

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 paper presents the results of a multi-faceted research and development effort that synergistically integrates artificial intelligence 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 strategic center of gravity analysis, and the deployment and evaluation of these agents in several courses at the US Army War College.

This paper presents the results of a multi-objective collaboration between the Learning Agent Laboratory of George Mason University, on one side, and the Center for Strategic Leadership and the Department of Military Strategy, Planning, and Operations of the US Army War College, on the other side. 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, as detailed in the following.

The artificial intelligence research objective is the development of the Disciple approach for building instructable knowledge-based systems or agents (Tecuci 1988; 1998). 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 that is 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, email 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 et al., 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 an intelligent assistant as easily as he 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

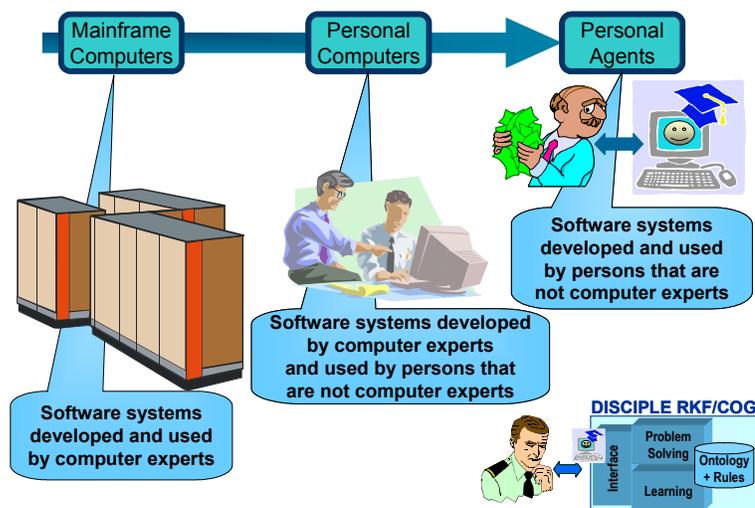


Figure 1: View on the evolution of the software development process

most recent predecessors: Disciple-WA (Tecuci et al. 1999) and Disciple-COA (Tecuci et al. 2001). All these three systems were developed as part of the “High Performance Knowledge Bases” and “Rapid Knowledge Formation” programs supported by DARPA and AFOSR (Burke 1999). 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 strategic center of gravity analysis. This 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 center of gravity of an entity (state, alliance, coalition, or group) was introduced by Karl von Clausewitz (1832) as “the foundation of capability, the hub of all power and movement, upon which everything depends, the point against which all the energies should be directed.”

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 US senior military service colleges, there is a great emphasis on the center of gravity analysis (Strange 1996). This introduces 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 US Army War College. In the “Case Studies in Center of Gravity Analysis” course (the COG 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 MAAI course), the students

use personal Disciple 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 in 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. After that, the Disciple agent can be used as a personal tutor, teaching the students in a way that is similar to how it was taught by the teacher (Hamburger and Tecuci, 1998; Tecuci and Keeling, 1999).

Each of the three objectives discussed above 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 in each of them. Moreover, it offers a new perspective on how to combine research in artificial intelligence, with research in a specialized domain, and with the development and deployment of prototype systems in education and practice.

The rest of the paper presents the current status of this research and development effort. The next section presents in more detail the center of gravity challenge problem. This 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 US 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 paper concludes with a summary of the synergistic aspects of this

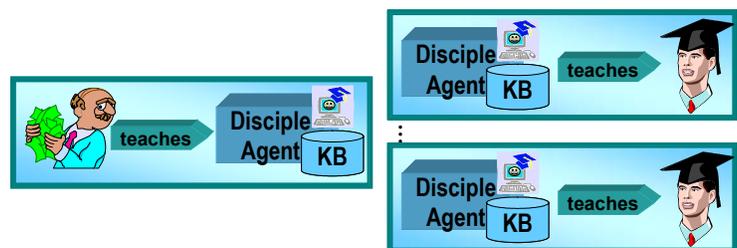


Figure 2: View on the future use of instructable agents in education

collaborative work and future research directions.

