INFORMATION AGE WARFARE – INTELLIGENT AGENTS IN THE CLASSROOM AND THE STRATEGIC ANALYSIS CENTER

COL Michael Bowman, U.S. Army, Fort Belvoir, VA; Gheorghe Tecuci, George Mason University, Fairfax VA; Mihai Boicu, George Mason University, Fairfax VA; COL (Ret) Jerome Comello, U.S. Army War College, Carlisle Barracks, PA

ABSTRACT

This paper presents a unique research and development project conducted over the past four years by the George Mason University (GMU) Learning Agents Center (LAC) and the U.S. Army War College. This research synergistically integrates cutting edge artificial intelligence research, military strategy research, and the practical use of agents in strategic analysis and education in the context of Information Age Warfare and Network-Centric Operations. This research has produced a collaborative assistant for rapid knowledge formation (RKF) and reasoning, called Disciple-RKF. Disciple-RKF enables a team of subject matter experts who do not have prior knowledge engineering experience to rapidly build and use intelligent knowledge-based agents, with limited assistance from a knowledge engineer. Each subject matter expert teaches a personal Disciple-RKF agent while collaborating with it in solving specific problems. During this process, each Disciple-RKF agent learns from its expert and builds its knowledge base. These knowledge bases are then merged into an integrated knowledge base for Disciple-RKF. Disciple-RKF has been successfully applied to develop knowledge-based agents for military center of gravity analysis which have been used in the curriculum of the US Army War College.

1. Introduction

Military historians and practitioners from Sun Tzu, to Carl Von Clausewitz, and now modern writers like Robert R. Leonard, have documented and examined the critical importance of information in warfare. Sun Tzu wrote 2500 years ago that to avoid peril in war, and to maximize the likelihood of victory, it was essential to know as much as possible about your enemy and your own forces. Clausewitz wrote almost 200 years ago that the confusion of events that make up combat, and the constant uncertainty that confounded a commander's understanding of what was going on around him, created what he described as the "fog of war," and that to be successful a commander must see past that fog.

We have entered a new era that some call the Information Age heralding our vastly increased capability to collect, process, disseminate, and utilize information. When modern writers describe Information Age Warfare they emphasize not only the technical aspects of this globally increased capacity to handle information, they place more emphasis than ever on the dramatic impact, from strategic down to tactical levels, that public access to information can produce (Alberts, et al., 2001). One hundred years ago genocide could occur in a remote corner of the world and go unnoticed and thereby ignored for generations. Today, a single dramatic event widely disseminated over the Internet, such as the beheading of a hostage in Iraq, can shift the foreign policy of a nation over night. Modern writers like Leonard acknowledge the current relevance of the classic principles and tenets of warfare described by Clausewitz and others, but they stress that in this modern age, access to information and rapid decision making are more important than ever (Leonard, R., 1998).

Modern military strategy and tactics are being transformed to take full advantage of Information Age technology in what are called "network-centric operations." Network-centric operations seek to maximize collection and dissemination of information and understanding throughout all levels of friendly forces, while attempting to disrupt enemy access to information and understanding. While access and understanding are important, it is still the commander's ability to rapidly make correct decisions that is most crucial to success in combat. Successful network-centric operations require that commanders and other decision makers have access to appropriate stores of organized information, tools to visualize the most important elements of that information, and automated, trust-worthy decision aids to support their most critical decisions.

The cooperative research, development, and experimentation described in this paper was designed to address key issues in Information Age Warfare and network-centric operations. It uses cutting-edge artificial intelligence, knowledge engineering, and machine learning technologies to acquire, organize, and present critical information about a classic military problem domain – strategic center of gravity analysis. It then goes several key steps further by providing an automated decision support tool that can assist the military decision maker, capture how and why critical decisions were reached, and provides a training mechanism for developing future military decision makers. These tools have been improved and extended through hands-on use and feedback from military officers over the past four years, and validated through formally documented experimentation. They are now an imbedded part of the curriculum of the US Army War College.

The more general objective of this research project was to develop and to experimentally validate a collaborative assistant for rapid knowledge formation and reasoning. This assistant enables a team of subject matter experts who do not have prior knowledge engineering experience to rapidly develop an integrated knowledge base for a complex application. The emphasis of this research was on acquiring expert problem solving knowledge that is not normally represented in written documents. This task is complementary to that of acquiring knowledge that has already been expressed in textbooks or other documents. This research has resulted in the development, experimental use, and transition of a complex knowledge engineering environment, called Disciple-RKF, and its application to the military center of gravity analysis domain, as described in the following, and in several publications.

Figure 1 introduces the general approach investigated. In this approach, each subject matter expert teaches a personal Disciple-RKF agent, while collaborating with it in solving specific problems. During this process, the Disciple-RKF agent learns from the expert, building, extending and improving its knowledge base. The resulting knowledge bases of all these Disciple-RKF agents are then integrated by a knowledge engineer. The Disciple-RKF agent with the integrated knowledge base can then be used in three ways. It can be used by a non-expert as a problem solver. It can be used by an expert as a problem solving assistant. Finally, it can be used by a student as a tutoring system.

