Training and Using DISCIPLE Agents

A Case Study in the Military Center of Gravity Analysis Domain

Gheorghe Tecuci, Mihai Boicu, Dorin Marcu, Bogdan Stanescu, Cristina Boicu, and Jerome Comello

■ This article presents the results of a multifaceted research and development effort that synergistically integrates AI research with military strategy research and practical deployment of agents into education. It describes recent advances in the DISCIPLE approach to agent development by subject-matter experts with limited assistance from knowledge engineers, the innovative application of DISCIPLE to the development of agents for the strategic center of gravity analysis, and the deployment and evaluation of these agents in several courses at the U.S. Army War College.

This article presents the results of a multiobjective collaboration between the Learning Agents Laboratory of George Mason University, on the one side, and the Center for Strategic Leadership and the Department of Military Strategy, Planning, and Operations of the U.S. Army War College, on the other side. A distinguishing feature of this collaboration is the synergistic integration of AI research with military strategy research and the practical use of agents in education, as detailed in the following.

The AI research objective is the development of the DISCIPLE approach for building instructable knowledge-based systems or agents (Tecuci 1998, 1988). The DISCIPLE approach advocates the creation of a powerful learning agent shell that can be taught by a person to solve problems in a way similar to how that person would teach a student or an assistant.

We think that the DISCIPLE approach contributes directly to a new age in the software systems development process, as illustrated in figure 1. In the mainframe computers age, the software systems were both built and used by computer science experts. In the current age of personal computers, these systems are still being built by computer science experts, but many of them (such as text processors, electronic-mail programs, or internet browsers) are now used by persons that have no formal computer education. Continuing this trend, we think that the next age will be that of the personal agents, where typical computer users will be able to both develop and use special types of software agents (Tecuci, Boicu, and Marcu 2000). The DISCIPLE approach attempts to change the way intelligent agents are built, from "being programmed" by a knowledge engineer to "being taught" by a user who does not have prior knowledge engineering or computer science experience. This approach would allow a typical computer user, who is not a trained knowledge engineer, to build by himself/herself an intelligent assistant as easily as he/she now uses a word processor to write a paper.

Over the years, we have developed a series of increasingly advanced learning agent shells forming the DISCIPLE family. The most recent family member, DISCIPLE-RKF, represents a significant advancement over its most recent predecessors: DISCIPLE-WA (Tecuci et al. 1999) and DISCIPLE-COA (Tecuci et al. 2001). All three systems were developed as part of the High Performance Knowledge Bases Program and

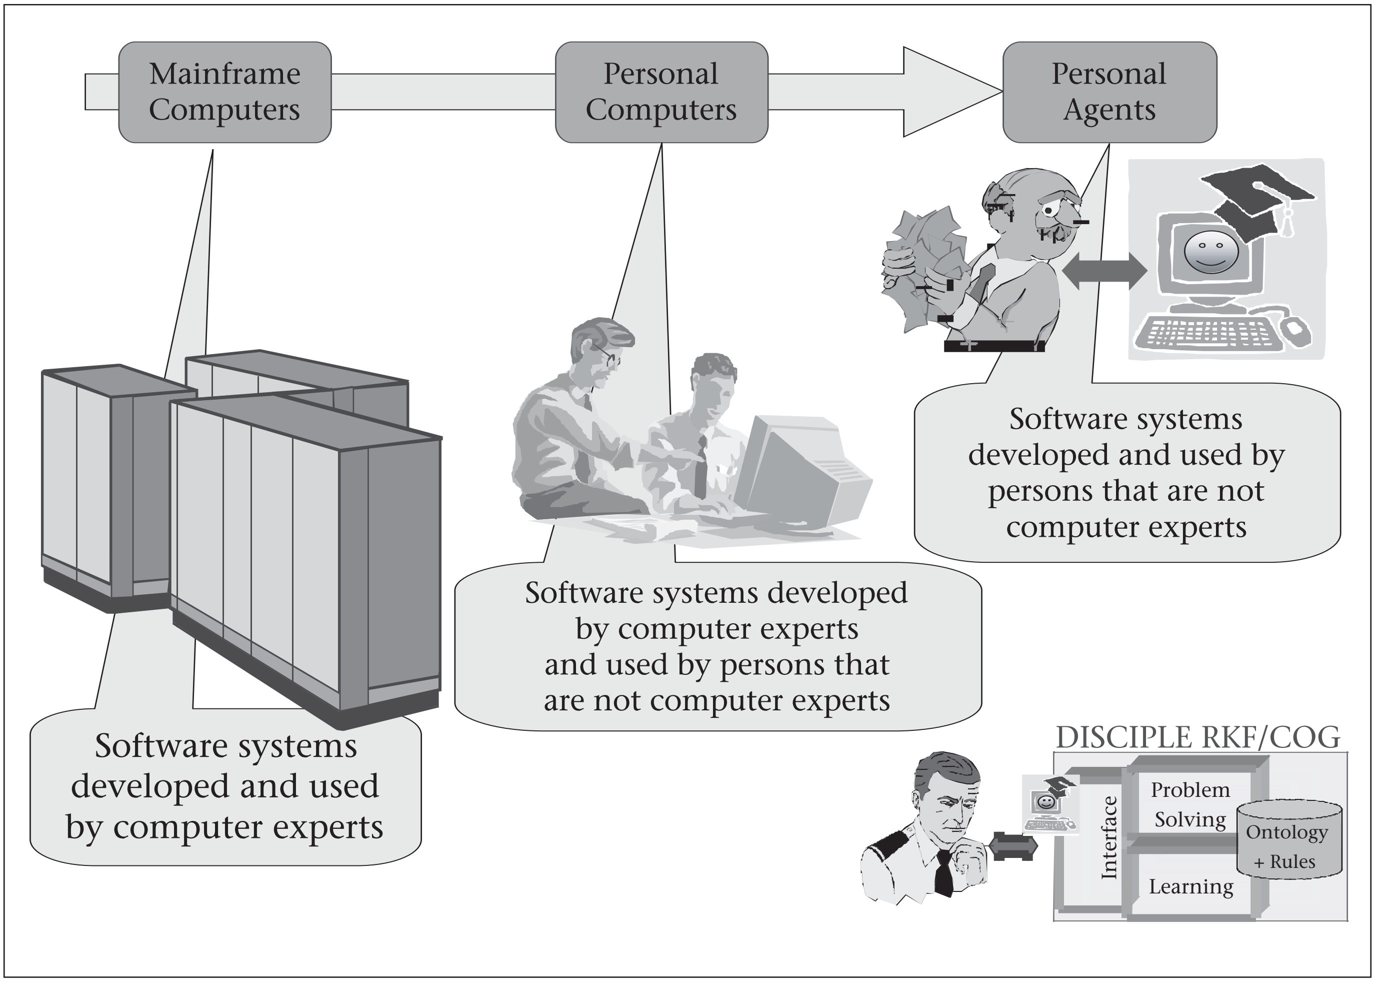

Figure 1. View on the Evolution of the Software Development Process.

the Rapid Knowledge Formation Program, supported by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Office of Scientific Research (AFOSR).[1] Both programs emphasized the use of innovative challenge problems to focus and evaluate the research and development efforts. The challenge problem for the DISCIPLE-RKF system is the strategic center of gravity analysis, which brings us to the second objective of this effort, the military strategy research objective of clarifying and formalizing the center of gravity analysis process by using the general task-reduction paradigm of problem solving. The concept of the center of gravity of an entity (state, alliance, coalition, or group) was introduced in the nineteenth century by Karl von Clausewitz (1976) as the foundation of capability, "the hub of all power and movement, on which everything depends,… the point against which all the energies should be directed" (595–596).

Correctly identifying the centers of gravity of the opposing forces is of highest importance in any conflict. Therefore, in the education of strategic leaders at all the United States senior military service colleges, there is a great emphasis on the center of gravity analysis (Strange 1996). Hence, we have the third objective of this research, the educational objective of enhancing the educational process of senior military officers through the use of intelligent agent technology. Using the DISCIPLE approach, we have developed intelligent agents for strategic center of gravity analysis that are used in several courses at the U.S. Army War College. In the Case Studies in Center of Gravity Analysis course, the students (who are high-ranking military officers, from lieutenant colonels to generals) use a DISCIPLE agent that was taught some of the instructor's expertise in center of gravity analysis. The students use DISCIPLE as an intelligent assistant that supports them both in learning about the center of gravity analysis concept and in developing a center of gravity analysis report for a war scenario. In the follow-on Military Applications of Artificial Intelligence course, the students use personal DISCI-

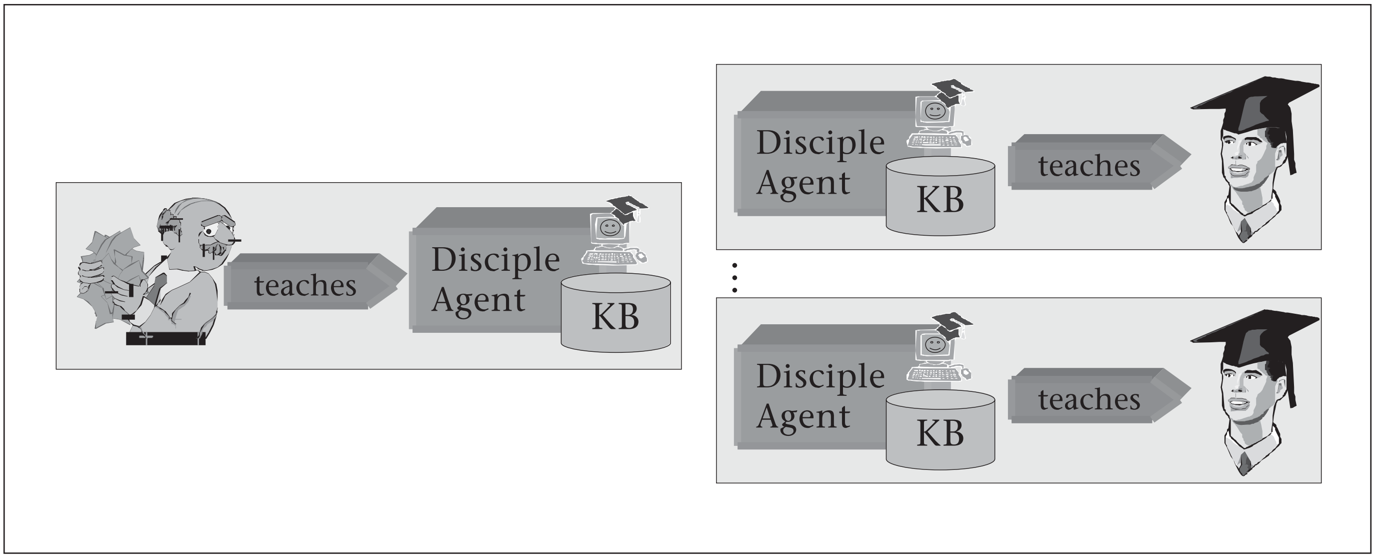

Figure 2. View on the Future Use of Instructable Agents in Education.

