there are two additional requirements that arise here. First, the automated pilots must conform to human limitations on aspects such as G forces, reaction times, and levels of accuracy. Second, the automated pilots need to reason with temporal intervals. This reasoning with temporal intervals also shows up in LW's tracking its missile's expected time of flight, as well as in other situations in this environment.

Normally, the scenario described here would continue beyond the snapshot depicted in Figure 9, and the pilots would continue to employ different maneuvers until either they themselves or their opponents get destroyed or run away. Following the scenario, they would participate in a debriefing process in which they would be asked, by an instructor, an evaluator, or a superior officer about what happened during the mission and why they did what they did at key junctures [Johnson 94a, Johnson 94b]. Debriefing is an important means by which others can validate and gain confidence in the pilot's judgement. With respect to capabilities, debriefing reinforces the need for communication, though this time in a more expansive and less time-critical setting. It also reveals the need for some form of episodic memory to enable the pilots to remember what transpired during their engagements, plus — and most critically — a self-explanation capability. Such a capability is essential for automated pilots, as it can help give military personnel confidence in the agents' knowledge and reasoning capabilities. Though debriefing can also potentially provide feedback from which the pilot could learn, such learning doesn't actually come up as a requirement in this scenario, because these are expert pilots performing a routine mission. Such learning is likely to become increasingly important in the future though.

The scenario that we have just gone through is only one of many possible scenarios, and a large number of variations are possible. Nonetheless, the snapshots and associated discussion bring out most of the key requirements for building automated pilots in this virtual environment, at least within the space of simulation parameters that we have encountered so far. These requirements are summarized below:

- Goal-driven behavior
- Knowledge-intensive behavior
- Reactivity
- Real-time performance
- Conforming to human reaction times and limitations
- Overlapping of performance of multiple high-level tasks
- Multi-agent coordination
- Communication
- Agent modeling (especially opponent modeling)
- Temporal reasoning (dealing with time intervals)
- Planning
- Maintaining episodic memory

- Explanation

Additional requirements are also expected to come up as these requirements are addressed, and as new types of vehicles (such as helicopters) and missions are investigated. For example, learning is eventually going to be required in order to improve performance both during and across engagements. However, because it doesn't come up here to any great extent, we won't focus on it further in this paper, other than to mention that using it effectively in this domain looks to be quite non-trivial.

### 3. Creating Intelligent Automated Pilots

To create effective automated pilots, it is necessary to address the list of requirements from the previous section in an integrated fashion. We have chosen to do so using Soar, a unified software architecture being developed as a basis for integrated intelligent systems [Laird, Newell, and Rosenbloom 87, Rosenbloom, et al. 91]. Soar is also a developing unified theory of cognition [Newell 90, Newell 92a, Newell 92b, Rosenbloom and Laird 94]. This theory entails a number of constraints (e.g., concerning human reaction times, knowledge representation, and learning) that bear directly on potential solutions to the requirements presented above. For readers unfamiliar with Soar, a brief description appears on the opposite page.

TacAir-Soar has been constructed from Soar through the addition of perceptual and motor interfaces (in the form of C code) that allow TacAir-Soar pilots to fly ModSAF planes in the DIS environment; specific knowledge about the tactical air combat domain (in the form of rules); and strategies for aiding in addressing those requirements not fully addressed by the Soar architecture itself (also in the form of rules). This interface and knowledge is currently maintained as a constant across all TacAir-Soar-based pilots. Individual pilots can then vary in terms of the mission and vehicle parameters provided to them during the briefing process, and to a small extent, the knowledge and tactics that they employ (by selectively disabling parts of less "competent" pilots). Allowing wider variations in structure, knowledge (including of tactics), temperament, and preferences is still a topic for future work.

The first two required capabilities — for goal-oriented and knowledge-intensive behavior — are addressed by the manner in which TacAir-Soar uses its architecture and knowledge in the process of dynamically expanding and contracting a goal hierarchy for this domain. To see this, let's look at a concrete example of LL's operation from a situation that arises shortly after it has decided to prosecute the intercept of DL and DW (Figure 5). Figure 10 depicts LL's hierarchy of problem spaces and operators. (Goals are represented implicitly in this diagram as the desires to apply operators that have reached an impasse.) This hierarchy gets generated as follows.

In the top-most problem space (**TOP-PS**), LL is attempting to accomplish its mission by applying the *EXECUTE-MISSION* operator. The termination condition of this operator is the completion of LL's mission (which is to protect its home base for a mission-specified time period). Since this has not yet been achieved, a subgoal is generated to complete the application of *EXECUTE-MISSION*. The **EXECUTE-MISSION** problem space is selected for use in this subgoal. In this problem space, LL selects the *INTERCEPT* operator to perform the intercept. The termination condition of this operator — that the opponents are either destroyed or chased away — is also not yet achieved, leading to a second subgoal of completing the intercept. The **INTERCEPT** problem space is selected for use in this goal. Within this problem space LL selects the *EMPLOY-MISSILE* operator, which uses missiles to either destroy the opponents or to force them to run away. However, since LL has not reached a position from which it can fire a missile, a third subgoal is generated. The **EMPLOY-MISSILE** problem space is selected for use in this subgoal. Within this problem space LL attempts to achieve a firing position for its



## Soar

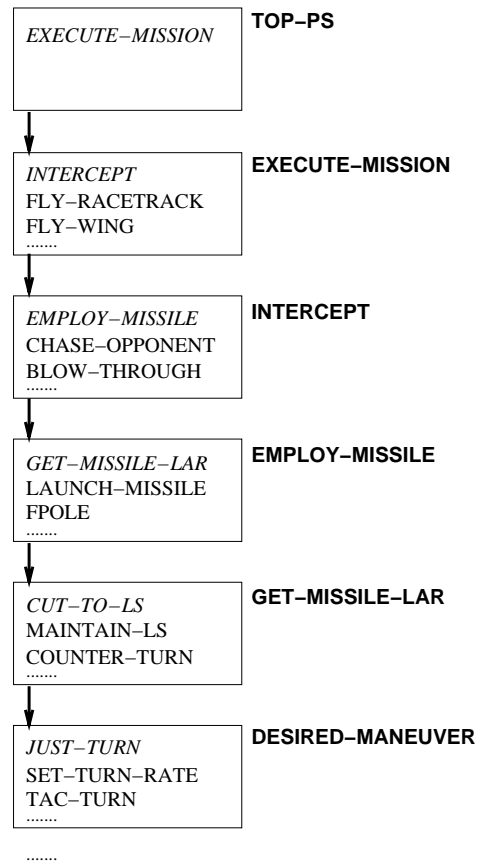
All tasks in Soar are formulated as attempts to achieve goals in problem spaces [Newell *et al* 91]. Each problem space consists of a set of states and a set of operators. States represent situations and operators represent actions. Operators perform the basic deliberate acts of the system. They can perform simple, primitive actions that modify the internal state (such as determining if all commit criteria are achieved) and/or generate primitive external actions (such as switching radar modes), or they can perform arbitrarily complex actions, such as executing a mission. The basic processing cycle is to repeatedly propose, select and apply operators of the problem space to a state, moving ahead one decision at a time. Operator selection, application, and termination are all dynamically determined by the system's knowledge.

