

An Innovative Application from the DARPA Knowledge Bases Programs Rapid Development of a Course of Action Critiquer

Gheorghe Tecuci, Mihai Boicu, Michael Bowman, and Dorin Marcu

Commentary by Murray Burke

■ This article presents a learning agent shell and methodology for building knowledge bases and agents and their innovative application to the development of a critiquing agent for military courses of action, a challenge problem set by the Defense Advanced Research Projects Agency's High-Performance Knowledge Bases Program. The learning agent shell includes a general problem-solving engine and a general learning engine for a generic knowledge base structured into two main components: (1) an ontology that defines the concepts from an application domain and (2) a set of task-reduction rules expressed with these concepts. The development of the critiquing agent was done by importing ontological knowledge from CYC and teaching the agent how an expert performs the critiquing task. The learning agent shell, the methodology, and the developed critiquer were evaluated in several intensive studies, demonstrating good results.

A great challenge for AI is the development of theories, methods, and tools that would allow users that do not have knowledge engineering or computer science

experience to build knowledge bases and agents by themselves. We believe success in this area will have an even greater impact on our society than the development of personal computers. Indeed, if personal computers allowed every person to become a computer user, without the need for special training in computer science, solutions to this AI challenge would allow any such person to become an agent developer. Agent development by typical computer users would lead to a large scale use of computers as personalized intelligent assistants, helping their users in a wide range of tasks. The key issue is that the development of such an agent should be as easy for the user as it currently is to use a word processor.

In this article, we present recent progress made in the George Mason University Learning Agents Laboratory toward this goal. First, we introduce the concept of a learning agent shell as a tool to be used directly by a subject-matter expert (SME) to develop an agent. Then, we present the DISCIPLE family of learning agent shells and the successful application

In his invited talk at the 1993 National Conference on Artificial Intelligence, Edward Feigenbaum compared the technology of a knowledge-based computer system with a tiger in a cage. Rarely does a technology arise that offers such a wide range of important benefits. However, this technology is still far from achieving its potential. This tiger is in a cage, and to free it, the AI research community must understand and remove the bars of the cage.

We now know that a knowledge-based system needs a great deal of knowledge to be truly useful. However, building a large and high-performance knowledge base is still a long, inefficient, and error-prone process. Responding to this problem, the Defense Advanced Research Projects Agency (DARPA) has sponsored a sequence of two programs for the development of the second generation of knowledge-based systems science and technology: (1) the High-Performance Knowledge Bases Program (HPKB FY97-99) and (2) the Rapid Knowledge Formation Program (RKF FY00-03).^{1,2} The goal of the HPKB Program was to produce the technology needed to enable system developers to rapidly construct large knowledge bases that provide comprehensive coverage of topics of interest, are reusable by multiple applications with diverse problem-solving strategies, and are maintainable in rapidly changing environments. The tasks of knowledge representation and acquisition are too difficult to be started from scratch each time a new knowledge base needs to be built. Therefore, this program supported the development of methods for creating knowledge bases by selecting, composing, extending, specializing, and modifying components from a library of reusable ontologies, common domain theories, and generic problem-solving

strategies. In addition, it supported the development of methods for rapidly extracting knowledge from natural language texts and the World Wide Web and for knowledge acquisition from subject matter experts (SMEs). An important emphasis of the HPKB Program was the use of *challenge problems*, which are complex, innovative military applications of AI that are intended to focus the research and development efforts and measure the effectiveness of alternative technical approaches. The participants collaborated in the development of knowledge-based systems for solving these challenge problems. These systems were the subject of intensive annual evaluations during which the completeness and correctness of the developed knowledge bases were measured as well as the time required to build the knowledge bases and the ease of modifying them to assimilate new or changed knowledge. The challenge problems for the first part of the HPKB Program and the corresponding evaluation results were presented in the Winter 1998 issue of *AI Magazine* (Cohen et al. 1998). The second part of the HPKB Program was based on even more complex challenge problems. One challenge problem was an extension of the year-one crisis-management problem, requiring rapid development of a large knowledge base containing broad but relatively shallow knowledge (such as general knowledge of countries, politics, geopolitical events, economics) necessary to discover and understand information about nascent and emerging crises from a wide range of potential information sources. Two teams, one composed of Teknowledge, Cycorp, Textwise, and Northwestern University, and the other composed of Science Applications International Corporation (SAIC), SRI International, the Massachusetts Institute of Technolo-

gy (MIT), Stanford University, and Northwestern University, developed two end-to-end integrated systems that were evaluated by Information Extraction and Transport Inc. (IET), the challenge problem developer, in the summer of 1999. Both systems demonstrated high performance through knowledge reuse and semantic integration and created a significant amount of reusable knowledge. The other challenge problem required the rapid development of knowledge bases containing comprehensive battlefield knowledge (for example, terrain characteristics, force structures, military organizations, troop movements, military strategy). The problem was to assess various aspects of military courses of action (COAs), such as their viability, their correctness, and their strengths and weaknesses with respect to the principles of war and the tenets of army operations; to justify the assessments made; and to propose improvements. This challenge problem was solved by developing a complex end-to-end system integrating complementary technologies developed by different research groups. Teknowledge, Northwestern University, and the University of Edinburgh's Artificial Intelligence Applications Institute (AIAI) developed the translation and fusion module that interpreted and combined the information included in the COA sketch and the COA statement and generated an internal representation of the input COA based on a *CYC* ontology. Four critiquers, each developed by a different team (a joint team from Teknowledge and Cycorp, the *DISCIPLE* team from George Mason University, the *EXPECT* team from University of Southern California Information Sciences Institute [USC/ISI], and the *LOOM* team from USC/ISI shared the generated representation of the COA and critiqued differ-

of one member of this family, *DISCIPLE-COA*, to the Defense Advanced Research Project Agency's (DARPA) High-Performance Knowledge Bases (HPKB) Program's course-of-action (COA) challenge problem. We describe the challenge problem and the *DISCIPLE-COA* shell and methodology used to build the COA critiquing agent. We then present the results of DARPA's evaluation of the developed tools and COA critiquers. We also briefly present the results of a separate knowledge-acquisition experiment with *DISCIPLE-COA*. We con-

clude the article with a discussion of this research.

Learning Agent Shell

We have introduced the concept of a *learning agent shell* as a new class of tools for rapid development of practical end-to-end knowledge-based agents by domain experts with limited assistance from knowledge engineers (Tecuci et al. 1999). A learning agent shell includes a general problem-solving engine and

ent aspects of it. The answers from each critiquer were displayed by a solution viewer developed by SAIC. The integrated system and its individual components were evaluated by Alphatech as exhibiting performance at the level of a subject-matter expert.

The HPKB Program emphasized the development of very large knowledge bases by knowledge engineers, demonstrating the utility of large knowledge bases and the feasibility of large-scale reuse. It has produced reusable knowledge libraries, including an upper ontology and middle theories for crisis and battlefield reasoning. The follow-on RKF Program emphasizes the development of knowledge bases directly by the domain experts. Its central objective is to enable distributed teams of SMEs to enter and modify knowledge directly and easily, without the need for prior knowledge engineering experience. The emphasis is on content and the means of rapidly acquiring this content from individuals who possess it, with the goal of gaining a scientific understanding of how ordinary people can work with formal representations of knowledge. Therefore, the program's primary requirement is the development of functions that enable SMEs to understand the contents of a knowledge base, enter new theories, augment and edit existing knowledge, test the adequacy of the knowledge base under development, receive explanations of theories contained in the knowledge base, and detect and repair errors in content. Because of the complexity of these tasks, the approaches developed in the RKF Program exploit the synergies among complementary AI technologies, such as natural language discourse processing, problem solving and learning by analogy, and commonsense reasoning.

RKF is organized in a manner similar

to HPKB, with challenge problems administered by an evaluation contractor (IET assisted by Veridian Pacific-Sierra Research (PSR) and George Mason University Institute for Biosciences, Bioinformatics, and Biotechnology) as a basis for formal evaluation of the technology provided by the developers. There are two integrated teams and several component technology developers that contribute to one or both of them. One team is coordinated by Cycorp and includes research groups from USC/ISI, University of Edinburgh, Teknowledge, SAIC, and Northwestern University. The other team is coordinated by SRI and includes research groups from the University of Texas at Austin, USC/ISI, Stanford University, Boeing, the University of Massachusetts at Amherst, Veridian/PSR, and Northwestern University. The component technology developers include Stanford University, MIT, Northwestern University, George Mason University, University of West Florida, and Pragati.

The RKF challenge problems are designed to test SMEs' abilities to use RKF tools to build knowledge bases in two main categories related to understanding biological weapons: (1) textbook knowledge and (2) expert knowledge. The textbook knowledge challenge problem covers standard textbooks in undergraduate biology and is designed to drive the development of the technology that will enable an SME to develop knowledge bases for domains in which knowledge is already relatively well organized, self contained, or comprehensive and for which there might exist recognized tests of subject understanding defined independently of the RKF Program. These tests will be based on questions appearing in the textbooks themselves or in associated academic tests for the relevant subject. The expert knowl-

edge challenge problem, however, covers practical, hands-on tasks for military analysis regarding defensive biological warfare. This problem is designed to drive the development of the technology that will enable the professional SME to develop knowledge that is task oriented, is not necessarily written down in books, and is germane for an application of national security importance. No appropriate, independently defined test exists now, so it will be necessary to define appropriate tests in the context of the RKF Program.

This article presents the DISCIPLE technology developed by the George Mason University Learning Agents Laboratory and its application to the COA challenge problem. The developed tool, called DISCIPLE-COA, is an important result of the HPKB Program that is representative of the many accomplishments of the program. DISCIPLE is now evolving from a tool for knowledge engineers to a tool that can be used directly by SMEs, as demonstrated by its successful knowledge-acquisition experiment performed at the U.S. Army Battle Command Battle Lab. In this sense, DISCIPLE is a good illustration of the transition from HPKB to RKF.

—Murray Burke

Notes

1. Murray Burke, Rapid Knowledge Formation (RKF) Program Description, dtsn.darpa.mil/iso/programtemp.asp?mode=3.
2. David Gunning and Murray Burke, High-Performance Knowledge Bases (HPKB) Program Description, dtsn.darpa.mil/iso/index2.asp?mode=9.

a general learning engine for building a knowledge base consisting of an ontology and a set of problem-solving rules. The process of developing a knowledge-based agent for a specific application relies on importing ontological knowledge from existing knowledge repositories and teaching the agent how to perform various tasks in a way that resembles how an expert would teach a human apprentice when solving problems in cooperation (figure 1). This process is based on mixed-initiative reasoning that integrates the complementary

knowledge and reasoning styles of the SME and the agent and on a division of responsibility for those elements of knowledge engineering for which they have the most aptitude, such that together they form a complete team for knowledge base development.

