Special Operations Forces IFORs

Glenn Taylor
Frank Koss
Paul Nielsen
Soar Technology, Inc.
317 N First Street
Ann Arbor, MI 48103
(734) 327-8000
{glenn, koss, nielsen}@soartech.com

Keywords:
Computer generated forces, dismounted infantry, artificial intelligence, Soar

ABSTRACT: *Intelligent Forces (IFORs) have advanced greatly over the past several years, especially in the tactical air domain. Recently, Soar Technology, Inc., applied our experience with IFORs to the domain of dismounted infantry. As a proof of concept, we developed an Army Ranger special operations force (SOF) team that performs a long-range reconnaissance mission. The team consists of dismounted infantry (DI) in the Joint Semi-automated Forces (JSAF) simulator. Each of the six team members is represented as a separate agent developed in the Soar cognitive architecture. While preliminary, this work shows the potential value of IFORs in the ground domain and serves as a basis for future development. This work will be included in the Joint Forces Command's Unified Vision 2001 (UV01) exercise.*

1. Overview

Much headway has been made in using Intelligent Forces (IFORs, distinguished as being fully autonomous successors to Semi-Automated Forces, or SAFs) to model human behavior in military simulation, particularly in the tactical air domain. Leveraging our prior experience in developing IFORs, we have developed SOF-Soar, the first application of Soar to modeling ground forces in military simulation. SOF-Soar is a prototype model of an Army Ranger Special Operations Forces (SOF) team performing long-range reconnaissance missions. This team consists of Dismounted Infantry (DI) entities in the JSAF simulator. SOF-Soar will be included in the Joint Forces Command Unified Vision 2001 (UV01) exercise.

2. Prior Work

SOF-Soar is third in a line of successful models of human behavior for military simulation built on the framework of the Soar cognitive architecture [1]. In 1992, under the DARPA STOW project, TacAir-Soar was developed in ModSAF as a broad model of fixed-wing aircraft pilot behaviors [2]. Later, these ideas were adopted to create RWA-Soar, its rotary-wing equivalent [3].

TacAir-Soar played important roles in exercises and demonstrations such as Roadrunner'98 and COYOTE'98 [4]. More recently, TacAir-Soar was indispensable in the Joint Forces Command's Joint Experiment '99 and Attack Operations '00, as well as many of the Navy's Fleet Battle Experiments. Additionally, TacAir-Soar was fielded as part of the Battle Force Tactical Trainer (BFTT) delivered to the Navy. TacAir-Soar's inclusion in these projects has demonstrated one of the advantages Soar-based systems have shown over conventional SAFs—autonomy. A single operator can control hundreds of Soar agents, with intervention required only when the operator wants to change their mission details [2].

Prior work has been done in developing behaviors for Dismounted Infantry. JSAF and ModSAF include Task Frame behaviors [5] for DIs, including limited direct action behaviors such as dropping C4 charges on vehicles. More extensively, DI-SAF [6] represents a vast overhaul of the ModSAF DI models and Task Frame-based behaviors. SOF-Soar, while using the underlying physical DI models in JSAF, represents a very different approach to human behavior modeling.

3. The Architecture

SOF-Soar uses the same overall architecture as TacAir-Soar, which has been discussed in detail previously [7]. We will provide here a brief overview of the architecture and its components for context. Figure 3.1 illustrates the interfaces between JSAF, Soar, and the SOF-Soar behaviors. The components are described individually below.

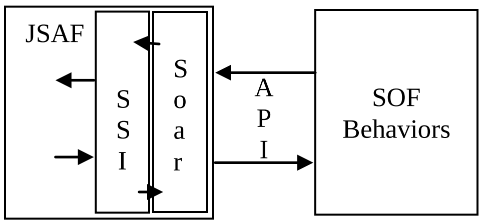

Figure 3.1: The SOF-Soar Architecture

3.1 JSAF

JSAF as the underlying simulation provides terrain information, information about other entities, sensor and weapons models, as well as the physical models controlled by SOF-Soar agents. The choice of JSAF for the SOF-Soar work was partly practical in nature—Soar was already integrated into JSAF, and a robust interface exists between them.

3.2 Soar

Soar is both a theory of human cognition and an embodiment of that theory in a programming architecture. Developed in 1982 at Carnegie Mellon University by Allan Newell and his students John Laird and Paul Rosenbloom, Soar has served worldwide as the basis of research in cognitive science, psychology and artificial intelligence, as well as the reasoning engine for some commercial applications. For more details Soar's history and architecture, see [1].

