

Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving

Gheorghe Tecuci, Mihai Boicu
Learning Agents Center
George Mason University
Fairfax, VA, 22030, USA
{tecuci, mboicu}@gmu.edu

Cindy Ayers, David Cammons
Center for Strategic Leadership
US Army War College
Carlisle, PA, 17013, USA
{cindy.ayers, david.cammons}@us.army.mil

Keywords: Structured Argumentation, Building Analysts, Terrorism, Knowledge Bases, Mixed-Initiative Reasoning

Abstract

This paper presents research on developing a new type of tool that can alleviate several systemic problems faced by the traditional intelligence analysis process. The tool is a personal cognitive assistant that can rapidly acquire expertise in intelligence analysis directly from intelligence analysts, can train new analysts, and can help analysts find solutions to complex problems through mixed-initiative reasoning, making possible the synergistic integration of a human's experience and creativity with an automated agent's knowledge and speed, and facilitating the collaboration with complementary experts and their agents.

1. Introduction

Traditional intelligence analysis suffers from several systemic problems including: information overload; intelligence sharing difficulties; lack of time, methods, and resources for analytic collaboration with area experts; limited capabilities in regard to the consideration of multiple hypotheses; socio-cultural and socio-psychological bias informing the analytic process; lack of time and resources for critical analysis and after-action review; "group-think" (a lack of diverse opinions informing the process) and "paralysis by analysis"; loss of analytic expertise due to downsizing and attrition; lack of time and resources needed to train new analysts; and limited availability and use of tools to improve the analytic process (Lowenthal, 1999; National Commission on Terrorist Attacks Upon the United States, 2004; Wheaton, 2001).

This paper presents joint research performed by the Learning Agents Center of George Mason University and the Center for Strategic Leadership of the US Army War College aimed at developing a new type of analytic tool that will help alleviate several of the above problems. This tool, called, Disciple-LTA (learner, tutor, and assis-

tant) is a personal cognitive assistant that can rapidly acquire expertise in intelligence analysis directly from intelligence analysts, can train new analysts, and can help analysts find solutions to complex problems through mixed-initiative reasoning, making possible the synergistic integration of a human's experience and creativity with an automated agent's knowledge and speed, and facilitating the collaboration with complementary experts and their agents. This new type of intelligent agent, capable of learning, tutoring and decision making assistance, is intended to act as a career-long aid to intelligence analysts. It will be used during classroom learning, for skills maintenance and growth after classroom learning, and for decision support in the field.

This research on creating Disciple-LTA builds on the Disciple theory, methodology, and family of agent shells for the development of knowledge-based agents by subject matter experts, with limited assistance from knowledge engineers (Tecuci, 1998). Previous versions of the Disciple agent shells were used to build agents for course of action critiquing and center of gravity analysis, which were successfully evaluated as part of DARPA's High Performance Knowledge Bases and Rapid Knowledge Formation programs (Tecuci et al. 2001). The Disciple agents for center of gravity analysis have been successfully used in several courses at the US Army War College since 2001 (Tecuci et al., 2002). In the "Case Studies in Center of Gravity Analysis" course the students use trained Disciple agents as intelligent assistants that help them develop a center of gravity analysis of a war scenario. At the same time, the students learn how to perform such an analysis. In the follow-on course, "Military Applications of Artificial Intelligence" the students, now experts in center of gravity analysis, teach personal Disciple agents their own expertise, and then evaluate both the developed agents and the development process.

Working closely with the users helped us identify the weaknesses and strengths of the Disciple prototype agents and guided the development of improved agents

for subsequent sessions of these courses. A similar strategy is being used for the Disciple-LTA project which is based on the latest version of the Disciple approach (Disciple-RKF) but extends it in several important directions related to its application to intelligence analysis, its capabilities to represent and reason with pieces of evidence, and its ability to act as a tutoring system and as a mixed-initiative reasoning system (particularly with respect to solution composition), as discussed in this paper.

2. Intelligence Analysis through Task Reduction and Solution Composition

We are developing a systematic approach to intelligence analysis which is both natural for the human analyst and appropriate for an automated agent. This approach is based on the general task-reduction/solution-composition paradigm of problem solving which proved to be suitable for a wide variety of domains (Durham, 2000; Lowrance et al., 2001; Powel and Schmidt, 1988; Tecuci et al., 2001). Our approach is illustrated by the reasoning tree in Figure 1. Such a tree is jointly developed by the analyst and his or her Disciple-LTA assistant and is intended to be a natural and explicit representation of the thread of logic of the analyst, as if he or she would be thinking aloud, as discussed in the following.