The Center of Gravity Problem

Military literature distinguishes between three levels of conflicts: a strategic level focusing on winning wars, an operational level focusing on winning campaigns, and 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” (1832), 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, while adequately protecting 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. Recognizing these difficulties, the Center for Strategic Leadership of the US 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 (Giles and Galvin 1996), 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 COG Analysis

Disciple-RKF/COG is an agent used in the US Army War College course titled “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 below.

First, Disciple guides the student to identify, study and describe the aspects of a campaign (such as the 1945 US 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 that aspect, or to 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 asks for the opposing forces. Once the student indicates Japan_1945 and US_1945 as opposing forces, Disciple includes them in the table of contents, together with general characteristics, which the student can specify (see the left hand side of Figure 4). The student may then click on any of these aspects (e.g. “Industrial capacity” under “Economic factors” of Japan_1945) and the agent guides the student in specifying it. The student’s specification may 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

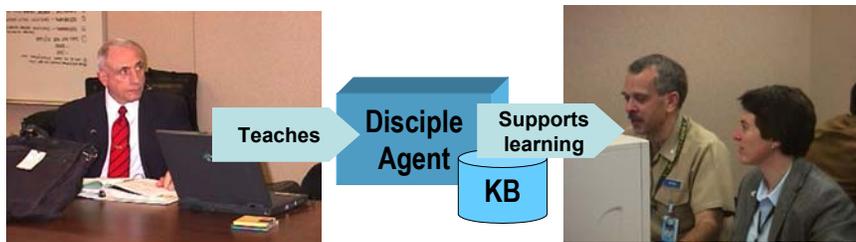


Figure 3: A step toward the vision of using instructable agents in education

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 COG 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 Staffs, 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 or/and for 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 COG 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 and testing a strategic center of gravity candidate. To perform this task,

Disciple asks itself a series of questions. The answer of each question allows Disciple to reduce the current task to simpler ones, until Disciple has enough information to first identify a strategic COG candidate, and then to test it, determining whether it should be eliminated or not.

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 COG 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 US. 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