For expert-level performance, sufficient knowledge is generally available so that the selection of the next appropriate operator is not problematic. However, when operator selection knowledge is insufficient to determine the next operator to apply, an *impasse* occurs, leading to the creation of a subgoal to determine which operator should be selected (possibly through some search or planning technique [Laird, Newell, and Rosenbloom 87, Rosenbloom, Lee, & Unruh 93]). Similarly, if an operator is too complex and/or unfamiliar for the available application knowledge to handle directly, an impasse occurs, leading to the creation of a subgoal to apply the operator. This type of operator and associated subgoal is ubiquitous in TacAir-Soar, where operators such as *INTERCEPT* are selected, and in turn lead to a goal in which operators specific to intercepts, such as *PLAN-INTERCEPT-GEOMETRY* or *EMPLOY-MISSILE* are proposed, selected and applied to carry out the intercept. This is a goal decomposition scheme in which the decomposition of a goal is not static, but is determined step by step as operators are dynamically selected and applied. Thus, subgoals arise during the process of selecting and/or applying operators, leading to the dynamic generation of a goal hierarchy. These subgoals disappear when the associated impasse is resolved, either because an operator can be selected or because a selected operator is terminated.

Each problem space for each goal in the hierarchy has its own state. Each such state includes a representation of all of its supergoals (and their states) — Soar's subgoals are functionally at the meta-level [Rosenbloom, Laird, & Newell 88] — plus possibly representations that are local to that particular goal and problem space. The top state includes all sensor data from the external environment, which is thus also available in all subgoals as well. At any time, states at any level of the goal hierarchy can change, usually through the changing of sensor values. Since Soar's knowledge is also active for all levels, an operator can terminate at any level of the hierarchy — including intermediate levels — at any time. Such a termination automatically flushes all lower levels of the hierarchy, and may lead to the selection of a new operator to replace the terminated one.

All of Soar's knowledge, be it for selection, application or termination of operators is stored in the form of productions (condition-action rules). Any changes in goals, states and perceptions can cause these productions to fire. There is no conflict resolution, so all applicable productions are allowed to fire in parallel. Operator selection knowledge consists of productions that test the current goal and/or state and generate symbolic preferences about the absolute or relative worth of operators. For instance, given a particular situation, a production might generate a preference that an operator to turn left is better than an operator to turn right. As the state changes, some preferences may be retracted (if the situation no longer matches the conditions of the rules that generated them) and others generated, so that *decisions* — which are made by a fixed, architectural decision procedure — can always be based on preferences relevant to the current situation. Operator application knowledge consists of productions that modify the state in accordance with the particular operator selected, and possibly generate output commands. Operator termination knowledge consists of productions that test the state and generate a termination signal if the state corresponds to what the operator is to accomplish.

Goals, and their results, form the basis of Soar's learning mechanism, *chunking*. Chunking acquires new productions, called *chunks*, that summarize the processing that leads to subgoal results; that is, to elements generated in subgoals that are accessible in supergoals. A chunk's actions are based on these results. Its conditions are based on those aspects of the supergoals that are relevant to the determination of the results. Once a chunk is learned, it can fire in relevantly similar future situations, directly producing the required result, and possibly avoiding the impasse that led to its formation. This chunking process is a form of explanation-based learning (EBL) [Mitchell *et al.* 86, Rosenbloom and Laird 86].



**Figure 10:** LL's problem space hierarchy. Boxes indicate problem spaces, with names in bold to their right. Names within boxes indicate operators available within problem spaces. Italicized names indicate currently selected operators.

missile by selecting the *GET-MISSILE-LAR* operator. (LAR stands for *launch acceptability region*, i.e., the region from which LL can effectively fire a missile at its opponents). Since such a position has not been achieved, a fourth subgoal is generated, and the **GET-MISSILE-LAR** problem space is selected for it. Here, the *CUT-TO-LS* operator is selected to achieve lateral separation. Finally, the *CUT-TO-LS* operator leads to a fifth subgoal being generated within which the **DESIRED-MANEUVER** problem space is selected. Within this space, the *JUST-TURN* operator is selected, and applies directly to issue a new heading command to the cockpit interface, causing the aircraft to turn in an attempt to achieve lateral separation.

As LL turns, if the termination conditions of any of the operators in the hierarchy is satisfied, then that operator is terminated, and all of the subgoals generated in response to that operator are removed — they are now irrelevant. For instance, if the desired lateral separation is achieved, then the termination condition of the *CUT-TO-LS* operator is achieved. That operator is then terminated and replaced by the *MAINTAIN-LS* operator — to maintain lateral separation — while its subgoal and associated problem space (**DESIRED-MANEUVER**) are flushed. Similarly, if a missile firing position is reached, the termination condition of LL's *GET-MISSILE-LAR* operator is satisfied, and hence it is terminated along with the two subgoals beneath it.

The goal-driven aspect of this behavior should be obvious from the example. The knowledge-intensive aspect arises because all of TacAir-Soar's behavior is driven by accessing relevant knowledge about action. The generation of alternative operators occurs via rules that examine the goal, problem space, and

state and create instantiated operators that may be relevant. Once these operators are generated, further rules fire to generate preferences that help select the operator most appropriate for the current situation. Once an operator is selected, additional rules fire to perform the appropriate actions.

The requirement for reactivity is addressed through a combination of Soar's use of productions to enforce ubiquitous context-sensitivity in the representation and use of knowledge, plus Soar's open and integrative decision procedure. TacAir-Soar's productions are constantly testing input from sensors (such as the heading of an opponent), inferences drawn by other knowledge about collections of sensor data (such as that an opponent is a threat), the current goals and operators being pursued (such as intercepting the opponent), and potential actions to take in this situation (such as achieving a position to fire a missile). This knowledge can react to changes in the situation (such as a sudden turn by its opponent) by suggesting new operators to pursue at any level in the goal hierarchy, generating preferences among suggested operators, terminating existing operators, or even directly suggesting new external actions. When the productions have finished making their suggestions (generating their preferences) the decision procedure integrates all available suggestions/preferences together to determine what changes should actually be made to the goal/operator hierarchy. Thus, at every step, TacAir-Soar's behavior is based on its assessment, given the current situation and its own goals, of what is the best thing to do.

While using rules to provide reactivity is similar to the approach taken in pure situated-action systems [Agre and Chapman 87], this domain demands, and TacAir-Soar enables, the use of persistent internal state. Relying only on externally available cues is simply insufficient for effective performance here. Simple examples of this phenomenon can be seen in situations where other aircraft are not visible on an agent's radar. For instance, LL loses contact with DL in Figure 8, yet needs to re-orient itself to fire a second missile at DL. If LL waits for DL to re-appear on its radar, it may have already lost its advantage. Similarly, DL loses contact with the other aircraft in Figure 7, yet needs to turn a full circle. As discussed earlier, DL cannot just keep turning until it sees *some* aircraft back on its radar, since there could be other aircraft in the vicinity. In such cases, TacAir-Soar can engage in a substantial effort to deliberately maintain internal state information.