2. Synergistic Collaboration and Transition to the Army War College

This project involved a multi-objective collaboration between the Learning Agents Center of George Mason University and the Center for Strategic Leadership along with the Department of Military Strategy, Planning, and Operations of the US Army War College (Tecuci et al., 2002). The US Army War College selected the problem domain (strategic center of gravity analysis), and provided extensive subject matter expertise (faculty and senior officers from all the military services), and experimentation support for both the developed technology and the resulting knowledge bases and agents.

A distinguishing feature of this collaboration is the synergistic integration of artificial intelligence research, with military strategy research, and the practical use of agents in education. The artificial intelligence research objective was the development of knowledge bases and agents by subject matter experts using learning agent technology. The military strategy research objective was the development of a systematic approach to center of gravity determination. The educational objective was the enhancement of the educational process of senior military officers and strategic leaders through the use of intelligent agent technology. This integration accelerated the development of the artificial intelligence technology and of intelligent agents for center of gravity analysis. It also facilitated the transition of this work to the US Army War College, where Disciple agents have been used since 2001 in a sequence of two joint warfare courses, “319jw Case Studies in Center of Gravity Analysis,” and “589jw Military Applications of Artificial Intelligence.”

3. The Center of Gravity Analysis Challenge Problem

Military center of gravity analysis was used as a challenge problem to test the knowledge acquisition, learning, and problem solving methods of Disciple-RKF. The concept of center of gravity (Clausewitz, 1832), introduced by Karl von Clausewitz, is fundamental to military strategy, denoting the primary source of moral or physical strength, power, or resistance of a force (Strange, J., 1996). The most important objective of a force (state, alliance, coalition, or group), in any type of conflict, is to protect its own center of gravity while attacking the center of gravity of its enemy. There is great emphasis on center of gravity analysis, in the education of strategic leaders at all U.S. senior military service colleges. This analysis requires a wide range of background knowledge not only from the military domain, but also from the political, psychosocial, economic, geographic, demographic, historic, international, and other domains. In addition, the situation, the adversaries involved, their goals, and their capabilities can vary in important ways from one scenario to another. Center of gravity analysis is a very good example of knowledge-intensive, expert problem-solving that a Disciple agent should be able to learn.

The approach to center of gravity analysis used in our experimentation is based on the work of Strange (Strange, J., 1996) and Giles and Galvin (Giles, P.,

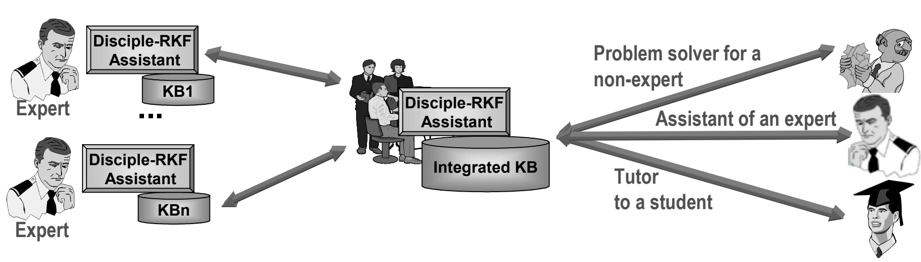

Figure 1: General approach to rapid knowledge formation by subject matter experts.

and Galvin, T.P, 1996), and developed with experts from the US Army War College. It consists of two main phases, *identification* and *testing*. During the identification phase, center of gravity candidates from different elements of power of a force (such as government, military, people, economy) are identified. For instance, a strong leader is a center of gravity candidate with respect to the government of a force. During the testing phase, each candidate is analyzed to determine whether it has all the critical capabilities that are necessary to be the center of gravity. For example, a leader must be protected, be informed, be able to communicate (with the government, the military, and the people), be influential (with the government, the military, and the people), be a driving force, have support (from the government, the military, and the people), and be irreplaceable. For each capability, one needs to determine the existence of the essential conditions, resources and means that are required by that capability to be fully operative, and which of these, if any, represent critical vulnerabilities.

4. The Architecture of the Disciple-RKF Agents

The architecture of Disciple-RKF includes the components from Figure 2 (Boicu et al., 2004). The core of the system is the learning agent shell, which has the following domain-independent components:

- A problem solving component based on the task reduction paradigm of problem solving, including the following modules:
  - Modeling assistant that helps the user to express his/her contributions to the problem solving process;
  - Interactive problem solving agent;
  - Autonomous problem solving agent.

