

## Virtual Humans for Team Training in Virtual Reality

Jeff Rickel and W. Lewis Johnson

*USC Information Sciences Institute & Computer Science Department*  
*4676 Admiralty Way, Marina del Rey, CA 90292-6695*  
*rickel@isi.edu, johnson@isi.edu, <http://www.isi.edu/isd/VET/vet.html>*

### Abstract

This paper describes the use of virtual humans and distributed virtual reality to support team training, where students must learn their individual role in the team as well as how to coordinate their actions with their teammates. Students, instructors, and virtual humans cohabit a 3D, simulated mock-up of their work environment, where they can practice together in realistic situations. The virtual humans can serve as instructors for individual students, and they can substitute for missing team members. The paper describes our learning environment, the issues that arise in developing virtual humans for team training, and our design for the virtual humans, which is an extension of our Steve agent previously used for one-on-one tutoring.

## 1 Introduction

Complex tasks often require the coordinated actions of multiple team members. Team tasks are ubiquitous in today's society; for example, teamwork is critical in manufacturing, in an emergency room, and on a battlefield. To perform effectively in a team, each member must master their individual role and learn to coordinate their actions with their teammates. There is no substitute for hands-on experience under a wide range of situations, yet such experience is often difficult to acquire; required equipment may be unavailable for training, important training situations may be difficult to re-create, and mistakes in the real world may be expensive or hazardous. In such cases, distributed virtual reality provides a promising alternative to real world training; students, possibly at different locations, cohabit a 3D, interactive, simulated mock-up of their work environment, where they can practice together in realistic situations.

However, the availability of a realistic virtual environment is not sufficient to ensure effective learning. Instructors are needed to demonstrate correct performance, guide students past impasses, and point out errors that students might miss. Yet requiring instructors to continually monitor student activities places a heavy burden on their time, and may severely limit students' training time. In addition, team training requires the availability of all appropriate team members, and may require adversaries as well. Thus, while virtual environments allow students to practice scenarios anywhere and anytime, the need for instructors and a full set of teammates and adversaries can provide a serious training bottleneck.

One solution to this problem is to complement the use of human instructors and teammates with intelligent agents that can take their place when they are unavailable. The intelligent agents cohabit the virtual world with human students and collaborate (or compete) with them on training scenarios. Intelligent agents have already proven valuable in this role as fighter pilots in large battlefield simulations [7, 10], but such agents have a limited ability to interact with students. Our work focuses on a different sort of agent: a virtual human that interacts with students through face-to-face collaboration in the virtual world, either as an instructor or teammate. We call our agent Steve (Soar Training Expert for Virtual Environments).

Our prior work focused on Steve's ability to provide one-on-one tutoring to students for individual tasks [14, 15]. Steve has a variety of pedagogical capabilities one would expect

of an intelligent tutoring system. However, because he has an animated body, and cohabits the virtual world with his student, he can provide more human-like assistance than previous disembodied tutors. For example, he can demonstrate actions, use gaze and gestures to direct the student's attention, and guide the student around the virtual world. This makes Steve particularly valuable for teaching tasks that require interaction with the physical world.

This paper describes our extensions to Steve to support team training. Steve's prior skills provided a solid foundation for his roles in team training, but several new issues had to be addressed. Each agent must be able to track the actions of multiple other agents and people, understand the role of each team member as well as the interdependencies, and communicate with both human and agent teammates for task coordination. In the remainder of this paper, we describe our learning environment for team training (Section 2), our solutions to these issues (Section 3), and related and future work (Section 4).

## 2 The Learning Environment

Our learning environment is designed to mimic the approach used at the naval training facility in Great Lakes, Illinois, where we observed team training exercises. The team to be trained is presented with a scenario, such as a loss of fuel oil pressure in one of the gas turbine engines that propels the ship. The team must work together, guided by standard procedures, to handle the casualty. At Great Lakes, the team trains on real, operational equipment. Because the equipment is in operation, the trainers have limited ability to simulate ship casualties; for example, they must mark gauges with grease pencils to indicate hypothetical readings. In our learning environment, the team, consisting of any combination of Steve agents and human students, is immersed in a simulated mock-up of the ship; the simulator creates the scenario conditions. As at Great Lakes, each student is accompanied by an instructor (human or agent) that coaches them on their role.