The basic real-time performance requirement is that simulated pilots need to be able to respond to environmental contingencies as fast as do human pilots; otherwise the behavior can be unrealistic. Although combat between jet aircraft sounds like the epitome of a high-speed, short-deadline task, it turns out that it mostly is not. In particular, at least in the initial phases of the combat, given the great distances involved, pilots do have seconds or tens of seconds available for their initial activities, such as identifying contacts on radar, or planning intercept geometry. However, even at these distances, there can be short bursts of intense activities, such as maneuvers to evade missiles, and unrealistically slow reaction times can be fatal. As the aircraft continue to maneuver and get closer to each other, slow reaction times become increasingly problematic. Such problems can arise either from too many decisions being required before action occurs, or from too much time being required for each decision. We have approached the former problem by focusing on expert-level performance in TacAir-Soar. In particular, we have tried to provide it enough knowledge so as to avoid the need for extensive meta-level reasoning and planning, especially during time-critical segments of an engagement. This expert-level knowledge allows operators and actions to be selected without impasses occurring, and therefore without extensive numbers of decisions. The latter problem — of individual decisions taking too long — can potentially occur because of slowness in the ModSAF simulator, the Soar architecture, the interface code between ModSAF and Soar (i.e., the cockpit abstraction), or the use of TacAir-Soar's knowledge. There is little action to be readily had from the first three sources: we have no direct control over the ModSAF simulator; the Soar architecture is already heavily optimized, with a new C-based version that is a factor of 15-20 faster than the previous Lisp-based version [Laird et al. 93]; and the interface is quite minimal. We have thus

focused our efforts on the fourth source — the knowledge (i.e., the rules). This has involved using rule coding styles that help ensure inexpensive rule match, detecting and eliminating unnecessary rule firings, and moving some low-level, bottom-up computations out of the rules and into the interface. As an example of moving computations from rules into the interface, consider what information the plane's radar provides to a TacAir-Soar-based pilot. Initially this consisted of only the position, heading and speed of detected planes. Productions would then compute relevant relations, such as the bearing to the plane, its lateral separation, and the bearing of the agent's plane from the other plane (these are both relevant for certain tactics). However, these values needed to be recomputed each time the planes moved (essentially every decision). So, the calculations were moved into the interface where they could be done more efficiently with dedicated C code, and the relations thus become part of the input from the sensor. This approach may seem inappropriate. However, it turns out that designers of fighter aircraft have also heavily utilized this strategy. They attempt to reduce the pilot's cognitive load as much as possible by off-loading many of the low-level, bottom-up computations onto the plane's instruments.

In addition to the basic real-time requirement, there is also a secondary real-time requirement here that stems from how the ModSAF simulator interacts with the synthetic pilots. ModSAF's control structure is based on the assumption that each agent can be trusted to relinquish control of the processor on its own, and therefore that each agent can be free to take complete control of the processor while it is active. Therefore, during each cycle of the simulation, each agent is called in turn to perform actions until it relinquishes control. Once all of the agents are done the environment is updated. ModSAF depends on each such cycle — of executing all of the agents plus updating the environment — completing in between 66 and 500 milliseconds total (the underlying simulation dynamically adjusts to the time required), though it can be longer for short periods of time without disrupting the simulation. To cope with these implicit soft deadlines, and thus to deserve the trust that ModSAF is implicitly placing in the TacAir-Soar-based pilots, we have limited each pilot to one decision per simulation cycle, and utilized the techniques described earlier to minimize the amount of time required per decision. The softness of the deadlines in this domain makes this an acceptable approach, as long as the system itself is fast enough, and the worst case behavior is not too far out of line.

The bottom line with respect to the real-time requirement is that on a single Silicon Graphics R4400-based Indigo Workstation, we are currently able to run the ModSAF simulator plus up to four TacAir-Soar-based pilots, while remaining within plausible cognitive limits 94% of the time and the simulator's expected time limits 100% of the time; and up to six pilots while remaining within plausible cognitive limits 75% of the time and the simulator's expected time limits 88% of the time. To go beyond six aircraft scenarios, additional aircraft can be added by distributing them across multiple workstations. As the need for synthetic pilots expands to hundreds (or thousands), the question of how much hardware is required per pilot will itself start to become critical. This suggests that additional speedups will always be useful.

The requirement of conforming to human reaction times is closely related to the real-time requirement. In fact, the basic real-time requirement is just that synthetic behavior be not too much slower than the corresponding human behavior. However, there is also a symmetric requirement that synthetic behavior not be too much quicker than the corresponding human behavior. To some extent, the Soar architecture starts us off in the right ball park in terms of how long it takes to do simple cognitive tasks [Newell 90]. And in practice, there has not been a significant problem with cognition running too quickly. What has proven to be more of a problem though is the rate at which human physical activities are simulated. Because we are not using a model of the human body, and in particular, of how fast it can perform physical activities such as moving eyes, scanning dials and radar screens, and turning knobs, TacAir-Soar pilots can react significantly more quickly than do human pilots. Pragmatically such problems can often

be ignored because they are effectively masked — that is, reduced to the level of noise — by the much longer response times of the plane or environment. For example, the time to respond to updated radar information may be dominated by the refresh time of the radar itself, which can require from two to ten seconds. In other cases we have introduced deliberate delays to slow down some of TacAir-Soar's responses. Similar ad hoc approaches have also been taken to modeling other human physical limitations; for example, maximum G forces sustainable by human pilots are modeled by restricting the plane so that it cannot exceed them (even if the plane theoretically could). The right long-term solution to these problems, of course, would be to utilize a better model of the physical body.

The remaining requirements from the list presented in Section 2 — overlapping of performance of multiple high-level tasks, agent modeling and coordination, communication, temporal reasoning, planning, maintaining episodic memory and explanation — are ones for which the underlying Soar architecture does not either directly provide a solution or dictate a specific approach to a solution. These requirements have thus presented interesting research challenges that have to be addressed, at least to some extent, in order to provide competent behavior in this environment. In many ways this research has been made easier by the generalized support Soar provides for intelligent behavior, as discussed under the capabilities already described. However, Soar also imposes a number of constraints on development — stemming from its being a theoretically motivated attempt at constructing a coherent architecture for intelligent behavior rather than simply a grabbag of tools — that don't always fit precisely with the demands of this environment. These mismatches either reflect ways in which the Soar architecture is inadequate, or they reflect constraints within which people do, and synthetic systems should, find ways of existing. Which of these is the case is an open question, though an actively investigated one.