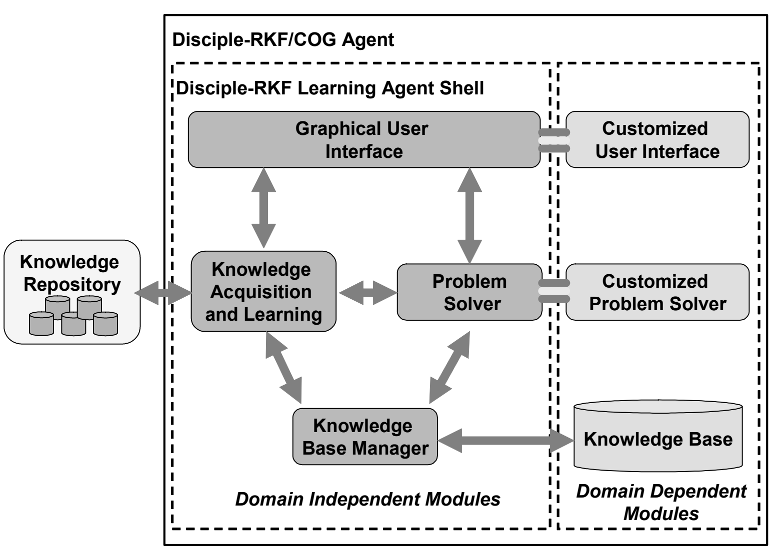

Figure 2: The Architecture of the Disciple-RKF agents.

- A knowledge acquisition and learning component for acquiring and refining the knowledge of the agent and allowing a wide range of operations including ontology import, user definition of knowledge base elements, ontology learning, and rule learning. This component includes the following modules:
  - Ontology development modules:
    - Tree-based browsers for objects and features;
    - Graph-based browsers for objects and features (association browser, hierarchical browser);
    - Viewers and editors for objects and features;
    - Ontology import module;
    - Knowledge base merging module.
  - Instances elicitation modules:
    - Scenario elicitation module;
    - Scripts editor.
  - Learning and refining modules:
    - Task formalization and learning module;
    - Explanation generation module;
    - Rule learning module;
    - Rule refinement with positive examples;
    - Rule refinement with negative examples;
    - Rule analysis module;
    - Rule regeneration module;
    - Exceptions-based knowledge base refinement;
    - Feature learning module.
- A knowledge base manager which controls the access to and the updates of the knowledge base. Each module of Disciple-RKF can access the knowledge base only through the functions of the knowledge base manager.
- A windows-based, domain-independent, graphical user interface.

The three components in the right hand side of Figure 2 are the typical domain dependent components of a Disciple-RKF agent that was customized for a specific application, such as center of gravity analysis. The Disciple-RKF center of gravity agent (also referred as Disciple-COG) includes the following components:

- A customized problem solving component that extends the basic task-reduction component in order to satisfy the specific problem solving requirements of the application domain.
- Customized graphical user interfaces which are built for the specific Disciple agent to allow the experts and the end users to communicate with the agent as close as possible to the way they communicate in their domains.
- The knowledge base of the Disciple agent containing knowledge specific to the center of gravity domain.

5. Knowledge Base Development Methodology

The Disciple approach covers all the phases of agent development and use (Tecuci 1998; Boicu et al., 2001; Tecuci et al., 2002). First, a knowledge engineer works with a subject matter expert to develop an ontology for the application domain. They use the ontology import module to extract relevant ontology elements from existing knowledge repositories, as well as the ontology editors and browsers of Disciple-RKF to create new elements. Figure 3 shows the interfaces of three of Disciple's ontology browsers: at top is the association browser which displays an object and its relationships with other objects; at bottom left is the tree browser which displays the hierarchical relationships between the objects in a tree structure; and at bottom right is the hierarchical browser which displays the hierarchical relationships between the objects in a graph structure.