The concept of a learning agent shell is an extension of the concept of expert system shell (Clancey 1984). As an expert system shell, it includes a general inference engine that can be reused for multiple applications. In addition, a learning agent shell exhibits an organization of

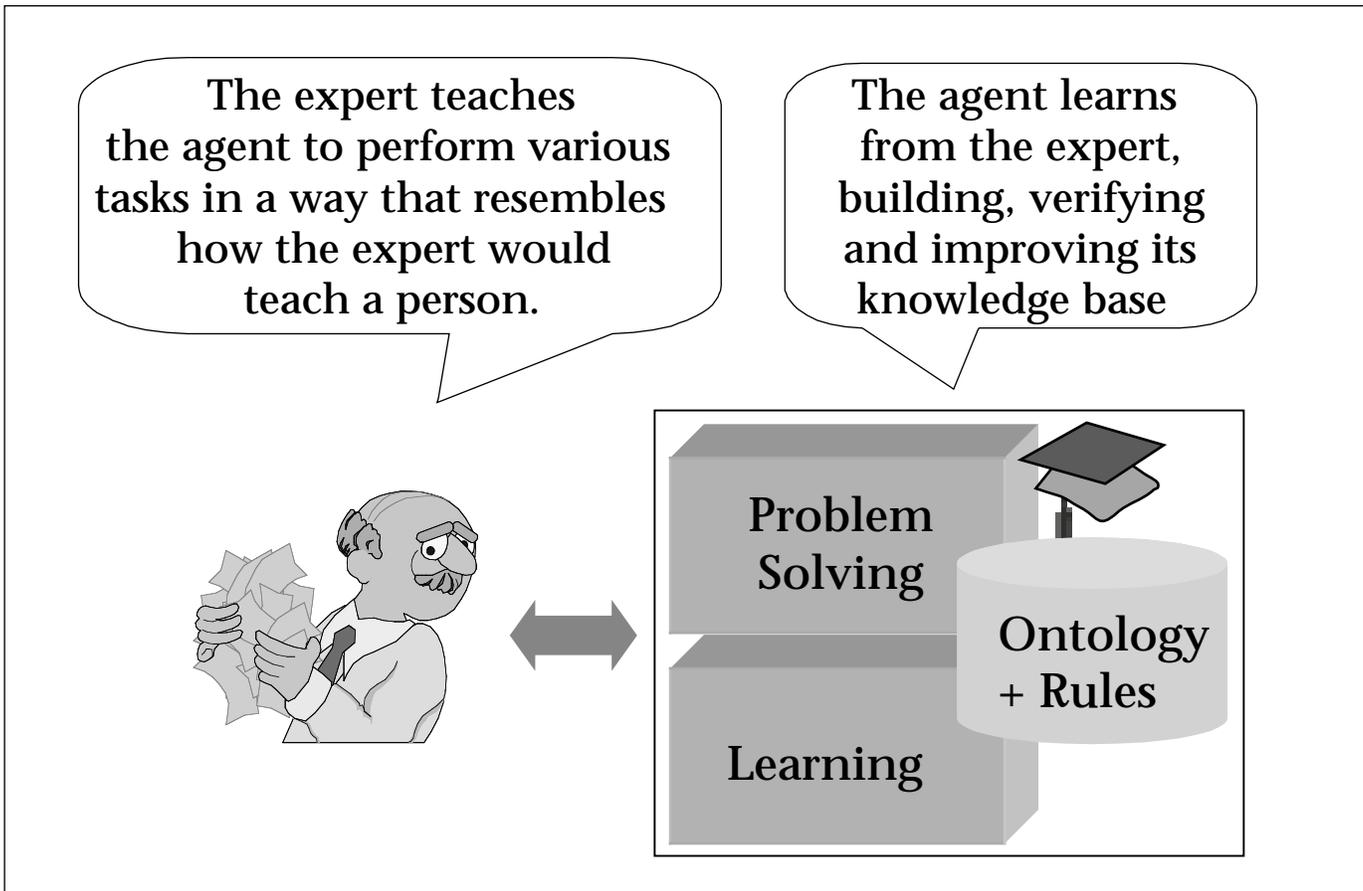


Figure 1. Interacting with a Learning Agent Shell.

the knowledge base into an ontology that specifies the terms from a particular domain and a set of problem-solving rules expressed with these terms.

The *ontology* is the more general component of the knowledge base, characteristic to an entire domain, such as medicine or the military. A *domain ontology* specifies terms that are useful in a wide range of different applications in a domain. For example, a military ontology would include specifications of military units and military equipment that are likely to be included in the knowledge base of any agent developed for a particular military application. Moreover, there is wide agreement in any mature domain on the basic terms of this domain, allowing one to reuse ontological knowledge that was previously developed to build a new knowledge base. As a consequence, a learning agent shell should include modules for importing ontological knowledge from existing knowledge bases.

The problem-solving rules represent the specific component of the knowledge base. The rules are not only specific to a particular appli-

cation in a given domain, but they are even specific to a particular SME. Consider, for example, an agent that assists a military commander in critiquing courses of action (COAs) with respect to the principles of war and the tenets of army operations (an agent that will be described in more detail in this article). The rules will identify strengths and weaknesses in a military COA and will obviously be domain specific. Moreover, they are likely to include subjective elements that are based on the experience of a specific military expert. Defining such problem-solving rules is a complex knowledge engineering task. Therefore, the learning agent shell should include mixed-initiative methods for learning such rules from a natural interaction with an SME.

The DISCIPLE Shell

Over the years, we have developed a series of increasingly more advanced learning agent shells from the DISCIPLE family. The problem-solving engine of a DISCIPLE agent is based on the general task-reduction paradigm. In this

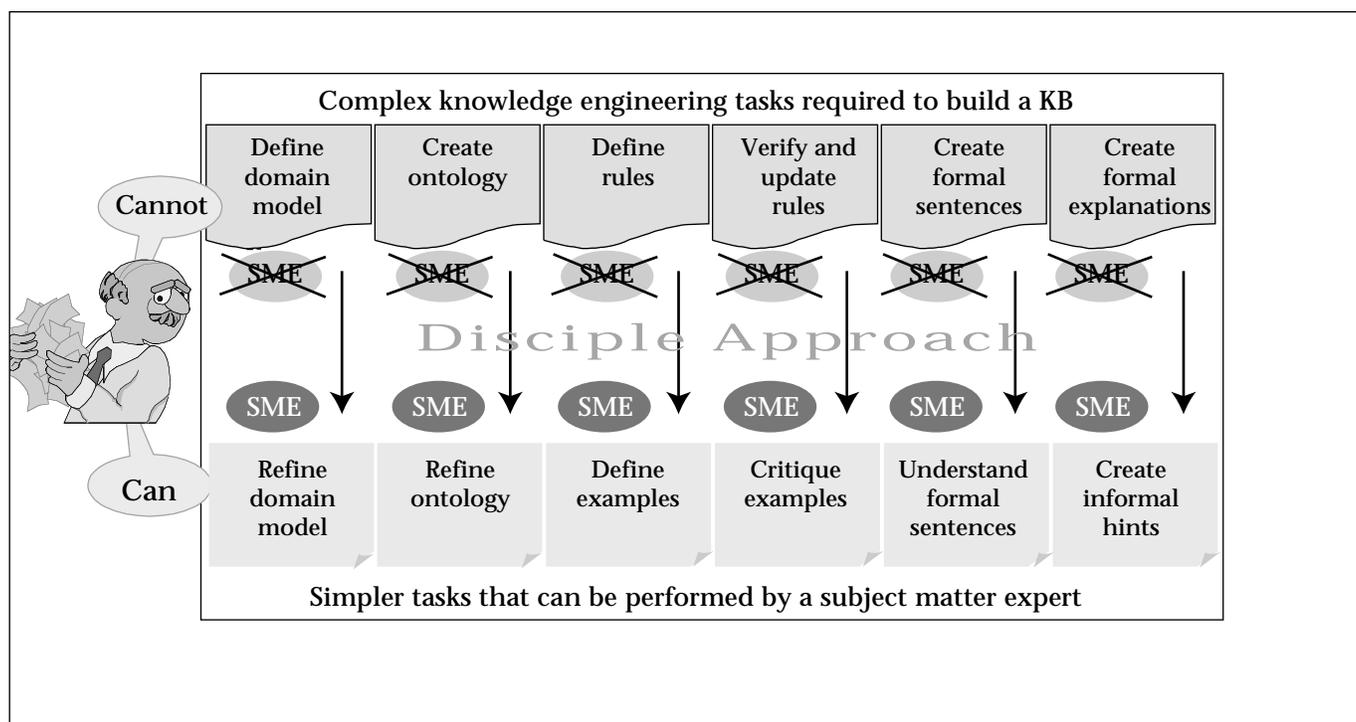


Figure 2. The General Strategy behind the DISCIPLÉ Approach.

paradigm, a task to be accomplished by the agent is successively reduced to simpler tasks until the initial task is reduced to a set of elementary tasks that can be performed immediately. The ontology of a DISCIPLÉ agent is based on the frame knowledge model of the open knowledge base connectivity (OKBC) protocol. OKBC has been developed as a standard for accessing knowledge bases stored in different frame-representation systems (Chaudhri et al. 1998). It provides a set of operations for a generic interface to such systems. There is also an ongoing effort to develop OKBC servers for various systems, such as CYC (Lenat 1995), ONTOLINGUA (Fikes, Farquhar, and Rice 1997), and LOOM (MacGregor 1991). These servers are becoming repositories of reusable ontologies and domain theories and can be accessed using the OKBC protocol. The use of the OKBC knowledge model for the DISCIPLÉ ontology facilitates the import of ontological knowledge from the OKBC compliant knowledge repositories. Because we are using the task-reduction paradigm, the problem-solving rules of a DISCIPLÉ agent are task-reduction rules. DISCIPLÉ uses an original representation of partially learned rules called *plausible version space* (PVS). Much of the power of the DISCIPLÉ approach comes from this PVS representation of the rules and from the multistrategy learning methods used to learn them. They allow DISCIPLÉ to incrementally learn a rule, starting from only one exam-

ple, and immediately use it in problem solving.

The DISCIPLÉ approach is based on several levels of synergism between the expert who has the knowledge to be formalized and the agent that incorporates knowledge engineering methods to formalize it. This multilevel synergism is achieved through mixed-initiative reasoning that integrates complementary human and automated reasoning to take advantage of their respective knowledge, reasoning styles, and computational strengths. The mixed-initiative reasoning is based on a division of responsibility between the expert and the agent for those elements of knowledge engineering for which they have the most aptitude, such that together they form a complete team for the development of the agent's knowledge base. From the point of view of the SME, this mixed-initiative approach results in the replacement of the difficult knowledge engineering tasks required to build a knowledge base, tasks that cannot be performed by an SME, with simpler tasks that can be performed by the expert, as shown in figure 2.