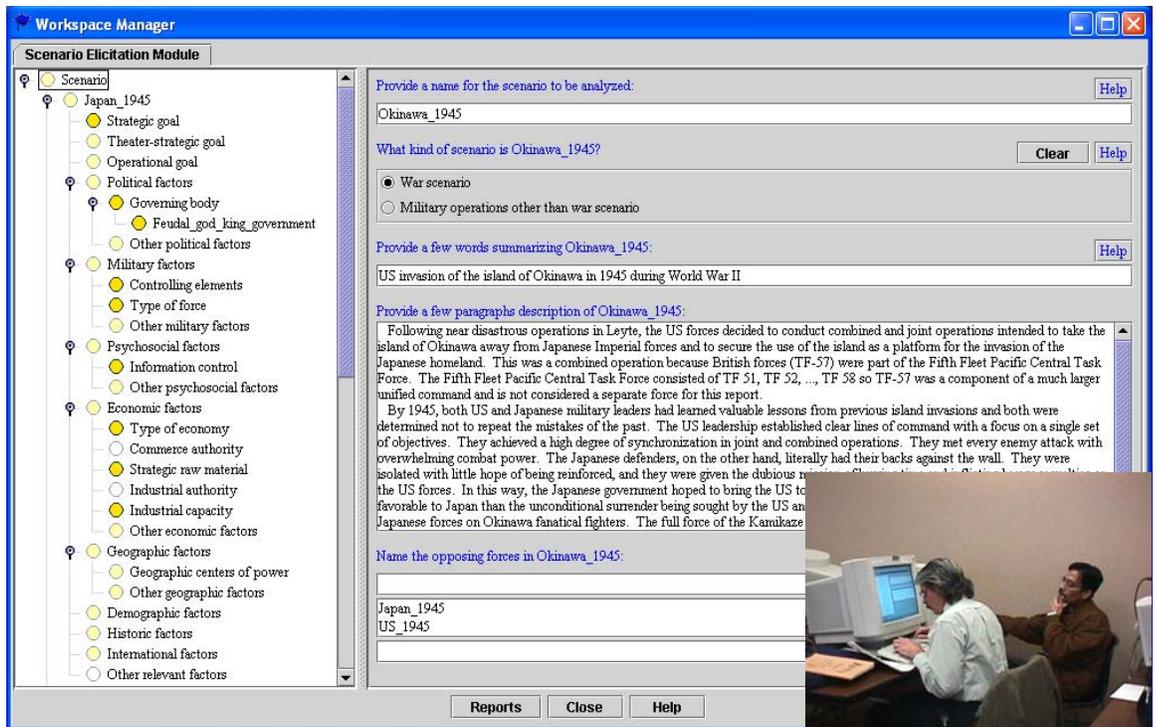


Figure 4: Scenario specification interface

the Clausewitz's trinity of power (government, military, and people), Emperor Hirohito is 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 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 Staffs -- 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 campaign. We recognize that subject matter experts often differ in their judgments as to the identification and analysis of center of gravity candidates for any particular scenario.

As illustrated above, 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 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. Secondly, center of gravity analysis is influenced by personal experiences and subjective judgments, and