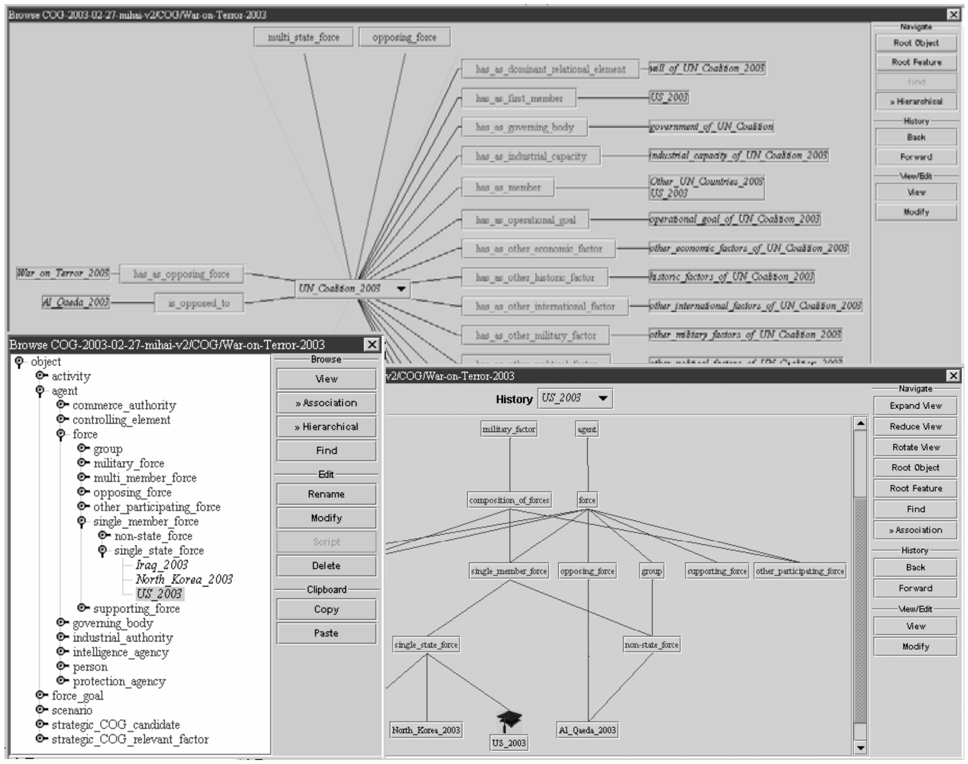

Figure 3. Three ontology browsers of Disciple-RKF

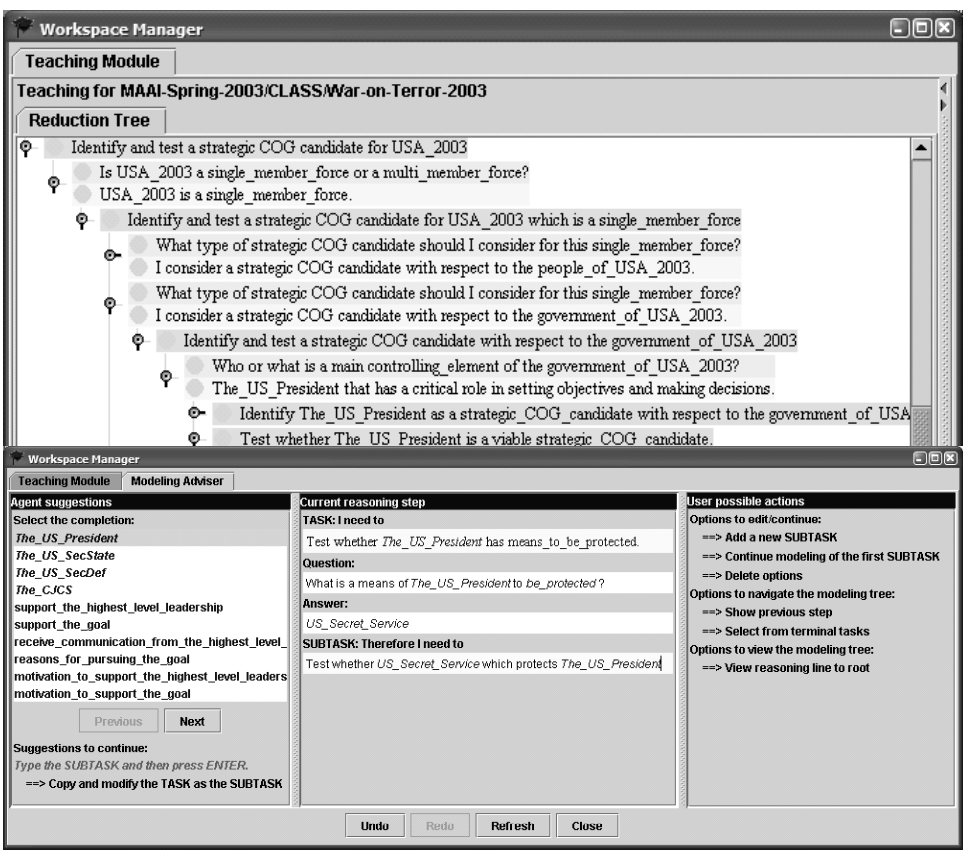

Figure 4: The reasoning tree and the modeling assistant interfaces.

The result of this knowledge base development phase is an object ontology which is complete enough to be used as a generalization hierarchy for learning, allowing the expert to train the Disciple agent on how to solve problems, with limited assistance from a knowledge engineer. The expert formulates a specific problem solving task and shows the agent the corresponding problem solving steps, helping the agent to understand them. Each problem solving step indicated by the expert consists of a task to be reduced, a question related to that task, the answer to the question, and one or several subtasks or solutions that reduce the task (Bowman, 2002). The top part of Figure 4 shows a fragment of a reasoning tree generated by the interactive problem solver consisting of a sequence of task reduction steps. The bottom part shows the interface of the Modeling Assistant that helps the expert express how to reduce the task at the bottom of the reasoning tree. This assistant may suggest the question to be asked, the answer to the question, and the knowledge base elements the expert may be referring to, since the expert's input is in natural language.