To build a knowledge base, we have first to develop a model of the application domain that will make explicit, at a qualitative and informal level, the way the expert solves problems. In the case of DISCIPLÉ, we have to model the process of solving a specific problem as a sequence of qualitative and informal problem-reduction steps, where a complex problem is

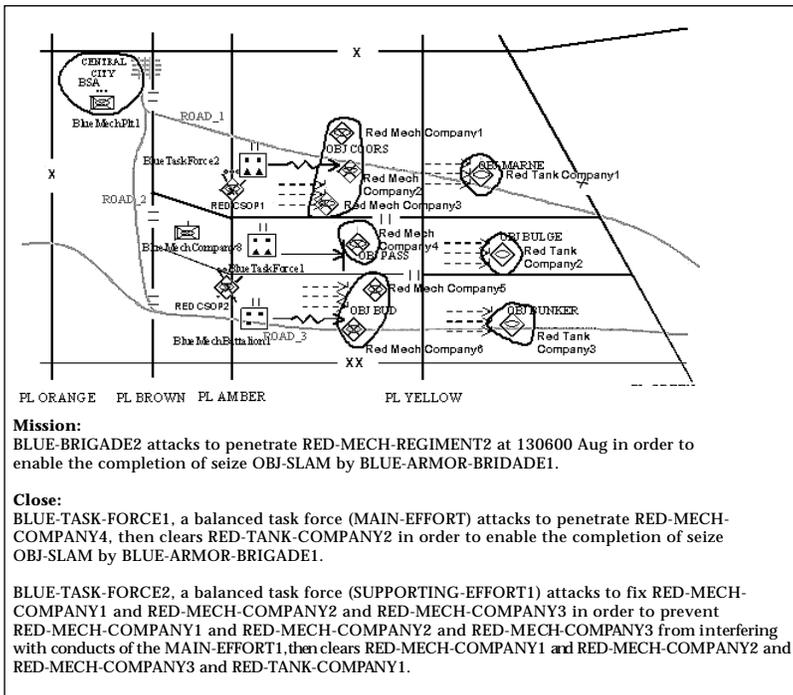


Figure 3. A Sample of a Course-of-Action (COA) Sketch and a Fragment of a COA Statement.

successively reduced to simpler problems. Then we have to build an ontology that will define the terms used to express the problems and their solutions. Finally, we have to define formal problem-solving rules, verify the rules, and update them. In general, these processes require the creation and modification of formal sentences and formal explanations. An SME cannot perform these tasks. In fact, they are very hard even for a knowledge engineer, leading to the well-known knowledge-acquisition bottleneck in the creation of knowledge-based agents.

In the DISCIPLE approach, the difficult knowledge engineering tasks required to develop a knowledge-based agent are replaced with simpler tasks that can be performed by an SME with limited assistance from a knowledge engineer. Instead of developing a complete model of the application domain, the expert would need to start with an initial model, extending and refining it with the help of the DISCIPLE agent. The use of the OKBC knowledge model for the DISCIPLE ontology facilitates the import of ontological knowledge from the OKBC compliant knowledge repositories. Therefore, instead of creating an ontology from scratch, the expert would only need to update and extend an imported ontology (Boicu et al. 1999). Instead of defining a complex problem-solving rule, the expert would only need to define a specific example of a problem-solving

episode because DISCIPLE will be able to learn a rule from this example. Instead of debugging a complex problem-solving rule, the expert would only need to critique specific examples of problem-solving episodes, and DISCIPLE will accordingly update the corresponding rule.

Most of the time, the expert would not need to create formal sentences or explanations but only to understand such sentences-explanations that were generated by DISCIPLE based on informal hints provided by the expert.

The Course-of-Action Challenge Problem

A *military COA* is a preliminary outline of a plan for how a military unit might attempt to accomplish a mission. A COA is not a complete plan; it omits many details of the operation, such as exact initial locations of friendly and enemy forces. After receiving orders to plan for a mission, a commander and staff complete a detailed and practiced process of analyzing the mission, conceiving and evaluating potential COAs, selecting a COA, and preparing detailed plans to accomplish the mission based on the selected COA. The general practice is for the staff to generate several COAs for a mission and then make a comparison of these COAs based on many factors, including the situation, the commander's guidance, the principles of war, and the tenets of army operations. The commander makes the final decision on which COA will be used to generate his/her plan based on the recommendations of the staff and his/her own experience with the same factors considered by the staff (Jones 1999).

The COA challenge problem consisted of rapidly developing a knowledge-based critiquing agent that can automatically critique COAs for ground force operations, systematically assess selected aspects of a COA, and suggest repairs to it. The role of this agent is to act as an assistant to the military commander, helping the commander to choose between several COAs under consideration for a certain mission. The agent could also help students learn to develop COAs.

The input to the COA critiquing agent consists of the description of a COA that includes the following aspects:

First is the *COA sketch*, such as the one in the top part of figure 3. It is a graphic depiction of the preliminary plan being considered. It includes enough of the high-level structure and maneuver aspects of the plan to show how the actions of each unit fit together to accomplish the overall purpose and omits much of the execution detail that will be included in

the eventual operational plan. The three primary elements included in a COA sketch are (1) control measures that limit and control interactions between units; (2) unit graphics that depict known, initial locations and composition of friendly and enemy units; and (3) mission graphics that depict actions and tasks assigned to friendly units. The COA sketch is drawn using a palette-based sketching utility.

Second is the *COA statement*, such as the partial one shown in the bottom part of figure 3. It clearly explains what the units in a COA will do to accomplish the assigned mission. This text includes a description of the mission and the desired end state as well as standard elements that describe purposes, operations, tasks, forms of maneuver, units, and resources to be used in the COA. The COA statement is expressed in a restricted but expressive subset of English.

Third are selected products of mission analysis, such as the areas of operations of the units, avenues of approach, key terrain, unit combat power, and enemy COAs.

Based on this input, the critiquing agent has to assess various aspects of the COA, such as its viability (its suitability, feasibility, acceptability, and completeness), its correctness (which considers the array of forces, the scheme of maneuver, and the command and control), and its strengths and weaknesses with respect to the principles of war and the tenets of army operations. The critiquing agent should also be able to clearly justify the assessments made and propose improvements to the COA.

General Presentation of DISCIPLE-COA

In the HPKB Program, the COA challenge problem was solved by developing an integrated system composed of four critiquers, each built by a different team, to solve a part of the overall problem: (1) a joint team from Teknowledge and Cycorp, (2) the EXPECT team from USC/ISI, (3) the LOOM team from USC/ISI, and (4) the DISCIPLE team from George Mason University. All these critiquers shared an input ontology that contains the terms needed to represent the COAs.

The COAs to be critiqued were provided by Alphatech. As presented in the previous section, each such COA is represented by a sketch and a textual description. A statement translator (developed by AIAI of the University of Edinburgh), a COA sketcher (developed by Teknowledge), and a geographic reasoner (developed by Northwestern University) transform and fuse these external representations

Assess COA411 with respect to the Principle of Mass

There is a major strength in COA411 with respect to mass because BLUE-TASK-FORCE1 is the MAIN-EFFORT1 and it acts on the decisive point of the COA (RED-MECH-COMPANY4) with a force ratio of 10.6, which exceeds a recommended force ratio of 3.0. Additionally, the main effort is assisted by supporting action SUPPRESS-MILITARY-TASK1 which also acts on the decisive point. This is good evidence of the allocation of significantly more than minimum combat power required at the decisive point and is indicative of the proper application of the principle of mass.

There is a strength in COA411 with respect to mass because BLUE-TASK-FORCE1 is the main effort of the COA and it has been allocated 33% of available combat power but this is considered just a medium level weighting of the main effort.

There is a strength in COA411 with respect to mass because BLUE-MECH-COMPANY8 is a COMPANY-UNIT-DESIGNATION level maneuver unit assigned to be the reserve. This is considered a strong reserve for a BRIGADE-UNIT-DESIGNATION level COA and would be available to continue the operation or exploit success.

Reference: FM 100-5 pg 2-4, KF 113.1, KF 113.2, KF 113.3, KF 113.4, KF 113.5 - To mass is to synchronize the effects of all elements of combat power at the proper point and time to achieve decisive results. Observance of the Principle of Mass may be evidenced by allocation to the main effort of significantly greater combat power than the minimum required throughout its mission, accounting for expected losses. Mass is evidenced by the allocation of significantly more than minimum combat power required at the decisive point.

Figure 4. Solutions Generated by DISCIPLE-COA.

into a description in the CYC language, according to the input ontology. This description is further used by all the critiquers.

SAIC completed the integrated system with a solution viewer that provided a uniform presentation of the critiques generated by the four developed critiquers.

Our critiquer, called DISCIPLE-COA, identifies the strengths and the weaknesses of a COA with respect to the principles of war and the tenets of army operations (U.S. Army 1993). There are nine principles of war: (1) objective, (2) offensive, (3) mass, (4) economy of force, (5) maneuver, (6) unity of command, (7) security, (8) surprise, and (9) simplicity. They provide general guidance for the conduct of war at the strategic, operational, and tactical levels. The tenets of army operations describe the charac-