The first of the remaining requirements — simultaneous performance of multiple high-level tasks — has turned out to be the most problematic of the set. In particular, as limited by the Soar architecture, TacAir-Soar cannot directly construct multiple independent goal hierarchies in service of multiple high-level tasks. For instance, consider a situation where LL needs to interpret an opponent's maneuver while it is simultaneously attempting to reach its own missile firing position. In this case, LL constructs a goal hierarchy for achieving missile firing position, as shown in Figure 10. The problem is that, though a second goal hierarchy is also required here — to interpret the opponent's maneuver — LL cannot independently construct it because of the constraint imposed by Soar that there be only a single active goal hierarchy. The approach to this problem that is currently implemented in TacAir-Soar is to commit to just one of the goal hierarchies, and to install operators for the remaining high-level goals, as needed, at the bottom of this goal hierarchy. Thus, in the above example, LL commits to a hierarchy for attaining a missile firing position, but then installs an *interpret-maneuver* operator as a subgoal of the the *just-turn* operator. Although, with sufficient care, this approach can be made to work, it also has a significant problem: any attempt to update the upper goal hierarchy will eliminate the lower one, because Soar automatically interprets "lower" as meaning "dependent on", even though the hierarchies really are independent here. We generally attempt to minimize this problem by ensuring that the lower hierarchy is present for only a brief time — in particular, it is usually limited to the application of a single operator within a single decision — however, a more general solution is clearly required. Alternative solutions that are currently under investigation include: changing the architecture to allow the explicit use of a forest of goal hierarchies [Jones et al. 94]; merging operators across goal hierarchies so that a single hierarchy can represent the state of all of the hierarchies simultaneously [Covrigaru 92]; and flushing the current goal hierarchy whenever a new one needs to be established, while depending on learning to compile goal hierarchies into single operators, and thus to reduce (or eliminate) the overhead of constantly flushing and rebuilding large goal hierarchies [Rubinoff and Lehman 94].

The next requirement — agent coordination — is an extremely difficult problem in general [Durfee and Montgomery 90]; however, the military simplifies it as much as possible by relying on shared knowledge of common doctrine and tactics, and by providing detailed mission specifications that outline ahead of time the roles of the coordinating pilots. Most of this knowledge is already encoded in TacAir-Soar agents so that they can correctly perform the missions, although some additional knowledge is required about the capabilities of other agents. However, there are still many circumstances in which dynamic coordination is necessary. For example, a more realistic version of the scenario described in Section 2 would include an E2C aircraft (which has a radar that can see in a 250 mile radius) aiding LL and LW, and a ground controller (also with a large radar) aiding DL and DW. These additional agents support the fighters by providing them with sightings of planes and sometimes directives for their missions. In other situations, a controller may dynamically change a plane's mission. Participation in these types of scenarios requires the TacAir-Soar agents to communicate while simultaneously maneuvering their aircraft (a capability that is presently achieved via the technique outlined in the previous paragraph). It also requires the agents to be *taskable*, that is, to be able to dynamically change their missions (and goals) in response to interactions with other agents. Current TacAir-Soar agents do participate in these types of scenarios, both as pilots and controllers [Laird, et al. 94].

The primary communication requirement in the scenario arises in the context of inter-agent coordination. This form of communication is currently based on a fixed set of message templates that are derived from the actual messages spoken by pilots (there are over seventy different types of messages that TacAir-Soar agents can send and receive). While this has proven to be adequate in the scenarios encountered so far, its lack of flexibility will likely prove problematic in the future. To provide for flexible natural-language communication, a research effort is underway to integrate NL-Soar [Lehman et al. 91], an independently developed, Soar-based, real-time natural language system into TacAir-Soar [Rubinoff and Lehman 94].

A secondary communication requirement in the scenario arises during post-mission debriefing [Johnson 94a, Johnson 94b]. Here, TacAir-Soar must accept questions and generate responses. Questions are currently input via a dynamically generated set of menus. Answers are generated via a combination of text (via a simple Functional-Unification-Grammar-based language generator implemented in Soar) and graphical depictions of aircraft configurations. The communication techniques used here, while more general than those presently used for inter-agent coordination, still need to be generalized. These two forms of communication capabilities also need to be merged into a single multi-functional capability.

The next three requirements from the list — agent modeling, planning, and temporal reasoning — are also at the top of our research agenda. For each of these capabilities, the real-time, dynamic, limited-information, multi-agent nature of this environment, poses the important research challenge. For instance, much of the existing work in agent modeling, as well as in the closely related area of plan-recognition, has focused on static, single-agent environments (except perhaps for a recent effort [Rao and Murray 94] also in the arena of air-combat simulation). The techniques that have emerged from this work are not directly applicable in our environment, given the agents' complex mix of goal-driven and reactive behavior. The key idea in the solution that we are investigating is that the mechanisms that a TacAir-Soar agent employs in generating its own flexible and reactive behaviors may be used to track other agents' flexible and reactive behaviors [Tambe and Rosenbloom 94]. Preliminary solutions from these investigations are being incorporated into TacAir-Soar.

The next requirement in the list is that of maintaining episodic memory. Episodic memory in TacAir-Soar is used primarily to support explanation — an agent cannot explain its decisions and actions if it cannot remember them. It is also employed to a limited extent during agent modeling, so that an

opponent's current actions are interpreted based upon its actions in the past. Episodic memory in general raises hard learning issues for Soar: since episodic memory is a basic characteristic of human cognition, a unified theory of cognition ought to provide a general mechanism for it. The episodic memory model in TacAir-Soar was designed more specifically to meet specific requirements for agent debriefing as well as general constraints imposed by the domain and the architecture. The general constraints are that episodic memory should add minimal processing overhead during mission execution, and for efficiency reasons should not substantially increase the size of Soar's working memory. Minimizing processing overhead means that episodic memory construction should not involve the explicit application of operators that can compete with and interfere with the other tasks that the agent is performing. The need to minimize working memory usage implies that simple solutions such as logging all events and state changes are not adequate. In order to provide a high-quality debriefing capability, it is important that the agent have an accurate memory of events and their sequence, perhaps more accurate than human episodic memory, which is prone to errors both in reconstructing events and in determining their proper order. The current implementation of episodic memory is a compromise among these requirements. The sequencing of events (i.e., significant operator applications), is recorded in working memory, so that it can be recalled accurately. The states in which those events occurred, on the other hand, are stored by committing state changes to long term memory. The long term memory mechanism employs chunking, using a variant of the data chunking mechanism used in other Soar systems to remember declarative information [Rosenbloom, Laird, & Newell 87].

The final requirement — the ability to explain ones own behavior — is both constrained and facilitated by the Soar architecture. The constraint arises from the inability of Soar-based agents to directly examine their own rules. This constraint is motivated by, among other things, comparable human limitations on self-reflection. It implies that understanding of ones own behavior must come from observation of, learning about, and possibly experimentation with, that behavior, rather than by direct examination of the "code" that generated it. The facilitation arises from the meta-level support Soar provides for exploring its own behavior. Such support enables Soar to reflect on both the actual decisions that were made and to reason hypothetically about alternative situations and decisions, in order to identify which situational factors are most significant for decision making. In contrast to the time-limited nature of the behavior required during missions, where extended meta-level activity can be problematic, such behavior is just what is needed during a debriefing session. Furthermore, the results of meta-level reasoning are summarized in chunks, which facilitate subsequent explanation. Currently, TacAir-Soar's explanation capability allows it to answer questions regarding the actions it performed, and the conclusions it came to, during task performance. In the future, additional capabilities may enable TacAir-Soar to obtain feedback from expert pilots during such an explanation session, providing TacAir-Soar an opportunity to improve its performance by learning from this feedback.

### **3.1. Current Status and Future Needs**

TacAir-Soar is at present capable of participating in simulated tactical air-combat scenarios involving fighters who are either on their own or flying in coordinated pairs (i.e., as a section). They may additionally be supported by ground or air controllers who can use their much larger radar coverage to provide information to the fighters on enemy fighters that are outside of their own radar range. These controllers can be realized either by humans, via a special graphical interface to ModSAF, or by specifying a ground or air-control mission to a TacAir-Soar agent.