Soar is a rule-based system that frames decision-making as the selection and application of operators to achieve goals. Soar represents long-term knowledge as production rules in the form of *if…then* statements whose *if* patterns match against a representation of the environment and the agent's own internal state. Actions serve to propose new operators, dynamically decompose more abstract operators, or send motor-control commands to the underlying simulation. Soar supports both goal-directed and reactive behavior, which makes it an ideal choice for implementing agents that must act in complex environments and realistic timeframes.

3.3 Soar-Simulation Interface

The Soar-Simulation Interface (SSI) [7] is the translation layer that exists between the environment (in this case, the JSAF simulator) and the agent. It converts information about the world into a form the agent can reason about, and converts agent actions into observable effects in the environment. While SOF-Soar represented a new domain for Soar development, some of the information needed by the SOF agents was already available via the TacAir-Soar SSI, including waypoint information for navigation and simulated radios for communication. Some additions were required, such as commands for moving along the ground and processing new sensor information.

The SSI currently provides three sensors to the SOF-Soar agents (naked-eye visual, binoculars, and night-vision binoculars) and three forms of explicit communication (radio, satellite communication, wireline, and hand signals).

4. Long Range Reconnaissance Patrol

As a proof of concept, we focused our development effort on implementing a six-man SOF team performing Long Range Reconnaissance Patrol (LRRP) missions, an example of which is illustrated in Figure 4.1.

The following information is briefed to the agents before the mission starts:

- waypoint names and locations
- directed routes between points
- teams (Radio and Observation)
- roles within teams (leads versus subordinates)
- names of team members
- observation criteria
- mission duration

All of this information would be briefed to real Special Forces in the mission rehearsal stage prior to insertion behind enemy lines.

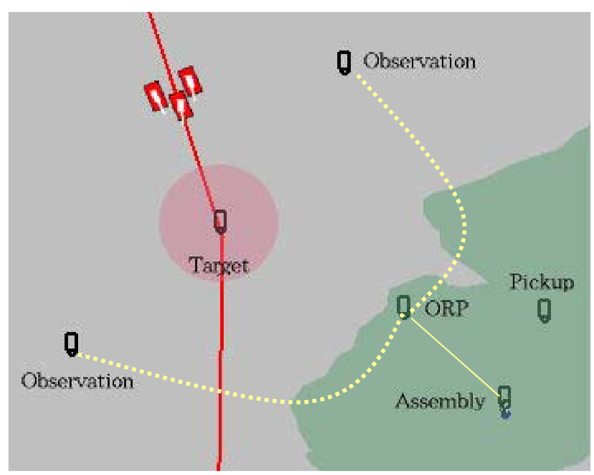

Figure 4.1: A Long Range Reconnaissance Mission

4.1 Patrol

The 6-man SOF team is instantiated near the Assembly point. They proceed to the assigned point and await the arrival of all other members of the team before marching on, via a pre-planned route, to the Objective Rally Point (ORP). The ORP is the last covered and concealed position before the Target location. When the whole team arrives at the ORP, they split up into three two-man teams: one Radio Team and two Observation Teams. The Radio Team stays at the ORP and uses it as a setup position to listen for reports from the lookouts. The Observation Teams move by concealed routes to their respective observation points around the Target area and observe that area for a set amount of time. When the mission duration has expired, the Observation Teams rendezvous with the Radio Team at the ORP, and the entire team heads to the pickup zone for extraction.

As noted earlier, some routes the SOF-Soar agents utilize are fully defined in the mission briefing as route segments to be followed precisely; others are denoted only by their endpoints, and are marked as requiring concealment. We have avoided the overhead of path planning in these cases by deferring to the concealed route routines native to JSAF, which take into account terrain line of sight calculations and the locations of suspected enemy positions. Given these criteria, JSAF returns a directed list of points, with the connecting segments marked as concealed or not. The agent then integrates this route into its knowledge base, and proceeds to follow it.

We have defined a subset of the Long Range Reconnaissance Patrol mission, termed Stationary Reconnaissance. The agents are created at the ORP or their observation points and do not move for the duration of the mission. In this case, they serve simply as immobile ground sensors that report what they observe. This helps sidestep issues related to mission length. In reality, Special Forces missions may take place deep behind enemy lines, may involve cross-country treks of 100 miles or more from assembly to the target, and can last for weeks at a time. Large-scale simulation exercises typically do not have the luxury of waiting for this sort of time-realistic infiltration to take place. Other than the lack of movement, the two missions are identical.