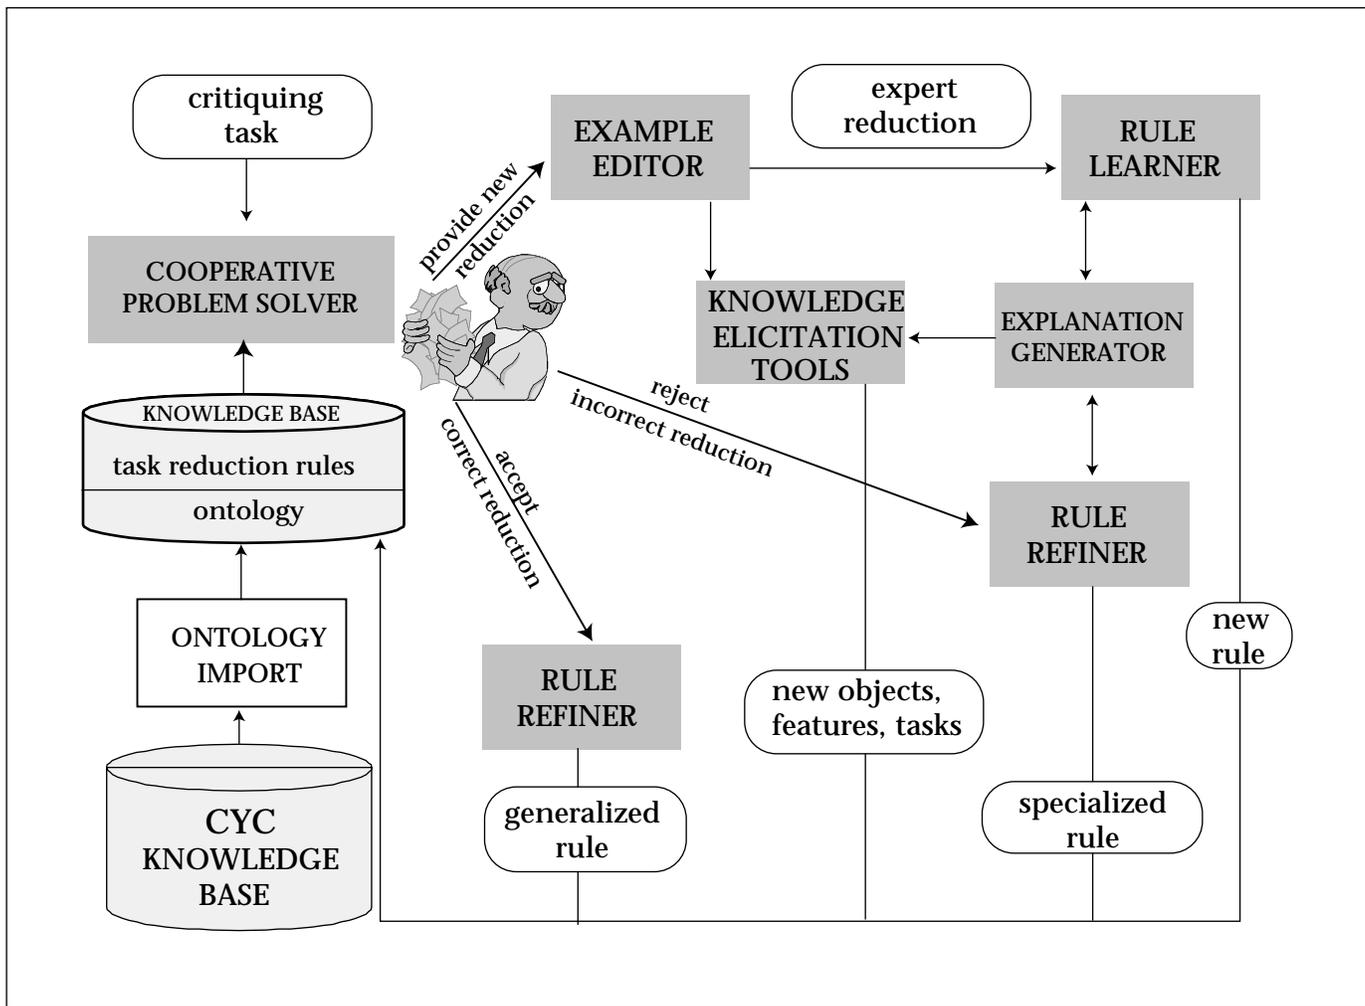


Figure 5. Expert Agent Interactions.

teristics of successful operations. There are five such tenets: (1) initiative, (2) agility, (3) depth, (4) synchronization, and (5) versatility. Figure 4, for example, shows some of the strengths of the COA from figure 3 with respect to the principle of mass, identified by DISCIPLÉ-COA.

In addition to generating answers in natural language, DISCIPLÉ also provides the reference material based on which answers are generated, as shown in the bottom part of figure 4. Also, the DISCIPLÉ-COA agent can provide justifications for the generated answers at three levels of detail, from a very abstract one (that shows the general line of reasoning followed) to a very detailed one (that indicates each of the knowledge pieces used in generating the answer).

Figure 5 shows the main modules of DISCIPLÉ-COA and their interactions. The ontology import module was used to integrate the input COA ontology into DISCIPLÉ's knowledge base and translate the CYC representation of specific

COAs into DISCIPLÉ's representation.

After the description of a specific COA, such as the one from figure 3, has been loaded into DISCIPLÉ's knowledge base, the SME can teach the DISCIPLÉ agent how to critique it. The expert invokes the cooperative problem solver, selects an initial critiquing task (such as "Assess COA with respect to the Principle of Surprise"), and asks the DISCIPLÉ agent to solve it. DISCIPLÉ uses its task-reduction rules to reduce the current task to simpler tasks, showing the expert the reductions found. The expert can accept a reduction proposed by the agent, can reject it, or can decide to define himself/herself a new reduction, as illustrated in figure 5.

To define a new reduction, the expert uses the *example editor*, which, in turn, can invoke several knowledge elicitation tools, such as the object editor, the feature editor, or the task editor, if the specification of the example involves new knowledge elements that are not present in the current knowledge base. Once the reduc-

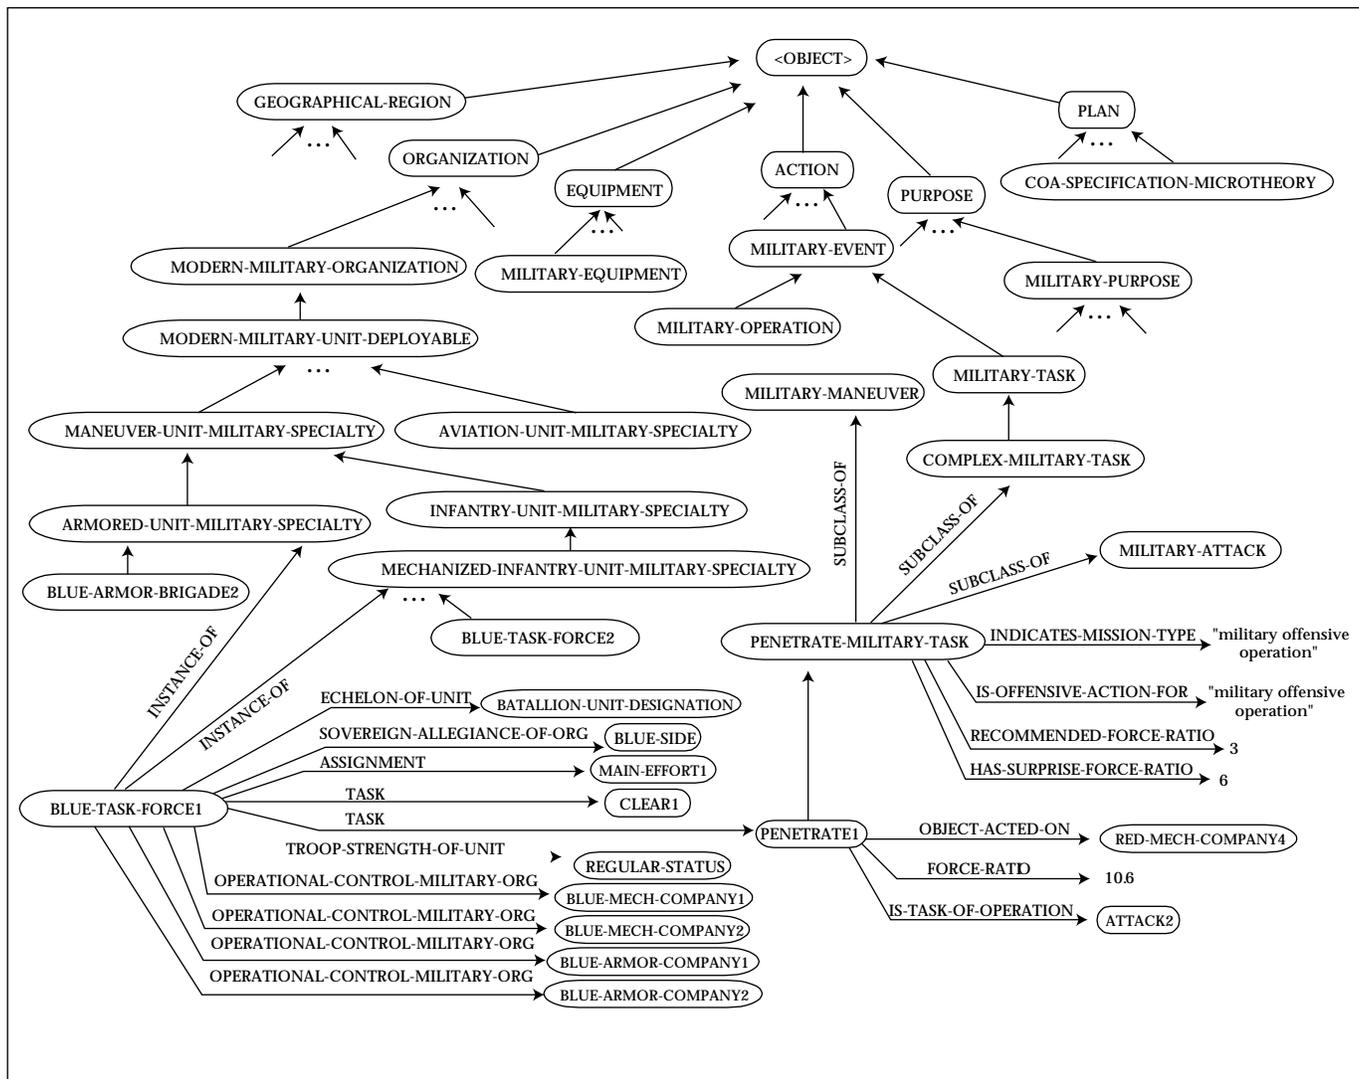


Figure 6. Fragment of the Ontology.

tion has been defined, the *rule learner* is invoked to generalize the example to a task-reduction rule. The rule learner automatically invokes the *explanation generator* that tries to find the explanation of why the reduction indicated by the expert is correct. The explanation generator proposes several plausible explanations from which the expert has to select the correct ones. The expert can help the agent find the correct explanations by providing a hint.

If the expert accepts a reduction proposed by the agent, then the *rule refiner* is invoked and can generalize the rule that has led to this reduction.

If the expert rejects a reduction proposed by the agent, then the agent attempts to find an explanation of why the reduction is not correct, with the explanation generator invoked

again, as described earlier. The explanation found is used by the rule refiner to specialize the rule.

After a new rule is learned or an existing rule is refined, the cooperative problem solver resumes the task-reduction process until a solution of the initial problem is found.

In addition to the cooperative problem solver, DISCIPLE-COA also includes an autonomous problem solver that is used after DISCIPLE-COA has been trained.

The next sections describe in more detail this process of developing the disciple-coa agent.

Ontology Development

For DISCIPLE-COA, an initial ontology was defined by importing the input ontology built by Teknowledge and Cycorp for the COA chal-

lenge problem. The imported ontology was further developed by using the ontology-building tools of DISCIPLE. These tools include specialized browsers and editors for the various knowledge pieces of DISCIPLE (for example, the object editor, the object feature editor, the task feature editor, the hierarchy browser, and the association browser).

DISCIPLE's ontology includes objects, features, and tasks, all represented as frames, according to the knowledge model of the OKBC protocol (Chaudhri et al. 1998).

The objects represent either specific individuals or sets of individuals. The objects are hierarchically organized according to the generalization relation (subclass-of/superclass-of and instance-of/type-of). Figure 6, for example, presents a fragment of the object ontology. The top part represents the upper level of the object ontology that identifies the types of concept represented in the ontology. They include GEOGRAPHICAL-REGION, ORGANIZATION, EQUIPMENT, and ACTION. Each of these concepts is the top of a specialized hierarchy, such as the hierarchy of organizations showed in the left part of figure 6. The leaves of this hierarchy are specific military units, corresponding to a specific COA to be critiqued by DISCIPLE. Each concept and instance of the object hierarchy is described by specific features and values. For example, the bottom part of figure 6 shows the description of the specific military unit called BLUE-TASK-FORCE1. BLUE-TASK-FORCE1 is described as both an ARMORED-UNIT-MILITARY-SPECIALTY and a MECHANIZED-INFANTRY-UNIT-MILITARY-SPECIALTY. The other features describe BLUE-TASK-FORCE1 as being at the battalion level; belonging to the blue side, designated as the main effort of the blue side; performing two tasks, PENETRATE1 and CLEAR1; having a regular strength; and having under its operational control four other units. The values of the features of BLUE-TASK-FORCE1 are themselves described in the same way. For example, one of the tasks performed by BLUE-TASK-FORCE1 is PENETRATE1. PENETRATE1 is defined as a penetration task and, therefore, inherits all the features of the penetration tasks in addition to the features that are directly associated with it.