Each student gets a 3D, immersive view of the virtual world through a head-mounted display (HMD) and interacts with the world via data gloves. Lockheed Martin's Vista Viewer software [17] uses data from a position and orientation sensor on the HMD to update the student's view as she moves around. Additional sensors on the gloves keep track of the student's hands, and Vista sends out messages when the student touches virtual objects. These messages are received and handled by the simulator, which controls the behavior of the virtual world. Our current implementation uses VIVIDS [13] for simulation authoring and execution. Separate audio software broadcasts environmental noises through headphones on the HMD based on the student's proximity to their source in the virtual world. Our current training environment simulates the interior of a ship, complete with gas turbine engines, a variety of consoles, and their surrounding pipes, platforms, stairs and walls. A course author can create a new environment by creating new graphical models, a simulation model, and the audio files for environmental sounds.

Our architecture for creating virtual worlds [9] allows any number of humans and agents to cohabit the virtual world. While the behavior of the world is controlled by a single simulator, each person interacts with the world through their own copy of Vista and the audio software, and each agent runs as a separate process. The separate software components communicate by passing messages via a central message dispatcher; our current implementation uses Sun's ToolTalk as the message dispatcher. This distributed architecture is modular and extensible, and it allows the various processes to run on different machines, possibly at different locations. It greatly facilitates team training, where arbitrary combinations of people and agents must cohabit the virtual world; our extension to team training would have been more difficult had we originally designed a more monolithic system geared towards a single student and tutor.

Humans and agents communicate through spoken dialogue. An agent speaks to a person (teammate or student) by sending a message to the person's text-to-speech software (Entropic's TrueTalk), which broadcasts the utterance through the person's headphones. When a person speaks, a microphone on their HMD sends their utterance to speech recognition soft-



Figure 1: (a) Steve describing an indicator light; (b) one Steve agent speaking to another; (c) one Steve agent watching another

ware (Entropic's GraphVite), which broadcasts a semantic representation of the utterance to all the agents. Currently, Vista provides no direct support for human-to-human conversation; if the humans are not located in the same room, they must use telephone or radio to converse.

For team training, teammates and instructors must track each other's activities. Each person sees each other person in the virtual world as a head and two hands. The head is simply a graphical model, so each person can have a distinct appearance. Each Vista tracks the position and orientation of its person's head and hands via the sensors, and it broadcasts the information to agents and the other Vistas. Each agent appears as a human upper body (see Figure 1). To distinguish different agents, each agent can be configured with its own shirt, hair, eye, and skin color, and its voice can be made distinct by setting its speech rate, base-line pitch, and vocal tract size (these parameters are supported by the TrueTalk software).

### 3 Agent Design

#### 3.1 Architecture

Each Steve agent consists of three main modules: perception, cognition, and motor control [15]. The perception module monitors messages from other software components, identifies relevant events, and maintains a snapshot of the state of the world. It tracks the following information: the simulation state (in terms of objects and their attributes), actions taken by students and other agents, the location of each student and agent, and human and agent speech (separate messages indicate the beginning of speech, the end, and a semantic representation of its content). In addition, if the agent is tutoring a student, it keeps track of the student's field of view; messages from the student's Vista indicate when objects enter or leave the field of view. The cognition module, implemented in Soar [11], interprets the input it receives from the perception module, chooses appropriate goals, constructs and executes plans to achieve those goals, and sends out motor commands to the motor control module. The motor control module accepts the following types of commands: move to an object, point at an object, manipulate an object (about ten types of manipulation are currently supported), look at someone or something, change facial expression, nod or shake the head, and speak. The motor control module decomposes these motor commands into a sequence of lower-level messages that are sent to the other software components (simulator, Vista Viewers, speech synthesizers, and other agents) to realize the desired effects. (See [15] for more details on this architecture.)

To allow Steve to operate in a variety of domains, his architecture has a clean separation between domain-independent capabilities and domain-specific knowledge. The code in the perception, cognition, and motor control modules provides a general set of capabilities that are independent of any particular domain. These capabilities include planning, replanning, and plan execution; mixed-initiative dialogue; assessment of student actions; question answering ("What should I do next?" and "Why?"); episodic memory; communication with teammates; and control of a human figure [15]. To allow Steve to operate in a new domain,

<b>Task</b> transfer-thrust-control-ccs	<b>Task</b> loss-of-fuel-oil-pressure
<b>Steps</b> press-pacc-ccs, press-scu-ccs	<b>Steps</b> transfer-thrust-control-ccs, ...
<b>Causal links</b>	<b>Causal links</b> ...
press-pacc-ccs achieves ccs-blinking for press-scu-ccs	<b>Ordering</b> ...
press-scu-ccs achieves thrust-at-ccs for end-task	<b>Roles</b>
<b>Ordering</b> press-pacc-ccs before press-scu-ccs	eoow: (transfer-thrust-control-ccs pacc), ... ;
<b>Roles</b> pacc: press-pacc-ccs; scu: press-scu-ccs	engrm: (transfer-thrust-control-ccs scu), ...