4.2 Observation

The two Observation Teams are given pre-briefed criteria that define what entities and activities to report. The teams can have multiple observation profiles.

Once the teams reach their respective observation points, they remain stationary and look for entities already in or passing through the target area defined by a point and a radius. The target profile contains the following information:

- Entity type and subtype of vehicles
- Size of the group
- Direction the entities are heading
- Group force (friendly, enemy, neutral)

When a member of an Observation Team sees an entity that matches a profile, that agent begins building a cluster associated with that profile. Since terrain features may block views of some of the elements, the agent waits a prescribed amount of time in case others appear. If more appear, and match the profile, those are added to the cluster, and the agent waits longer. Once the time from the last observed vehicle exceeds the clustering time, and the total number of elements in the cluster equals or exceeds the group size defined in the profile, the agent sends a report to its local group.

Besides the target profile, other features are used to cluster elements into a single group, including heading and speed variance between elements, and time and space interval between sightings. If the variances or intervals are too great, or the size of the group is too small, the cluster is dropped from consideration, and the agent forgets about it.

Observation reports are broadcast to everyone in the SOF team, and agents that receive these reports build up their own conceptual representations of the target area. The Radio Team must correlate the reports of multiple Observation Teams into a single "mental picture" before reporting to the Forward Operations Base (FOB). These clusters represent the largest group observed; for instance, if one report mentions a group of 3 and another a group of 4, the correlated cluster will denote 4 vehicles observed. Just as Observation Teams wait for more vehicles to cluster, the Radio Team waits a prescribed amount of time for other reports to come in that might elaborate a known cluster. Once that time has expired, and the Radio Team has constructed a consistent mental picture, it will report that observation to the FOB.

4.3 Reporting

The Radio Team sends messages to the FOB in the form of SALUTE reports, a common format used by military and civilian observers alike for describing scenes of activity. SALUTE reports are made up of six components:

S – size of the group or event
A – the activity of the observed group
L – location of the event
U – unit markings
T – time of sighting
E – equipment carried by the group

To avoid the complexity having to implement full natural language generation behaviors, SOF-Soar agents use a simple template mechanism for reporting what they've seen. A message is an English sentence with slots for each of the SALUTE components listed above. For instance, if a group of red T-72s tanks is driving south through the observation area, the agents might report the following:

> *"At 1304 observed 3 t-72s tank moving at 56 degrees 34 minutes 34 seconds by 64 degrees 43 minutes 22 seconds heading 174 at 24 miles per hour with red markings and unknown equipment"*

Due to the current limitations of JSAF, some of the SALUTE components contain minimal or default information. For instance, equipment carried by other entities is not visible to observers, units have no particular markings other than their respective force colors, and units can only be observed as moving or stationary. To fill out these aspects of the report, the agents will always report "unknown" equipment, the force color for markings, and either "moving" or "stopped" for activity.

5. Communication with SOF-Soar Agents

Communication is an important aspect of agent behavior, enabling them to coordinate with other agents to solve problems, particularly in the military simulation domain [8]. As we have described, SOF-Soar agents use simulated radios as their primary form of communication. Each radio is assigned a particular frequency, and all agents that have a radio tuned to that frequency, within the range of that transmitter, can "hear" the broadcast. SOF-Soar agents can direct their message to individual recipients by prefacing the message with the name of the intended recipient. It is up to the receiver to process or disregard the message based on the named recipient.

SOF-Soar agents also communicate non-verbally by *moving*. That is, a subordinate agent will follow his lead when he sees his lead move. However, this is the limit to the resolution of physical articulation afforded by JSAF, so other useful forms of non-verbal communication such as hand signals have to be simulated as radio messages.

So far, we have described only inter-agent communication. However, human operators must also be able to communicate with the agents to convey mission specifications and commands. Two such tools for accomplishing this are the Communications Panel and the SOF Exercise Editor.

5.1 Communications Panel

For the purposes of running exercises, we allow human simulation operators to communicate directly with individual SOF-Soar agents. A tool called the Communications Panel (or Comm-Panel) enables an operator to tell the agents to change their mission parameters during mission execution. The agents receive the commands as normal text messages on their radios. Currently, there are commands to redirect agents to different observation points, a different ORP, and a different Pickup Zone. The agents can also be told to terminate their mission immediately, which causes them all to rendezvous at the ORP, and then proceed to the Pickup Zone. Lastly, Observation Team members can be told to change their observation criteria, or to add new criteria.