By taking advantage of the networked environment (see Figure 1), we have been able to run scenarios involving up to eight fighters and two controllers. TacAir-Soar's main area of expertise is Beyond-

Visual-Range combat, where pilots only have radar and communication information about enemy aircraft, as opposed to Within-Visual-Range combat, where pilots can also directly see the enemy aircraft. TacAir-Soar pilots can fly three different types of aircraft, three different types of "aircraft vs aircraft" combat missions, and a substantial variety of maneuvers and tactics. TacAir-Soar includes approximately 200 operators spread over 24 problem spaces and approximately 1700 rules. It has consumed about two years worth of our team effort to reach this point (including a large ramp-up time for relevant knowledge acquisition and infrastructure development).

More recently, we have been extending TacAir-Soar to fly types of aircraft and perform types of missions that are quite different from those used in "aircraft versus aircraft" scenarios. With respect to aircraft, the key new type is an *attack helicopter* [Gunston 86], which acts much like a hybrid between a fixed-wing aircraft and a ground vehicle (such as a tank). Its missions, tactics and weapons, differ in a number of significant ways from fixed-wing air-to-air combat, most importantly in the way it needs to reason about, plan with respect to, navigate within, and exploit terrain features such as hills and trees. With respect to fixed-wing aircraft missions, the key new one is *close air support*, which is an "aircraft versus ground targets" mission conducted in support of ground troops. A key additional capability that TacAir-Soar pilots require for close air support is the ability to accept new missions dynamically and plan the details of its attack. This involves determining factors such as the timing, ordnance, altitude, and method of attack. This planning is based on information regarding the weather, terrain, target size, enemy anti-aircraft capabilities, etc. Thus, the behavior of the system may be very different if one of these parameters is changed (e.g., if there are medium-level clouds instead of it being a clear day). This planning capability is a prototype for some of the types of planning that will be required for other missions.

During the next three years, we will be expanding the scope (breadth) and competence (depth) of TacAir-Soar in preparation for the STOW-97 exercise. In terms of scope, TacAir-Soar needs to cover a total of sixteen different mission types, including suppression of electronic defenses, strategic attack, interdiction, surface attack, troop transport, reconnaissance, and re-fueling. Some of these are natural extensions of its air-combat and close air support missions, while others are quite different. Providing competence in these missions is not just a matter of developing each mission for an individual plane to fly. Many of these missions, such as close air support, require a section or division (two sections) of planes to check in with various controllers who can dynamically modify various aspects of the mission, such as the general route the planes are to take, or the specific target of the attack. Moreover, there are integrated missions, such as an offensive strike, that involve up to twenty to thirty aircraft, each with its own mission, but coordinating its behavior with other aircraft. In support of the various missions to be performed, TacAir-Soar will need to know how to fly up to as many as fifty different aircraft types. While many of these aircraft are quite similar to the aircraft it can already fly — at least at the level they are modeled for our simulations — some, such as the helicopters, are quite different. We anticipate that these new types of missions and aircraft will drive much of our future research and development effort in this area.

While the requirements with regard to scope are thus at least somewhat clear, the requirements with regard to competence are not as clear. In particular, the goal of the TacAir-Soar system is not to win a simulated air-combat via any and all means (including simulator quirks). Rather, the goal involves remaining faithful to the specified constraints, including the mission parameters and continued display of human-like behavior. So, in some scenarios, TacAir-Soar may fail at performing its mission in a tactical sense, but still be completely successful in terms of the requirements for intelligent automated forces. Our plan is therefore to use constant feedback from, and comparisons with, expert human pilots to drive both the development and evaluation of TacAir-Soar's competence (complemented with more abstract



evaluations of individual system components, where feasible). Generally this involves having experts watch simulated engagements involving TacAir-Soar and/or human pilots, and then asking the experts to analyze and comment on the behavior of the participants. So far this has led to enthusiastic comments about TacAir-Soar's behavior, along with some specific suggestions about how its behavior is inappropriate. However, there is still a long way to go.

#### 4. Related Work

There are at least four classes of systems to which TacAir-Soar needs to be compared. The first class includes systems that provide intelligent forces for battlefield simulations in general, and air-combat simulations in particular. Semi-automated forces (introduced in Section 1) tend to be developed as combinations of finite state machines (FSMs) and arbitrary code [Calder et al. 93]. This strategy has proven adequate for encoding fragments of well-defined behavior, but has so far fallen considerably short of autonomously competent behavior (thus requiring frequent intervention by knowledgeable humans). FSM languages are simply too restrictive to support the representation of human-like intelligence. The ability to use arbitrary code does provide significant improvements in flexibility; however, this flexibility of *language* unfortunately does not by itself translate into flexibility of *behavior* by the system that is programmed within the language. What is missing is the higher-level support necessary to enable such behavioral flexibility, such as the support for dynamic integration of knowledge during behavior, and goal-driven problem solving and planning. Most of the systems that go beyond the simple FSM strategy are structured as simple rule-based expert systems for tactical decision making [Zytkow and Erickson 87, Kornell 87]. Given TacAir-Soar's own knowledge-intensive (rule-based) approach to decision making, it does share some similarities with such systems. However, while expert systems have some of the strengths required for realistic simulation, they are weak in other areas. For example, in a standard rule-based approach it is difficult to capture the complexity of the multiple dynamic goals about which pilots must reason. In addition, these systems typically rely only on high-level tactical knowledge, and as a result prove to be rather rigid — unless they can be preprogrammed for every possible contingency, their performance degrades greatly when faced with unexpected situations. Finally, expert systems generally ignore the other cognitive aspects of the task, such as modeling of other agents in the environment, and behaving in a human-like manner. Soar was specifically constructed so as to overcome such limitations of standard rule-based expert systems. Besides Soar, one other system in this class that attempts to provide for flexible behavior is a recent effort at the Australian AI institute [Rao et al. 93], that attempts to use the Procedural Reasoning System (PRS) [Georgeff and Lansky 86] as a basis for modeling pilots in air-combat simulation. However, this effort does not yet appear to be far enough along to evaluate the extent to which it can actually provide such flexibility in this demanding domain.

The second class of systems includes those based on other AI agent architectures that attempt the integration of a variety of component capabilities. This is a much broader category of systems. For example, consider the large number of systems presented at the 1991 and 1994 AAAI spring symposia on agents and agent architectures (and the planned 1995 symposium). Unfortunately, the characteristics of the space of agent designs and architectures is not well enough understood to enable a clear understanding of the relationship between TacAir-Soar and these other systems. About the best we can do at this point is to compare the number and types of capabilities exhibited by these systems. Based on such an evaluation, TacAir-Soar appears to be one of a very few systems that integrate together so many distinct capabilities (although clearly not all of its capabilities exist in as general a form as they ultimately ought to). Another strong effort in the arena of highly integrated agents would appear to be Vere and Bickmore's *Basic Agent*, which combines limited natural language understanding and generation, planning, temporal reasoning, plan execution, episodic memory, simulated symbolic perception and some

general world knowledge [Vere and Bickmore 90].