Each problem solving step indicated by the expert is an example from which the agent learns a problem solving rule. First, the expert and Disciple collaborate in formalizing the tasks from the expert's examples and Disciple learns general task patterns. Next the expert helps Disciple find an explanation of why the task reduction step is correct and Disciple learns a general task reduction rule. The top part of Figure 5 shows the interface of the task formalization module, with the tasks in natural language in the middle and their formalizations on the right. The bottom part of Figure 5 shows the explanation generation and selection interface. The task reduction rule learned from the example specified in the modeling assistant (bottom of Figure 4) is shown in the rule viewer in Figure 6.

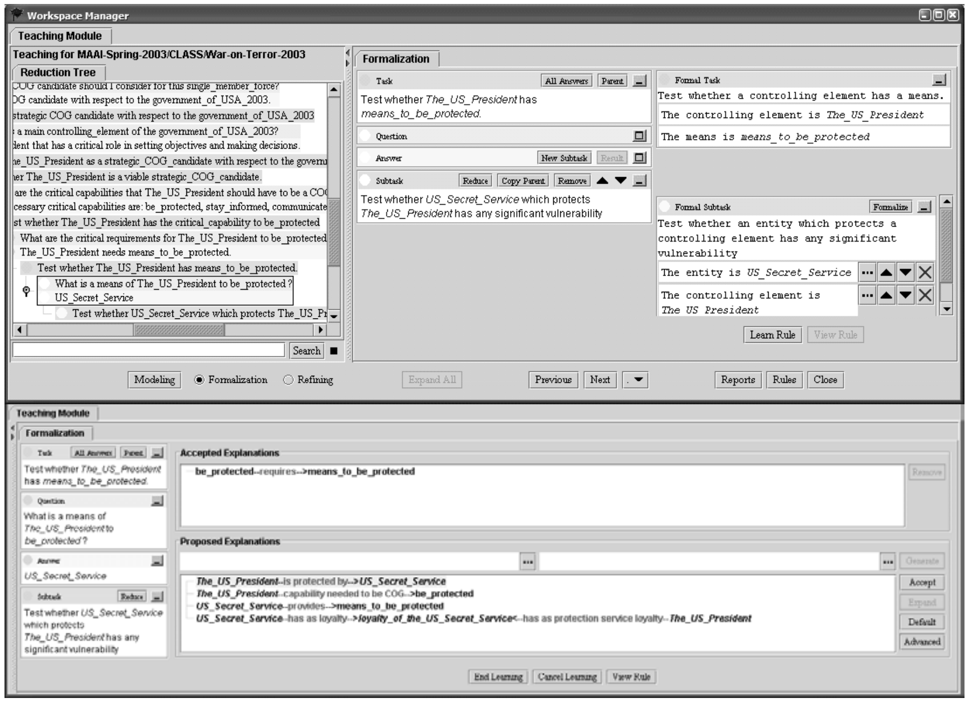

Figure 5: Task formalization and example explanation.

As Disciple learns new rules from the expert, the interaction between the expert and Disciple evolves from a teacher-student interaction toward an interaction where both collaborate in solving a problem. During this mixed-initiative problem solving phase, Disciple learns not only from the contributions of the expert, but also from its own successful or unsuccessful problem solving attempts, which leads to refinement of learned rules. At the same time, Disciple extends the object ontology with new objects and features.

Copies of Disciple agents may also be trained in parallel by different experts. In this case, the individual knowledge bases have to be merged into an integrated knowledge base, as discussed in section 8.

6. Use of Disciple-COG in Scenario Elicitation Experiments

Successive versions of the customized Disciple-RKF/COG agent were used in the "Case Studies in Center of Gravity Analysis" course at the US Army War College beginning in 2001, becoming part of the course syllabus. Figure 7 shows how the agent was used. First Disciple was taught how to analyze a scenario, based on the expertise of the course's instructor. The students then

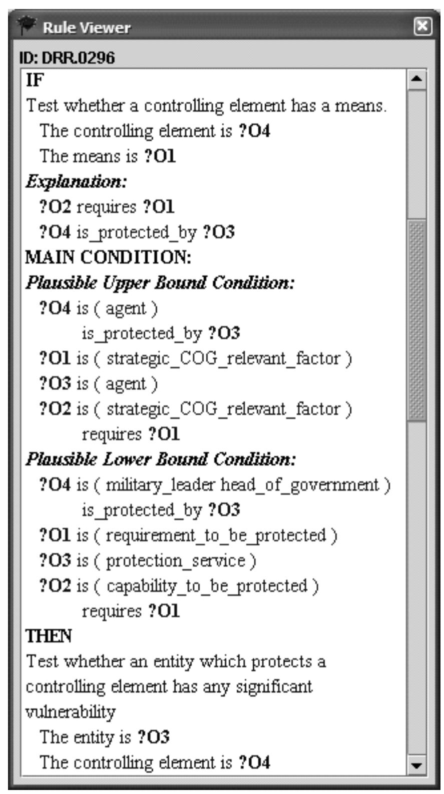

Figure 6: Learned rule.

used Disciple as an intelligent assistant that helps them develop a center of gravity analysis of a war scenario. Each session of this course was an experiment in scenario elicitation from subject matter experts. These experiments demonstrated that the Disciple approach can be used to develop agents that have been found to be useful for a complex military domain (Tecuci et al., 2004b).