PLE agents as subject-matter experts, teaching them their own problem-solving expertise in center of gravity analysis.

The DISCIPLE approach is particularly relevant to education, figure 2 illustrating our long-term research vision in this area. As shown on the left-hand side of figure 2, a teacher teaches a DISCIPLE agent through examples and explanations, in a way that is similar to how the teacher would teach a student. The DISCIPLE agent can then be used as a personal tutor, teaching the students in a way similar to how it was taught by the teacher (Hamburger and Tecuci 1998; Tecuci and Keeling 1999).

Each of the three objectives discussed earlier is recognized as important and difficult in its own right. Our experience with addressing them together in a synergistic manner has resulted in faster progress for each of them. Moreover, it offers a new perspective on how to combine research in AI with research in a specialized domain and with the development and deployment of prototype systems in education and practice.

The rest of this article presents the current status of this research and development effort. The next section presents in more detail the center of gravity challenge problem. This discussion is followed by an end user perspective on a developed DISCIPLE agent for center of gravity analysis, called DISCIPLE-RKF/COG, which is used in the Case Studies in Center of Gravity Analysis course at the U.S. Army War College. The following section presents an overview of the DISCIPLE-RKF shell and its use to build the DISCIPLE-RKF/COG agent, emphasizing its new capabilities with respect to the previous DISCIPLE shells. This section also discusses the deployment and evaluation of DISCIPLE in the Military Applications of Artificial Intelligence course. The article concludes with a summary of the synergistic aspects of this collaborative work and future research directions.

The Center of Gravity Problem

Military literature distinguishes between three levels of conflicts: (1) a strategic level focusing on winning wars, (2) an operational level focusing on winning campaigns, and (3) a tactical level focusing on winning battles. One of the most difficult problems that senior military leaders face at the strategic level is the determination and analysis of the centers of gravity for friendly and opposing forces. Originally introduced by Clausewitz in his classical work *On War* (1976), the center of gravity is now understood as representing "those characteristics, capabilities, or localities from which a military force derives its freedom of action, physical strength, or will to fight" (Joint Chiefs of Staff 2001). The force's goal should be to eliminate or influence the enemy's strategic center of gravity yet adequately protect its own.

Center of gravity determination 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. Therefore, when performing center of gravity analysis, experts rely on their own professional experience and intuitions, without following a rigorous approach. Recog-

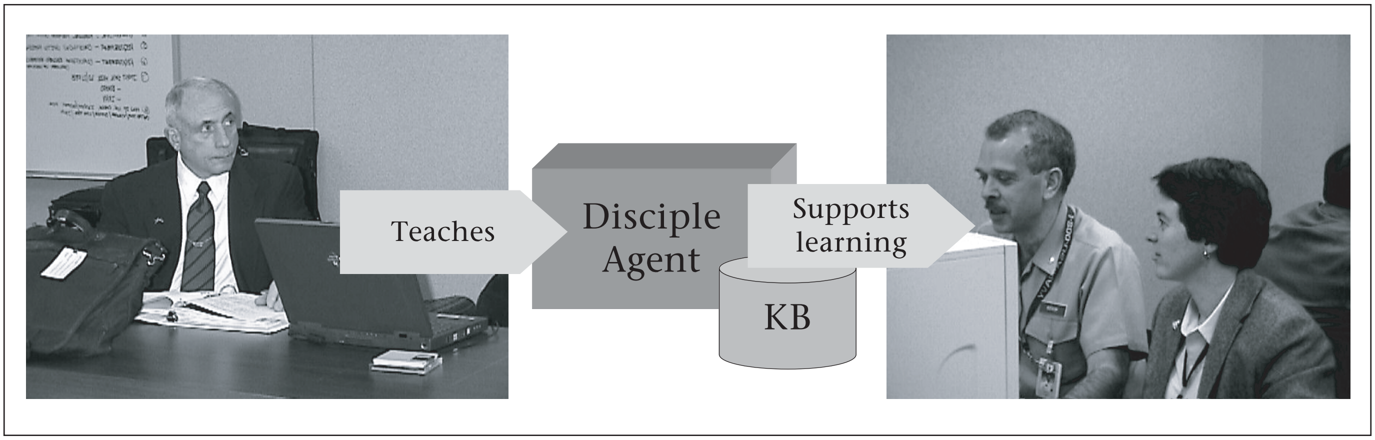

Figure 3. A Step toward the Vision of Using Instructable Agents in Education.

nizing these difficulties, the Center for Strategic Leadership of the U.S. Army War College started, in 1993, an effort to elicit and formalize the knowledge of a number of experts in center of gravity analysis. This research resulted in a monograph on center of gravity analysis,[2] which provided a basis for the application of DISCIPLE to this high-value application domain and for the development of the DISCIPLE-RKF/COG instructable agent presented in the next section.

A DISCIPLE Agent for Center of Gravity Analysis

DISCIPLE-RKF/COG is an agent used in the U.S. Army War College course entitled Case Studies in Center of Gravity Analysis. The use of DISCIPLE in this course is a step toward the vision illustrated in figure 2 on the use of instructable agents in education. Indeed, as shown in figure 3, we have worked with the course's instructor to teach a DISCIPLE agent some of his expertise in center of gravity analysis. Then, DISCIPLE helped the students learn to perform a center of gravity analysis of an assigned war scenario, as discussed next.

First, DISCIPLE guides the student to identify, study, and describe the aspects of a campaign (such as the 1945 U.S. invasion of the island of Okinawa) that are relevant for center of gravity analysis. The student-agent interaction takes place as illustrated in figure 4. The left part of the window is a table of contents, whose elements indicate various aspects of the scenario. When the student selects one such aspect, DISCIPLE asks specific questions intended to acquire from the student a description of this aspect or update a previously specified description. All the answers are in natural language.

Taking the Okinawa_1945 scenario as our example, DISCIPLE starts by asking for a name and a description of the scenario and then the names of the opposing forces. Once the student indicates Japan_1945 and US_1945 as the opposing forces, DISCIPLE includes them in the table of contents, together with general characteristics that the student can specify (see the left-hand side of figure 4). The student can then click on any of these aspects (for example, "industrial capacity" under "economic factors" of Japan_1945), and the agent guides the student in specifying it. The student's specification can prompt additional questions from DISCIPLE and a further expansion of the table of contents. An orange, yellow, or white circle marks each title in the table of contents, indicating respectively that all, some, or none of the corresponding questions of DISCIPLE have been answered. However, the student is not required to answer all the questions.

DISCIPLE can be asked, at any time, to identify and test the strategic center of gravity candidates for the current specification of the scenario. Figure 5 shows the center of gravity solution viewer. Its left-hand side contains the list of the center of gravity candidates identified by DISCIPLE for each of the opposing forces in the Okinawa_1945 scenario. For Japan_1945, they are the will of the people of Japan, Emperor Hirohito, the Japanese Imperial General Staff, the military of Japan, and the industrial capacity of Japan. When a candidate is selected in the left-hand side of the viewer, its (abstract or detailed) justification for identification and testing will be displayed in the right-hand side of the viewer. The top part of figure 5 shows the abstract justification for the identification of Emperor Hirohito as a strategic center of gravity candidate. The bottom part of the figure shows the testing of this candidate. DISCIPLE uses the task-reduction paradigm to generate these justifications. It starts with the top-level problem-solving task of identifying

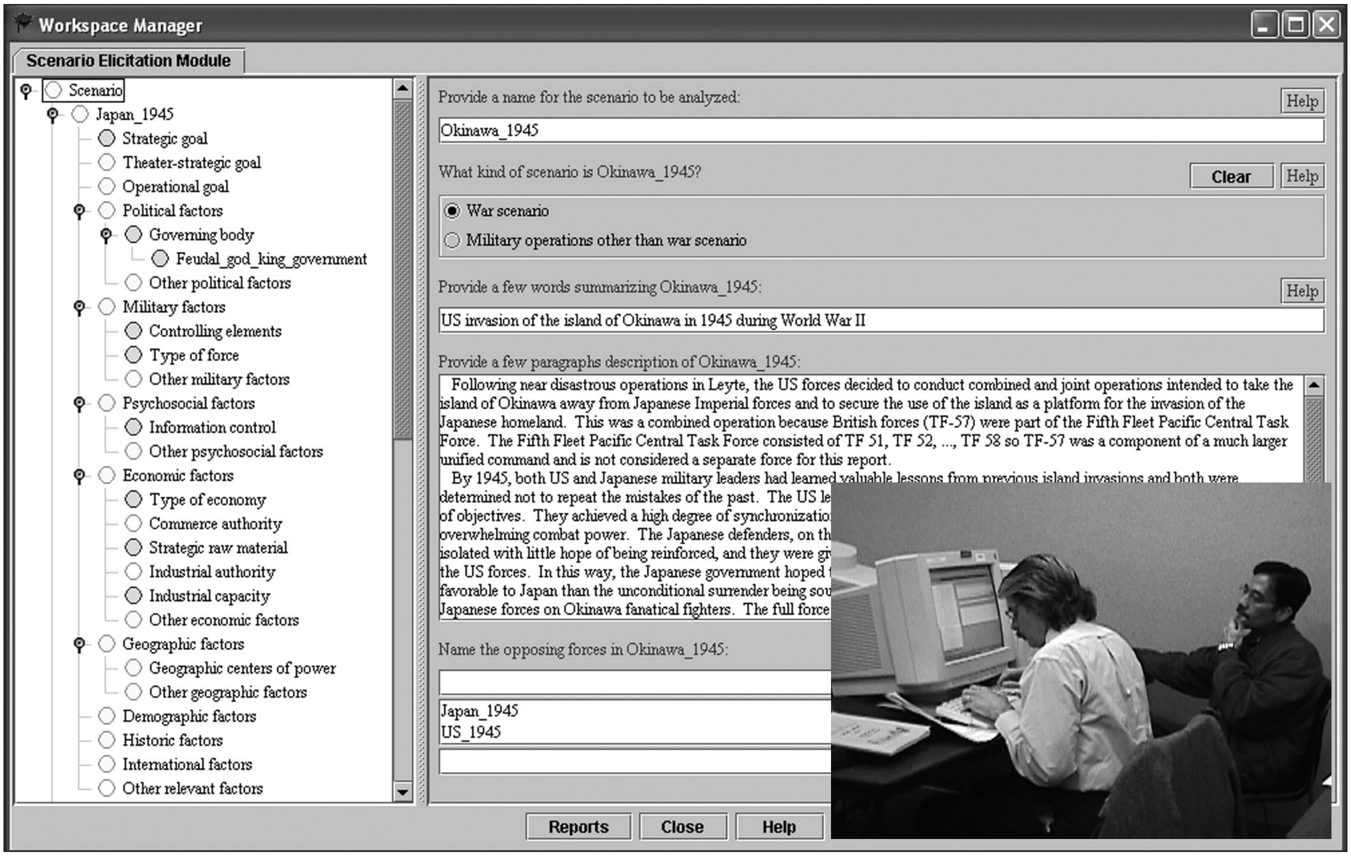

Figure 4. Scenario Specification Interface.

and testing a strategic center of gravity candidate. To perform this task, DISCIPLE asks itself a series of questions. The answer to each question allows DISCIPLE to reduce the current task to simpler ones, until DISCIPLE has enough information to first identify a strategic center of gravity candidate and then to test it, determining whether it should be eliminated.