Figure 2: Example team task descriptions

a course author simply specifies the appropriate domain knowledge in a declarative language. (Recent work has focused on acquiring the knowledge from an author's demonstrations and the agent's experimentation [2]). The knowledge falls in two categories: perceptual knowledge (knowledge about objects in the virtual world, their relevant simulator attributes, and their spatial properties) and task knowledge (procedures for accomplishing domain tasks and text fragments for talking about them). For details about Steve's perceptual knowledge, see [15]; the remainder of the paper will focus on Steve's representation and use of task knowledge.

### 3.2 Representing Task Knowledge

Most of Steve's abilities to collaborate with students on tasks, either as a teammate or tutor, stem from his understanding of those tasks. As the scenario unfolds, Steve must always know which steps are required, how they contribute to the task goals, and who is responsible for their execution. In order to handle dynamic environments containing other people and agents, he must understand the tasks well enough to adapt them to unexpected events; he cannot assume that the task will follow a pre-specified sequence of steps. Moreover, our goal was to support a declarative representation that would allow course authors to easily specify task knowledge and update it when necessary [14].

Our representation for individual tasks, used in our previous work for one-on-one tutoring, satisfies these design criteria. The course author describes each task using a standard plan representation [16]. First, each task consists of a set of steps, each of which is either a primitive action (e.g., press a button) or a composite action (i.e., itself a task). Composite actions give tasks a hierarchical structure. Second, there may be ordering constraints among the steps; these constraints define a partial order over the steps. Finally, the role of the steps in the task is represented by a set of causal links; each causal link specifies that one step in the plan achieves a goal that is a precondition for another step in the plan or for termination of the task. For example, pulling out a dipstick achieves the goal of exposing the level indicator, which is a precondition for checking the oil level.

This task representation is suitable for structured tasks based on standard procedures. It would not be suitable for tasks that require creative problem solving, such as design tasks. Fortunately, many tasks in industry and the military have this type of structure, including operation and maintenance of equipment, trauma care, and surgical procedures. Moreover, this representation need not be viewed as a fixed sequence of steps; rather, it is a general causal network of steps and goals, and can be used by a planning algorithm to dynamically order the steps even in the face of unexpected events, as described in Section 3.3.

To extend Steve to team training, we had to decide how to assign team members to task steps. Research on multi-agent teams has addressed methods by which teammates dynamically negotiate responsibility for task steps. However, supporting such negotiation among a team of agents *and* people would require more sophisticated natural language dialogue capabilities than Steve currently has. Fortunately, many team tasks have well-defined roles that are maintained throughout task execution, so we focus on this class of tasks.

Extending Steve to support such team tasks required one simple addition to each task description: a mapping of task steps to team roles. For example, Figure 2 shows a simplified task model for transferring thrust control to the central control station of a ship. Two roles

must be filled: one operator mans the propulsion and auxiliary control console (PACC) in the central control station (CCS), and another operator mans the shaft control unit console (SCU) in the engine room. The PACC operator requests the transfer by pressing the CCS button on her console, which results in the CCS button blinking on both consoles. When the CCS button is blinking on the SCU, the SCU operator presses it to finalize the transfer. This last action achieves the end goal of the task, which is indicated in the task description by specifying its effect as a precondition of the dummy step "end-task."<sup>1</sup>

If a step in the task is itself a team task, it will have its own roles to be filled, and these may differ from the roles in the parent task. Therefore, the parent task specifies which of its roles plays each role in the subtask. For example, the first task in Figure 2 is a subtask of the second (loss-of-fuel-oil-pressure). The task description calls for the executive officer of the watch (EOOW) to play the role of the PACC operator and for the engine room officer (ENGRM) to play the role of the SCU operator for the transfer of thrust control.