We need to “Assess whether Location_A is a training base for terrorist operations.” In order to perform this assessment task, the analyst and the agent will ask themselves a series of questions. The answer to each question will lead to the reduction of the current assessment task to simpler assessment tasks. The first question asked is: “What type of factors should be considered to assess the presence of a terrorist training base?” The answer is “Political environment, physical structures, flow of suspected terrorists, weapons and weapons technology, other suspected bases in the region, and terrorist sympathetic population.” This answer leads to the reduction of the above top level task to 6 simpler assessment tasks, one for each identified factor. Each such task is further reduced in a similar manner, guided by a corresponding question and answer. For instance, the fourth task is reduced to 8 simpler tasks. The second of these 8 tasks is reduced to a simpler task, and this simpler task is further reduced to four even simpler tasks, one of which is “Assess whether there are explosive experts in the vicinity of Location-A.” The purpose of this successive task-reduction process is to reduce a complex intelligence analysis task T to a set of simpler intelligence analysis tasks T_i which could be performed through evidence analysis.

The next step is to search for and analyze pieces of evidence that are relevant to each of the tasks T_i . We are developing a systematic approach to evidence analysis which identifies different types of evidence and defines analyses procedures that are specific to each type. This approach is inspired by the theory of evidence developed by (Schum, 2001) which distinguishes between the following types of evidence: *tangible* (objects, documents, images, measurements, charts), *unequivocal testimonial*

(direct observation, second hand, or opinion), *equivocal testimonial* (complete equivocation or probabilistic), *missing tangible or testimonial*, and *authoritative records* (accepted facts).

To illustrate our approach, let us consider a report from Person-Z who claims to have repeatedly seen Person-E, a known explosive expert, in the vicinity of Location-A. This is a piece of evidence which is potentially relevant to the tasks “Assess whether there are explosive experts in the vicinity of Location-A.” In Schum’s terminology, this is an unequivocal testimonial evidence on a direct observation of Person-Z. Consequently, one has to assess three aspects: 1) the relevance of this evidence with respect to the assessment of whether there are explosive experts in the vicinity of Location-A; 2) the competence of Person-Z with respect to providing this kind of evidence; and 3) the credibility of Person-Z.

To assess the credibility of Person-Z one has to assess his veracity, objectivity and observational accuracy. The *veracity* of an observer refers to the degree to which that observer believes that the event actually occurred (i.e. is Person-Z lying or is he telling the truth?). The *objectivity* of an observer refers to the degree to which one attends to the evidence of his or her senses and does not let personal motivations or expectations determine what he or she will believe. The *observational sensitivity or accuracy* of an observer refers to the degree to which one’s senses (as well as the conditions of observations and the observer’s physical condition at the time of observation) gives evidence to reported observation (Schum, 2001).

Once the veracity, objectivity, and observational sensitivity of Person-Z are assessed, they are combined into an assessment of the credibility of Person-Z, as illustrated at the bottom of Figure 1. Person-Z’s credibility is further combined with his competence and with the relevance of his testimony, to obtain a partial solution of the task “Assess whether there are explosive experts in the vicinity of Location-A.” This partial solution is subsequently composed with the partial solutions corresponding to other pieces of evidence, to obtain the following solution to the above task: “There is very strong evidence that there are explosive experts in the vicinity of Location-A.”

Solutions for the other tasks shown at the middle of Figure 1 are found in a similar way. These solutions are successively combined, from bottom up, to produce the following solution for the assessment task from the top of the tree: “There is strong evidence that Location-A is a training base for terrorist operations.” Notice that the solution composition process is also guided by questions and answers.

To summarize, the analysis illustrated in Figure 1 consists of the following phases: 1) A complex assessment task T is successively reduced to simple assessment tasks T_i that can be performed through evidence analysis; 2) Potentially relevant pieces of evidence E_j for each task T_i are identified; 3) Available evidence is analyzed and a solution for each task T_i is obtained; and 4) The solutions of the tasks T_i are successively combined to obtain the solution for the initial task T.