The abstract justifications shown in the right-hand side of figure 5 are obtained by keeping only the sequence of questions and answers from the detailed justification (that is, by eliminating the task names). Notice that Emperor Hirohito is identified as a strategic center of gravity candidate for Japan_1945 in the Okinawa_1945 scenario because he is the main controlling element of the government of Japan, having a critical role in setting objectives and making decisions. After being identified as a candidate, Emperor Hirohito is analyzed based on various elimination tests, but he passes all of them. Because Japan_1945 has a feudal god-king government and Emperor Hirohito is its god-king, he could make the government accept the unconditional surrender of Japan, which is the main strategic goal of the United States. As commander in chief of the military, he can also impose his will on the military of Japan. Finally, he could also make the people of Japan accept unconditional surrender. Being able to impose his will on the Clausewitz's trinity of power (government, military, and people), Emperor Hirohito was very likely to be the strategic center of gravity of Japan in 1945.

As another example, consider the industrial capacity of Japan_1945, which is another source of strength, power, and resistance because it produces the war materiel and transports of Japan. DISCIPLE, however, eliminates this strategic center of gravity candidate because the military and the people of Japan_1945 are determined to fight to the death and not surrender even with diminished war materiel and transports.

In the example scenario portrayed here, DISCIPLE eliminates all but two candidates for Japan—Emperor Hirohito and the Japanese Imperial General Staff—and suggests that the student should select one of them as the strategic center of gravity of Japan in 1945. It is important to point out that this example is only one possible approach to the analysis of Japan's center of gravity for the Okinawa cam-

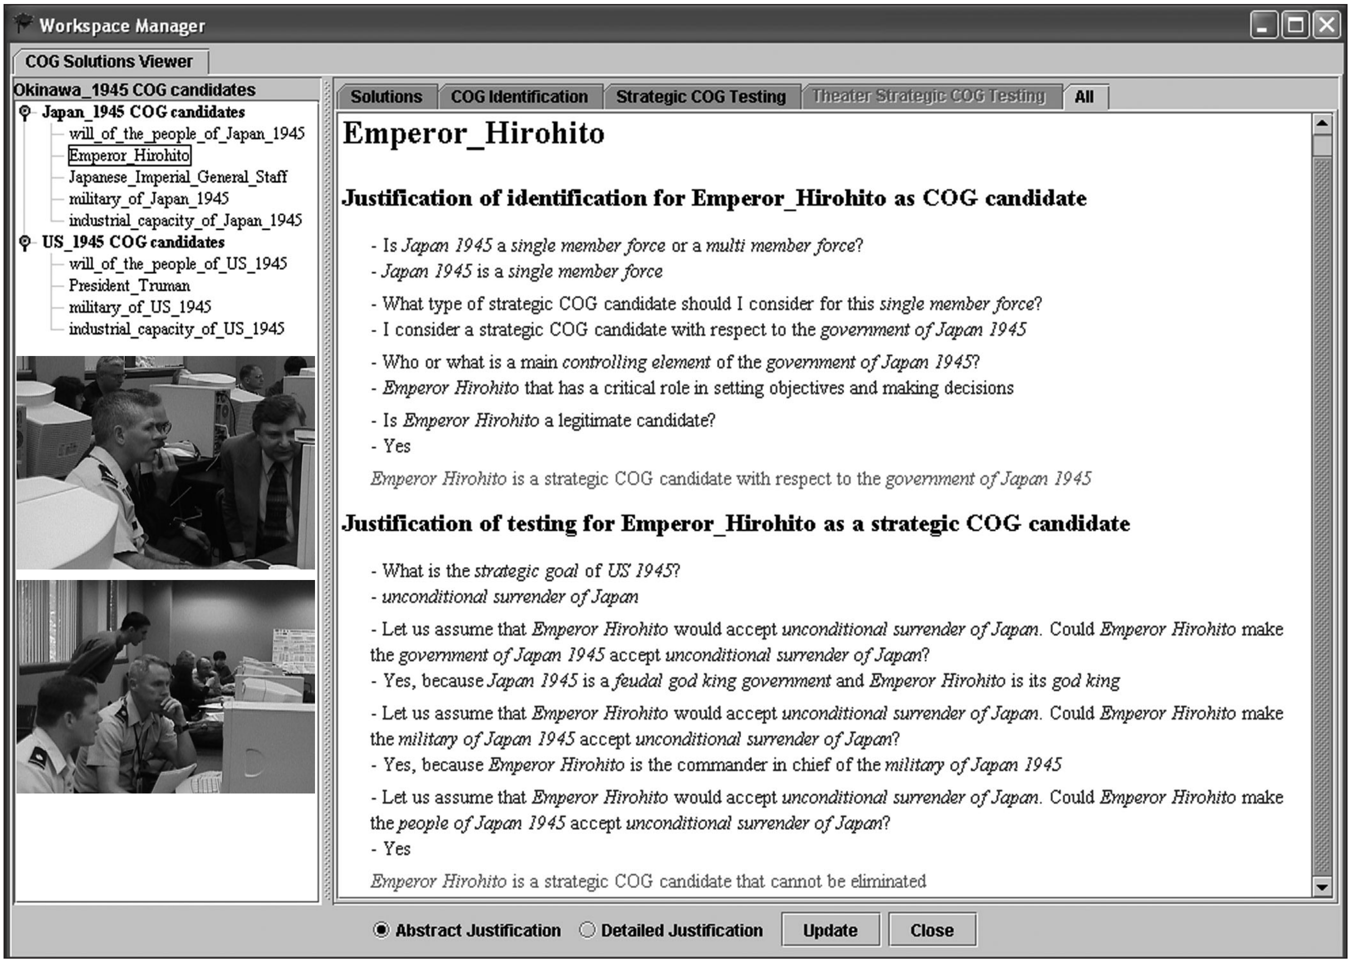

Figure 5. The Problem Solving Interface of DISCIPLE-RKF/COG.

paign. We recognize that subject-matter experts often differ in their judgments about the identification and analysis of center of gravity candidates for any particular scenario.

As illustrated, DISCIPLE guides the student to identify, study, and describe the relevant aspects of the opposing forces in a particular scenario. Then DISCIPLE identifies and tests the strategic center of gravity candidates, as illustrated in figure 5. After that, DISCIPLE generates a draft analysis report, a fragment of which is shown in figure 6. The first part of this report contains a description of the scenario, being generated by DISCIPLE based on the information elicited from the student. The second part of the report includes all the center of gravity candidates identified by DISCIPLE, together with their justifications for identification and testing. The student must now finalize this report by examining each of the center of gravity candidates and their justifications, completing, correcting, or even rejecting DISCIPLE's reasoning and providing an alternative line of reasoning. This process is productive for several reasons. First, the agent generates its proposed solutions by applying general reasoning rules and heuristics learned previously from the course's instructor to a new scenario described by the student. Second, center of gravity analysis is influenced by personal experiences and subjective judgments, and the student (who has unique military experience and biases) might have a different interpretation of certain facts.

This requirement for the critical analysis of the solutions generated by the agent is an important educational component of military commanders that mimics military practice. Commanders have to critically investigate several courses of action proposed by their staff and make the final decision on which one to use.

Use of DISCIPLE in the Center of Gravity Course

Successive versions of DISCIPLE have been used in both the winter and spring sessions of the Case Studies in Center of Gravity Analysis

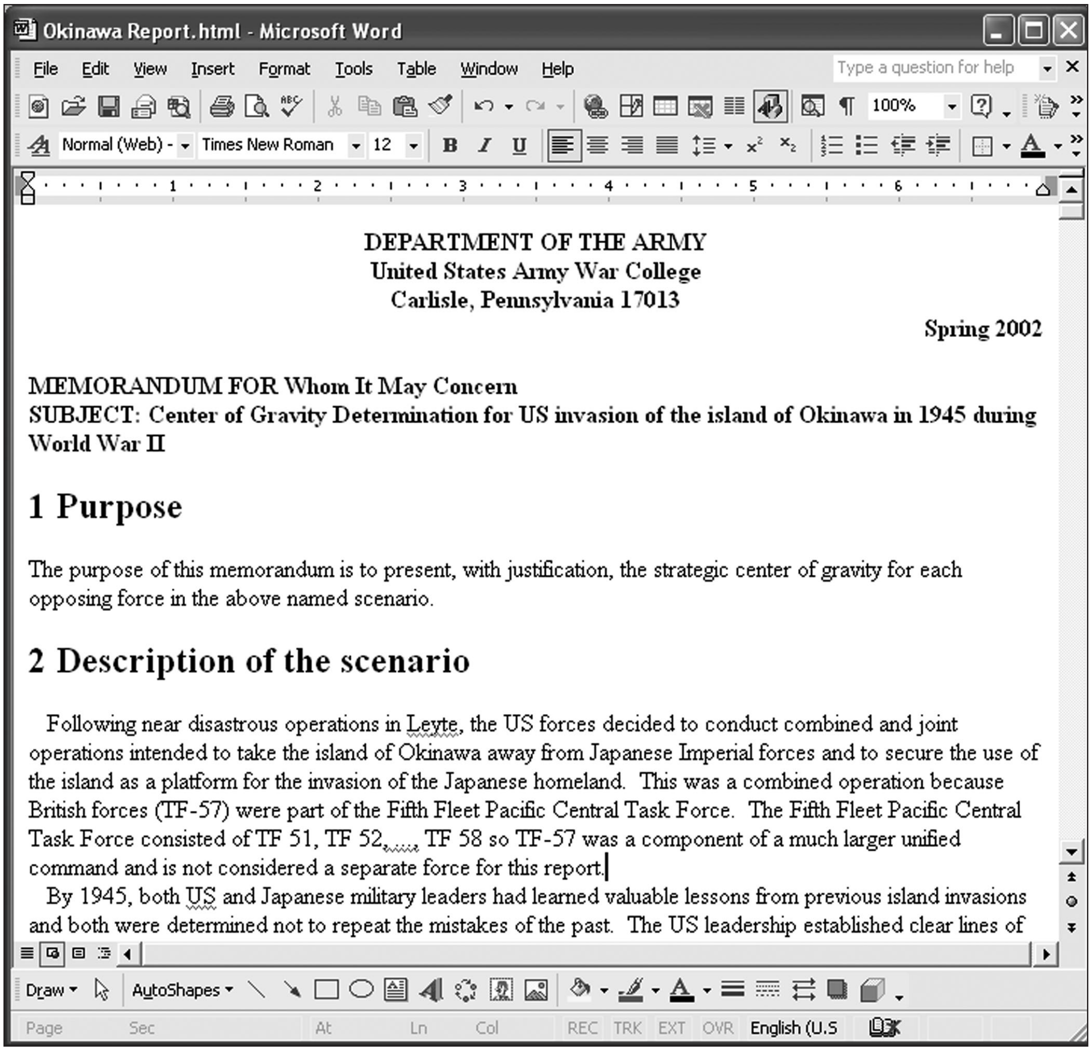

DEPARTMENT OF THE ARMY
United States Army War College
Carlisle, Pennsylvania 17013

Spring 2002

MEMORANDUM FOR Whom It May Concern
SUBJECT: Center of Gravity Determination for US invasion of the island of Okinawa in 1945 during World War II

1 Purpose

The purpose of this memorandum is to present, with justification, the strategic center of gravity for each opposing force in the above named scenario.