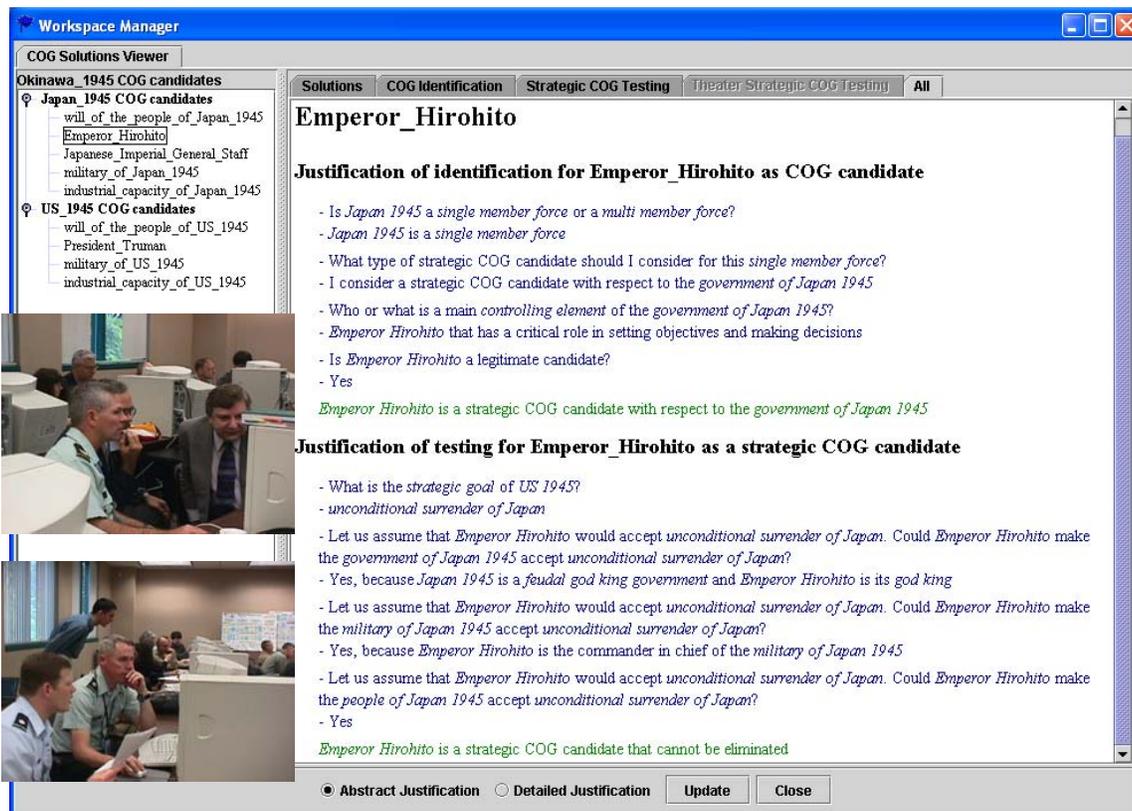


Figure 5: The problem solving interface of Disciple-RKF/COG

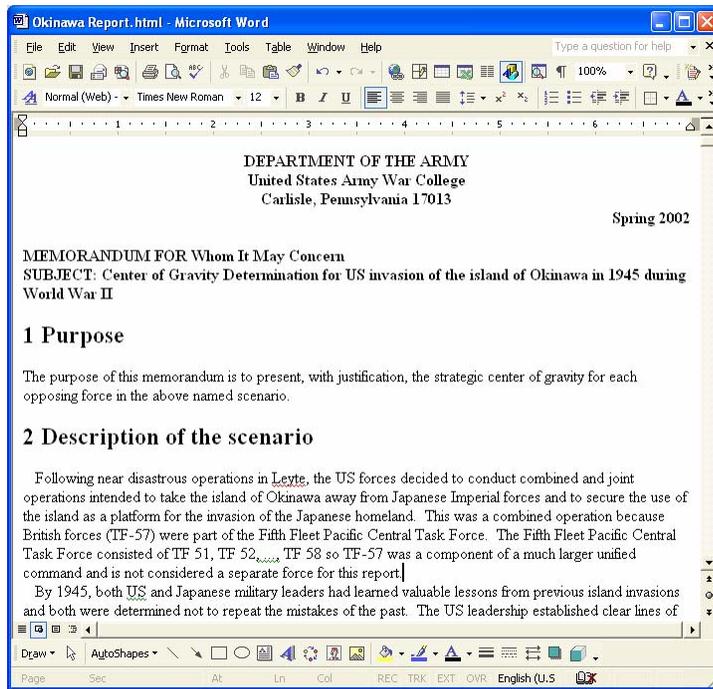


Figure 6: The report generated by Disciple-RKF/COG

the student (who has unique military experience and biases) may 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 actions proposed by their staff and to make the final decision on which one to use.

Use of Disciple in the COG course

Successive versions of Disciple have been used in both the Winter and Spring sessions of the “Case Studies in Center of Gravity Analysis” 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 US officers and 3 international fellows), 3 students in the Spring-2001 session (1 US officer and 2 international fellows), 13 students in the Winter-2002 session (11 US officers and 2 international fellows), and 10 students in the Spring-2002 session (2 US 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, and 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 instance, are some of the evaluations of the 13 students from the Winter 2002 session, being generally representative of the evaluations from all the other sessions. On a 5-point scale (strongly disagree, disagree, neutral, agree, strongly agree), 9 students agreed and the other 4 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 9 agreed, and the other 3 strongly agreed with the statement “Disciple helped me to learn to perform a strategic center of gravity analysis of a scenario.” One student disagreed, but 4 students agreed and the other 8 strongly agreed that “The use of Disciple was a useful learning experience.” Finally, one student disagreed, but 9 students agreed and the other 3 strongly agreed that “Disciple should be used in future versions of this course.”

To our knowledge, this is the first time that 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 US Army War College.

Agent Development with Disciple-RKF

The Disciple-RKF/COG agent presented in the previous section was developed by using the Disciple-RKF learning agent shell, as will be described 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: an object ontology that defines the terms from a specific application domain, and a set of task reduction rules expressed with these terms. Disciple-RKF represents a significant evolution as 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 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, to clarify how these problems could be solved using Disciple's task reduction paradigm, and to 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

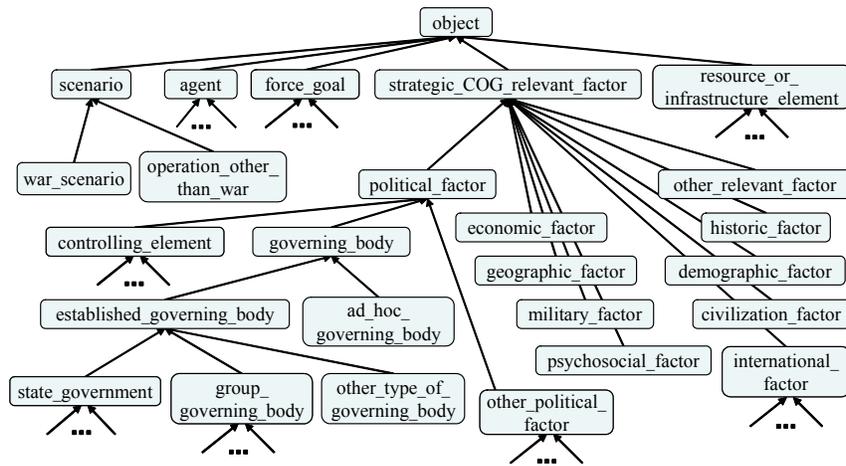


Figure 7: A fragment of the COG object ontology

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).