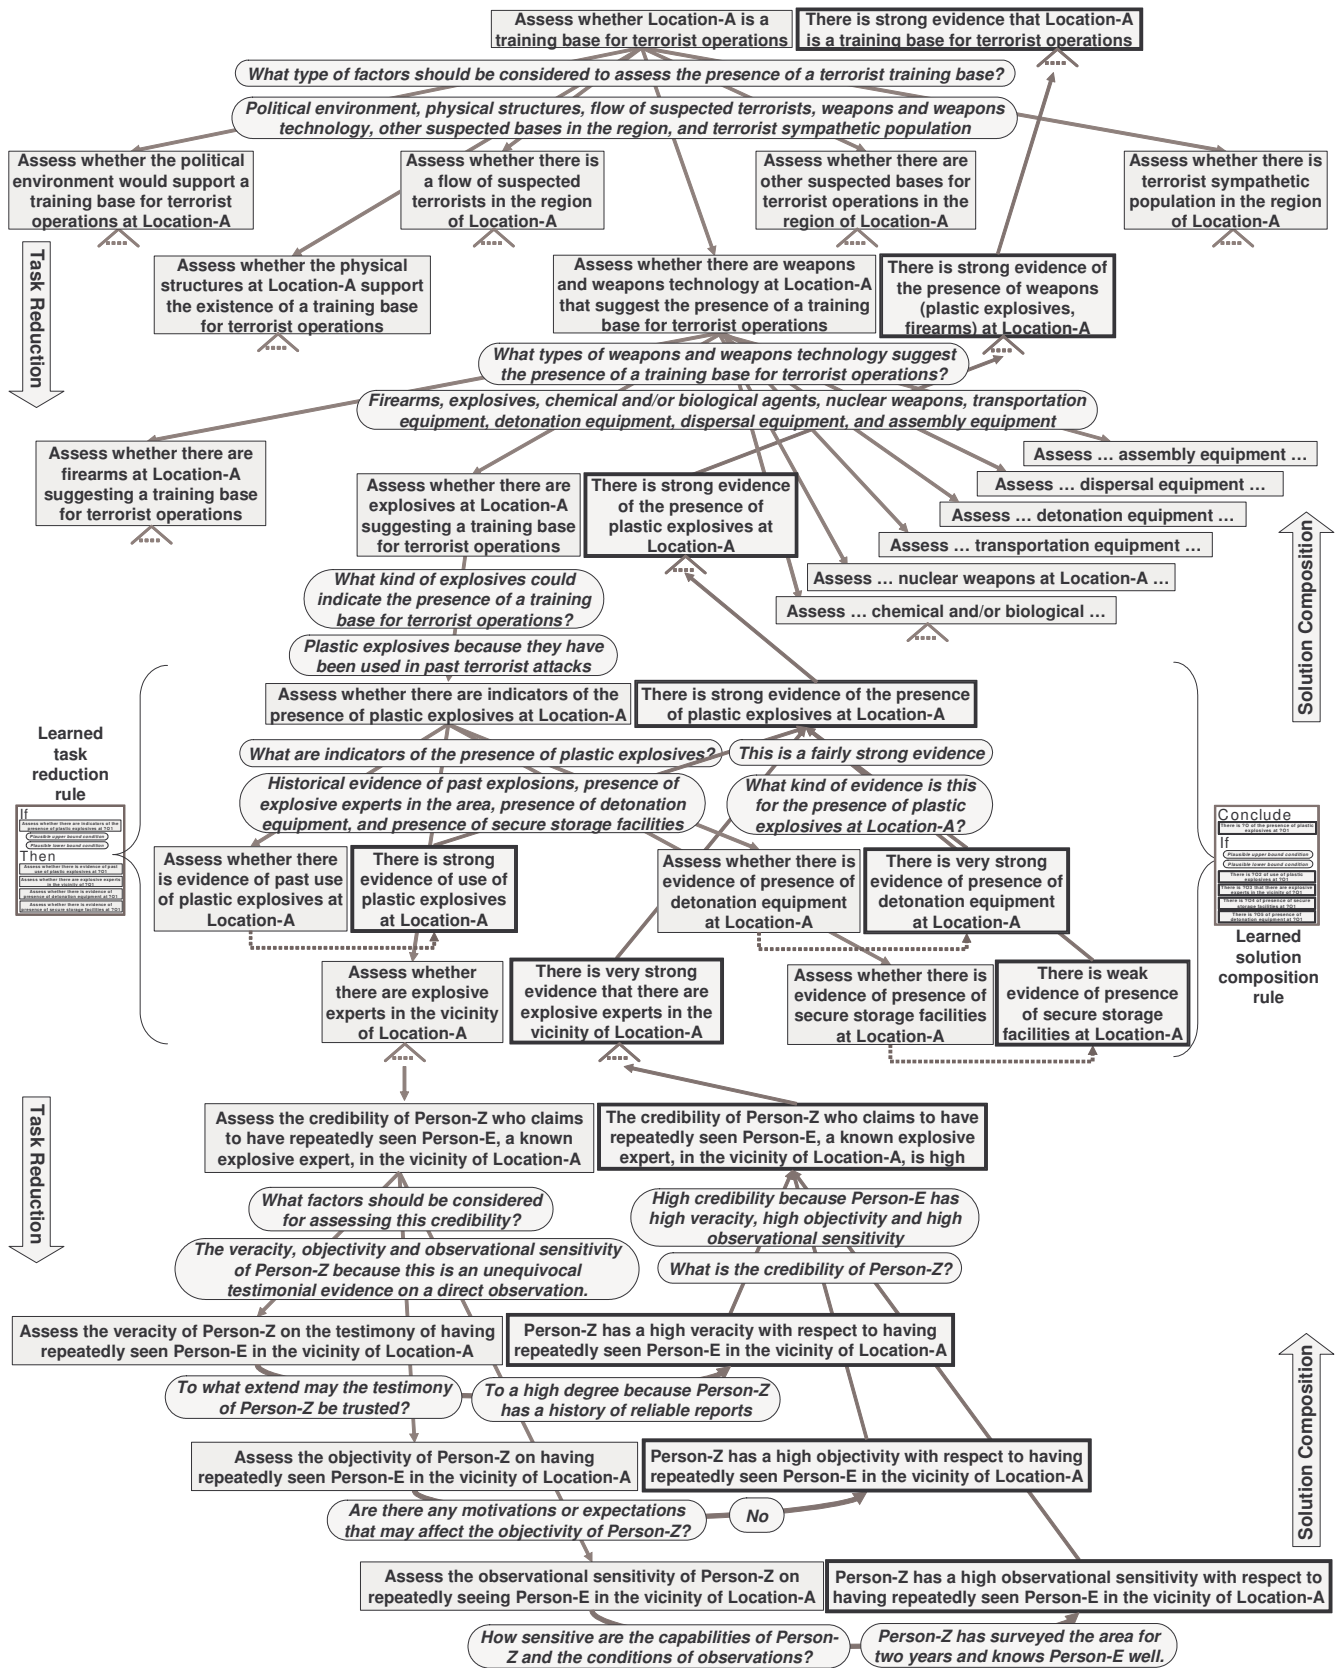


Figure 1: Intelligence analysis through task reduction and solution combination

3. Disciple-LTA

The overall architecture of Disciple-LTA, which supports the above analysis, is shown in the center of Figure 2. As a tool, Disciple-LTA is a general knowledge-based agent which has no specific knowledge in its knowledge base, but can be taught by an intelligence analyst, and can develop its knowledge base to become an analyst's assistant. Disciple-LTA has a multi-agent architecture composed of three groups of cooperating agents: problem solving agents, learning agents, and tutoring agents.

The problem solving agents support various intelligence analysis tasks, such as problem definition, hypotheses generation, information collection, hypotheses evaluation, hypothesis selection, and report generation. The main problem-solving engine of Disciple-LTA is based on the general task-reduction/solution-composition paradigm illustrated in the previous section. To be able to generate a reasoning tree like the one from Figure 1, the knowledge base of a Disciple agent is structured into an object ontology and a set of if-then problem solving rules. The object ontology is a hierarchical representation of the objects from the application domain, together with their properties and relationships (Fensel 2000). The objects to be represented include different types of intelligence resources, such as HUMINT (e.g. agents, informers, observants) or OSINT (e.g. books, webpages, newspaper articles), as well as descriptions of domain-specific objects such as different types of explosives, locations, etc. The if-then problem solving rules are expressed using the objects from the ontology. Each rule indicates how and under what conditions a complex task can be

reduced to simpler tasks, or how and under what conditions the solutions of the simpler tasks can be combined into the solution of the complex task (Boicu, 2002).