Task descriptions (e.g., Figure 2) specify the structure of tasks, but they leave the goals (e.g., ccs-blinking) and primitive steps (e.g., press-pacc-ccs) undefined. The course author defines each primitive step as an instance of some action in Steve's extensible action library. For example, the step press-pacc-ccs would be defined as an instance of press-button in which the particular button to be pressed is pacc-ccs, the name of an object in the virtual world. The course author defines each goal by the conditions in the simulated world under which it is satisfied. For example, ccs-blinking is satisfied when the simulator attribute scu-ccs-state has the value "blinking." Thus, Steve is able to relate his task knowledge to objects and attributes in the virtual world.

### 3.3 Using Task Knowledge

When someone (e.g., a human or agent instructor) requests a team task to be performed, each Steve agent involved in the task as a team member or instructor uses his task knowledge to construct a complete task model. The request specifies the name of the task to be performed and assigns a person or agent to each role in that task. Starting with the task description for the specified task, each agent recursively expands any composite step with its task description, until the agent has a fully-decomposed, hierarchical task model. Role assignments in the request are propagated down to subtasks until the task model specifies which team member is responsible for each step. For example, if the task is loss-of-fuel-oil-pressure, with Joe as the EOOW, then Joe will play the role of the PACC for the subtask transfer-thrust-control-ccs, and hence he is responsible for the step press-pacc-ccs. Since all agents have the same task knowledge, each agent will construct the same hierarchical task model with the same assignment of responsibilities.

For simulation-based training, especially for team tasks, agents must be able to robustly handle unexpected events. Scripting an agent's behavior for all possible contingencies in a dynamic virtual world is difficult enough, but the problem is compounded when each agent must be scripted to handle unexpected actions by any human team member. One option is to simply prevent human students from deviating from standard procedures, but this robs the team of any ability to learn about the consequences of mistakes and how to recover from them. Instead, we have designed Steve to use his task knowledge to adapt task execution to the unfolding scenario.

To do this, each agent maintains a plan for how to complete the task from the current state of the world. The task model specifies all steps that *might* be required to complete the task; it can be viewed as a worst case plan. Agents continually monitor the state of the virtual world, identify which goals in the task model are already satisfied, and use a partial-order planning algorithm to construct a plan for completing the task [15]. This plan is a subset of the task model, consisting of the steps relevant to completing the task, the ordering constraints among them, and the causal links that indicate the role of each step in achieving the end goals. In

<sup>1</sup>This representation for end goals is standard in AI planners [16].

our prior work, this plan would specify how an agent intended to complete a task; for team training, the plan specifies how the agent intends for the team to collectively complete the task, with some causal links specifying the interdependencies among team members (i.e., how one team member's action depends on a precondition that must be achieved by a teammate). Thus, agents dynamically interleave construction, revision, and execution of plans to adapt to the unfolding scenario.

If an agent is serving only as a missing team member, it simply performs its role in the task, waiting when appropriate for the actions of its teammates, and communicating with them when necessary. In contrast, an agent serving as an instructor for a human student interacts with that student in a manner similar to one-on-one tutoring. The agent can demonstrate the student's role in the task, explaining each action it takes, or it can monitor the student as she performs the task, answering questions when the student needs help. Moreover, the agent instructor can easily shift between these two modes as the task proceeds; the student can always interrupt the agent's demonstration and ask to finish the task herself, and she can always ask the agent to demonstrate a step when she gets stuck.

### 3.4 Team Communication

In team tasks, coordination among team members is critical. Although team members can sometimes coordinate their actions by simply observing the actions of their teammates, spoken communication is typically required. Team leaders need to issue commands. Team members often need to inform their teammates when a goal has been achieved, when they are starting an activity, and when they detect an abnormal condition. Because team communication is so important, it must be taught and practiced in team training.

We model team communication as explicit speech acts in the task descriptions. For the sort of structured tasks we have studied, this is natural; all the documented team procedures given to us specified when one team member should say something to another and how it should be said. To support this, we extended Steve's action library to include a new type of action: a speech act from one team member to another. The specification of the act requires four components: (1) the name of the task role to which the speech act is directed (e.g., SCU), (2) the name of the attribute being communicated (e.g., thrust-location), (3) the value being communicated for that attribute (e.g., ccs), and (4) the appropriate text string (e.g., "Thrust control is now at the central control station"). Each speech act appears as a primitive action in the task description, allowing us to explicitly model its relationship to the task, including the role responsible for performing it, ordering constraints on when it should be said, and causal links that specify how its effect contributes to completing the task (i.e., which other steps depend on that result).