2 Description of the scenario

Following near disastrous operations in Leyte, the US forces decided to conduct combined and joint operations intended to take the island of Okinawa away from Japanese Imperial forces and to secure the use of the island as a platform for the invasion of the Japanese homeland. This was a combined operation because British forces (TF-57) were part of the Fifth Fleet Pacific Central Task Force. The Fifth Fleet Pacific Central Task Force consisted of TF 51, TF 52,…, TF 58 so TF-57 was a component of a much larger unified command and is not considered a separate force for this report.

By 1945, both US and Japanese military leaders had learned valuable lessons from previous island invasions and both were determined not to repeat the mistakes of the past. The US leadership established clear lines of

Figure 6. The Report Generated by DISCIPLE-RKF/COG.

course during the past two academic years and will continue to be used in the future. The attendance of these courses was as follows: 10 students in the winter 2001 session (7 U.S. officers and 3 international fellows), 3 students in the spring 2001 session (1 U.S. officer and 2 international fellows), 13 students in the winter 2002 session (11 U.S. officers and 2 international fellows), and 10 students in the spring 2002 session (2 U.S. officers and 8 international fellows). The students were lieutenant colonels, colonels, or generals from all the military services. At the end of each course, the students completed detailed evaluation forms about DISCIPLE and its modules, addressing many issues ranging from judging its usefulness in achieving course's objectives to judging its methodological approach to problem solving to judging the ease of use and other aspects of various modules. As the capabilities of the used DISCIPLE agents evolved, the evaluation questions also evolved. The following, for example, are some of the evaluations of the 13 students from the Winter 2002 session, which are generally representative of the evaluations from all the other sessions. On a five-point scale (strongly disagree, disagree, neutral, agree, strongly agree), nine students agreed and four strongly agreed that "the use of DISCIPLE is an assignment that is well suited to the course's learning objectives." One student was neutral, but nine agreed, and three strongly agreed with the statement that "DISCIPLE helped me to learn to perform a strategic center of gravity analysis of a scenario." One student disagreed, but four students agreed and eight strongly agreed that "the use of DISCIPLE was a useful learning experience." Finally, one student disagreed, nine students agreed, and three strongly agreed that "DISCIPLE should be used in future versions of this course."

To our knowledge, this is the first time that

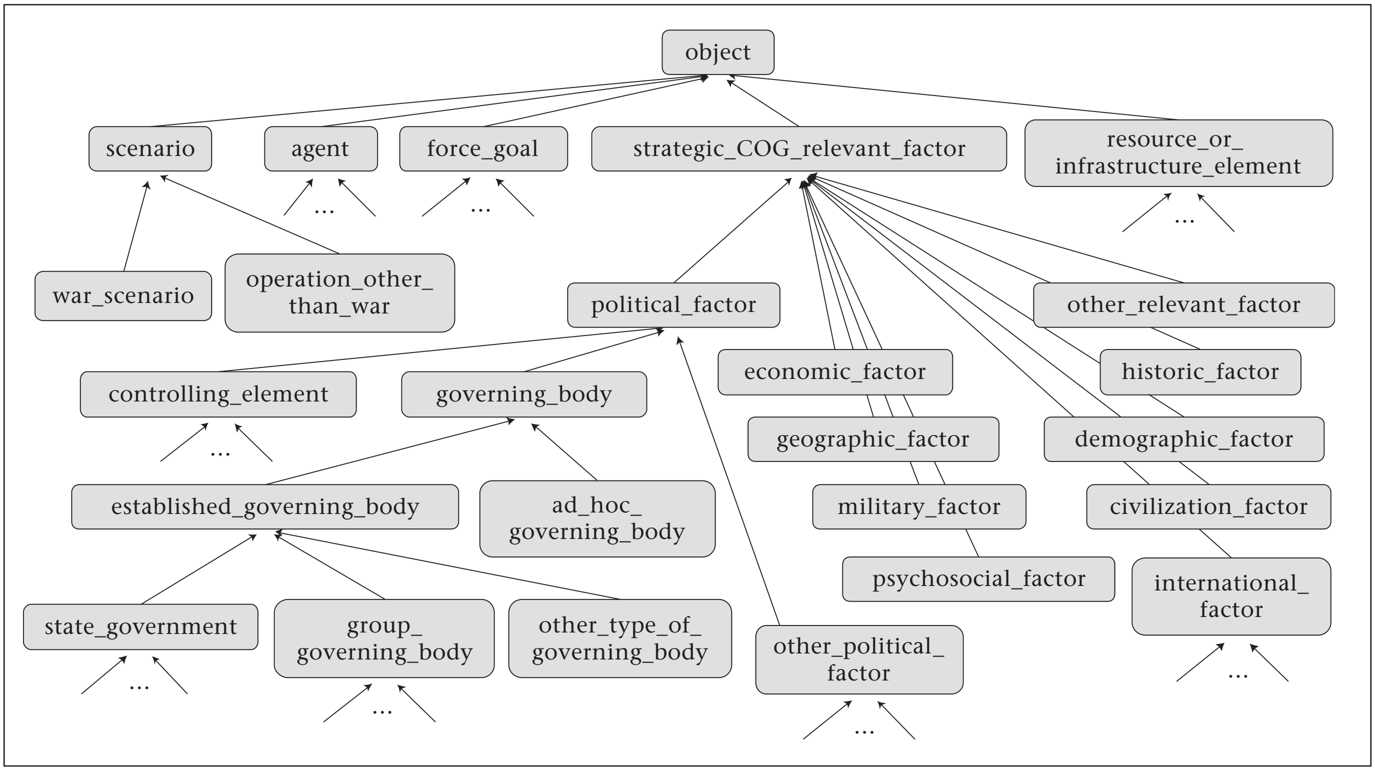

Figure 7. A Fragment of the Center of Gravity Object Ontology.

intelligent agents for the strategic center of gravity identification and testing have been developed and used. The next section discusses the development of these agents and their use in the Military Applications of Artificial Intelligence courses at the U.S. Army War College.

Agent Development with DISCIPLE-RKF

The DISCIPLE-RKF/COG agent presented in the previous section was developed using the DISCIPLE-RKF learning agent shell, as we describe in this section. DISCIPLE-RKF consists of an integrated set of knowledge acquisition, learning, and problem-solving modules for a generic knowledge base having two main components: (1) an object ontology that defines the terms from a specific application domain and (2) a set of task-reduction rules expressed with these terms. DISCIPLE-RKF represents a significant evolution compared to the previous DISCIPLE shells. It implements more powerful knowledge representation and reasoning mechanisms and has an improved interface that facilitates mixed-initiative reasoning. Even more significantly, DISCIPLE-RKF incorporates new modules that allow a subject-matter expert to perform additional knowledge engineering tasks, such as scenario specification, modeling of his/her problem-solving process, and task formalization.

In general, the process of developing a specific knowledge-based agent with DISCIPLE-RKF consists of two major stages: (1) the development of the object ontology by the knowledge engineer and the subject-matter expert and (2) the training of DISCIPLE by the subject-matter expert.

In the first development stage, a knowledge engineer works with a subject-matter expert to specify the type of problems to be solved by the DISCIPLE agent, clarify how these problems could be solved using DISCIPLE's task-reduction paradigm, and develop an object ontology.

The object ontology consists of hierarchical descriptions of objects and features, represented as frames, as in the knowledge model of the open knowledge base connectivity protocol (Chaudhri et al. 1998). An object hierarchy fragment from the center of gravity domain is shown in figure 7, and a feature hierarchy fragment is shown in figure 8. The careful design and development of the object ontology is of utmost importance because it is used by DISCIPLE as its generalization hierarchy for learning. DISCIPLE-RKF includes a suite of ontology modules, such as tree-based and graph-based browsers and viewers (that allow an easy and intuitive navigation of the ontology) and editors (used to develop and maintain the ontology).

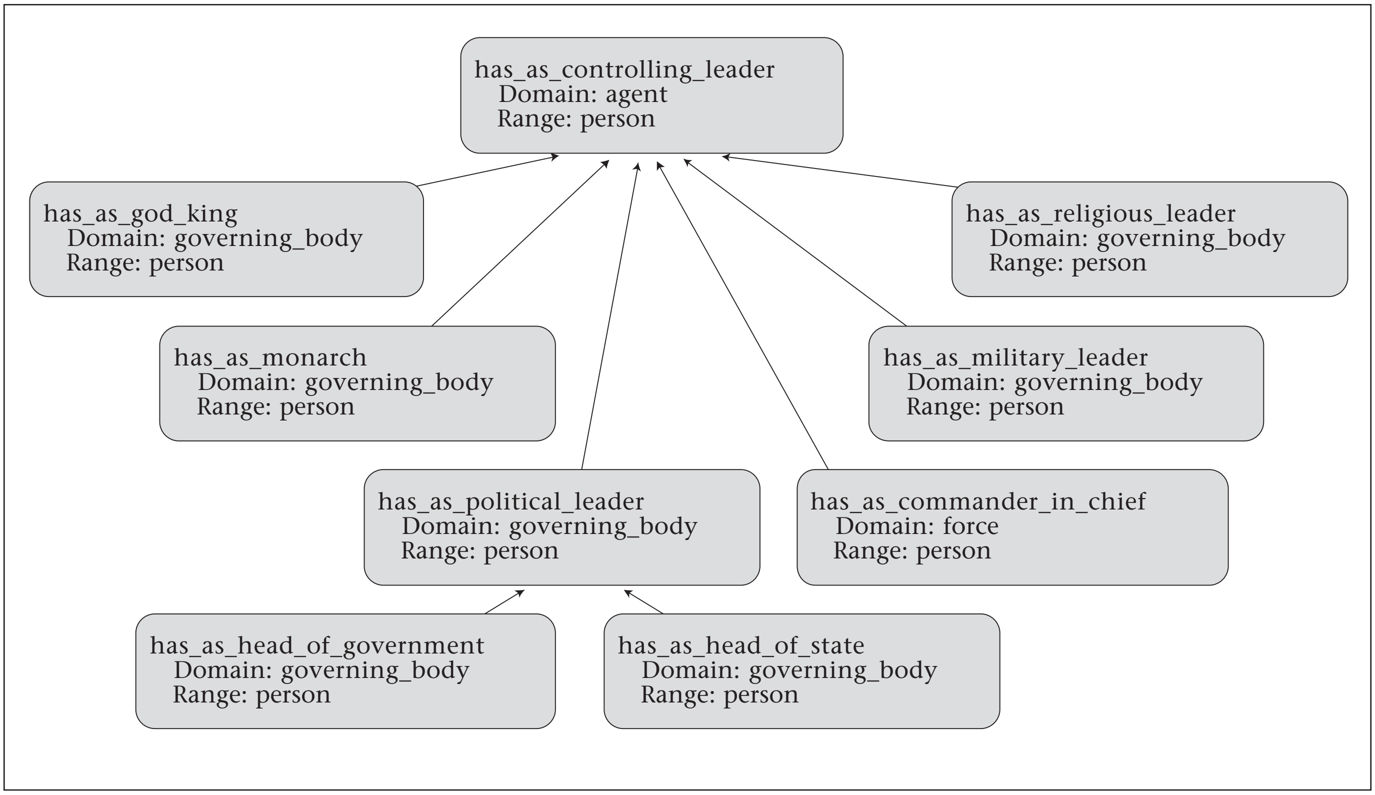

Figure 8. A Fragment of the Center of Gravity Feature Ontology.