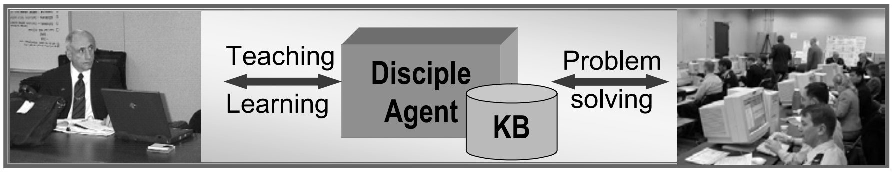

Figure 7: Training and using the Disciple agent (the COG course).

7. Use of Disciple-RKF/COG in Agent Teaching Experiments

Figure 8 shows the use of Disciple in the "589jw Military Applications of Artificial Intelligence" course at the US Army War College. In this course, the students teach personal Disciple agents their own expertise in center of gravity analysis and then evaluate both the developed agents and the development process. In the 2001 experiments, the students used historic scenarios (such as the 1943 WWII Okinawa campaign) with state actors to teach personal Disciple agents how to identify center of gravity candidates. In the 2002 experiment, the students used historic scenarios and a hypothetical scenario with state actors to teach personal Disciple agents how to identify center of gravity candidates and to eliminate those candidates that do not pass certain tests. In the 2003 experiment, the students used historic, current and hypothetical scenarios, with both state and non-state actors, to teach personal Disciple agents how to test center of gravity candidates based on the concepts of critical capabilities, critical requirements, and critical vulnerabilities. A total of 38 US and international officers from all military branches and the Reserve Components have attended these courses. At the end of these three experiments, 10 of them strongly agreed, 20 agreed, 7 were neutral and only one disagreed with the statement *"I think that a subject matter expert can use Disciple to build an agent, with limited assistance from a knowledge engineer."* This result shows that significant progress has been made in developing the technology that will allow subject matter experts to build their own intelligent assistants.

8. Experiment in Parallel Knowledge Base Development by Subject Matter Experts

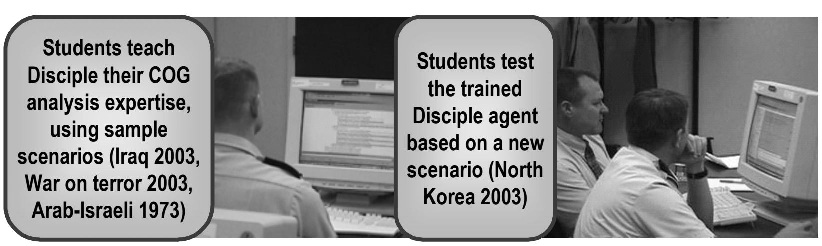

Figure 8: Teaching and testing a personal Disciple agent (the MAAI course).

The Spring 2003 session of the "Military Applications of Artificial Intelligence" course included an experiment in parallel knowledge base development by subject matter experts, which is illustrated in Figure 9 (Tecuci et al., 2004a). Before starting the experiment, the Disciple-RKF agent was trained to identify leaders as center of gravity candidates. The knowledge base of this agent contained the definitions of 432 concepts and features, 29 tasks and 18 task reduction rules. However, the agent had no knowledge of how to test the identified candidates. A domain analysis and ontology development was then performed, involving all the subject matter experts. This considered the example of testing whether Saddam Hussein, in the Iraq 2003 scenario, had all the required critical capabilities to be the center of gravity for Iraq. Based on this domain analysis, the ontology of Disciple-RKF was extended with the definition of 37 new concepts and features identified with the help of the subject matter experts.

The 13 subject matter experts were next grouped into five teams of 2 or 3 experts each, and each team was given a copy of the extended Disciple-RKF agent. Each team then trained its agent to test whether a leader had one or two critical capabilities, as indicated in Figure 1. For example, Team 1 trained its agent how to test whether a leader has the critical capabilities of staying informed and being irreplaceable. Three scenarios (Iraq 2003, Arab-Israeli War 1973, and War on Terror 2003) were used for the training, with the experts teaching Disciple-RKF how to test each strategic leader from these scenarios. As a result of the training performed by the experts, the knowledge base of each Disciple-RKF agent was extended with new features, tasks, and rules, as indicated in Figure 11. For example, the knowledge base of the agent trained by Team 1 was extended with 5 features, 10 tasks and 10 rules. The average training time per team was 5 hours and 28 minutes, and the average rule learning rate per team was 3.53 rules/hour. This included the time spent in all the agent training activities, such as scenario specification, modeling expert reasoning, task formalization, rule learning, problem solving, and rule refinement.