The third class is other systems implemented within the Soar architecture. Though there are many such systems — descriptions of a number of them can be found in [Rosenbloom, Laird, & Newell 93] — there are two important characteristics of TacAir-Soar that set it apart. First, TacAir-Soar is to actually be fielded and used in a real-world environment. This simulation environment is real world in the sense that it has been developed by outside (commercial and government) groups, it is of significant direct value to these groups, and they have a direct need for participation in it by intelligent agents. Second, TacAir-Soar requires the integration of a more diverse set of capabilities than any previous Soar-based system. Most of these capabilities have already been exhibited by individual Soar-based systems — knowledge-intensive and goal-driven behavior [Prietula et al. 91, Rosenbloom et.al. 85], reactivity and real-time performance [Laird and Rosenbloom 90, Pearson et al. 93], planning [Rosenbloom, Lee, & Unruh 93, Stobie et al. 92], agent modeling [Hill and Johnson 94, Ward 91] and natural language communication [Lehman et al. 91] — and these earlier systems have provided the strong foundations necessary for the work on TacAir-Soar. However, most of these systems, in reflecting the research interests of their developers, have targeted individual capabilities in isolation, or a small number of capabilities in combination. There have been initial promising steps towards integration of multiple capabilities, e.g., [Polk et al. 89, Laird and Rosenbloom 90, Huffman 94, Nelson, Lehman & John 94], but these have been exceptions. We have thus realized for some time now, as have some others, a need for Soar systems that provide a tighter integration of a wider range of capabilities. For instance, Hayes-Roth in her comments on the work on integration in Soar states that:

Let me presume to offer a friendly challenge to the Soar team. Give us "Agent-Soar", integrating all of the "content" of the existing Soar applications in a graceful, principled fashion within a single instance of the Soar architecture. Show us the resulting economies of storage and Agent-Soar's global perspective on its knowledge. Show us how Agent-Soar can exploit all of its content to meet several of the component requirements for intelligence simultaneously... [Hayes-Roth 93]

TacAir-Soar takes some concrete steps to address this friendly challenge. As outlined in the previous section, TacAir-Soar already integrates a diverse set of capabilities and it promises to integrate even more, including natural language and learning.

The fourth class is systems that specialize in TacAir-Soar's individual capabilities, such as planning or agent modeling. These capabilities represent core research issues in AI, which our project must address, although within the confines of creating integrated pilot agents; and we necessarily look to these specialized systems to aid us in in this endeavor. However, there are just too many of them to even begin trying to list them here.

## 5. Conclusion

Interactive simulation environments provide rich virtual worlds for research in building integrated intelligent agents that can interact with humans, each other, and their environment. Our current project is based on one such virtual world: battlefield simulations. We are building automated pilot agents that can participate in simulated combat, both with and against humans. The challenge this task poses is the creation of intelligent pilot agents that integrate a broad range of capabilities in a human-like manner.

We are building a system called TacAir-Soar to address this challenge. Given the support it obtains from the Soar architecture, TacAir-Soar is easily able to exhibit a number of the required capabilities — to the extent that it is able to be the first AI system to directly participate in an operational military exercise. However, there are other capabilities for which the Soar architecture provides only a framework for their development, or indeed proves somewhat of a hindrance to their development. These individual

capabilities, and their integration in TacAir-Soar, have presented an interesting set of issues for both present and future research. They will also likely contribute to further evolution of the Soar architecture.

A second important set of issues for future work concerns the evaluation of the agents. Clearly, one important and large-scale test for the agents is the simulated theater of war (STOW) demonstration in 1997. Between ten and fifty thousand automated agents, and up to one thousand humans, are to participate in this exercise. This will provide a gross evaluation of the viability, realism and usefulness of our agents. However, additional forms of evaluation are also required. These include a detailed analysis of agent performance in terms of the individual capabilities and components and the interactions resulting from their integration, as well as detailed comparisons of the behavior of the agents to humans, and debriefings of agents by human subject matter experts. Lessons learned from such evaluations could help extend this intelligent agent technology beyond battlefield simulations, into other collaborative and competitive virtual worlds for education, entertainment, training, and manufacturing.

### **Acknowledgment**

We gratefully acknowledge the contribution of the other members of the team involved in the creation of the TacAir-Soar system, particularly, Jill Lehman, Paul Nielsen, and Robert Rubinoff.

This research was supported under contract N00014-92-K-2015 from the Advanced Systems Technology Office (ASTO) of the Advanced Research Projects Agency (ARPA) and the Naval Research Laboratory (NRL) to the University of Michigan (and a subcontract to the University of Southern California Information Sciences Institute from the University of Michigan). This project would not have been possible without the visionary support of Cmdr. Dennis McBride of ARPA/ASTO; and the technical support by Ed Harvey, Tom Brandt, Bob Richards, and Craig Petersen of BMH Inc.; Andy Ceranowicz and Joshua Smith of Loral Inc.; and David Keirse of Hughes Aircraft Co.

## References

- [Agre and Chapman 87] Agre, P. E., and Chapman, D.  
Pengi: An implementation of a theory of activity.  
In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, Calif.: The AAAI Press, August, 1987.
- [Anderson, et al. 90] Anderson, J. R., Boyle, C. F., Corbett, A. T., and Lewis, M. W.  
Cognitive modeling and intelligent tutoring.  
*Artificial Intelligence* 42:7-49, 1990.
- [Bates et al. 92] Bates, J., Loyall, A. B., and Reilly, W. S.  
*Integrating reactivity, goals and emotions in a broad agent*.  
Technical Report CMU-CS-92-142, School of Computer Science, Carnegie Mellon University, May, 1992.
- [Bratman, et al. 88] Bratman, M. E., Israel, D. J., and Pollack, M. E.  
Plans and resource-bounded practical reasoning.  
*Computational Intelligence* 4(4):349-355, 1988.
- [Brooks 91] Brooks, R. A.  
Intelligence without representation.  
*Artificial Intelligence* 47:139-160, 1991.
- [Calder et al. 93] Calder, R. B., Smith, J. E., Courtemanche, A. J., Mar, J. M. F., Ceranowicz, A. Z.  
ModSAF behavior simulation and control.  
In *Proceedings of the Third Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, March, 1993.
- [Covrigaru 92] Covrigaru, A.  
*Emergence of meta-level control in multi-tasking autonomous agents*.  
PhD thesis, University of Michigan, 1992.
- [Cremer et al. 94] Cremer, J., Kearney, J., Papelis, Y., and Romano, R.  
The software architecture for scenario control in the Iowa driving simulator.  
In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Crooks et al. 92] Crooks, W. H., Fraser, R. E., Jacobs, R. S., McDonough, J. G.  
*Functional description: WISSARD aircraft simulators*.  
Technical Report TR-11100-15600-08-92, Illusion Engineering Inc, Westlake Village, CA, August, 1992.
- [Durfee and Montgomery 90] Durfee, E., and Montgomery, T. A.  
A hierarchical protocol for coordinating multi-agent behavior.  
In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, Calif.: The AAAI Press, August, 1990.
- [Etzioni 93] Etzioni, O.  
Intelligence without robots: A reply to Brooks.  
*AI Magazine* 14:7-13, 1993.