5.2 SOF Exercise Editor

For rapid exercise development, we are currently developing a mission specification tool for the SOF-Soar agents, called the Special Operations Forces Exercise Editor (SOFEE). SOFEE is a Java-based graphical tool that lets the user specify all the pre-briefed information the agents expect or require, facilitating the generation of large exercises in short periods of time. As the number of missions performed by the SOF-Soar agents grows, the SOFEE will expand to accommodate those mission types. We also plan on folding in FWA and RWA elements and missions to integrate previous work done in TacAir-Soar and RWA-Soar.

6. Situational Awareness

SOF-Soar entities, like TacAir-Soar and RWA-Soar entities, maintain their own situational awareness. They obtain information about the environment from sensors, communications, and from the pre-briefed mission specification. When they become aware of an entity in the simulation, they can deliberately seek more information by focusing their attention on that specific target. If the entity has gone out of the range of the model's sensors, information about that entity remains in attentional memory for a set amount of time before being dropped; the agent will completely forget about the entity unless the agent has deliberately decided to remember it. (This same forgetting mechanism holds for communication as well.) Similarly, if an agent has lost contact with an entity it is aware of, the agent projects the location of the agent while it is not directly sensed. If the agent has lost contact with the entity for an extended period of time, the agent deliberately forgets about it.

A tool called the Situational Awareness Panel was developed for TacAir-Soar to visualize this information

[9]. It includes information about what the agent senses, as well as internal state information, the current goals the agent is pursuing, and important milestones achieved during the course of the mission. This can be used to display similar information about SOF-Soar agents, to allow human inspection of the agent's reasoning process.

7. Domain-Specific Issues

While we were able to utilize some prior work developed for Soar behaviors in the air domain, there are clearly many issues particular to the ground domain that have until now been ignored.

For example, Dismounted Infantry entities in JSAF experience fatigue at the physical level if they exceed a certain speed, and gradually recover once they have slowed down below a certain threshold or stop altogether. Our SOF-Soar agents are not cognizant of this fact, so are not able to use movement optimally. Likewise, terrain considerations are not nearly as important in the air domain as with ground forces. We have so far relied upon the human operator and JSAF to provide our SOF agents information about route planning. Ideally, SOF-Soar agents would be capable of terrain reasoning at some abstract level. Physical articulation, team movement, and team coordination is potentially much more complex with DI entities than we have had to address in the air domain.

8. Future Work

Our current model of SOF behaviors is only in its infancy, so there are many short and long-term additions planned.

A first priority is expanding the breadth and fidelity of the LRRP missions. The agents' movement model is fairly simple: each subordinate follows the leader in a single-file line across terrain. Adding patrol tactics such as bounding overwatch, listening halts, and hasty ambushes would greatly improve the visual realism of the model. And while the SOF agents can recognize any vehicle native to JSAF, the observation criteria are too quantitative in nature, and even a small variance outside the bounds might result in incorrect clusters being formed when common sense would dictate a more reasonable clustering.

The next step is expanding the set of missions the SOF-Soar agents perform. Of pressing interest is that of Direct Action (DA) missions—those in which enemy forces are engaged in small, precision strikes. This might include demolition of specific targets such as radar dishes or command centers by placing explosives, or by calling in air support and guiding munitions with laser designators. Along these lines, SOF-Soar agents currently lack behaviors for interacting with or reacting to other entities in the simulation. Military Operations in Urban Terrain (MOUT) is of great interest lately in the simulation community. We feel that SOF-Soar can contribute to this effort, given the foundation we have based on the results of this preliminary work.

There has been much domain-independent work done in the context of TacAir-Soar and RWA-Soar that could readily be applied to SOF-Soar. This includes a more realistic perceptual attention model [10], fatigue effects on decision-making [11], an explicit teamwork model [12], and speech-based communication with the agents [8]. Incorporating this work will serve to improve the realism of the SOF agents.

Finally, the DI-SAF project represents an overhaul of ModSAF to support Dismounted Infantry behaviors at a higher level of fidelity than was previously available [6]. Included in the DI-SAF overhaul were features such as DI entities with freely exchangeable equipment, a model of sound so nearby explosions or gunfire could be heard, and more articulate physical models. All of these improvements to the simulation greatly contributed to the efficacy of the DI-SAF project. One might assume, by its lineage from ModSAF, that JSAF would benefit from a similar overhaul to support higher fidelity DI behaviors, independent of the behavior system.