After the 5 Disciple-RKF agents were trained, their knowledge bases were merged by a knowledge engineer using the knowledge base merging tool of Disciple-RKF. The knowledge engineer also performed a general test of the integrated knowledge base, which included 10 new features, 102 new tasks, and 99 new rules. During this process two semantically equivalent features were unified, 4 rules were deleted, and 12 other rules were refined by the knowledge engineer. The other 8 features and 83 rules acquired from the experts were not changed. Most of the modifications were done to remove rule redundancies, or to specialize overly general rules.

Next, each expert team tested the integrated agent on a new scenario, North Korea 2003. Each team was asked to judge the correctness of each reasoning step performed by the agent for only the capabilities which that SME team had performed agent training. The result was 98.15% correctness.

9. Conclusions and Future Research Directions

The experiment described in section 8 and illustrated in Figure 9 is the first of its kind ever performed. It demonstrates Disciple's capability for rapid and parallel development of knowledge bases by subject matter experts, with limited assistance from knowledge engineers, and the integration of the developed knowledge bases into a functioning agent. However, while significant progress has been made, much work remains to be done to improve the developed methods. For instance, while the subject matter expert has an increased role and independence in the agent development process, the knowledge engineer still has a critical role to play. The knowledge engineer must develop and provide a relatively complete and correct object ontology. The knowledge engineer also has to develop a generic modeling of the expert's problem solving process based on the task reduction paradigm. Even guided by this generic modeling, and using natural language, the subject matter expert has difficulty expressing his reasoning process. Therefore, more work is necessary to develop methods for helping the expert in this task, along the path opened by the Modeling Advisor. The experimentations performed to date have also revealed that the mixed-initiative reasoning methods of Disciple-RKF could be significantly empowered by enhancing the natural language processing capabilities of the system.

The development of Disciple-RKF has been focused on the initial acquisition of problem solving knowledge from a subject matter expert, which results in an initial knowledge base. Further work is required to develop