The hierarchy of objects is used as a generalization hierarchy for learning by the DISCIPLE agent. One way to generalize an expression is to replace an object with a more general one from such a hierarchy. For example, PENETRATE-MILITARY-TASK from the bottom right side of figure 6 can be generalized to COMPLEX-MILITARY-TASK, MILITARY-MANEUVER, or MILITARY-ATTACK. The goal of the learning process is to select the correct generalization.

The features used to describe the objects and

the tasks are themselves represented in the feature hierarchy.

Training the DISCIPLE-COA Agent

The next step in the development of the DISCIPLE-COA critiquer was to teach DISCIPLE to critique COAs with respect to the principles of war and the tenets of army operations. The expert loads the description of a specific COA, such as COA411 represented in figure 3, and then invokes the cooperative problem solver with a task of critiquing the COA with respect to a certain principle or tenet. DISCIPLE uses its task-reduction rules to reduce the current task to simpler tasks, showing the expert the reductions found. The expert can accept a reduction proposed by the agent, reject it, or decide to define a new reduction. From each such interaction, DISCIPLE will either refine a previously learned rule or will learn a new task-reduction rule. After a new rule is learned or an existing rule is refined, the cooperative problem solver resumes the task-reduction process until a solution of the initial problem is found.

Initially, DISCIPLE does not contain any rules. Therefore, all the problem-solving steps (that is, task reductions) must be provided by the expert, as illustrated in figure 7 and explained in the following.

To assess COA411 with respect to the principle of security, the expert and DISCIPLE need a certain amount of information, which is obtained by asking a series of questions. The answer to each question allows one to reduce the current assessment task to a more detailed one. This process continues until the expert and DISCIPLE have enough information about COA411 to make the assessment. As shown in figure 7, the initial task is reduced to that of assessing the security of COA411 with respect to the countering of enemy reconnaissance. Then one asks whether there is any enemy reconnaissance unit present in COA411. The answer identifies RED-CSOP1 as such a unit because it is performing the task SCREEN1. Therefore, the task of assessing security for COA411 with respect to countering enemy reconnaissance is now reduced to the better-defined task of assessing security when enemy reconnaissance is present. The next question to ask is whether the enemy reconnaissance unit is destroyed. In the case of COA411, RED-CSOP1 is destroyed by the task DESTROY1. Therefore, one can conclude that there is a strength in COA411 with respect to the principle of security because the enemy reconnaissance unit is countered.

Figure 8 illustrates the process of teaching

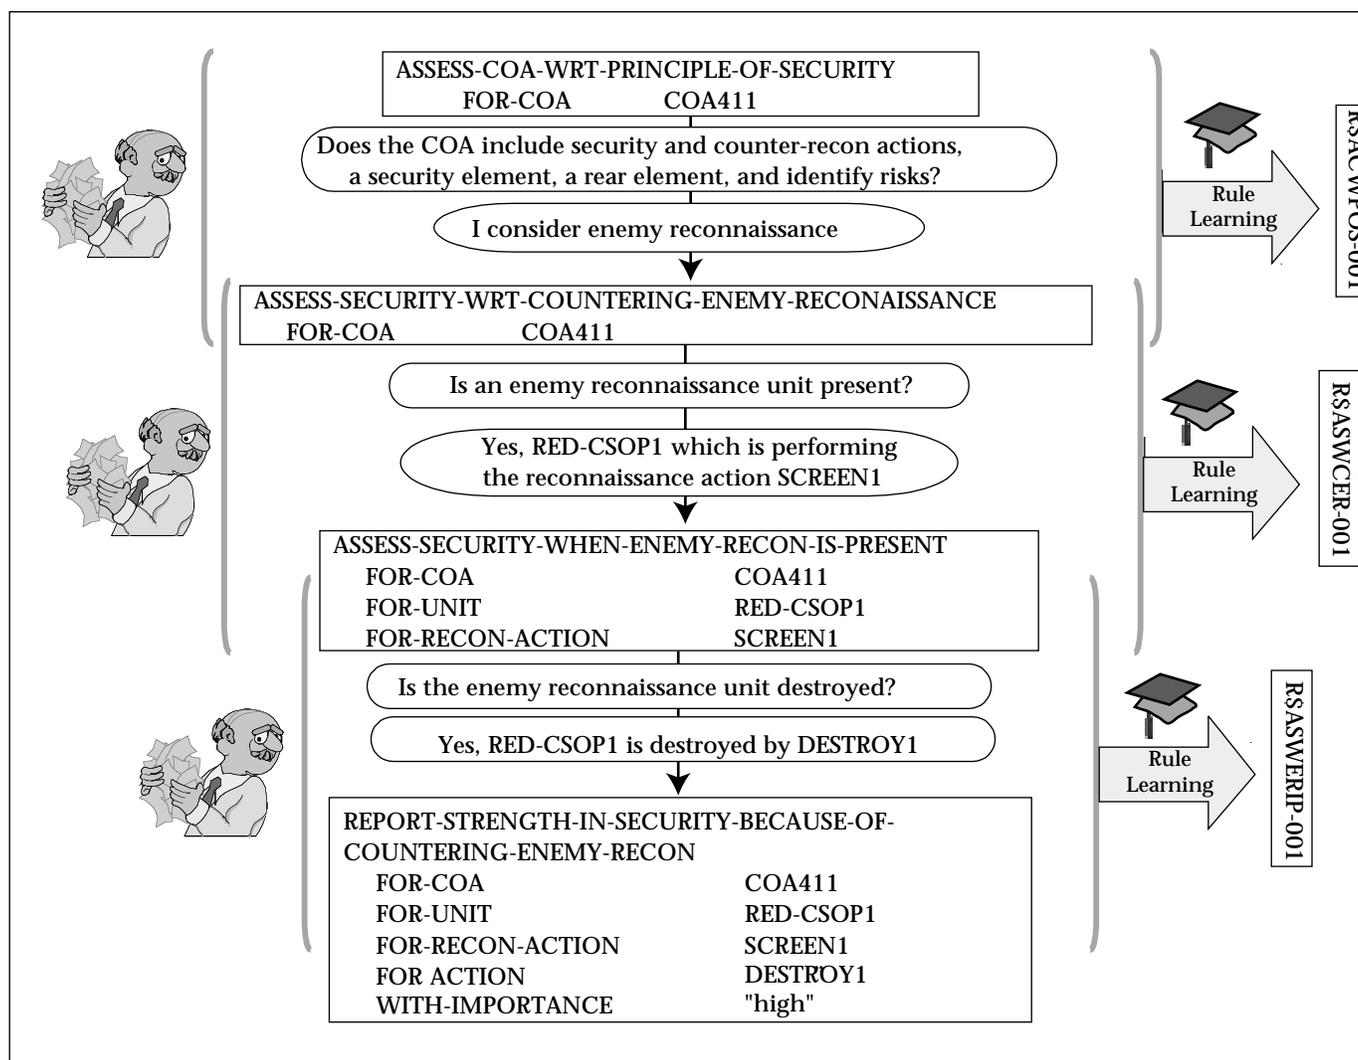


Figure 7. Task Reductions Indicated by the Expert.

DISCIPLE. The left side represents the reasoning process of the expert, where the question and the answer are in free natural language format. Although this line of reasoning is natural to a human expert, a learning agent cannot understand it. The explanation that would be understood by the agent is represented in the upper right part of figure 8 and consists of various relations between certain elements from its ontology. The first explanation piece states, in DISCIPLE's formal language, that RED-CSOP1 is an enemy unit. The second explanation piece expresses the fact that RED-CSOP1 is performing the action SCREEN1. Finally, the last explanation piece expresses the fact that SCREEN1 is a reconnaissance action. Although an expert can understand the meaning of these formal expressions, he/she cannot define them because he/she is not a knowledge engineer.

For one thing, he/she would need to use the formal language of the agent. Moreover, he/she would also need to know the names of the potentially many thousands of concepts and features from the agent's ontology.

Although defining the formal explanations of this task-reduction step is beyond the individual capabilities of the expert and the agent, it is not beyond their joint capabilities. Finding these explanation pieces is a mixed-initiative process of searching the agent's ontology because an explanation piece is a path of objects and relations in this ontology. In essence, the agent will use analogical reasoning and help from the expert to identify and propose a set of plausible explanation pieces from which the expert will have to select the correct ones. One explanation-generation strategy is based on an ordered set of heuristics for analog-