A new capability of DISCIPLE-RKF is the ability to define elicitation scripts for objects and features. These scripts guide the expert to define the instances that occur in a scenario (such as Okinawa_1945 or Emperor Hirohito, as illustrated before). Figure 9 shows three elicitation scripts associated with the "scenario" object. The top script specifies what the question is to be asked by DISCIPLE to elicit the name of the scenario, how the user's answer should be used to update the ontology, what other scripts should be called after updating the ontology, and even what the appearance of the interface is. The use of the elicitation scripts allows a knowledge engineer to rapidly build customized interfaces for DISCIPLE agents, such as the one illustrated in figure 4, thus effectively transforming this software development task into a knowledge engineering one.

The result of the first development stage is a customized DISCIPLE agent. In the second major stage of agent development, this agent is trained to solve problems by a subject-matter expert, with limited assistance from a knowledge engineer. Figure 10 shows the main phases of the agent training process, which starts with a knowledge base that contains only a general object ontology (but no instances, no problem-solving tasks, and no task-reduction rules) and ends with a knowledge base that incorporates expert problem-solving knowledge.

During the *scenario specification* phase, the scenario specification module (which is a new module of DISCIPLE-RKF) guides the expert in describing the objects that define a specific strategic scenario (for example, the U.S. invasion of the island of Okinawa in 1945). The expert does not work directly with the object ontology to specify the scenario. Instead, the expert-agent interaction takes place as illustrated in figure 4, being directed by the execution of the elicitation scripts. Experimental results show that the experts can easily perform this task.

After the expert has specified the Okinawa_-1945 scenario, he/she can start the *modeling* of his/her center of gravity reasoning for this particular scenario as a sequence of task-reduction steps. The expert expresses his/her reasoning in English, similarly to how he/she would think aloud as he/she would a problem, as illustrated in table 1. First, the expert formulates the top-level problem-solving task. To perform this task, the expert asks himself/herself a series of questions. The answer to each question allows the expert to reduce the current task to a simpler one. This process continues until the expert has enough information to first identify a strategic center of gravity candidate and then to determine whether it should be eliminated.

Experimental results show that this agent

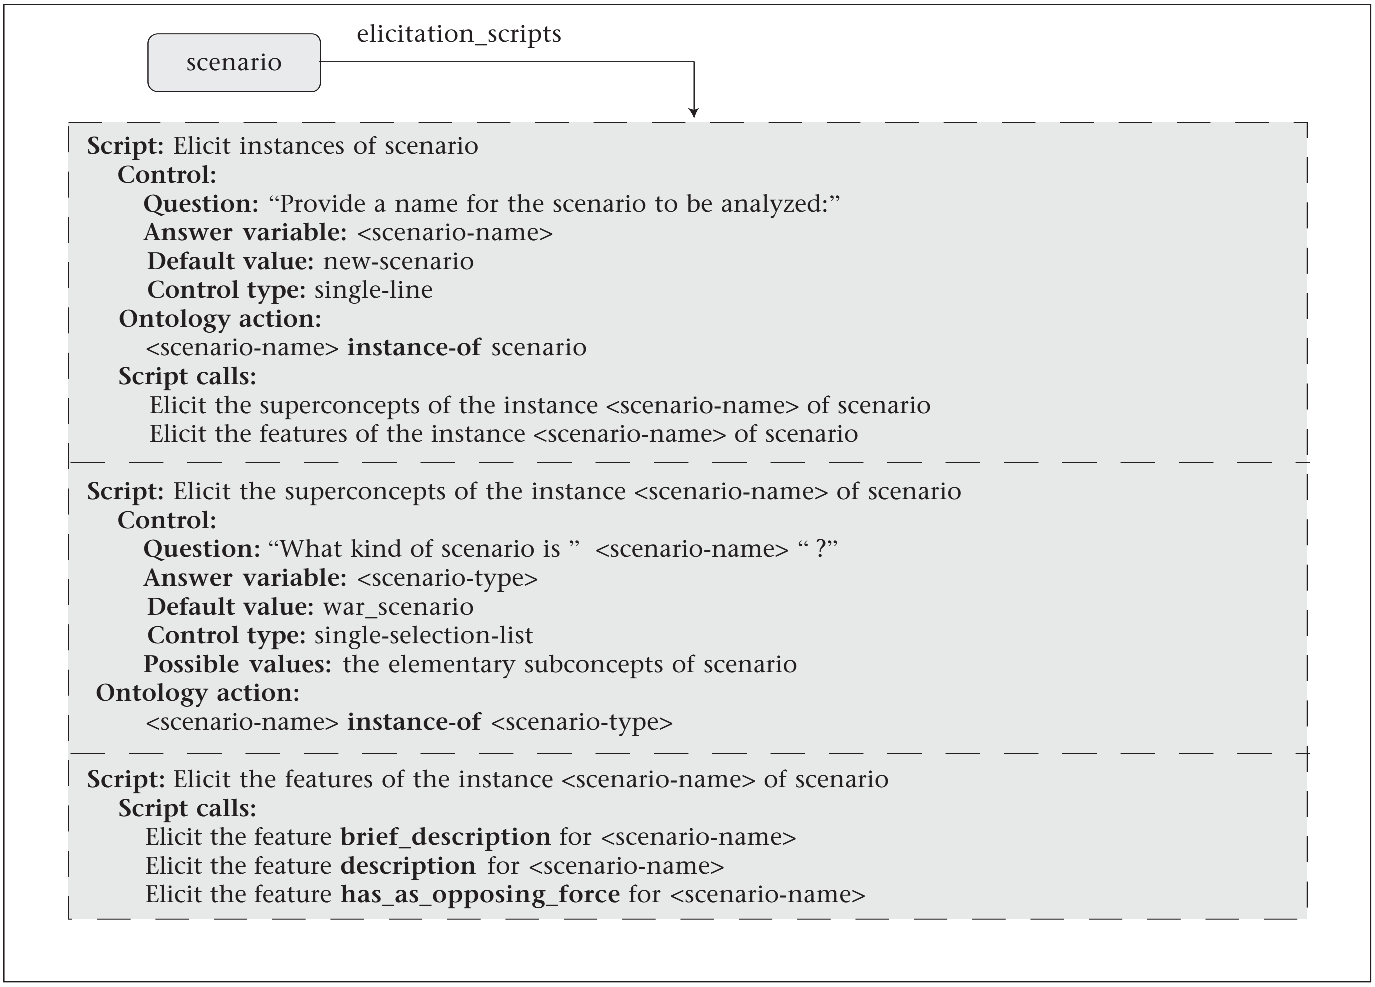

Figure 9. Sample Elicitation Scripts.

I need to
Identify and test a strategic COG candidate for the Okinawa_1945 scenario.
What kind of scenario is Okinawa 1945?
Okinawa 1945 is a war scenario.
Therefore I need to
Identify and test a strategic COG candidate for the Okinawa_1945 which is a war scenario.
Which is an opposing force in the Okinawa 1945 scenario?
Japan 1945
Therefore I need to
Identify and test a strategic COG candidate for Japan_1945.
…

Table 1. Sample Modeling of the Center of Gravity Analysis Process for a Specific Scenario.

training activity is the most challenging for the expert. We have therefore recently developed a modeling adviser to help the expert in this activity. Figure 11 shows the interface of this new module of DISCIPLE-RKF. The middle part of the screen contains the current task-reduction step that the expert is composing. At each state in this process, the right-hand side of the screen shows all the actions that could be performed in this state, and the left-hand side shows the action that the modeling adviser is actually recommending. For example, to specify the current subtask, the adviser suggested that the expert copy and modify the task. The modeling adviser can also suggest the question to be asked or the answer of the question. As mentioned, the expert expresses his/her reasoning in English. However, each time he/she starts to type a word, the agent lists on the left-hand side of the screen all the instances and concepts from the knowledge base that are consistent with the characters typed so far. This approach is useful for two different reasons: (1) it