Given this representation for team communication, Steve agents can both generate and comprehend such utterances during task execution. When an agent's plan calls for it to execute one of these speech acts, it sends the text string to appropriate speech synthesizers for its human teammates to hear, and it broadcasts a semantic representation of the speech act for its agent teammates to "hear." When a human says the appropriate utterance, her speech recognizer identifies it as a path through its domain-specific grammar, maps it to an appropriate semantic representation, and broadcasts it to the agents. Each agent checks its plan to see if it expects such a speech act from that person at that time. If so, it updates the specified attribute in its mental state with the specified value, and it nods to the student in acknowledgment. If the speech recognizer fails to understand the student's utterance, or the utterance is not appropriate at the current time, the student's instructor agent is responsible for giving the student appropriate feedback.

There are several important points about this approach. First, it only applies to structured tasks for which the required team communications can be specified in the task description; it will not suffice for tasks that require more arbitrary communication. Fortunately, many well structured team tasks, particularly in the military, include such a prescribed set of utterances.

Second, since Steve does not include any natural language understanding abilities, all valid variations of the utterances must be added to the grammar for the speech recognizer. Again, this is reasonable for tasks with prescribed utterances. Third, note the difference between our approach and communication messages in a purely multi-agent system; a speech recognizer cannot tell to whom the utterance is intended, so agents use the task model to determine whether the speaker is addressing them. Finally, each agent must treat a human student and their instructor as jointly performing a role; if either of them generates the speech act, it must be treated as coming from that role.

Although spoken communication is typically required for team tasks, nonverbal communication is also important. Human students can observe the actions of their nearby agent and human teammates, which is often required for proper team coordination. Agents look at a teammate when expecting them to do something, which can cue a student that she is responsible for the next step. Agents look at the teammate to whom they are speaking (Figure 1b), allowing students to follow the flow of communication and recognize when they are being addressed. Finally, agents react to their teammates' actions; they look at objects being manipulated by teammates (Figure 1c), and they nod in acknowledgment when they understand something a teammate says to them. For tasks that require face-to-face collaboration among team members, such nonverbal communication is critical.

#### 4 Discussion

Steve has been tested on a variety of naval operating procedures. In our most complicated team scenario, five team members must work together to handle a loss of fuel oil pressure in one of the gas turbine engines. This task involves a number of subtasks, some of which are individual tasks while others involve subteams. All together, the task consists of about three dozen actions by the various team members. Steve agents can perform this task themselves as well as in concert with human team members.

Several other recent systems have applied intelligent tutoring methods to team training, although none provides virtual humans like Steve. The PuppetMaster [12] serves as an automated assistant to a human instructor for large-scale simulation-based training, providing high-level interpretation and assessment to guide the instructor's interventions. It models team tasks at a coarser level than Steve, and is particularly suited to tracking large teams in very dynamic situations. AETS [19] uses detailed cognitive models to monitor a team of human students as they run through a mission simulation using the actual tactical workstations aboard a ship, rather than a virtual mockup. However, the system does not use these models to provide surrogate team members, and the automated tutor provides feedback to students only through a limited display window and highlighting of console display elements. AVATAR [5] supports simulation-based training for air traffic controllers and provides feedback on their verbal commands to simulated pilots and their console panel actions. However, the automated tutor and pilots have no planning capabilities, so scenarios must be more tightly scripted than in our approach. The Cardiac Tutor [6] trains a medical student to lead cardiac resuscitation teams. The system includes simulated doctors, nurses and technicians that play designated team roles, but they do not appear as virtual humans; they are heard but not seen. The representation of task knowledge appears less general than ours, since it is tailored particularly to trauma care. Unlike our system, where any team member could be a student or agent, their system is limited to a single student playing the role of the team leader.

Some limitations in our system could be alleviated by recent research results from related areas. To go beyond tasks with prescribed utterances, we could leverage ongoing research on robust spoken dialogue [1]. To handle tasks with shifting roles and unstructured communication among teammates, we could incorporate a more general theory of teamwork [8, 18]. To handle tasks that involve simultaneous physical collaboration (e.g., two people jointly lifting a heavy object), we will need a tighter coupling of Steve's perception and body control [3]. Although research in these areas is still incomplete, many useful methods have been developed.