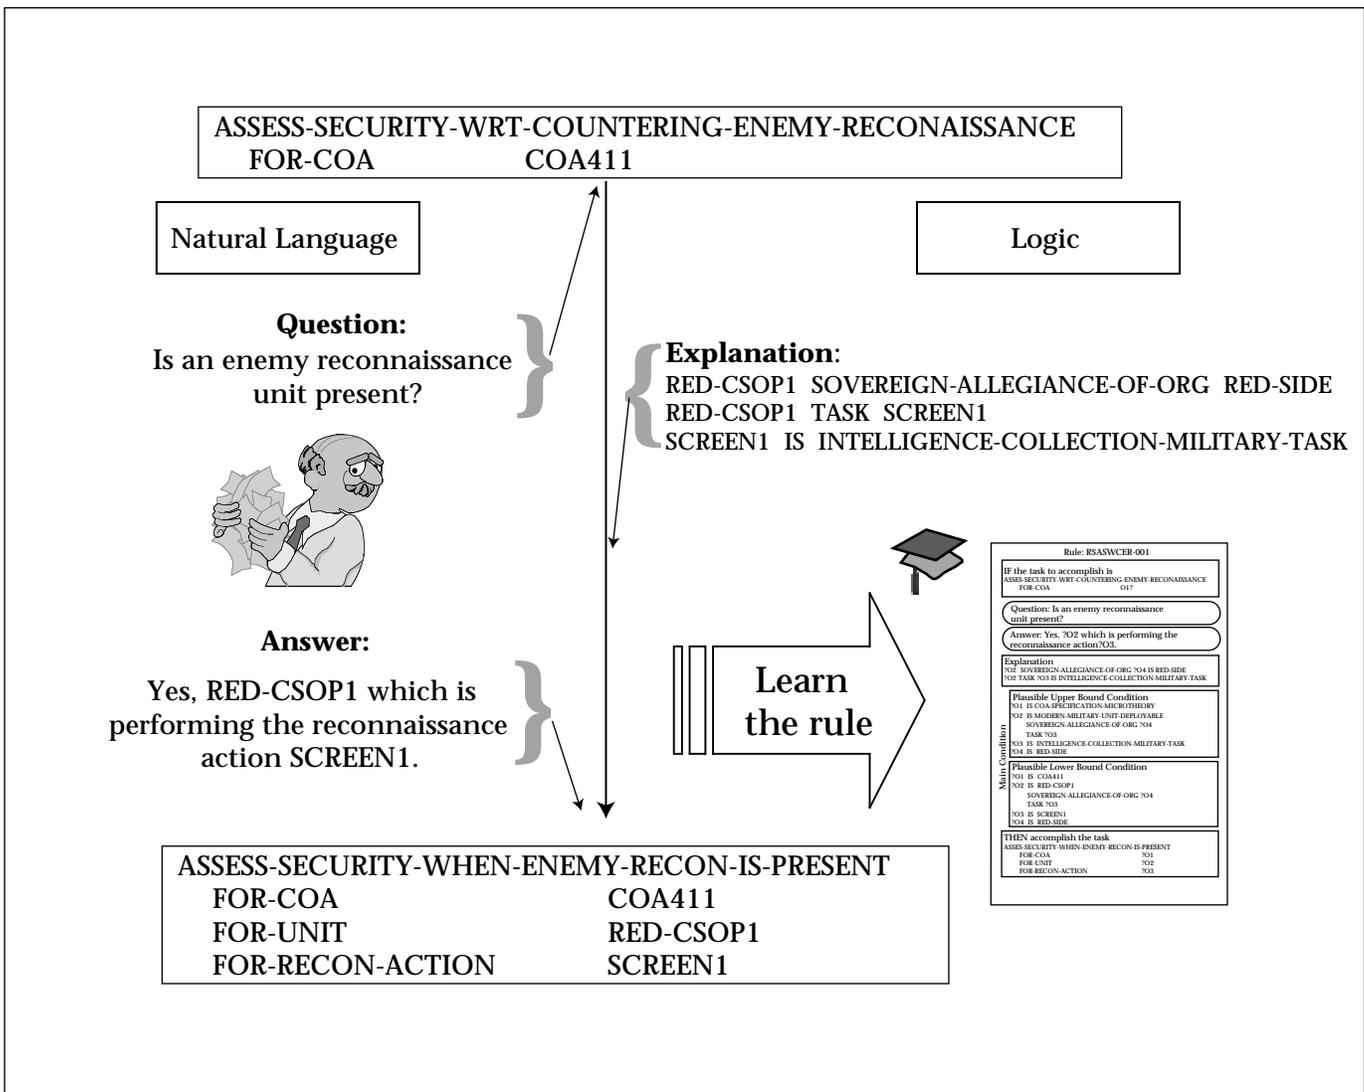


Figure 8. Teaching DISCIPLINE to Reduce a Task.

ical reasoning. These heuristics exploit the hierarchies of objects, features, and tasks to identify the rules that are similar to the current reduction and use their explanations as a guide to search for similar explanations of the current example. This cooperative explanation-generation process proved to be effective, as demonstrated by the successful knowledge-acquisition experiment described in this article.

From the example reduction and its explanation in figure 8, DISCIPLINE automatically generated the PVS rule in figure 9, which is an if-then rule, the components of which are generalizations of the elements of the example in figure 8. In addition, the rule contains two conditions for its applicability: (1) a plausible lower-bound condition and (2) a plausible upper-bound condition. These conditions approximate an

exact applicability condition that DISCIPLINE attempts to learn. Initially, the plausible lower-bound condition covers only the example in figure 8, restricting the variables from the rule to take only the values from this example. It also includes the relations between these variables that have been identified as relevant in the explanation of the example. The plausible upper-bound condition is the most general generalization of the plausible lower-bound condition. It is obtained by taking into account the domains and the ranges of the features from the plausible lower-bound conditions and the tasks to determine the possible values of the variables. The domain of a feature is the set of objects that might have this feature. The range is the set of possible values of this feature. For example, ?O2 is the value of the task feature FOR-UNIT and has as features SOVEREIGN-

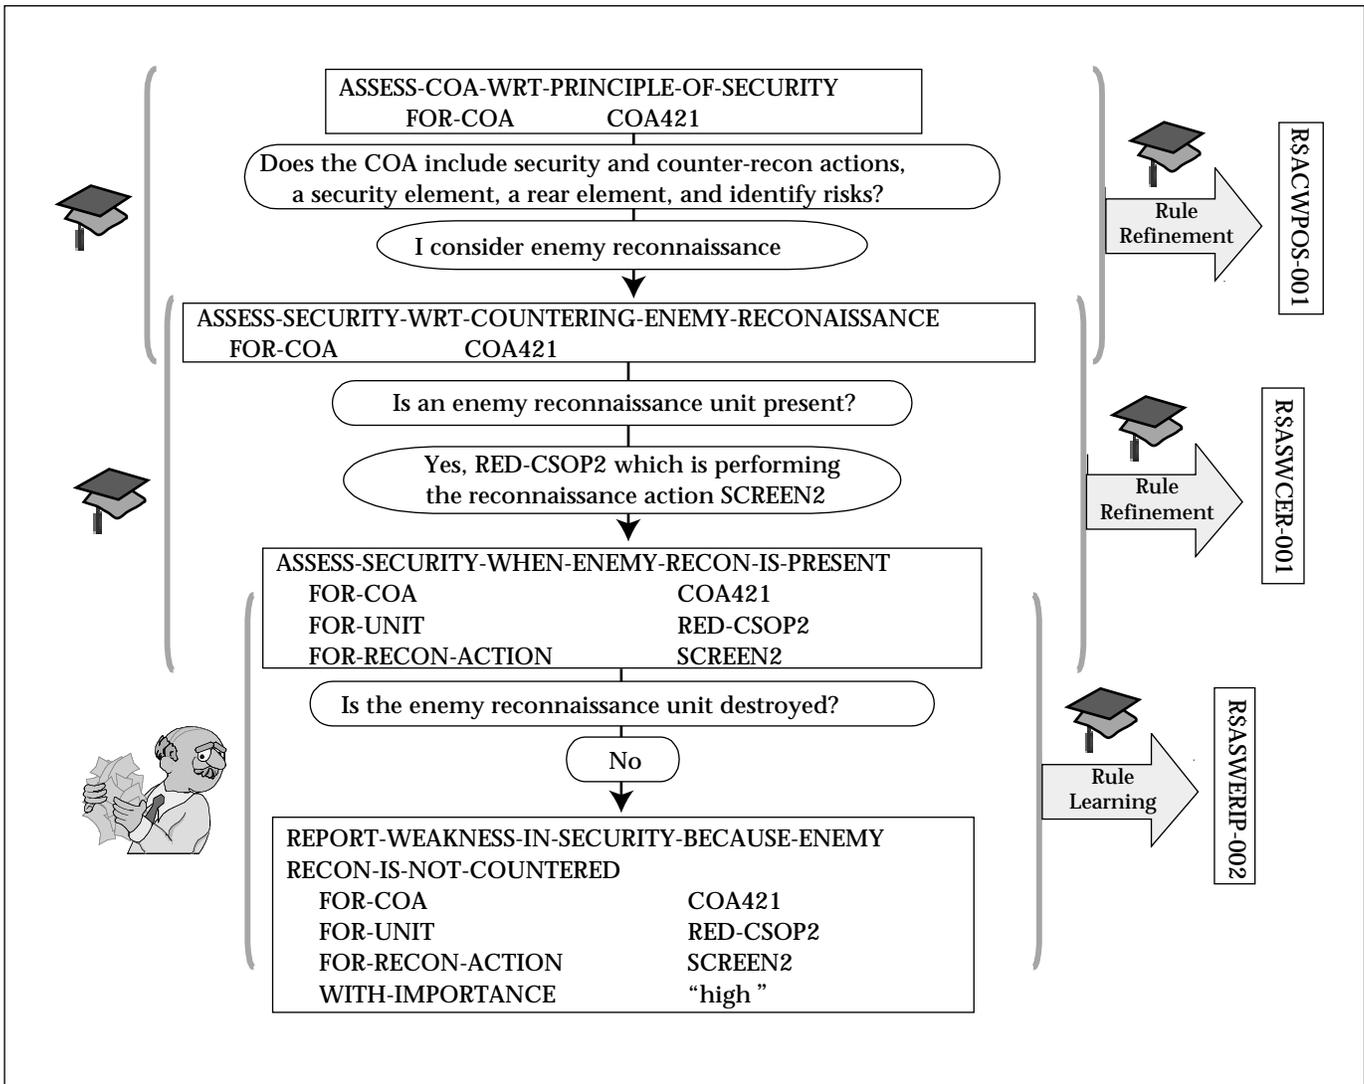


Figure 10. Cooperative Problem Solving and Learning.

Through refinement, the task-reduction rules can become significantly more complex than the rule in figure 9 (Boicu et al. 2000). For example, when a reduction proposed by DISCIPLE is rejected by the expert, DISCIPLE attempts to find an explanation for why the reduction is wrong. Then the rule can be refined with an except-when plausible version-space condition. The bounds of this version space are generalizations of the explanations that should not hold for the reduction rule to be applicable.

In any case, comparing the left-hand side of figure 8 (which is defined by the domain expert) with the rule from figure 9 (which is learned by disciple) suggests the usefulness of disciple for knowledge acquisition. In the traditional knowledge engineering approach, a knowledge engineer would need to manually define and debug a rule such as the one in fig-

ure 9. With disciple, the domain expert needs only to define an example reduction because disciple learns and refines the corresponding rule. This approach works well, as is demonstrated by the intense experimental studies conducted with disciple that are reported in the next section.

Evaluation of the Course-of-Action Critiquers

In addition to George Mason University, other three research groups have developed COA critiquers as part of the HPKB Program. Teknowledge and Cycorp have developed a critiquer based on the CYC system (Lenat 1995). The other two critiquers have been developed at USC/ISI, one based on the EXPECT system (Kim