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 the question 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 the appearance of the interface. The use of the elicitation scripts

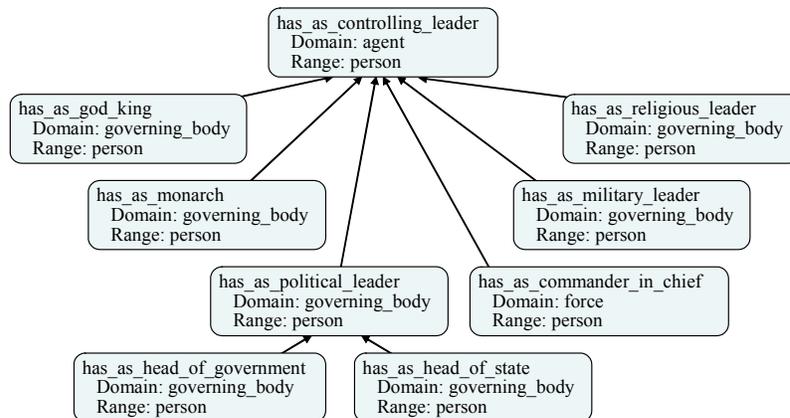


Figure 8: A fragment of the COG feature ontology

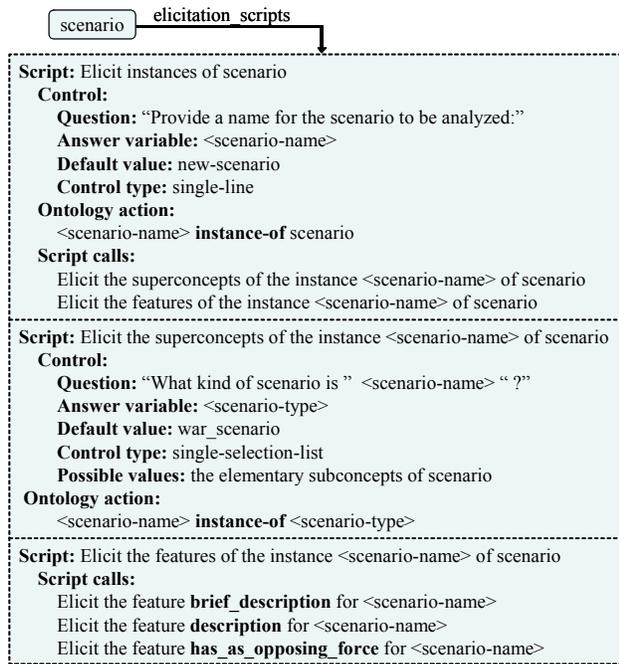


Figure 9: Sample elicitation scripts

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

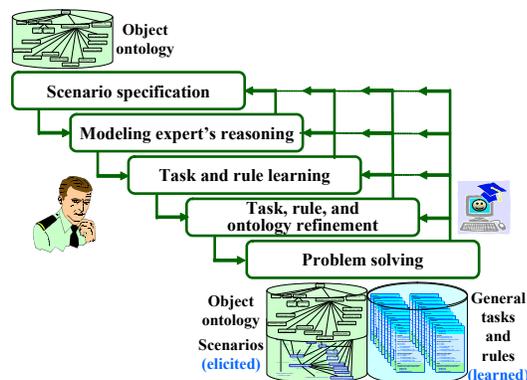


Figure 10: The main phases of the agent training process

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 (e.g. the US invasion of the island of Okinawa in 1945). The expert does not work directly with the object ontology in order 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 can start the *Modeling* of his COG reasoning for this particular scenario, as a sequence of task reduction steps. The expert expresses his reasoning in English, similarly to how he would think aloud while solving 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 a series of questions. The answer of 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 or not.

