Army Acquisition, Logistics and Technology (ALT), pp.34-35, Nov-Dec 2001.

Intelligent Agents, Tools for the Command Post and Commander

Michael Bowman, Gheorghe Tecuci, Mihai Boicu
Learning Agents Laboratory, Department of Computer Science, MSN 4A5,
George Mason University, 4400 University Drive, Fairfax, VA 22030, USA
{mbowman3, tecuci, mboicu}@gmu.edu

Abstract
A commander's ability to make good decisions under adverse conditions can be decisive in battle. 72 hours into an operation, he or she is likely to be running on catnaps and caffeine. Decision support systems based on Artificial Intelligence have been more hype than help up to this point, playing a very limited role in military command and control or support systems. New research and experimental results presented in this paper suggest this may soon change. Supported by DARPA, the US Air Force and US Army, the George Mason University Learning Agents Laboratory has taken a novel approach to the creation and use of intelligent agents to solve complex military problems. Our research goal is to develop methods and tools that allow users with minimal computer skills to easily build, teach, and maintain intelligent software agents. In this approach, the user teaches the agent to perform various tasks in a way that resembles how they would teach an apprentice or student - providing examples and explanations, supervising and correcting behavior. Created in this manner, an intelligent agent can combine by-the-book doctrine and tactics, with the lessons a commander learns throughout a professional military career. The agent would not replace the commander or staff, but with no loss of capacity due to stress or other environmental factors, it would provide concise, relevant and explainable considerations that the commander can take into account when making decisions.

Introduction

You are a brigade commander, 72 hours into the fight and running on caffeine and cat naps. Your operations officer is briefing you on possible courses of action (COA) for a new mission. He recommends a COA and explains the advantages. Something about the COA bothers you, but you just can't put your finger on it. Fatigue, stress or some other distraction is keeping you from recalling something that would make a difference in this decision. You make your best judgement and drive on, but your gut feeling leaves you thinking that there was a better way – if you had just had more time or a clearer head!

Our decisions are a function of our education, training, experience and personal preference. It is well documented that the decisions we make under stress are generally not as good as those we make when we are well rested, comfortable and relaxed, which drives the US military to select commanders based on proven ability to make good decisions under adverse conditions.

Revisit the opening scenario. The situation and the environment are the same, but this time you have another tool to assist you. An intelligent agent, trained by you to remember the lessons of a lifetime, will help you decide. The intelligent agent is computer software that runs on common computers, and accesses data from your battle command systems and planning tools regardless of whether they are powerful networked computer systems or hand written notes and sketches. The intelligent agent does not care how cold it is, or how much sleep you have had. In seconds, it evaluates the COAs and presents you with a list of strengths, weaknesses, and issues for each of them. You quickly scan through the list, discarding some, nodding agreement with others, until you come to the one or more jewels that you immediately recognize as being critical to the decision. Based on your own judgement and the recommendations of your staff, but now armed with a few additional key considerations, you make your command decision.

These considerations might be based on planning guidelines you learned in a classroom; an after action review from an exercise you participated in; or on new enemy tactics that you observed in battle yesterday. The intelligent agent combines the things that a computer does best, sorting and sifting through data, with the things a human does best, learning from a lifetime of experience. It provides concise, relevant and explainable considerations that the commander can take into account when making decisions. This is our vision for the use of intelligent agents in the command post of the near future!

George Mason University, Learning Agents Laboratory

Decision support and expert systems have been around for a while. To date, they have produced more hype than service, and they have played a very limited role in military systems. Even with today's rapid growth in computing power, most software products claiming to be intelligent don't solve complex, real-world problems.

The George Mason University (GMU) Learning Agents Laboratory (LALAB) is taking a novel approach to the creation and use of intelligent agents to solve complex problems. Our research goal is to develop methods and tools that allow users with minimal computer skills to easily build, teach, and maintain intelligent software agents.

This research was part of the DARPA High Performance Knowledge Base (HPKB) project with additional support from the US Air Force Office of Scientific Research and the US Army Battle Command Battlelab (BCBL). The work continues in the DARPA Rapid Knowledge Formation (RKF) project, still supported by the Air Force and now with cooperation from the US Army War College. The goal of HPKB was to test the claim that with the latest artificial intelligence technology large knowledge bases could be quickly built and updated. Our research indicates that with the right approach, intelligent agents can meet this intent.