9. Conclusion

Our work with SOF-Soar, however nascent in its development, clearly shows the promise of applying Soar-based IFORs to ground forces. Our agents currently perform Long Range Reconnaissance Patrol missions autonomously, and can accept mission-modification commands from an operator while performing their tasks. SOF-Soar demonstrates the same advantages shown in Soar-based systems such as TacAir-Soar, including goal-directed behavior and reduced operator requirements.

10. Acknowledgements

This work was funded under MSIAC contract SPO 700-99-D-0300. Mr Rae Dehncke at IDA was instrumental in providing programmatic support for modeling ground entities in Soar. Mr Phil Tuttle at AB Tech and Cmdr Art Galpin at SOCOM provided background knowledge about Special Operations Forces.

References

For brevity, we use "CGF&BR" for "Computer Generated Forces and Behavior Representation", in Orlando, FL.

[1] Laird, J., Newell, A., Rosembloom, P.: "Soar: An Architecture for General Intelligence" *Artificial Intelligence 33(1)*, pp. 1-64, 1987.
[2] Jones, R.M, Laird, J.E., & Nielsen, P.E., Coulter, K.J., Kenny, P.K., and Koss, F.V.: "Automated Intelligent Pilots for Combat Flight Simulation" *AI Magazine* 20(1), pp. 27-42, 1999.
[3] Hill, R., Chen, J., Gratch, J., Rosembloom, P., Tambe, M: "Intelligent Agents for the Synthetic Battlefield: A Company of Rotary Wing Aircraft" *Proceedings of Innovative Applications of Artificial Intelligence (IAAI-97)*, 1997.
[4] Nielsen, P.E., Smoot, D., Martinez, R., and Dennison, J.D. "Participation of TacAir-Soar in RoadRunner and Coyote Exercises at AFRL, Mesa, AZ" *Proceedings of 9th CGF&BR*, pp. 173-180, 2000.
[5] Calder, R., Smith, J., Courtemanche, A., Mar, J., Ceranowicz, A.: "ModSAF Behavioral Simulation and Control" in *Proceedings of 3rd CGF&BR*, 1994.
[6] Reece, D., Qurston, D., Kraus, M., Carbia, I.: "Updating ModSAF for Individual Combatants: the DI SAF Program" *Proceedings of the 7th CGF&BR*, 1998.
[7] Schwamb, K., Koss, F., Keirsey, D.: "Working with ModSAF: Interfaces for Programs and Users" *Proceedings of the 4th CGF&BR*, 1994.
[8] Nielsen, P.E., Koss, F., Taylor, G., Jones, R.M.: "Communication with Intelligent Agents" I/ITSEC 2000.
[9] Jones, R.: "Graphical Visualization of Situational Awareness and Mental State for Intelligent Computer Generated Forces" *Proceedings of the 8th CGF&BR*, 1999.
[10] Hill, R.: "Modeling Perceptual Attention in Virtual Humans" in Proceedings of the 8th CGF&BH, 1999.
[11] Jones, R., Neville, K., Laird, J.: "Modeling Pilot Fatigue with a Synthetic Behavior Model" *Proceedings of the 7th CGF&BR*, 1998.
[12] Tambe, M.: "Towards Flexible Teamwork" *Journal of Artificial Intelligence Research*, Vol. 7., pp 83-124, 1997.

Author Biographies

GLENN TAYLOR is a Research Engineer at Soar Technology, Inc. He is the lead developer of SOF-Soar and is also involved with continuing development on TacAir-Soar and several other Soar-based agent systems. He received his BS in computer science and his MS in computer science and engineering from the University of Michigan in 1994 and 1996, respectively.

FRANK KOSS is a Senior Systems Engineer at Soar Technology, Inc., where he continues work related to Soar, JSAF, and simulation. Previously, while at the University of Michigan, he developed the interface between the Soar cognitive architecture and the ModSAF simulator, and extended ModSAF itself. He received his BS in computer engineering from Carnegie Mellow University in 1991 and his MSE in computer science and engineering from the University of Michigan in 1993.

PAUL NIELSEN is a Vice President and one of the co-founders of Soar Technology, Inc. He received his Ph.D. in computer science from the University of Illinois in 1988. Prior to joining Soar Technology, he worked at the GE Corporate Research and Development Center and the University of Michigan Artificial Intelligence Laboratory.