Table 1: Sample modeling of the COG analysis process for a specific scenario

<p><i>I need to</i> 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. <i>Therefore I need to</i> 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 <i>Therefore I need to</i> Identify and test a strategic COG candidate for Japan_1945. ...</p>

Experimental results show that this is the most challenging agent training activity for the expert. We have therefore recently

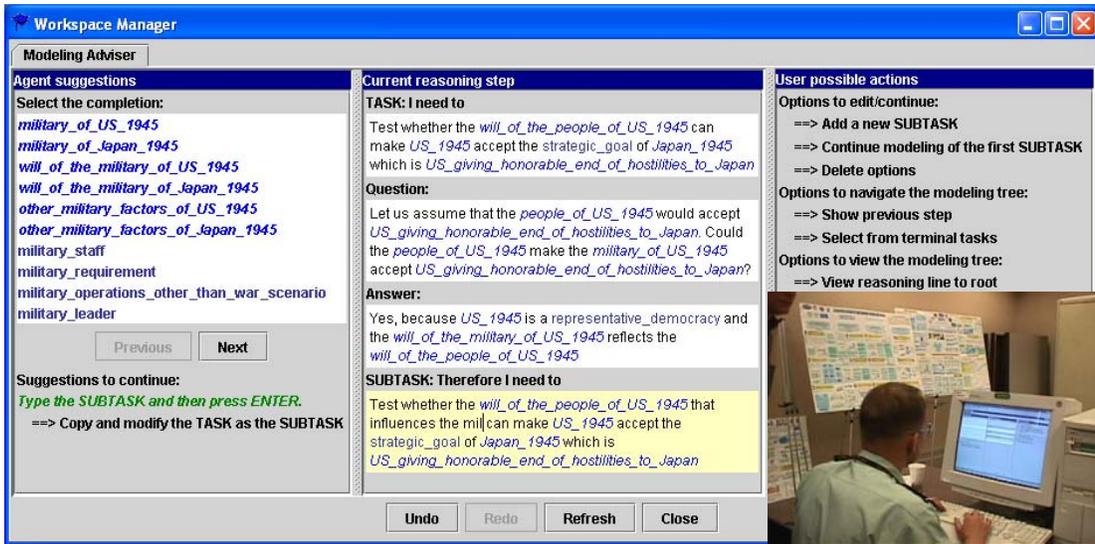


Figure 11: The modeling advisor interface

developed a Modeling Advisor 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 that state, and the left hand side shows the action that the Modeling Advisor is actually recommending. For instance, to specify the current subtask, the advisor suggested the expert to copy and modify the task. The Modeling Advisor may also suggest the question to be asked, or the answer of the question. As mentioned, the

expert expresses his reasoning in English. However, each time he starts to type a word, the agent lists in 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 is useful for two different reasons: it facilitates the user's input, and helps the agent to "understand" his phrases.

In the *Task and rule learning* phase, Disciple learns general tasks and general rules from the task reduction steps defined in the modeling phase. For instance, consider the reduction step from the middle of Figure 11, shown again in the left hand side of Figure 12. It consists of a task, a

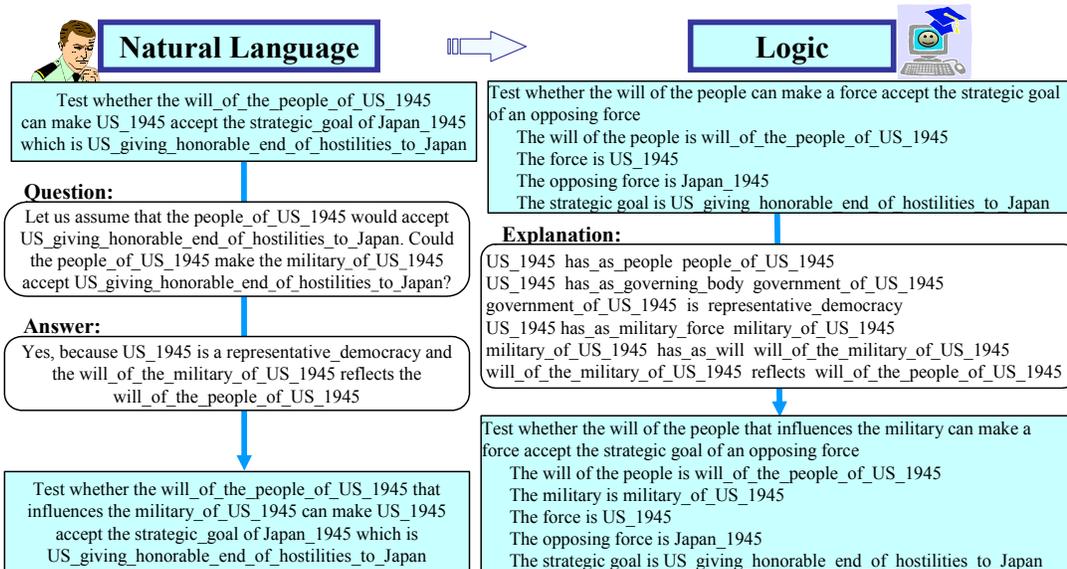


Figure 12: Mixed-initiative language to logic translation

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 may 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 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 above. 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