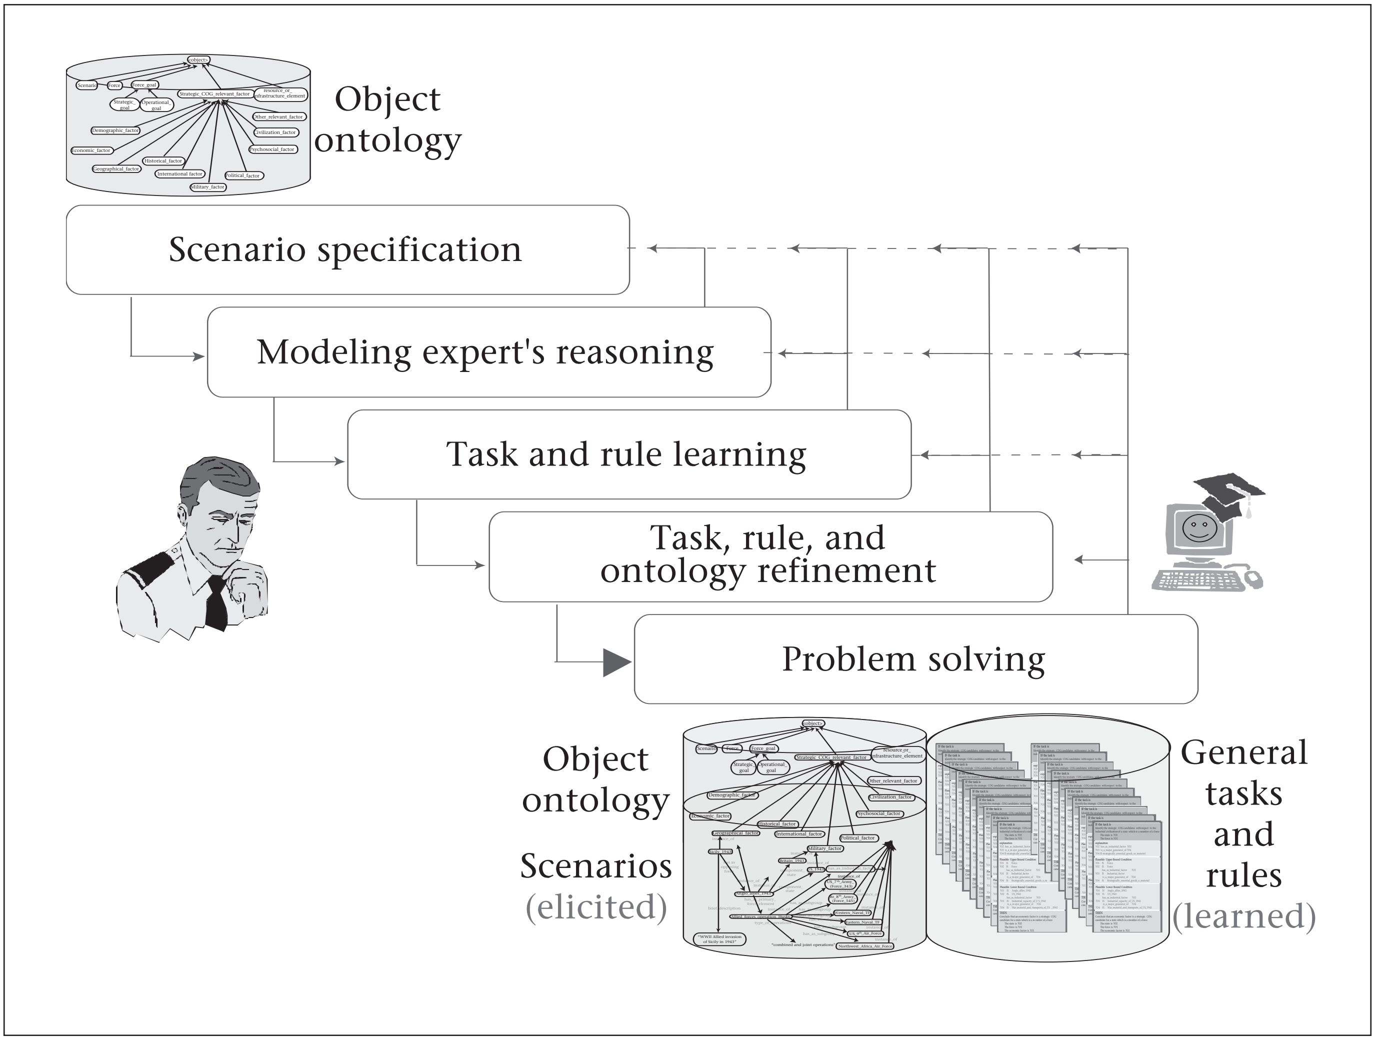

Figure 10. The Main Phases of the Agent Training Process.

facilitates the user's input and (2) it helps the agent to "understand" his/her phrases.

In the *task-learning and rule-learning* phase, DISCIPLE learns general tasks and general rules from the task-reduction steps defined in the modeling phase. For example, consider the reduction step from the middle of figure 11, shown again on the left-hand side of figure 12. It consists of a task, a question, an answer, and a subtask. Because all these expressions are in natural language, the expert and the agent collaborate to translate them into the formal logical expressions from the right-hand side of figure 12. First, the natural language expression of each task is structured into an abstract phrase called the *task name*, which does not contain any instance or constant, and several specific phrases representing the task's features. The formalization is proposed by the agent and can be modified by the expert. Next, the expert and the agent collaborate to also formalize the question and the answer from the left-hand side of figure 12 into the explanation from the right-hand side of figure 12. This explanation represents the best approximation of the meaning of the question-answer pair that can be formed with elements from the object ontology. In essence, the agent will use analogical reasoning and guidance from the expert to propose a set of plausible explanation pieces from which the expert will select the most appropriate ones (Tecuci et al. 2001).

Based on the formalizations from figure 12 and the object ontology from figures 7 and 8, the DISCIPLE agent learns the general task shown in figure 13 and the general rule shown in figure 14. Both the learned task and the learned rule have an informal structure, shown at the top of figure 13 and figure 14, respectively. They also have a formal structure, shown at the bottom of figure 13 and figure 14, respectively. The informal structure preserves the natural

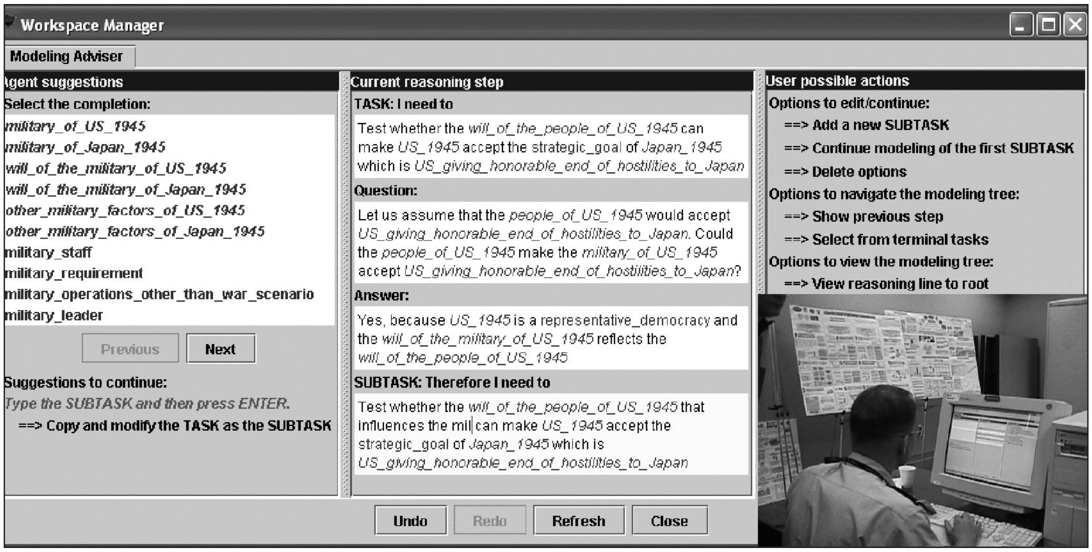

Figure 11. The Modeling Adviser Interface.

language of the expert and is used in agent-user communication. The formal structure is used in the actual reasoning of the agent.

Initially, when the agent has no rules and no tasks, the expert teaches DISCIPLE how to solve problems, and DISCIPLE generates partially learned tasks and rules, as indicated earlier. As DISCIPLE learns 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.

The learned formal rule in figure 14 includes two applicability conditions, a plausible upper-bound condition, and a plausible lower-bound condition. The plausible upper-bound condition results from a maximal generalization of the example and its explanation from figure 12. This condition allows the rule to be applicable in many analogous situations, but the result might not be correct. The plausible lower-bound condition results from a minimal generalization of the example and its explanation. This condition allows the rule to be applicable only in situations that are very similar to the one from which the rule was learned. Therefore, the corresponding reasoning is much more likely to be correct than the one corresponding to the upper-bound condition. The agent will apply the learned rule to solve new problems, and the feedback received from the expert will be used to further refine the rule. In essence, the two conditions will converge toward one another (usually through the specialization of the plausible upper-bound condition and the generalization of the plausible lower-bound condition), both approaching the exact applicability condition of the rule. *Rule refinement* could lead to a complex task-reduction rule, with additional except-when conditions that should not be satisfied for the rule to be applicable. The tasks are refined in a similar way (Boicu et al. 2000).

It is important to stress that the expert does not deal directly with the learned tasks and rules but only with their examples used in problem solving. Therefore, the complex knowledge engineering operations of defining and debugging problem-solving rules are replaced in the DISCIPLE approach with the much simpler operations of defining and critiquing specific examples.

After the DISCIPLE agent has been trained, it can be used in the autonomous problem-solving mode to identify and test the strategic center of gravity candidates for a new scenario, as was illustrated before.

Use of DISCIPLE in the Military Applications of Artificial Intelligence Course

Many of the students that take the Center of Gravity Analysis course in the winter session, together with additional students, take the Military Applications of Artificial Intelligence course in the spring session. The spring 2001 session was attended by 10 students (7 U.S. officers and 3 international fellows). The spring

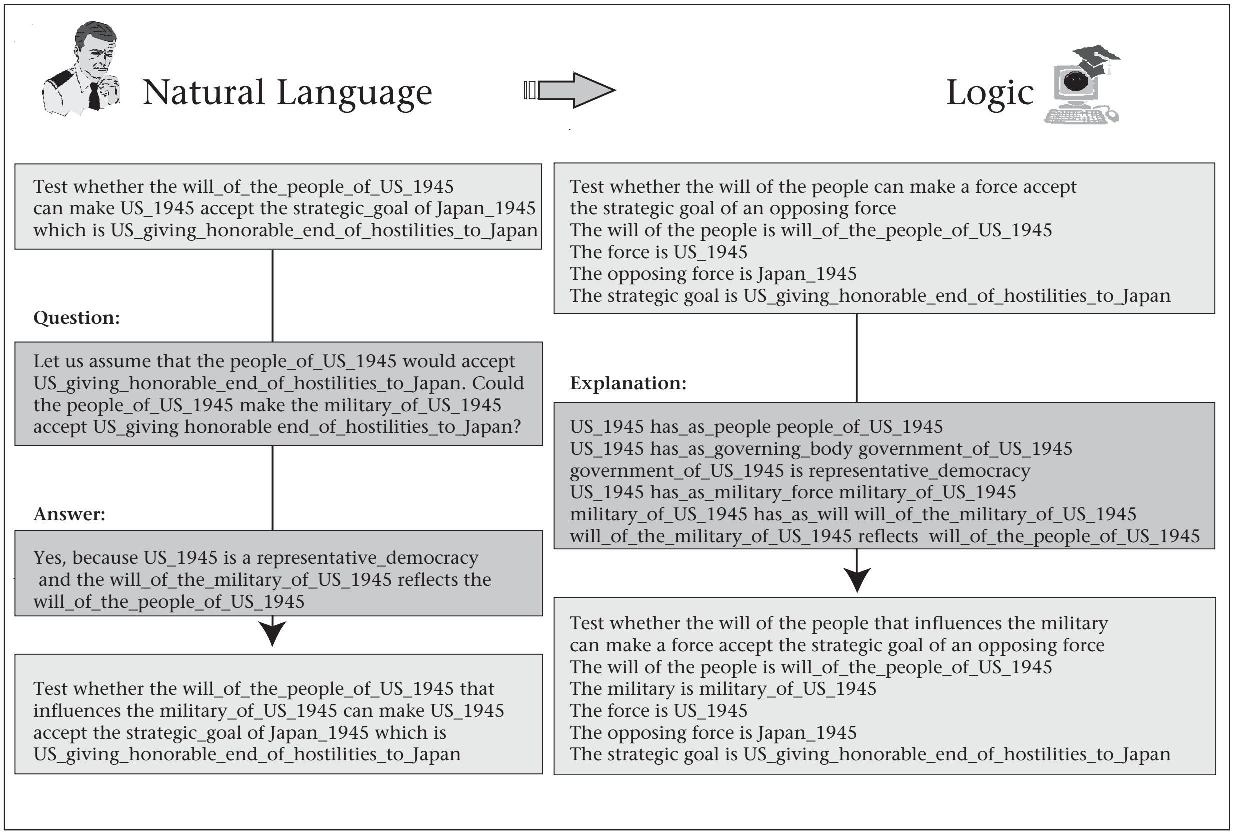

Figure 12. Mixed-Initiative Language to Logic Translation.