Acquiring Knowledge for Expert Decision Making: The Traditional Bottleneck and a Solution

A major stumbling block in building intelligent systems that solve problems on par with a human expert is known as the "knowledge acquisition bottleneck." This bottleneck comes from the requirement to transfer knowledge from an expert, through a knowledge engineer, to the computer (Figure 1). The knowledge engineer must learn what the expert knows and how the expert uses that knowledge. The engineer then uses various tools and techniques to build a knowledge base. This is a long, painful and inefficient process.

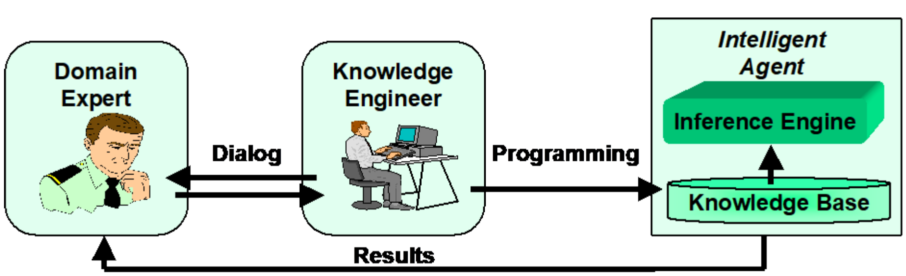

Figure 1: Traditional Approach to Building Intelligent Agents

Disciple is a theory, methodology and tool set, in which subject matter experts directly construct intelligent agents. In this approach, the expert directly teaches the agent to perform various tasks in a way that resembles how they would teach an apprentice or student, by giving the agent examples and explanations, and by supervising and correcting its behavior.

The knowledge base of a Disciple agent consists of an ontology that defines relevant terms and relationships from a problem domain, and a set of problem solving rules expressed with these terms. The general strategy behind the Disciple approach is to replace difficult knowledge engineering tasks required to build a knowledge base, with simpler tasks that can be performed by an expert, as shown in Figure 2.

The traditional approach to creation of a useful knowledge base requires very complex steps, including the creation of an ontology, the definition of problem solving rules and the validation and update of these rules. In general, these tasks require the creation of formal computer representations, a task that only a knowledge engineer can accomplish.

In the Disciple approach, these very complex tasks are replaced with simpler ones. Instead of creating an ontology, the expert updates and extends an initial ontology imported from existing sources of knowledge. Instead of defining a complex problem-solving rule, the expert identifies and explains an example solution from which Disciple learns a general rule. Instead of debugging a complex problem-solving rule, the expert critiques specific examples of agent problem solving from which Disciple updates corresponding rules. The expert will not need to create formal computer representations, just understand information generated by Disciple. Finally, the expert will not need to provide formal explanations, just informal hints that will guide Disciple in generating possible explanations from which the expert will choose.

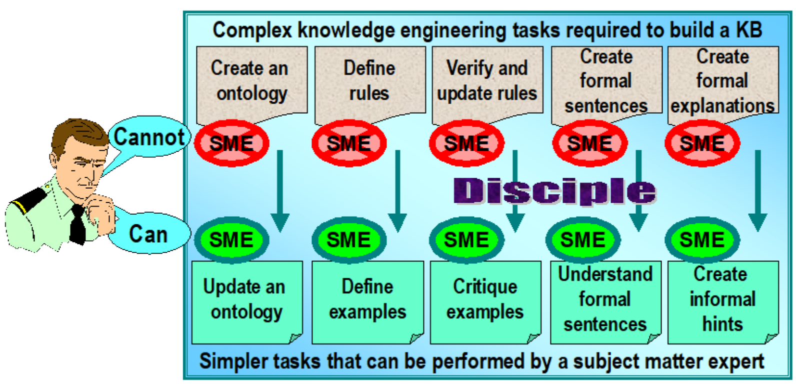

Figure 2: Disciple Approach for Building Intelligent Agents

Disciple's history, basic capabilities and inner workings are described in detail in Tecuci G., *BUILDING INTELLIGENT AGENTS: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, Academic Press, 1998. Recent papers describing improved capabilities are available on the GMU LALAB web page (http://www.lalab.gmu.edu).