Test whether the ?O1 can make ?O2 accept the strategic_goal of ?O3 which is ?O4	
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	
Plausible Upper Bound Condition	Plausible Lower Bound Condition
?O1 is strategic_COG_relevant_factor	?O1 is will_of_people
?O2 is agent	?O2 is opposing_force
is strategic_COG_relevant_factor	is single_state_force
?O3 is agent	?O3 is opposing_force
is strategic_COG_relevant_factor	is single_state_force
?O4 is force_goal	?O4 is strategic_goal

Figure 13: Task learned from the example in Figure 12

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 may not be correct. On the other hand, 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

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 the ?O7 reflects 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	?O2 is opposing_force
has_as_people ?O5	is single_state_force
has_as_military_force ?O6	has_as_people ?O5
has_as_governing_body ?O8	has_as_military_force ?O6
	has_as_governing_body ?O8
?O3 is strategic_COG_relevant_factor	?O3 is opposing_force
is agent	is single_state_force
?O4 is force_goal	?O4 is strategic_goal
?O5 is people	?O5 is people
?O6 is military_force	?O6 is military_force
has_as_will ?O7	has_as_will ?O7
?O7 is will_of_agent	?O7 is will_of_military
reflects ?O1	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

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 which should not be satisfied in order 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 MAAI 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 US officers and 3 international fellows). The Spring 2002 session was attended by 15 US officers. In this course

the students are given a general overview of Artificial Intelligence, 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 COG course where they are end-users of Disciple.

As far as agent development is concerned, the MAAI course is organized in two parts, a learning part during which the students (who are military experts) learn to use Disciple, and 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 7-page report describing a new scenario (the Okinawa scenario described in this paper), 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 (i.e. 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 five agents trained by the five teams acquired 179.2 facts to specify the

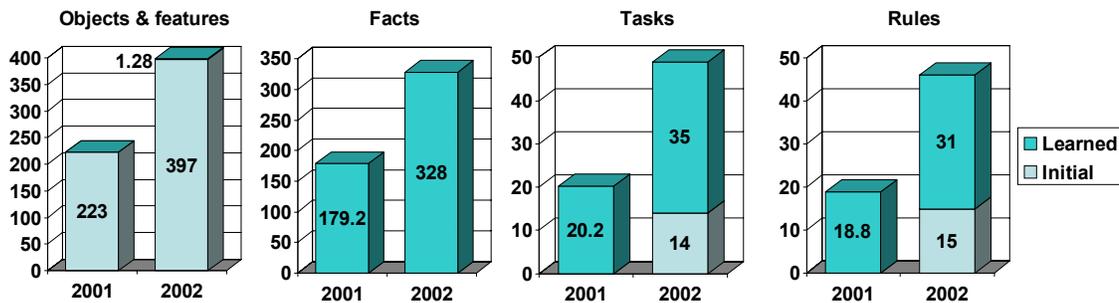


Figure 15: Knowledge base development during Spring-2001 and Spring-2002 experiments

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 that 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, containing questions about the main components of Disciple. One of the most significant results was that 7 out 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. Then, as part of their hands-on experience with Disciple, each of the 7 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 summarize the average characteristics of the knowledge bases developed by the 7 teams. First of all, 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 the 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 instance, 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 with respect to 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 (i.e. both because of the use of a single training scenario, and because of incomplete training for that 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 out of the 15 experts agreed, 2 experts strongly agreed and 2 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 the one required during the Spring-2001 experiment.

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

Conclusions

This paper presented the current status of a multi-faceted research and development effort that synergistically integrates research in artificial intelligence, research in center of gravity analysis, and the practical application to education.