Disciple-LTA allows the analyst to act as the orchestrator of the problem solving process, guiding the high-level exploration, while Disciple-LTA implements this guidance by taking into account analyst's assumptions, preferences and biases. To illustrate this, let us consider again the reasoning tree in Figure 1, this time developed by the trained Disciple-LTA agent. The rightmost assessment task from the upper part of Figure 1 is "Assess whether there is terrorist sympathetic population in the region of Location-A." The agent may reason under the analyst's assumption that there is such a terrorist sympathetic population, and may no longer investigate this issue. However, knowing that this is an assumption, the agent may also attempt to challenge it in the background by actually trying to solve this assessment task, and alerting the analyst if, for instance, evidence is found that recent events have turned the population against the terrorists.

Sometimes there is not enough information to find an answer to some question. In such a case the agent may explore what-if scenarios, each corresponding to a different plausible answer to the question.

Other times the tasks that an analyst has to perform may require the help of different types of experts. Consider, for instance, the task from the top left of Figure 1: "Assess whether the political environment would support a training base for terrorist operations at Location-A." Through its external expertise agent, Disciple-LTA may request the help of the agent belonging to an appropriate political analyst to perform this assessment task.




Figure 2: Disciple-LTA - Integration of intelligence analysis education and operations.

The learning agents of Disciple-LTA (see center of Figure 2) facilitate the rapid development of the knowledge base by capturing the problem solving expertise of experienced analysts, including their problem solving strategies, prior and tacit knowledge. Many of these learning agents are developments of the corresponding learning agents of Disciple-RKF (Tecuci et al., 2002). They include browsers and editors for ontology development and scenario elicitation. They also include agents for learning task reduction rules, and for refining the object ontology. New agents that are developed for Disciple-LTA include a modeling assistant that helps the user to express her reasoning using the task reduction paradigm, a learning agent for learning and refining solution composition rules, and a specialized editor for representing pieces of evidence.

The Disciple-LTA shell is used to rapidly develop a Disciple-LTA agent for a specific intelligence analysis domain by following a two phase process: 1) The development of an initial object ontology for that domain, which is performed jointly by a knowledge engineer and an expert analyst, and 2) The teaching of Disciple-LTA, which is performed by the intelligence analyst, with limited assistance from the knowledge engineer.

During the teaching process, the analyst considers typical intelligence analysis tasks, such as the one from the top of Figure 1, builds the reasoning tree, and helps the agent to understand each problem solving step. From each problem solving step the agent learns a general reasoning rule. Consider, for instance, the task "Assess whether there are indicators of the presence of plastic explosives at Location-A" from the middle of Figure 1, which is reduced to four subtasks. From this specific task reduction step Disciple-LTA learns a general if-then task reduction rule which will allow it to make a similar reduction in a future situation. Consider also the composition of the solutions of the four subtasks into the solution of the above task. From this solutions composition example Disciple-LTA learns a general solution composition rule.

As Disciple-LTA is trained by an analyst, its knowledge base evolves to represent better and better the analyst's expertise, factual and tacit problem solving knowledge, assumptions and biases. Therefore, in time, the interaction between the analyst and Disciple-LTA evolves from a teacher-student interaction toward an interaction where both collaborate in performing an intelligence analysis task. During this interaction Disciple-LTA learns not only from the contributions of the analyst, but also from its own successful or unsuccessful problem solving attempts, which lead to the refinement of the learned rules. At the same time, Disciple-LTA extends the object ontology with new objects and features.

The tutoring agents of Disciple-LTA (see center of Figure 2) enable it to teach new analysts how to perform intelligence analysis. These agents include an agent for learning and refining tutoring knowledge, an agent for learning a student's model, a test generator, and a student evaluator. The main idea is to teach new analysts in a

way that is similar to how Disciple-LTA was itself taught by an expert analyst. Thus the roles are now reversed, with the agent being the expert and the human the learner. The agent will now consider typical intelligence analysis tasks, such as the one from the top of Figure 1, and will explain to the student analyst how to solve them.