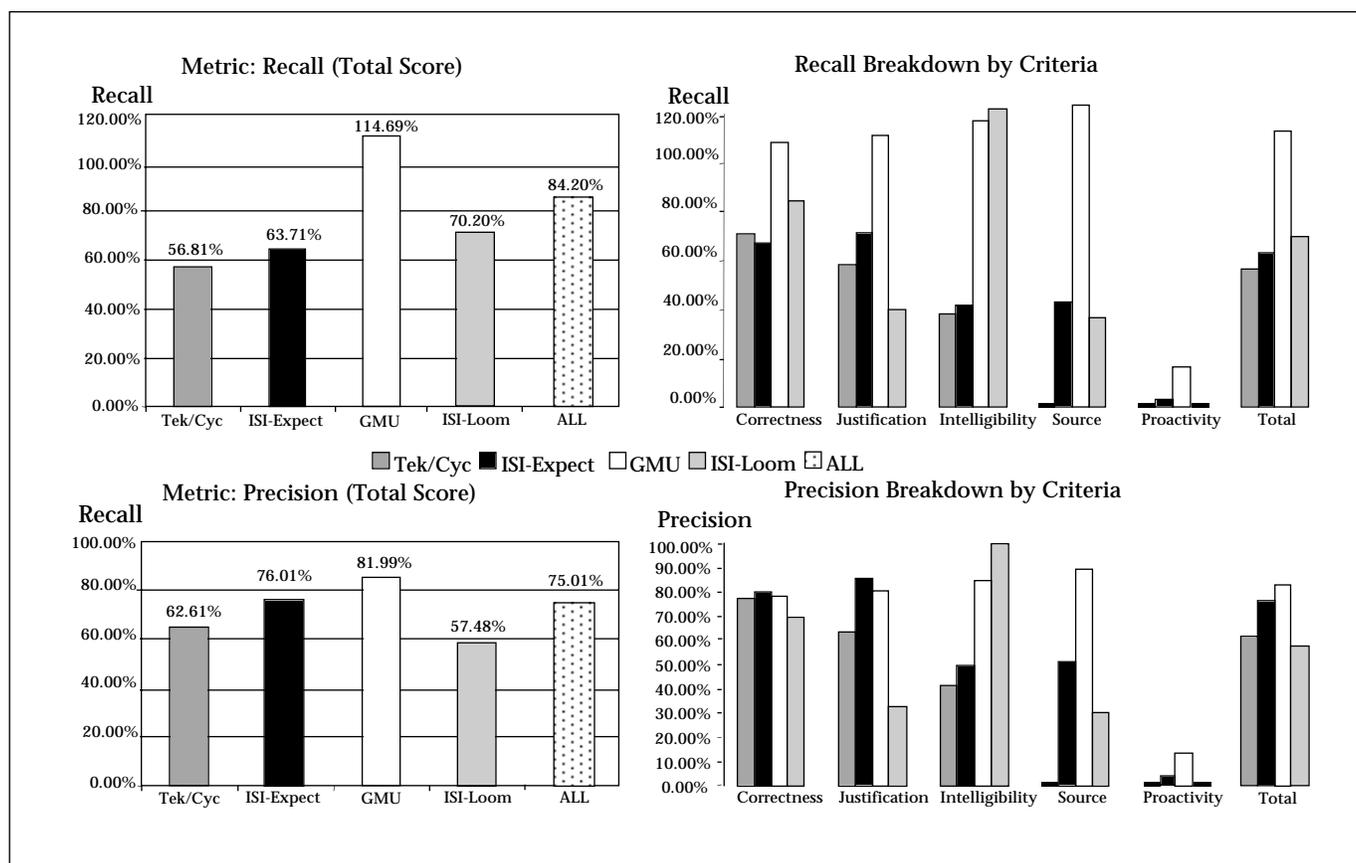


Figure 11. The Performance of the Developed Course-of-Action Critiquers and the Integrated System.

and Gil 1999) and the other based on the LOOM system (MacGregor 1999).¹ All the critiquers were evaluated as part of the HPKB's annual evaluation that took place from 6 to 16 July 1999 and included five evaluation items of increasing difficulty. Each item consisted of descriptions of various COAs and a set of questions to be answered about each of them. Item 1 consisted of COAs and questions that were previously provided by DARPA to guide the development of the COA critiquing agents. Item 2 included new test questions about the same COAs. Items 3, 4, and 5 consisted of new COAs that were increasingly more complex and required further development of the COA agents to properly answer the asked questions. Each of items 3, 4, and 5 consisted of two phases. In the first phase, each team had to provide initial system responses. Then the evaluator issued the model answers, and each team had a limited amount of time to repair its system, perform further knowledge acquisition, and generate revised system responses.

The responses of each system were scored by a team of domain experts along the following dimensions and associated weights: correctness—50 percent (matches model answer

or is otherwise judged to be correct), justification—30 percent (scored on presence, soundness, and level of detail), lay intelligibility—10 percent (degree to which a lay observer can understand the answer and the justification), sources—10 percent (degree to which appropriate sources are noted), and proactivity—10-percent extra credit (appropriate corrective actions or other information suggested to address the critique). Based on these scores, several classes of metrics have been computed, including recall and precision. *Recall* is obtained by dividing the score for all answers provided by a critiquer to the total number of model answers for the asked questions. This number was over 100 percent with our critiquer, primarily because of the extra credit received for generating additional critiques that were not among the model answers provided by the evaluator. *Precision* is obtained by dividing the same score by the total number of answers provided by the system (both the model answers provided by the evaluator and the new answers provided by the critiquer). The results obtained by the four evaluated critiquers are presented in figure 11. These graphs also show the averages of these results,

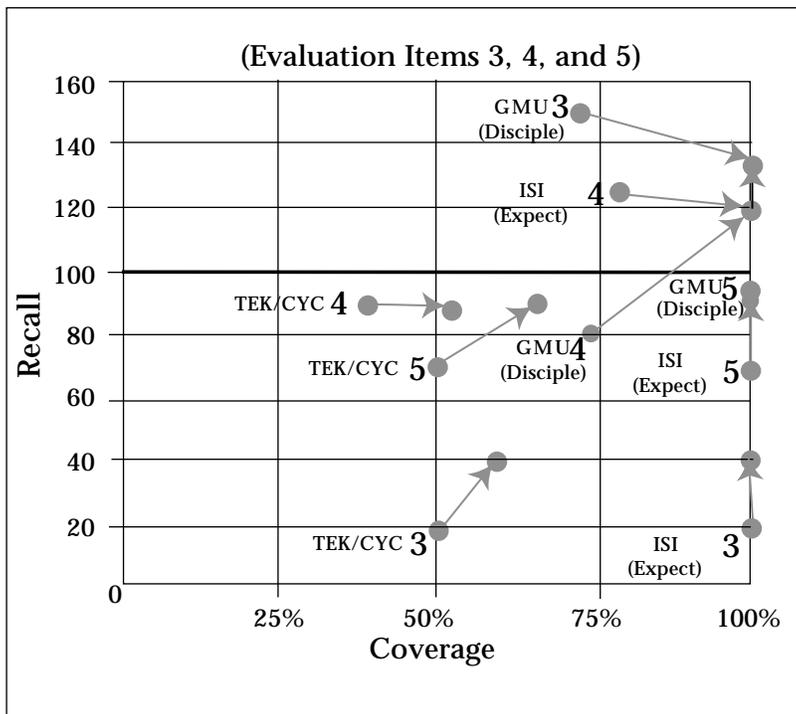


Figure 12. Coverage versus Recall, Prerepair, and Postrepair.

which represent the recall and the precision of the integrated system.

Figure 12 compares the recall and the coverage of the developed critiquers for the last three most complex items of the evaluation. For each item, the beginning of each arrow shows the coverage and recall for the initial testing phase, and the end of the arrow shows the same data for the modification phase. In this graph, the results that are above and to the right are superior to the other results. This graph also shows that all the systems increased their coverage during the evaluation. In particular, the knowledge base of DISCIPLE increased by 46 percent (from the equivalent of 6229 simple axioms to 9092 simple axioms), which represents a very high rate of knowledge acquisition of 286 simple axioms a day.

Direct Knowledge Acquisition from Subject-Matter Experts

During August 1999, we conducted a one-week knowledge-acquisition experiment with DISCIPLE-COA at the U.S. Army Battle Command Battle Lab in Fort Leavenworth, Kansas, to test the claim that domain experts that do not have prior knowledge engineering experience can teach DISCIPLE-COA (Tecuci et al. 2000). The experiment involved four such military experts

and had three phases: (1) joint training (days 1 to 3), (2) an individual teaching experiment (day 4), and (3) a joint discussion of the experiment (day 5). The entire experiment was videotaped. The training for the experiment included a detailed presentation of DISCIPLE's knowledge representation, problem-solving and learning methods, and tools. For the teaching experiment, each expert received a copy of DISCIPLE-COA with a partial knowledge base. This knowledge base was obtained by removing the tasks and the rules from the complete knowledge base of DISCIPLE-COA. That is, the knowledge base contained the complete ontology of objects, object features, and task features. We also provided the experts with the descriptions of three COAs—(1) COA411, (2) COA421, and (3) COA51—to be used for training DISCIPLE. These were the COAs used in the final phases of DARPA's evaluation of all the critiquers. Finally, we provided and discussed with the experts the modeling of critiquing of these COAs with respect to the principles of offensive and security. That is, we provided the experts with specific task reductions such as the one from figure 7 to guide them in teaching DISCIPLE-COA. Each expert then taught DISCIPLE-COA independently while he/she was supervised by a knowledge engineer, whose role was to help the expert if he/she reached an impasse while he/she was using DISCIPLE.

Figure 13 shows the evolution of the knowledge base during the teaching process for one of the experts, results that are representative for all four experts. In the morning the expert taught DISCIPLE to critique COAs with respect to the principle of offensive and in the afternoon he taught it to critique COAs with respect to the principle of security. In both cases, the expert first used COA411, then COA422 and then COA51. As one can see from figure 13, DISCIPLE initially learned more rules, and then the emphasis shifted to rule refinement. Therefore, the increase in the size of the knowledge base is greater toward the beginning of the training process for each principle. The teaching for the principle of offensive took 101 minutes. During this time, DISCIPLE learned 14 tasks and 14 rules (147 simple axioms equivalent). The teaching for security took place in the afternoon and consisted of 72 minutes of expert-DISCIPLE interactions. During this time, DISCIPLE learned 14 tasks and 12 rules (136 simple axioms equivalent). There was no or very limited assistance from the knowledge engineer with respect to teaching. The knowledge-acquisition rate obtained during the experiment was very high (9 tasks and 8 rules an hour, or 98 simple axioms equivalent an hour). At the end

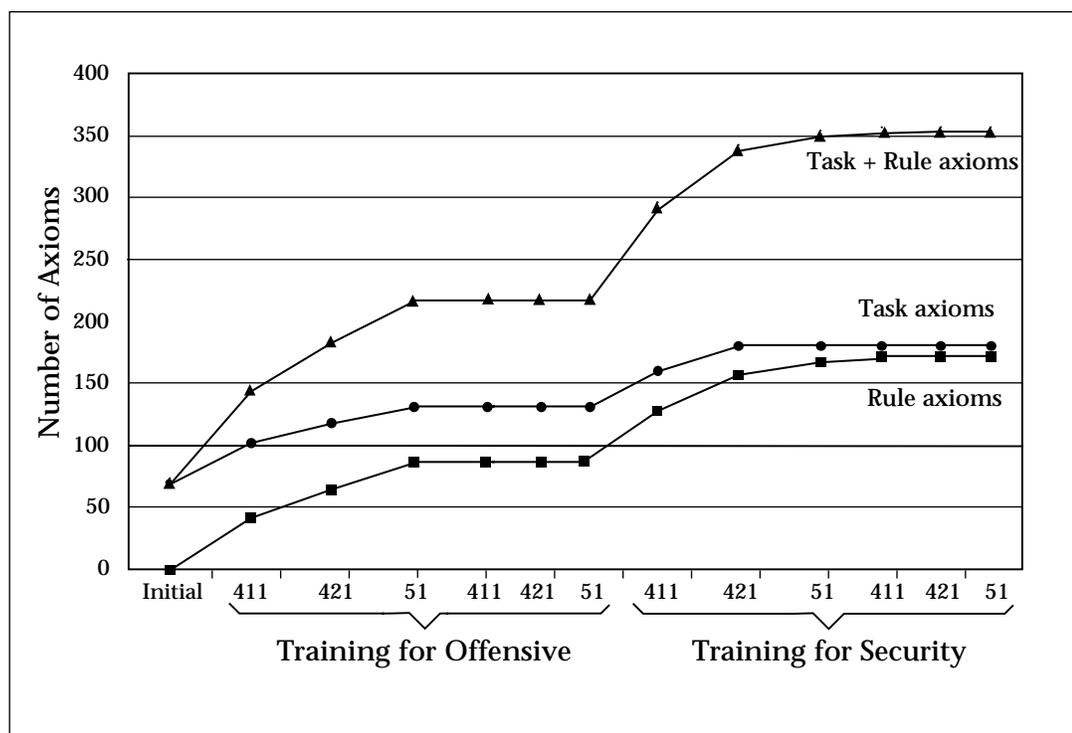


Figure 13. The Evolution of the Knowledge Base during the Knowledge-Acquisition Experiment.

of this training process, DISCIPLE-COA was able to correctly identify 17 strengths and weaknesses of the 3 COAs with respect to the principles of offensive and security.