Sample Application of Disciple – COA Critiquing

As part of HPKB we developed a Disciple agent to critique COAs for ground combat operations. The COAs were provided by the US Army and came in a standard format consisting of a multi-paragraph description (Figure 3) and a tactical sketch (Figure 4).

Mission:	BLUE-BRIGADE2 attacks to penetrate RED-MECH-REGIMENT2 at 130600 Aug in order to enable the completion of seize OBJ-SLAM by BLUE-ARMOR-BRIGADE1.
Close:	BLUE-TASK-FORCE1, a balanced task force (MAIN EFFORT1) attacks to penetrate RED-MECH-COMPANY4, then clears RED-TANK-COMPANY2 in order to enable the completion of seize OBJ-SLAM by BLUE-ARMOR-BRIGADE1. BLUE-TASK-FORCE2, a balanced task force (SUPPORTING EFFORT1) attacks to fix RED-MECH-COMPANY1 and RED-MECH-COMPANY2 and RED-MECH-COMPANY3 in order to prevent RED-MECH-COMPANY1 and RED-MECH-COMPANY2 and RED-MECH-COMPANY3 from interfering with conducts of the MAIN EFFORT1, then clears RED-MECH-COMPANY1 and RED-MECH-COMPANY2 and RED-MECH-COMPANY3 and RED-TANK-COMPANY1.