- [Georgeff and Lansky 86] Georgeff, M. P. and Lansky, A. L.  
Procedural knowledge.  
*Proceedings of the IEEE special issue on knowledge representation* 74:1383-1398, 1986.
- [Gunston 86] Gunston, B.  
*AH-64 Apache*.  
Osprey, London, 1986.
- [Hanks et al. 93] Hanks, S., Pollack, M. E., and Cohen, P. R.  
Benchmarks, test beds, controlled experimentation, and the design of agent architectures.  
*AI Magazine* 14(4):17-42, 1993.
- [Hayes-Roth 93] Hayes-Roth, B.  
On building integrated cognitive agents: a review of Allen Newell's Unified Theories of Cognition.  
*Artificial Intelligence* 59:329-341, 1993.
- [Hill and Johnson 94] Hill, R., and Johnson, W. L.  
Situated plan attribution for intelligent tutoring.  
In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, Calif.: The AAAI Press, 1994.
- [Huffman 94] Huffman, Scott B.  
*Instructable Autonomous Agents*.  
PhD thesis, The University of Michigan, Department of Computer Science and Engineering, January, 1994.
- [IEEE 93] The IEEE standards board.  
*IEEE Standard for information technology -- protocols for Distributed Interactive Simulation Applications*.  
Technical Report IEEE-Std-1278-1993, Institute for Electrical and Electronic Engineers, May, 1993.
- [Johnson 94a] Johnson, W. L.  
Agents that explain their own actions.  
In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*, pages 87-95. Orlando, FL, 1994.
- [Johnson 94b] Johnson, W. L.  
Agents that learn to explain themselves.  
In *Proceedings of the National Conference on Artificial Intelligence*. American Association for Artificial Intelligence, Seattle, WA, August, 1994.
- [Johnson et al. 94] Johnson, W. L., Jones, R. M., Keirse, D., Koss, F. V., Laird, J. E., Lehman, J. F., Nielsen, P. E., Rosenbloom, P. S., Rubinoff, R., Schwamb, K., Tambe, M., van Lent, M., and Wray, R.  
*Collected Papers of the Soar/IFOR project, Spring 1994*.  
Technical Report CSE-TR-207-94, University of Michigan, Department of Electrical Engineering and Computer Science, 1994.  
Also available as Technical Reports ISI/SR-94-379 from the University of Southern California Information Sciences Institute and CMU-CS-94-134 from Carnegie Mellon University; WWW access <http://ai.eecs.umich.edu/ifor/papers/index.html>.

- [Jones 94] Jones, R. M.  
Dynamic generation of complex behavior.  
In *Proceedings of the National Conference on Artificial Intelligence*, pages 1504.  
Menlo Park, Calif.: The AAAI Press, 1994.  
Video presented at AAAI-94.
- [Jones et al. 93] Jones, R. M., Tambe, M., Laird, J. E., and Rosenbloom, P.  
Intelligent automated agents for flight training simulators.  
In *Proceedings of the Third Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, March, 1993.
- [Jones et al. 94] Jones, R., Laird, J. E., Tambe, M., and Rosenbloom, P. S.  
Generating behavior in response to interacting goals.  
In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Kautz and Allen 86] Kautz, A., and Allen J. F.  
Generalized plan recognition.  
In *Proceedings of the National Conference on Artificial Intelligence*, pages 32-37.  
Menlo Park, Calif.: The AAAI press, 1986.
- [Kornell 87] Kornell, J.  
Reflections on using knowledge based systems for military simulation.  
*Simulation* 48:144-148, 1987.
- [Laird and Rosenbloom 90] Laird, J.E. and Rosenbloom, P.S.  
Integrating execution, planning, and learning in Soar for external environments.  
In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, Calif.: The AAAI press, July, 1990.
- [Laird et al. 93] Laird, J. E., Congdon, C. B., Altmann, E., and Doorenbos, R.  
Soar User's Manual, version 6.  
Available from [soar-requests@cs.cmu.edu](mailto:soar-requests@cs.cmu.edu).  
1993
- [Laird, et al. 94] Laird, J. E., Jones, R. M., and Nielsen, P. E.  
Coordinated behavior of computer generated forces in TacAir-Soar.  
In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Laird, Newell, and Rosenbloom 87] Laird, J. E., Newell, A. and Rosenbloom, P. S.  
Soar: An architecture for general intelligence.  
*Artificial Intelligence* 33(1):1-64, 1987.
- [Lehman et al. 91] Lehman, J. F., Lewis, R. L., and Newell, A.  
*Natural language comprehension in Soar: Spring 1991*.  
Technical Report CMU-CS-91-117, School of Computer Science, Carnegie Mellon University, 1991.

- [Loper et al. 94] The DIS steering committee.  
*The DIS vision: A map to the future of distributed simulation.*  
 Technical Report IST-SP-94-01, Institute for simulation and training, University of Central Florida, May, 1994.
- [Maes et al. 94] Maes, P., Darrell, T., Blumberg, B., and Pentland, S.  
 Interacting with Animated Autonomous Agents.  
 In J. Bates (editor), *Proceedings of the AAAI Spring Symposium on Believable Agents.*  
 1994.
- [Mitchell et al. 86] Mitchell, T. M., Keller R. M., and Kedar-Cabelli, S. T.  
 Explanation-based generalization: A unifying view.  
*Machine Learning* 1(1):47-80, 1986.
- [Moravec 90] Moravec, H.  
*Mind Children.*  
 Harvard University Press, Cambridge, Massachusetts, 1990.
- [Nelson, Lehman & John 94] Nelson, G., Lehman, J. F. & John, B. E.  
 Integrating cognitive capabilities in a real-time task.  
 In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society.*  
 1994.
- [Newell 90] Newell, A.  
*Unified Theories of Cognition.*  
 Harvard University Press, Cambridge, Massachusetts, 1990.
- [Newell 92a] Newell, A.  
 Unified Theories of Cognition and the role of Soar.  
 In Michon, J. and Akyurek, A. (editor), *Soar: A Cognitive Architecture in Perspective.*  
 Kluwer Academic, Cambridge, Mass., 1992.
- [Newell 92b] Newell, A.  
 Precis of *Unified Theories of Cognition.*  
*Behavioral and Brain Sciences* 15:425-492, 1992.
- [Newell et al 91] Newell, A., Yost, G. R., Laird, J. E., Rosenbloom, P. S., & Altmann, E.  
 Formulating the problem space computational model.  
 In R. F. Rashid (editor), *CMU Computer Science: A 25th Anniversary Commemorative*, pages 255-293. ACM Press/Addison-Wesley, New York, 1991.
- [Pearson et al. 93] Pearson, D.J., Huffman, S.B., Willis, M.B., Laird, J.E., and Jones, R.M.  
 A Symbolic Solution to Intelligent Real-Time Control.  
*IEEE Robotics and Autonomous Systems* 11:279-291, 1993.
- [Polk et al. 89] Polk, T.A., Newell, A., and Lewis, R.L.  
 Toward a unified theory of immediate reasoning in Soar.  
 In *Proceedings of the Annual Conference of the Cognitive Science Society*, pages  
 506-513. 1989.
- [Prietula et al. 91] Prietula, M., Hsu, W-L., Steier, D., and Newell, A.  
 Applying an architecture for general intelligence to scheduling.  
*ORSA Journal on Computing* 5(3), 1991.