The artificial intelligence research in knowledge bases and agent development by subject matter experts has benefited from the center of gravity analysis domain which 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 artificial intelligence research in that the 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 US Army War College courses, and are leading to a significant extension of the center of gravity monograph of Giles and Galvin (1996).

Finally, the innovative application of the artificial intelligence and center of gravity research to education, through the use of the Disciple agents, has had a significant impact on improving the COG and MAAI 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 COG and MAAI courses have also revealed several limitations of this approach and have provided numerous ideas for improvement. For instance, while the subject matter expert has an increased role and independence in agent development, the knowledge engineer still has a critical role to play. He has to assure the development of a fairly complete and correct object ontology. He 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 reasoning process. Therefore more work is needed to develop methods for helping the expert in this task, along the path opened by the Modeling Advisor.

The experimentations revealed that the

mixed-initiative reasoning methods of Disciple could be significantly 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 optimally represented, 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: research in artificial intelligence, research in a specialized domain, and 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, and none of them alone would have achieved its own goals to the same extent.

We will therefore continue this multi-objective 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 non-state 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 paper was sponsored by DARPA, AFRL, AFMC, USAF, under agreement number F30602-00-2-0546, by AFOSR under grant no. F49620-00-1-0072, and by the US 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 COG and MAAI 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 paper.

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 *Proc. of the Seventeenth International Conference on Machine Learning (ICML)*. Stanford, California: Morgan Kaufman.
- Boicu, M.; Tecuci, G.; Stanescu, B.; Marcu, D.; and Cascaval, C.E. 2001. Automatic Knowledge Acquisition from Subject Matter Experts. In *Proc. of the Thirteenth International Conference on Tools with Artificial Intelligence (ICTAI) 7-9 November 2001, Dallas, Texas*, 69-78. Los Alamitos, California: IEEE Computer Society.
- Burke, M. 1999. Rapid Knowledge Formation Program Description. At http://reliant.technowledge.com/RKF/projects/Darpa_RKF_PIP.htm, August 4, 2002.
- 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, California: AAAI Press.
- Clancey, W. J. 1984. NEOMYCIN: Reconfiguring a rule-based system with application to teaching. In Clancey W. J. and Shortliffe, E. H. eds. *Readings in Medical Artificial Intelligence*, 361-381. Reading, MA: Addison-Wesley.
- Clausewitz, C.V. 1832. *On War*, translated and edited by M. Howard and P. Paret. Princeton, NJ: Princeton University Press, 1976.
- Department of the Army 2001. *Field Manual 3-0, Operations*. Washington, D.C.: U.S. Gov. Printing Office.
- Giles, P.K.; and Galvin, T.P. 1996. *Center of Gravity: Determination, Analysis and Application*. CSL, U.S. Army War College, PA: Carlisle Barracks.
- Hamburger, H.; and Tecuci, G. 1998. Toward a Unification of Human-Computer Learning and Tutoring, In Goettl, B.P.; Half, H.M.; Redfield, C.L.; and Shute, V.J. (eds), *Intelligent Tutoring Systems*, 444-453, Berlin: Springer-Verlag.
- Joint Chiefs of Staff, 2001. Doctrine for Joint Operations, *Joint Publication 3-0*, Washington, D.C.: U.S. Joint Chiefs of Staff, 10 September 2001, III-22.
- Strange, J. 1996. *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*. Quantico, VA: Marine Corps University Foundation.
- Tecuci, G. 1988. *DISCIPLINE: A Theory, Methodology and System for Learning Expert Knowledge*, 197 pages, Thèse de Docteur en Science, University of Paris-South.
- Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London, England: Academic Press.
- Tecuci, G.; and Keeling, H. 1999. Developing an Intelligent Educational Agent with Disciple, *International Journal of Artificial Intelligence in Education*, vol. 10, no.3-4.
- 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, California: AAAI Press.
- Tecuci, G.; Boicu, M.; Marcu, D. 2000. Learning Agents Teachable by Typical Computer Users. In *Proc. of the AAAI-2000 Workshop on New Research Problems for Machine Learning*, Austin, Texas.
- Tecuci, G.; Boicu, M.; Bowman, M.; and Marcu, D. with a commentary by 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* Vol. 22, No. 2: 43-61.