There is a growing understanding of the principles behind effective team training [4]. Empirical experiments are beginning to tease out the skills that make teams effective (e.g., task skills vs. team skills), the basis for team cohesion (e.g., shared mental models), the best types of feedback (e.g., outcome vs. process), and the best sources of feedback (e.g., instructor vs. teammate). Because our approach allows us to model face-to-face interaction among human instructors and students and their agent counterparts, we are now in an excellent position to incorporate and experiment with a variety of these new ideas in team training.

## 5 Acknowledgments

This work was funded by the Office of Naval Research, grant N00014-95-C-0179. We are grateful for the contributions of our collaborators: R. Stiles and colleagues at Lockheed Martin; A. Munro and colleagues at Behavioral Technologies Laboratory; and R. Angros, B. Moore, and M. Thiebaut at ISI.

## References

- [1] J.F. Allen, E.W. Miller, E.K. Ringger, and T. Sikorski. Robust understanding in a dialogue system. In *Proc. 34th Annual Meeting of the Assoc. for Comp. Ling.*, pp. 62-70, 1996.
- [2] R. Angros, Jr., W.L. Johnson, and J. Rickel. Agents that learn to instruct. In *AAAI Fall Symposium on ITS Authoring Tools*, 1997. AAAI Press. AAAI Tech. Report FS-97-01.
- [3] N.I. Badler, C.B. Phillips, and B.L. Webber. *Simulating Humans*. Oxford University Press, 1993.
- [4] E. Blickensderfer, J.A. Cannon-Bowers, and E. Salas. Theoretical bases for team self-correction: Postering shared mental models. *Adv. in Interdiscipl. Studies of Work Teams*, 4:249-279, 1997.
- [5] C.A. Connolly, J. Johnson, and C. Lexa. AVATAR: An intelligent air traffic control simulator and trainer. In *Proc. ITS '98*, pp. 534-543. Springer, 1998.
- [6] C. Eliot and B.P. Woolf. An adaptive student centered curriculum for an intelligent training system. *User Modeling and User-Adapted Instruction*, 5:87-86, 1995.
- [7] R.W. Hill, Jr., J. Chen, J. Gratch, P. Rosenbloom, and M. Tambe. Intelligent agents for the synthetic battlefield: A company of rotary wing aircraft. In *Proc. IAAI-97*, pp. 1006-1012, 1997. AAAI Press.
- [8] N. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75, 1995.
- [9] W.L. Johnson, J. Rickel, R. Stiles, and A. Munro. Integrating pedagogical agents into virtual environments. *Presence: Teleoperators and Virtual Environments*, 7(6):523-546, December 1998.
- [10] R.M. Jones, J.E. Laird, and P.E. Nielsen. Automated intelligent pilots for combat flight simulation. In *Proc. IAAI-98*, pp. 1047-1054, 1998. AAAI Press.
- [11] J.E. Laird, A. Newell, and P.S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1):1-64, 1987.
- [12] S.C. Marsella and W.L. Johnson. An instructor's assistant for team-training in dynamic multi-agent virtual worlds. In *Proc. ITS '98*, pp. 464-473. Springer, 1998.
- [13] A. Munro and D.S. Surmon. Primitive simulation-centered tutor services. In *Proc. AI-ED Workshop on Arch. for Intell. Simulation-Based Learning Environments*, Kobe, Japan, 1997.
- [14] J. Rickel and W.L. Johnson. Intelligent tutoring in virtual reality: A preliminary report. In *Proc. Eighth World Conference on Artificial Intelligence in Education*, pp. 294-301. IOS Press, 1997.
- [15] J. Rickel and W.L. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence*, 1999. Forthcoming.
- [16] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [17] R. Stiles, L. McCarthy, and M. Pontecorvo. Training studio: A virtual environment for training. In *Workshop on Simul. and Interaction in Virtual Environments (SIVE-95)*, 1995. ACM Press.
- [18] M. Tambe. Towards flexible teamwork. *Jour. of Artificial Intelligence Research*, 7:83-124, 1997.
- [19] W. Zachary, J. Cannon-Bowers, J. Burns, P. Bilazarian, and D. Kreckler. An advanced embedded training system (AETS) for tactical team training. In *Proc. ITS '98*, pp. 544-553. Springer, 1998.