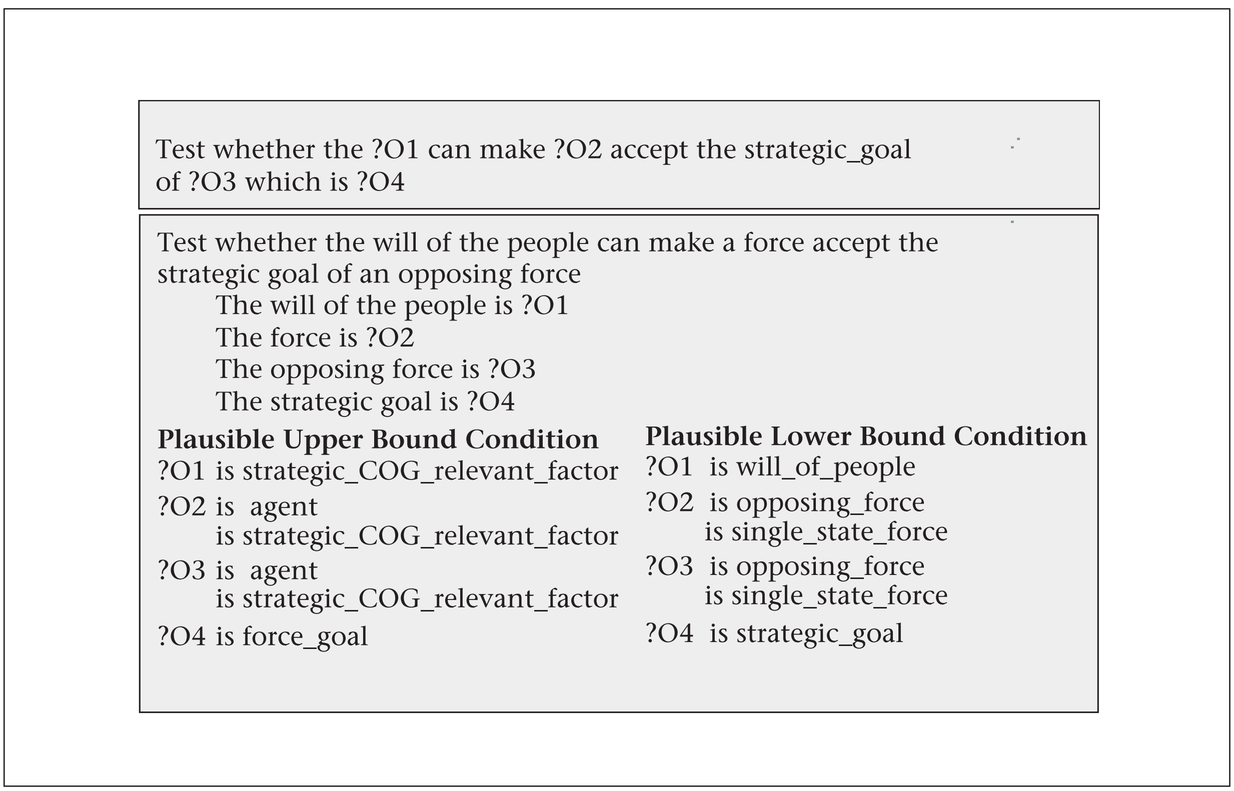

Figure 13. Task Learned from the Example in Figure 12.

IF: Test whether the ?O1 can make ?O2 accept the strategic_goal of ?O3 which is ?O4

Question: Let us assume that the ?O5 would accept ?O4.
Could the ?O5 make the ?O6 accept ?O4?
Answer: Yes, because ?O2 is a representative_democracy and thes ?O7 reflect the ?O1

THEN: Test whether the ?O1 that influences the ?O6 can make ?O2 accept the strategic_goal of ?O3 which is ?O4

IF: Test whether the will of the people can make a force accept the strategic goal of an opposing force

The will of the people is ?O1
The force is ?O2
The opposing force is ?O3
The strategic goal is ?O4

Explanation
?O2 has_as_people ?O5
?O2 has_as_governing_body ?O8
?O8 is representative_democracy
?O2 has_as_military_force ?O6 has_as_will ?O7
?O7 reflects ?O1

Plausible Upper Bound Condition		**Plausible Lower Bound Condition**	
?O1	is will_of_agent	?O1	is will _of_people
?O2	is force has_as_people ?O5 has_as_military_force ?O6 has_as_governing_body ?O8	?O2	is opposing_force is single_state_force has_as_people ?O5 has_as_military_force ?O6 has_as_governing_body ?O8
?O3	is strategic_COG_relevant_factor is agent	?O3	is opposing_force is single_state_force
?O4	is force_goal	?O4	is strategic_goal
?O5	is people	?O5	is people
?O6	is military_force has_as_will ?O7	?O6	is military_force has_as_will ?O7
?O7	is will_of_agent reflects ?O1	?O7	is will_of_military reflects ?O1
?O8	is representative_democracy	?O8	is representative_democracy

THEN: Test whether the will of the people that influences the military can make a force accept the strategic goal of an opposing force

The will of the people is ?O1
The military is ?O6
The force is ?O2
The opposing force is ?O3
The strategic goal is ?O4

Figure 14. Rule Learned from the Example in Figure 12.

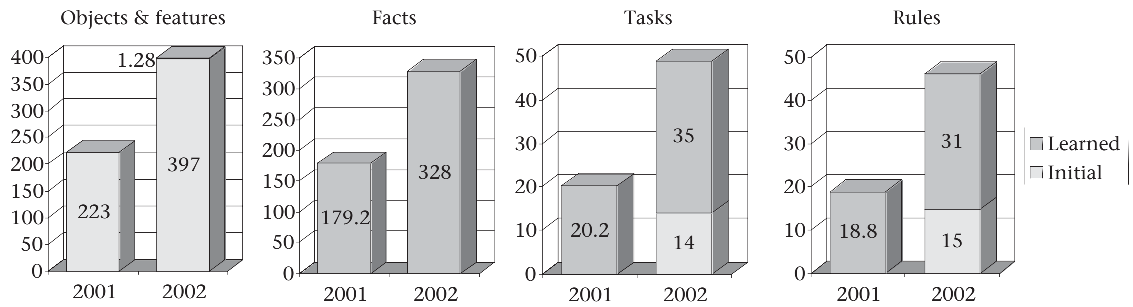

Figure 15. Knowledge Base Development during Spring 2001 and Spring 2002 Experiments.

2002 session was attended by 15 U.S. officers. In this course, the students are given a general overview of AI as well as an introduction to DISCIPLE-RKF. They are generally organized in two-person teams. Each team is given the project to train a personal DISCIPLE-RKF agent according to its own reasoning in center of gravity analysis for a certain historical scenario. That is, the students use DISCIPLE-RKF as subject-matter experts as opposed to the center of gravity course where they are end users of DISCIPLE.

As far as agent development is concerned, the Military Applications of Artificial Intelligence course is organized in two parts: (1) a learning part, during which the students (who are military experts) learn to use DISCIPLE, and (2) an experimentation part, during which each team trains its own agent.

In the spring 2001 session, each of the five teams learned to train its own DISCIPLE agent by using a different scenario. Then, in the last two 3-hour class sessions, the teams participated in a controlled agent training experiment that was videotaped in its entirety. Each team was provided with a copy of DISCIPLE-RKF that contained a generic object ontology but no specific instances, no tasks, and no rules. It received a seven-page report describing a new scenario (the Okinawa scenario described in this article) and was asked to train its DISCIPLE agent to identify center of gravity candidates for that scenario. After each significant phase of agent training and knowledge base development (that is, scenario specification, modeling, rule learning, and rule refinement), a knowledge engineer reviewed the team's work, and the team then made any necessary corrections under the supervision of the knowledge engineer. The left-hand side of the graphs in figure 15 summarize the average characteristics of the knowledge bases developed during the spring 2001 experiment. Notice that on average, the 5 agents trained by the 5 teams acquired 179.2 facts to specify the Okinawa scenario. They also learned 20.2 tasks and 18.8 rules for the identification of strategic center of gravity candidates. Although obviously incomplete (both because of the use of a single training scenario and because of incomplete training for this scenario), the knowledge bases were good enough for identifying correct center of gravity candidates not only for the Okinawa (training) scenario but also for the scenarios used for the class projects. At the end of this final experiment, the students completed a detailed questionnaire about the main components of DISCIPLE. One of the most significant results was that 7 of the 10 experts agreed, 1 expert strongly agreed, and 2 experts were neutral with respect to the statement: "I think that a subject matter expert can use DISCIPLE to build an agent, with limited assistance from a knowledge engineer." This experiment was conducted using a previous version of DISCIPLE-RKF that is described in Boicu et al. (2001).

The spring 2002 session of the Military Applications of Artificial Intelligence course was organized in a slightly different manner. All the students learned to use DISCIPLE during the lectures, using the World War II invasion of Sicily by the Allied Forces as a training scenario. As part of their hands-on experience with DISCIPLE, each of the seven teams trained its own DISCIPLE agent using a different scenario. In all but one case, the scenarios were those from the winter 2002 session of the Case Studies in Center of Gravity Analysis course.

The right-hand side of the graphs in figure 15 summarizes the average characteristics of the knowledge bases developed by the seven teams. First, it should be emphasized that this time the experts trained their agents not only to identify strategic center of gravity candidates for the given scenario but also to test them, which involves a more complex reasoning.

Notice that the size of the initial object ontology in spring 2002 was almost twice the size of the ontology from the spring 2001 experiment (397 versus 223 object and feature types). Moreover, this ontology was slightly extended during experimentation, with an average of 1.28 features, hinting to DISCIPLE's capability of learning with an evolving representation language. This increase in the size of the ontology, from spring 2001 to spring 2002, was required by the additional reasoning for testing the center of gravity candidates.

Notice also that the DISCIPLE agents from the spring 2001 experiment did not have any initial reasoning tasks or rules. The DISCIPLE agents from the spring 2002 experiment had 14 initial tasks and 15 initial rules that allowed the agents to perform the top-level reasoning illustrated in table 1. For example, these tasks and rules allowed DISCIPLE to reduce the task

> Identify and test a strategic COG candidate for the Sicily_1943 scenario.

to the task

> Identify and test a strategic COG candidate with respect to the people of US_1943.

Then the team had to teach its agent how to identify and test the strategic center of gravity candidates of an opposing force with respect to the people of US_1943 (as well as consider other aspects, such as the government, the military, or the economy). On average, each team taught its agent 35 tasks and 31 rules. Nevertheless, the developed knowledge bases were still incomplete for the same reasons as in the spring 2001 experiment (that is, the use of a single training scenario and incomplete training for this scenario). Again, however, the knowledge bases were good enough to allow each agent to (incompletely) analyze the scenarios of the other teams.