Figure 3: Segments of a COA description

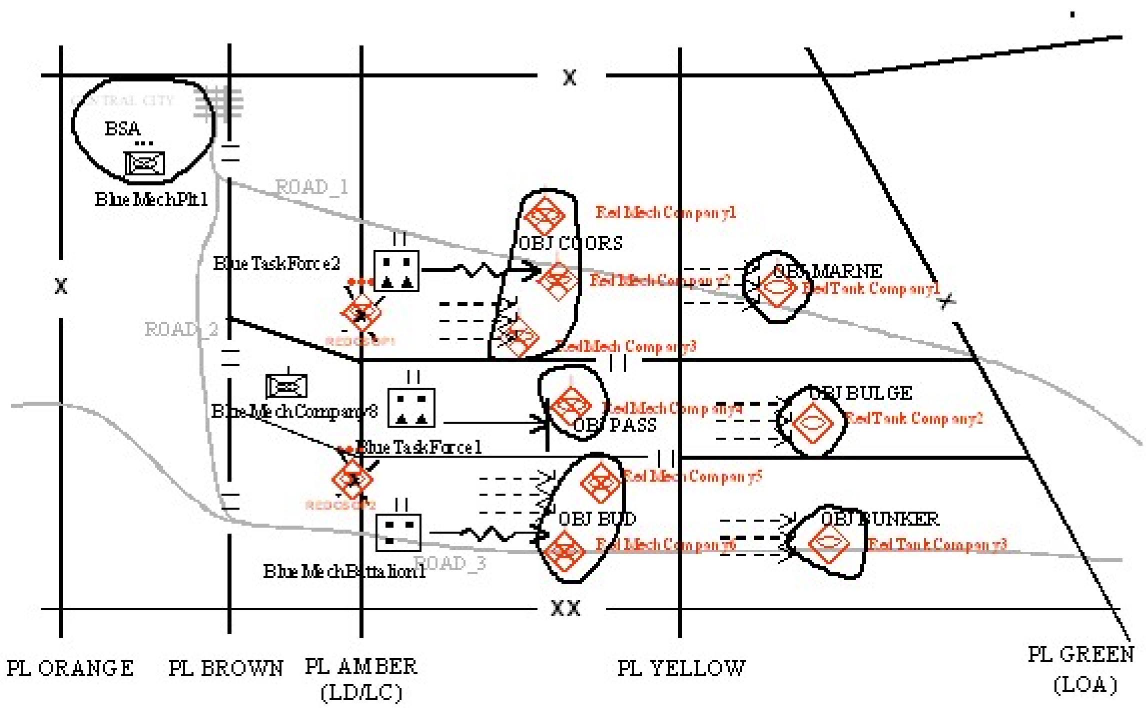

Figure 4: COA sketch

The Disciple COA-agent identifies strengths and weaknesses of a COA with respect to the principles of war and the tenets of army operations as described in US Army Field Manual 100-5. A general understanding of the principles and tenets exists, but military experts disagree on their application. Our goal was to create a tool that contained this common understanding while being flexible enough to allow rapid personalization by the experts training and using the agent. Figure 5 shows a strength identified by Disciple in the COA represented in Figures 3 and 4 for the principle of surprise.

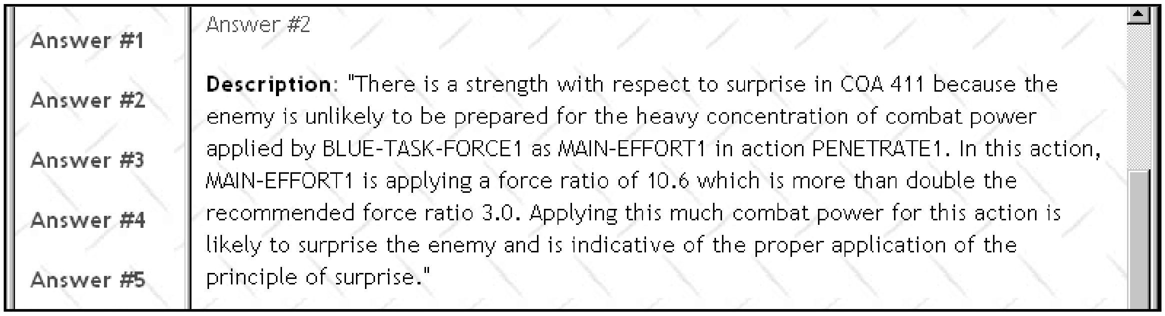

Figure 5: Solution generated by Disciple-COA.

A Methodology for Building Disciple Agents

The development of a specific Disciple agent includes two main processes: ontology development and agent training.

Building the domain ontology begins with importing background military knowledge, such as unit echelons and capabilities, from existing sources of knowledge. Additional terms and relationships identified by the expert are added as necessary. The Disciple-COA ontology was built by importing many terms needed to model the COA domain from a research knowledge base developed by Cycorp, called CYC.

Training a Disciple agent is an iterative process of showing it how to solve problems based on examples, letting the agent attempt to solve other problems, and providing the agent explanations for why these solutions are right or wrong. A strength of this approach is that the expert does not have to be perfect or comprehensive when he conducts agent training. Flaws in training show up naturely when Disciple tries to solve problems on its own. The expert merely has to examine Disciple' generated solutions and provide explanations for where it went wrong.

Experimental Results

The Disciple methodology and agents have been tested with other systems, as part of DARPA annual HPKB program evaluations. In summary, the experimental results show that Disciple based agents were highly effective in knowledge acquisition and complex problem solving, outperforming the other systems developed to solve similar problems.

In August 1999, we conducted a knowledge acquisition experiment to demonstrate that it is possible for military experts to directly train Disciple agents. Four US Army officers with 16 to 22 years of service successfully trained Disciple agents that critiqued COAs. Commenting on the usefulness and usability of Disciple, one officer, LTC John N. Duquette, stated: "The potential use of this tool by domain experts is only limited by their imagination—not their AI programming skills." To our knowledge, this is the first time subject matter experts with no prior knowledge engineering experience successfully directly trained intelligent agents to solve complex problems.

Conclusion

We have briefly presented a vision for the use of intelligent agents in a military command post, described some of the challenges faced, and presented the Disciple approach to overcoming those challenges. Our long-term goal is to develop the technology that allows typical computer users to directly build intelligent agents and knowledge bases as easily as they use personal computers for text processing. This will change the way intelligent agents will be built, from being programmed by a knowledge engineer, to being taught by a subject matter expert, and will contribute to a generalized application of agent technology in all areas of human activity.