After the experiment, each expert was asked to fill out a detailed questionnaire designed to collect subjective data for usability evaluation. All the answers took into account that DISCIPLE-COA was a research prototype and not a commercial product and were rated based on a scale of agreement with the question from 1 to 5, with 1 denoting not at all and 5 denoting very. For illustration, table 1 shows three questions and the answers provided by the four experts.

Conclusions

We introduced the concept of a learning agent shell and a methodology for rapid development of knowledge bases and agents based on the DISCIPLE learning agent shell. The DISCIPLE shell and methodology have been applied to the development of a critiquing agent that acts as an assistant to a military commander. This approach and the developed agent have been evaluated in two intensive studies. The first study concentrated on the quality of the developed critiquer and the ability to rapidly extend it by its developers and SMEs. The second study concentrated on the ability of domain experts to extend the knowledge base of the critiquer

with limited assistance from knowledge engineers. Both studies have shown that DISCIPLE has reached a significant level of maturity, being usable to rapidly develop complex knowledge-based agents.

The two main factors that contributed to the success of DISCIPLE-COA are (1) the synergistic collaboration between the SME and DISCIPLE in developing the knowledge base and (2) the multistrategy learning method of DISCIPLE that is based on the plausible version-space representation. Our research on PVSs has its origins in Mitchell's (1997) influential work on version spaces and his candidate elimination algorithm. We extended them along several dimensions, which led to a powerful and practical mixed-initiative multistrategy learning approach that synergistically integrates a wide range of knowledge-acquisition and machine learning strategies, including apprenticeship learning, empirical inductive learning from examples and explanations, and analogical learning. This method is based on a powerful knowledge representation language that includes the frame-based OKBC knowledge model for the representation of the ontological knowledge and complex task-reduction rules with multiple conditions. Moreover, we do not make the assumption that the representation space for learning needs to be completely defined before learning can take place. On the

Questions	Answers
Do you think that Disciple is a useful tool for knowledge acquisition?	<ul style="list-style-type: none"> • Rating 5. Absolutely! The potential use of this tool by domain experts is only limited by their imagination—not their AI programming skills. • 5 • 4 • Yes, it allowed me to be consistent with logical thought.
Do you think that Disciple is a useful tool for Problem Solving?	<ul style="list-style-type: none"> • Rating 5. Yes. • 5 (absolutely) • 4 • Yes. As it develops and becomes tailored to the user, it will simplify the tedious tasks.
Were the procedures / processes used in Disciple compatible with Army doctrine and/or decision making processes?	<ul style="list-style-type: none"> • Rating 5. As a minimum yes, as a maximum—better! • This again was done very well. • 4 • 4

Table 1. Sample Questions Answered by the Experts.

contrary, the representation language is assumed to be incomplete and partially incorrect and is itself evolving during rule learning through the improvement of the ontology. Because the learning process takes place in an evolving representation language, the various plausible bounds of a rule are subject to heuristic transformations that involve both generalization and specialization operations for each bound. These transformations are guided by hints, explanations, and analogies. Therefore, the learning process is efficient and does not suffer from any combinatorial explosion. Also, learning can take place even in the presence of exceptions when there is no rule that discriminates between the positive examples and the negative examples.

To conclude, our long-term vision for DISCIPLE, which guides our future work, is to evolve it to a point where it will allow normal computer users to build and maintain knowledge bases and knowledge-based agents as easily as they now use personal computers for text processing.

Acknowledgments

This research was done in the George Mason University Learning Agents Laboratory (LALAB). Research of the LALAB is sponsored by the Defense Advanced Research Projects Agency (DARPA), the Air Force Office of Scientific Research (AFOSR), and Air Force Research Lab-

oratory, Air Force Material Command, USAF, under agreement F30602-00-2-0546, grant F49620-97-1-0188 and grant F49620-00-1-0072. The U.S. government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. 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, AFOSR, the Air Force Research Laboratory, or the U.S. government. The evaluation of the COA critiquers in the HPKB program was conducted by Alphatech. The experts that participated in the BCBL knowledge-acquisition experiment were Lt. Col. John N. Duquette, Lt. Col. Jay E. Farwell, Maj. Michael P. Bowman, and Maj. Dwayne E. Ptaschek. Florin Ciucu, Cristian Levcovici, Cristina Cascaval, Ping Shyr, and Kathryn Wright contributed to DISCIPLE-COA.

Notes

1. R. MacGregor, 1999. Retrospective on LOOM. Available online at www.isi.edu/isd/LOOM/papers/macgregor/Loom_Retrospective.html.

References

Boicu, M.; Wright, K.; Marcu, D.; Lee, S. W.; Bowman, M.; and Tecuci, G. 1999. The DISCIPLE Integrated Shell and Methodology for Rapid Development of Knowl-

edge-Based Agents. In Proceedings of the Sixteenth National Conference on Artificial Intelligence, 900-901. Menlo Park, Calif.: American Association for Artificial Intelligence.

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*, 73-80. San Francisco, Calif.: Morgan Kaufmann.

Chaudhri, V. K.; Farquhar, A.; Fikes, R.; Karp, 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.

Clancey, W. J. 1984. NEOMYCIN: Reconfiguring a Rule-Based System with Application to Teaching. In *Readings in Medical Artificial Intelligence*, eds. W. J. Clancey and E. H. Shortliffe, 361-381. Reading, Mass.: Addison-Wesley.

Cohen, P.; Schrag, R.; Jones, E.; Pease, A.; Lin, A.; Starr, B.; Gunning, D.; and Burke, M. 1998. The DARPA High-Performance Knowledge Bases Project. *AI Magazine* 19(4): 25-49.

Fikes, R.; Farquhar, A.; and Rice, J. 1997. Tools for Assembling Modular Ontologies in ONTOLOGIA. In Proceedings of the Fourteenth National Conference on Artificial Intelligence, 436-441. Menlo Park, Calif.: American Association for Artificial Intelligence.

Jones, E. 1999. HPKB Course of Action Challenge Problem Specification, Alphatech, Inc., Burlington, Massachusetts.

Kim, J., and Gil, Y. 1999. Deriving Expectations to Guide Knowledge Base Creation. In Proceedings of the Sixteenth National Conference on Artificial Intelligence, 235-241. Menlo Park, Calif.: American Association for Artificial Intelligence.

Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM* 38(11): 33-38.

MacGregor, R. 1991. The Evolving Technology of Classification-Based Knowledge Representation Systems. In *Principles of Semantic Networks: Explorations in the Representations of Knowledge*, ed. J. Sowa, 385-400. San Francisco, Calif.: Morgan Kaufmann.

Mitchell, T. M. 1997. *Machine Learning*. New York: McGraw-Hill.

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

Tecuci, G.; Boicu, M.; Bowman, M.; Marcu,

D.; Shyr, P.; and Cascaval, C. 2000. An Experiment in Agent Teaching by Subject Matter Experts. *International Journal of Human-Computer Studies* 53(10): 583-610.

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.

U.S. Army. 1993. U.S. Army Field Manual, 100-5, Operations, Headquarters, Department of the Army, Washington, D.C.



Gheorghe Tecuci is professor of computer science and director of the Learning Agents Laboratory at George Mason University 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*, *Machine Learning: A Multistrategy Approach*, and *Machine Learning and Knowledge Acquisition: Integrated Approaches*. His e-mail address is tecuci@gmu.edu.



Mihai Boicu is a Ph.D. student in the School of Information Technology and Engineering at George Mason University and a research assistant in the Learning Agents Laboratory. Previously he was a tenured teacher in the Computer Science Department of the National School of Computer Science, Bucharest, Romania. His research interests include AI, machine learning, intelligent agents, knowledge acquisition, knowledge representation, problem solving, theory of algorithms, and education. He is a member of the American Association for Artificial Intelligence. His e-mail address is mboicu@gmu.edu.

Lt. Col. Michael Bowman is an active duty Army officer and a Ph.D. candidate at George Mason University. He received a B.S. from Ouachita Baptist University and



an M.S. from the Naval Postgraduate School. He was the U.S. Army product manager for Communications and Intelligence Support Systems and has had a variety of acquisition, automation, and tactical assignments at the Defense Intelligence Agency and the U.S. Military Academy and in several Army Field Artillery battalions. His e-mail address is mbowman3@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 intelligent user interfaces and machine learning. He is a member of the American Association for Artificial Intelligence. His e-mail address is dmarcu@cs.gmu.edu.



Murray Burke has served as the program manager of the High-Performance Knowledge Base Project and is currently the program manager of the Rapid Knowledge Formation Project at the Defense Advanced Research Projects Agency. He has over 30 years in command and control information systems research and development. He was cofounder and executive vice president of Knowledge Systems Concepts, Inc., a software company specializing in defense applications of advanced information technologies. Burke is a member of the American Association for Artificial Intelligence.

Expertise in Context: Human and Machine

COMPUTERIZED “expert systems” are among the best known applications of artificial intelligence. But what is expertise? The nature of knowledge and expertise, and their relation to context, is the focus of active discussion — even controversy — among psychologists, philosophers, computer scientists, and other cognitive scientists. The questions reach to the very foundations of cognitive theory — with new perspectives contributed by the social sciences. These debates about the status and nature of expert knowledge are of interest to and informed by the artificial intelligence community — with new perspectives contributed by “constructivists” and “situationalists.”

The twenty-three essays in this volume discuss the essential nature of expert knowledge, as well as such questions such as how “expertise” differs from mere “knowledge,” the relation between the individual and group processes involved in knowledge in general and expertise in particular, the social and other contexts of expertise, how expertise can be assessed, and the relation between human and computer expertise.

690 pp., ISBN 0-262-56110-7

Published by the AAAI Press / The MIT Press

**To order call 800-356-0343 (US and Canada)
or (617) 625-8569.**

**Distributed by The MIT Press,
5 Cambridge Center, Cambridge, MA 02142**



The cover is a reproduction of "Clean Three" by Marc Popoff. It is reproduced with permission, courtesy of Nathan Galanter, New York, New York.

**Paul J. Feltovich
Kenneth M. Ford
Robert R. Hoffman**