- [Rao and Murray 94]  
 Rao, A. S., and Murray, G.  
 Multi-agent mental-State recognition and its application to air-Combat modelling.  
 In *Proceedings of the Workshop on Distributed Artificial Intelligence (DAI-94)*. 1994.
- [Rao et al. 93] Rao, A. S., Lucas, A., Morley, D., Selvestrel, M., Murray, G.  
*Agent-oriented architecture for air-combat simulation*.  
 Technical Report Technical Note 42, The Australian Artificial Intelligence Institute,  
 1993.
- [Rosenbloom and Laird 86]  
 Rosenbloom, P. S. and Laird, J. E.  
 Mapping explanation-based generalization onto Soar.  
 In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages  
 561-567. 1986.
- [Rosenbloom and Laird 94]  
 Rosenbloom, P. S. and Laird, J. E.  
 On *Unified Theories of Cognition*: A response to the reviews.  
 In W. J. Clancey, S. W. Smoliar, & M. J. Stefik (editors), *Contemplating Minds: A  
 Forum for Artificial Intelligence*, pages 141-165. MIT Press, Cambridge, MA,  
 1994.  
 (Reprint of Rosenbloom & Laird, 1993, in *Artificial Intelligence*, vol. 59, pp. 389-413).
- [Rosenbloom et.al. 85]  
 Rosenbloom, P. S., Laird, J. E., McDermott, J., Newell, A., and Orciuch, E.  
 R1-Soar: An experiment in knowledge-intensive programming in a problem-solving  
 architecture.  
*Pattern Analysis and Machine Intelligence* 7:561-569, 1985.
- [Rosenbloom, et al. 91]  
 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):289-325, 1991.
- [Rosenbloom, Laird, & Newell 87]  
 Rosenbloom, P. S., Laird, J. E., & Newell, A.  
 Knowledge level learning in Soar.  
 In *Proceedings of Sixth National Conference on Artificial Intelligence*, pages 499-504.  
 AAAI, Seattle, 1987.
- [Rosenbloom, Laird, & Newell 88]  
 Rosenbloom, P. S., Laird, J. E., & Newell, A.  
 Meta-levels in Soar.  
 In P. Maes & D. Nardi (editors), *Meta-Level Architectures and Reflection*, pages  
 227-240. North Holland, Amsterdam, 1988.
- [Rosenbloom, Laird, & Newell 93]  
 Rosenbloom, P. S., Laird, J. E., & Newell, A. (editors).  
*The Soar Papers: Research on Integrated Intelligence*.  
 MIT Press, Cambridge, MA, 1993.
- [Rosenbloom, Lee, & Unruh 93]  
 Rosenbloom, P. S., Lee, S., & Unruh, A.  
 Bias in planning and explanation-based learning.  
 In S. Minton (editor), *Machine Learning Methods for Planning*, pages 197-232.  
 Morgan Kaufmann, San Mateo, CA, 1993.



- [Rubinoff and Lehman 94]  
 Rubinoff, R., and Lehman, J.  
 Natural language processing in an IFOR pilot.  
 In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Schwamb et al. 94]  
 Schwamb, K., Koss, F., and Keirse, D.  
 Working with ModSAF: Interfaces for programs and users.  
 In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Stobie et al. 92] Stobie, I., Tambe, M., and Rosenbloom, P.  
 Flexible integration of path-planning capabilities.  
 In *Proceedings of the SPIE conference on Mobile Robots*, pages (in press).  
 November, 1992.
- [Tambe et al 94] Tambe, M., Jones, R. M., Laird, J. E., & Rosenbloom, P. S.  
 Building believable agents for simulation environments.  
 In J. Bates (editor), *Working Notes of the AAAI Spring Symposium on Believable Agents*, pages 82-85. AAAI, Stanford, CA, 1994.
- [Tambe and Rosenbloom 94]  
 Tambe M., and Rosenbloom, P. S.  
 Event tracking in complex multi-agent environments.  
 In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Thorpe et al. 89] Thorpe, J. A., Shiflett, J. E., Bloedorn, G. W., Hayes, M. F., Miller, D. C.  
*The SIMNET network and protocols*.  
 Technical Report 7102, BBN systems and technologies corporation, July, 1989.
- [van Lent and Wray 94]  
 van Lent, M., and Wray, B.  
 A very low cost system for direct human control of simulated vehicles.  
 In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1994.
- [Vere and Bickmore 90]  
 Vere, S., and Bickmore, T.  
 A basic agent.  
*Computational Intelligence* 6:41-59, 1990.
- [Ward 91]  
 Ward, B.  
*ET-Soar: Toward an ITS for Theory-Based Representations*.  
 PhD thesis, School of Computer Science, Carnegie Mellon University, May, 1991.
- [Webber and Badler 93]  
 Webber, B., and Badler, N.  
 Virtual interactive collaborators for simulation and training.  
 In *Proceedings of the Third Conference on Computer Generated Forces and Behavioral Representation*. Orlando, Florida: Institute for Simulation and Training, University of Central Florida, May, 1993.

[Zytkow and Erickson 87]

Zytkow, J. M., and Erickson, M. D.  
Tactical manager in a simulated environment.  
*Methodologies for intelligent systems.*  
Amsterdam: Elsevier Science, 1987.

## Table of Contents

<b>1. Introduction</b>	<b>2</b>
<b>2. The Air-Combat Simulation Domain</b>	<b>5</b>
<b>3. Creating Intelligent Automated Pilots</b>	<b>15</b>
<b>3.1. Current Status and Future Needs</b>	<b>22</b>
<b>4. Related Work</b>	<b>24</b>
<b>5. Conclusion</b>	<b>25</b>
<b>Acknowledgment</b>	<b>26</b>

### List of Figures

<b>Figure 1: Human and automated pilots interact with the DIS environment via distributed simulators.</b>	<b>5</b>
<b>Figure 2: A snapshot of ModSAF's simulation of an air-combat situation.</b>	<b>7</b>
<b>Figure 3: A portion of an automated pilot's forest of goal hierarchies.</b>	<b>8</b>
<b>Figure 4: Snapshot 1: The initial positions of the aircraft in the air-combat simulation scenario.</b>	<b>8</b>
<b>Figure 5: Snapshot 2: LL attempts to achieve lateral separation.</b>	<b>10</b>
<b>Figure 6: Snapshot 3: LL fires a missile at DL.</b>	<b>11</b>
<b>Figure 7: Snapshot 4: DL and DW begin a Post-hole maneuver.</b>	<b>11</b>
<b>Figure 8: Snapshot 5: DL's turn causes it to lose radar contact with all other aircraft.</b>	<b>13</b>
<b>Figure 9: Snapshot 6: All four fighter aircraft have re-positioned themselves for a new attack.</b>	<b>13</b>
<b>Figure 10: LL's problem space hierarchy. Boxes indicate problem spaces, with names in bold to their right. Names within boxes indicate operators available within problem spaces. Italicized names indicate currently selected operators.</b>	<b>17</b>