At the end of the spring 2002 experiment, 9 of the 15 experts agreed, 2 experts strongly agreed, and 4 were neutral with respect to the statement, "I think that a subject matter expert can use DISCIPLE to build an agent, with limited assistance from a knowledge engineer," in spite of the fact that the training required this time was significantly more complex than that required during the spring 2001 experiment.

We consider these experiments to be a significant success, demonstrating that subject-matter experts can train personal agents their own problem-solving expertise with limited assistance from knowledge engineers.

Conclusions

This article presented the current status of a multifaceted research and development effort that synergistically integrates research in AI, research in center of gravity analysis, and the practical application to education.

The AI research in knowledge bases and agent development by subject-matter experts has benefited from the center of gravity analysis domain that provided a complex challenge problem. The identification and testing of strategic center of gravity candidates exemplifies expert problem solving that relies on a wide range of domain knowledge, a significant part of which is tacit. This research has also benefited from its practical application to education. Both the Case Studies in Center of Gravity Analysis course and the Military Applications of Artificial Intelligence course allowed us to perform thorough experimentations with real experts, resulting in the validation of our methods and providing many ideas for improvements.

The research in center of gravity analysis has benefited from the AI research in that agent development has helped clarify and formalize the center of gravity identification and testing process. The developed center of gravity reasoning models were validated in the U.S. Army War College courses and are leading to a significant extension of the center of gravity monograph of Giles and Galvin.[3]

Finally, the innovative application of the AI and center of gravity research to education, through the use of the DISCIPLE agents, has had a significant impact on improving the Center of Gravity Analysis and Military Applications of Artificial Intelligence courses. Done as a very successful experiment in 2001, it was made a regular part of the syllabi for 2002, to be continued in the following years.

The deployment and evaluation of DISCIPLE in the Case Studies in Center of Gravity Analysis and Military Applications of Artificial Intelligence courses have also revealed several limitations of this approach and have provided numerous ideas for improvement. For example, although the subject-matter expert has an increased role and independence in agent development, the knowledge engineer still has a critical role to play. He/she has to assure the development of a fairly complete and correct object ontology. He/she also has to develop a generic modeling of the problem-solving process based on the task-reduction paradigm. Even guided by this generic modeling, and using natural language, the subject-matter expert has difficulties in expressing his/her reasoning process. Therefore, more work is needed to develop methods for helping the expert in this task along the path opened by the modeling adviser.

The experimentations revealed that the mixed-initiative reasoning methods of DISCIPLE could significantly be empowered by developing the natural language–processing capabilities of the system.

Finally, because the expert who teaches DISCIPLE has no formal training in knowledge engineering or computer science, the knowledge pieces learned by the agent and the knowledge base itself will not be represented optimally and will require periodic revisions by a knowledge engineer. Examples of encountered problems with the knowledge base are semantic inconsistencies within a rule, proliferation of semantically equivalent tasks, and the violation of certain knowledge engineering principles. It is therefore necessary to develop mixed-initiative knowledge base reformulation and optimization methods to identify and correct such problems in the knowledge base.

The single most important lesson from this effort is the significant benefit resulted from the synergistic integration of the three complementary activities: (1) research in AI, (2) research in a specialized domain, and (3) development and deployment of prototype systems in education and practice. Each of these three activities contributed to the achievement of the goals of the other two; none of them alone would have achieved its own goals to the same extent.

We will therefore continue this multiobjective activity. We plan to improve the DISCIPLE approach by addressing the limitations revealed by the performed experimentations. We also plan to extend the formal treatment of the center of gravity analysis by addressing operations other than wars and nonstate combatants. Finally, we plan not only to maintain the developed DISCIPLE-RKF/COG agent but also to accordingly extend and improve its capabilities. Therefore, the maintenance of this application will actually be a by-product of this integrated effort.

Acknowledgments

The research described in this article was sponsored by DARPA, the Air Force Research Laboratory (AFRL), the Air Force Materiel Command (AFMC), and the United States Air Force under agreement F30602-00-2-0546; by the AFOSR under grant F49620-00-1-0072; and by the U.S. Army War College, benefiting from the direction of Murray Burke, Robert Herklotz, William Rzepka, Douglas Campbell, David Cammons, and David Brooks. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, AFMC, AFOSR, USAWC, or the U.S. government. Michael Bowman, Marcel Barbulescu, Xianjun Hao, Tyrus Berry, and other members of the Learning Agents Laboratory contributed to the development of DISCIPLE-RKF. William Cleckner, James Donlon, and Antonio Lopez helped with the organization of the Case Studies in Center of Gravity Analysis and the Military Applications of Artificial Intelligence courses. The students from these courses helped significantly with the advancement of this research, not only through the experimentations performed but also through the numerous suggestions provided. We are also grateful to Steve Chien and John Riedl for their helpful comments on an early version of this article.

Notes

1. M. Burke. 1999. Rapid Knowledge Formation Program Description. /reliant.teknowledge.com/RKF - projects/Darpa_RKF_PIP.htm, August 4, 2002.

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

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

References

Boicu, M.; Tecuci, G.; Marcu, D.; Bowman, M.; Shyr, P.; Ciucu, F.; and Levcovici, C. 2000. DISCIPLE-COA: From Agent Programming to Agent Teaching. In Proceedings of the Seventeenth International Conference on Machine Learning (ICML), 73–80. San Francisco, Calif.: Morgan Kaufman.

Boicu, M.; Tecuci, G.; Stanescu, B.; Marcu, D.; and Cascaval, C. E. 2001. Automatic Knowledge Acquisition from Subject-Matter Experts. In Proceedings of the Thirteenth International Conference on Tools with Artificial Intelligence (ICTAI), 69–78. Washington, D.C.: IEEE Computer Society.

Chaudhri, V. K.; Farquhar, A.; Fikes, R.; Park, P. D.; and Rice, J. P. 1998. OKBC. A Programmatic Foundation for Knowledge Base Interoperability. In Proceedings of the Fifteenth National Conference on Artificial Intelligence, 600–607. Menlo Park, Calif.: American Association for Artificial Intelligence.

Clausewitz, C. V. 1976. *On War*. Translated and edited by M. Howard and P. Paret. Princeton, N.J.: Princeton University Press.

Hamburger, H., and Tecuci, G. 1998. Toward a Unification of Human-Computer Learning and Tutoring, In *Intelligent Tutoring Systems,* eds. B. P. Goettl, H. M. Halff, C. L. Redfield, and V. J. Shute, 444–453. Berlin: Springer-Verlag.

Joint Chiefs of Staff. 2001. Doctrine for Joint Operations, Joint Publication 3-0, III-22, 10 September 2001. Washington, D.C.: U.S. Joint Chiefs of Staff.

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

Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. San Diego, Calif.: Academic.

Tecuci, G. 1988. DISCIPLE: A Theory, Methodology, and System for Learning Expert Knowledge, Thèse de Docteur en Science, Centre d'Orsay, Laboratoire de Recherche en Informatique, University of Paris-South.

Tecuci, G., and Keeling, H. 1999. Developing an Intelligent Educational Agent with DISCIPLE. *International Journal of Artificial Intelligence in Education* 10(3–4): 221–237.

Tecuci, G.; Boicu, M.; and Marcu, D. 2000. Learning Agents Teachable by Typical Computer Users. Paper presented at the Seventeenth National Conference on Artificial Intelligence (AAAI-00) Workshop on New Research Problems for Machine Learning, 30 July–3 August, Austin, Texas.

Tecuci, G.; Boicu, M.; Bowman, M.; Marcu, D.; and Burke, M. 2001. An Innovative Application from the DARPA Knowledge Bases Programs: Rapid Development of a High-Performance Knowledge Base for Course-of-Action Critiquing. *AI Magazine* 22(2): 43–-61.

Tecuci, G.; Boicu, M.; Wright, K.; Lee, S. W.; Marcu, D.; and Bowman, M. 1999. An Integrated Shell and Methodology for Rapid Development of Knowledge-Based Agents. In Proceedings of the Sixteenth National Conference on Artificial Intelligence, 250–257. Menlo Park, Calif.: American Association for Artificial Intelligence.

Gheorghe Tecuci is professor of computer science and director of the Learning Agents Laboratory at George Mason University, chair of AI at the U.S. Army War College, and a member of the Romanian Academy. He received two Ph.D.s in computer science, one from the University of Paris-South, Orsay, France, and the other from the Polytechnic University of Bucharest, Romania. He has published over 100 scientific papers and 5 books, including *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool, and Case Studies* (Academic Press, 1998); *Machine Learning: A Multistrategy Approach* (Morgan Kaufmann, 1994); and *Machine Learning and Knowledge Acquisition: Integrated Approaches* (Academic Press, 1995). His e-mail address is tecuci@gmu.edu.

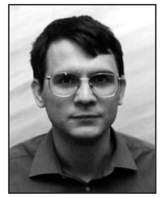

Mihai Boicu is research assistant professor of computer science and associate director of the Learning Agents Laboratory at George Mason University. He recently received the Outstanding Graduate Student in Information Technology Award, the IAAI-2002 Deployed Application Award, and the Centennial Coin of the U.S. Army War College. His domains of interest are knowledge representation, knowledge acquisition, multistrategy learning, and mixed-initiative reasoning with applications in instructable agents. He is a member of the American Association for Artificial Intelligence. His e-mail address is mboicu@gmu.edu.

Dorin Marcu is a Ph.D. candidate in computer science and a research assistant in the Learning Agents Laboratory at George Mason University. His current research interests are mixed-initiative human-computer interaction and adaptive user interfaces. He received the IAAI-2002 Deployed Application Award and the Centennial Coin of the U.S. Army War College. He is a member of the American Association for Artificial Intelligence. His e-mail address is dmarcu@cs.gmu.edu.

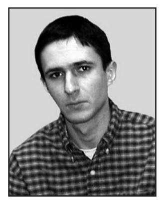

Bogdan Stanescu is a Ph.D. candidate in computer science and a research assistant in the Learning Agents Laboratory at George Mason University. His research interests are AI, machine learning, knowledge acquisition, and intelligent agents. He received the IAAI-2002 Deployed Application Award and the Centennial Coin of the U.S. Army War College. He is a member of the American Association for Artificial Intelligence. His e-mail address is bstanesc@cs.gmu.edu.

Cristina Boicu is a Ph.D. student in computer science at George Mason University. She has been a member of the Learning Agents Laboratory at George Mason University since 1999. Her research interests include intelligent agents, knowledge acquisition, knowledge discovery, and machine learning. She recently received a Computer Science Certificate of Achievement from George Mason University, the IAAI-2002 Deployed Application Award, and the Centennial Coin of the U.S. Army War College. Her e-mail address is ccascava@gmu.-edu.

Jerome J. Comello, Ph.D. and colonel infantry U.S. army retired, is currently professor of military studies in the Department of Military Strategy Planning and Operations at the U.S. Army War College. Formerly, he was associate professor of strategic military operations and planning, Center for Strategic Leadership, U.S. Army War College, and associate professor of national security, U.S. Marine Corps Command and Staff College.