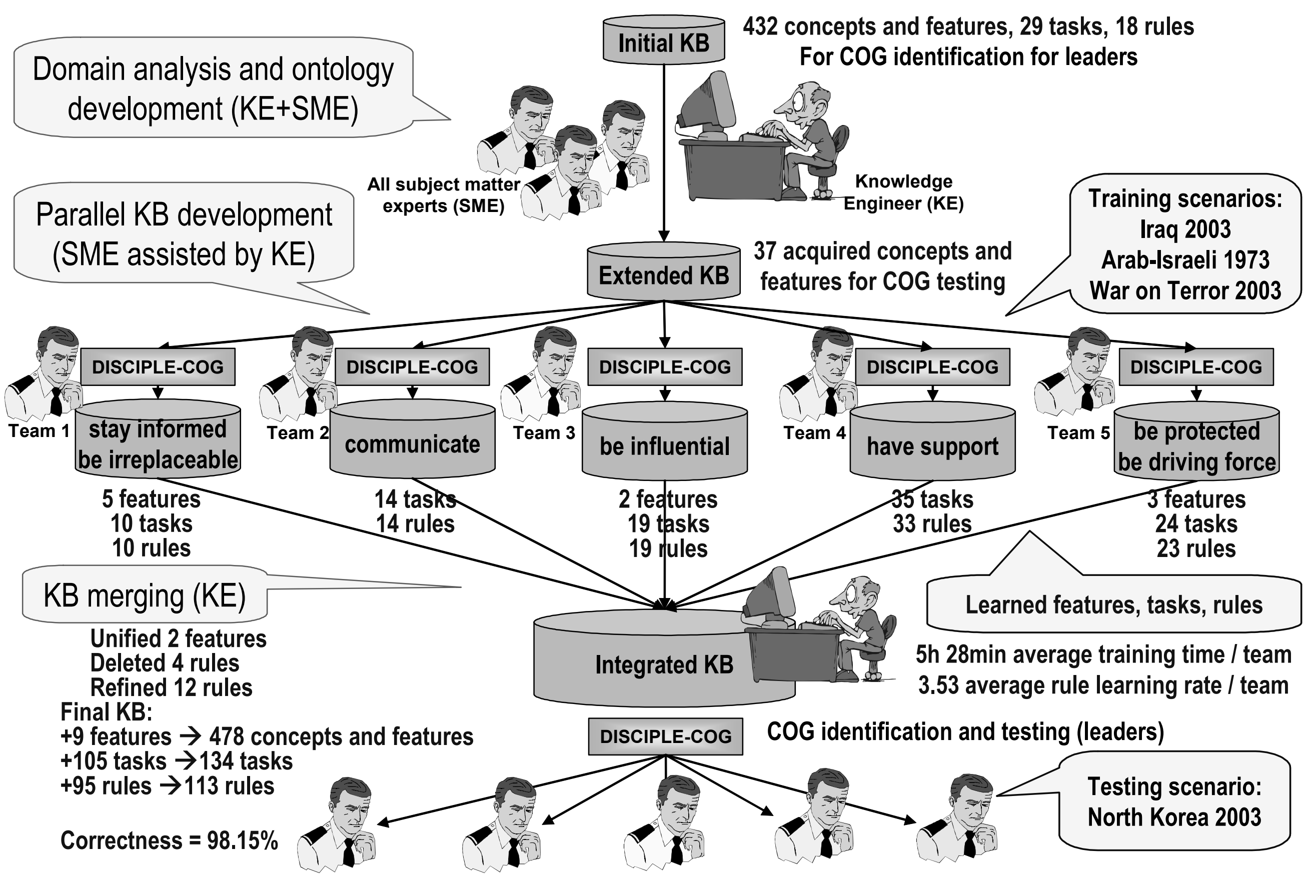

Figure 9. Experiment of rapid knowledge base development by subject matter experts.

knowledge acquisition and learning methods for knowledge base extension and refinement, as well as the integration, validation, and maintenance of the knowledge acquired from different subject matter experts. In particular, the Disciple approach can naturally be extended with methods and tools for:

- Knowledge acquisition for rules and ontology refinement, using mixed-initiative and multi-strategy techniques that will exploit the complementary elements of human and automated reasoning, and between several learning strategies (such as learning from examples, from explanations, by analogy, by abduction, or by abstraction).
- Acquisition of meta-rules that will capture an expert's rational for choosing among different ways of performing a problem solving task.
- Knowledge bases integration, validation, and maintenance (which will address problems such as ontology merging, inconsistencies within a knowledge piece, redundancies and inconsistencies among knowledge pieces acquired from different experts, refinement and reorganization of the object ontology for increased performance, restructuring and refinement of the acquired problem solving rules, etc.).

ACKNOWLEDGEMENTS

This research was performed in the Learning Agents Center of George Mason University and was sponsored by several US Government agencies, including the Defense Advanced Research Projects Agency, the Air Force Research Laboratory, the Air Force Office of Scientific Research, and the US Army War College, under the agreement numbers F30602-00-2-0546, F49620-00-1-0072, FA8750-04-1-0257. The views and conclusions contained therein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the sponsoring US Government agencies. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. Dorin Marcu, Bogdan Stanescu, Cristina Boicu, Marcel Barbulescu, Hao Xianjun and Catalin Balan have contributed to the Disciple RKF system.

REFERENCES

Alberts, D., Garstka, J., Hayes, R., Signori, D., *Understanding Information Age Warfare*, Washington D.C., CCRP Publications, 2001. 35.

Boicu M., Tecuci G., Stanescu B., Marcu M. and Cascaval C.E., Automatic Knowledge Acquisition from Subject Matter Experts, in *Proceedings of the Thirteenth International Conference on Tools with Artificial Intelligence* (ICTAI), 7-9 November 2001, Dallas, Texas, IEEE Computer Society, Los Alamitos CA, 2001, pp 69-78.

Boicu M., Tecuci G., Stanescu B., Marcu D., Barbulescu M., Boicu C., *Design Principles for Learning Agents*, in *Proceedings of AAAI-2004 Workshop on Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems,* July 26, San Jose, AAAI Press, Menlo Park, CA, 2004.

Bowman, M., A Methodology for Modeling Expert Knowledge that Supports Teaching-Based Development of Agents," *PhD Dissertation in Information Technology*, School of Information Technology and Engineering, George Mason University, Fairfax Virginia, 2002, 66-110.

Clausewitz, K.V. 1832. *On War,* translated and edited by. Howard M. and Paret P., Princeton NJ, Princeton University Press, 1976, 595-596.

Giles, P., and Galvin, T., *Center of Gravity: Determination, Analysis and Application*. CSL, U.S. Army War College, PA: Carlisle Barracks, 1996.

Leonard, R., *The Principles of War for the Information Age*, Novato CA, Presidio Press, 1998, 22-34.

Strange, J., *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*, Quantico VA, US Marine Corps University Foundation, 1996.

Tecuci, G. 1998. *Building intelligent agents: an apprenticeship multistrategy learning theory, methodology, tool and case studies*. London: Academic Press.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., *Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain*, in *AI Magazine,* 24, 4, 2002, AAAI Press, Menlo Park CA, 2002, pp.51-68.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Barbulescu M., *Parallel Knowledge Base Development by Subject Matter Experts* in *Proceedings of the 14th International Conference on Knowledge Engineering and Knowledge Management,* EKAW 2004, 5-8 October 2004, Whittlebury Hall, Northamptonshire, UK, Springer-Verlag, 2004a.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Barbulescu M., *A University Research Group Experience with Deploying an Artificial Intelligence Application in the Center of Gravity Analysis Domain*, in *Proceedings of AAAI-2004 Workshop on Fielding Applications of Artificial Intelligence,* July 25, San Jose, AAAI Press, Menlo Park, CA, 2004b.