4. Experimentation Environment

We are developing Disciple-LTA using an approach similar to the User-Centered Systems Engineering Process (DeBellis and Haapala, 1995) which encourages the developers and the users to collaborate during software design. Furthermore, because there is a strong correlation between a system's success and the way it fits within an existing organization, our approach will take into account the organization of the military, its cycles of military education and practice, and the new challenges that it faces in the current war on terror. Disciple prototypes are developed iteratively and incrementally, and are evaluated in periodic formal experiments, to obtain crucial feedback. Successive prototypes have increasing functionality and approximation of user needs. This approach identifies risks and problems early, making corrections less expensive and more effective.

We are using an experimentation environment which is similar to the one which was very successfully used for Disciple-RKF, as part of the DARPA's Rapid Knowledge Formation Program (Tecuci et al., 2002). The experimentation is to be conducted in the Military Applications of Artificial Intelligence (MAAI) course, taught at the US Army War College, in Spring 2005. This is a 10-week, 3 hours/week course, attended by military intelligence analysts and other military personnel with interest in artificial intelligence. The students, who have no prior knowledge engineering experience, will be introduced to the Disciple-LTA agent which they will use as a tutoring system, problem solving assistant, and learner.

First the students will use Disciple-LTA as an intelligent tutoring system by considering typical intelligence analysis tasks and understanding how Disciple-LTA solves them. The main goal of this phase is to teach the students how to systematically solve intelligence analysis tasks by using task reduction and solution composition, as illustrated in section 2. Then the students will use Disciple-LTA as a problem solving and learning agent, considering tasks for which Disciple-LTA was only partially trained, such as "Assess whether Al Qaeda has nuclear weapons." During this phase the students will understand how an agent can be used as a problem solving assistant and how it can be taught by its expert user.

During the MAAI course, the modules of Disciple-LTA will be extensively evaluated and feedback from the users will be collected to inform the development of an improved Disciple-LTA agent, which will be again evaluated during the next session of the MAAI course. As will be discussed next, this experimentation environment is also intended to validate a proposed integration of intelligence analysis education and operations.

5. Integration of Intelligence Analysis Education and Operations

Figure 2 presents a long-term vision on the life cycle of a Disciple-LTA cognitive assistant which integrates intelligence analysis education and operations. The starting point (Phase 1) is the development of a Disciple-LTA agent shell customized for intelligence analysis tasks.

In Phase 2, the Disciple-LTA agent shell is trained by expert analysts (with limited assistance from a knowledge engineer) to perform intelligence analysis tasks. The subject matter experts are Army War College professors.

In Phase 3 the trained agents take the role of intelligent tutors, teaching the AWC students in a way that is similar to how they were taught by the expert analysts.

The AWC students will take their agents with them, to be used as expert collaborators, during Phase 4. During the agent's normal use in operations the analyst and the agent will encounter novel situations which are new opportunities for learning. However, the analyst will be primarily concerned with the current problem solving process and will have neither the time nor the incentive to support the learning of the agent. Therefore, the agent will have to learn by employing resource bounded, non-disruptive learning techniques, storing relevant experiences for mixed-initiative learning to occur during periodic after-action reviews, represented as Phase 5. During the after-action reviews, the agent will learn not only new reasoning patterns, but also a model of its analyst, incorporating his/her preferences, biases, and assumptions. After the completion of Phase 5, the agent will reenter Phase 4 of its life cycle.

During its use, the agent will acquire a significant amount of problem solving knowledge. However, because its user is not a knowledge engineer, this knowledge will not always be optimally represented and organized. Therefore, during Phase 6, a knowledge engineer will interact with the agent to optimize its knowledge base, with limited assistance from an expert analyst. The optimized agent will be returned to its user (see the link from Phase 6 to Phase 5). Notice also that the knowledge bases of different agents will incorporate the expertise of their analysts. These knowledge bases could be kept to represent the knowledge of various expert analysts, allowing the preservation and future access to this knowledge. A comparative analysis of these knowledge bases (performed during Phase 6) will also reveal valuable new knowledge to be integrated into an improved agent to be used in Phase 2 of a new cycle (see the link from Phase 6 to Phase 2), therefore injecting the learned knowledge back into the process. Phase 6 will also provide Phase 1 with a specification on how to improve the learning agent shell, based on the lessons learned in the previous phases (see the link from Phase 6 to Phase 1).

6. Conclusions

This paper has presented current research on developing a new type of cognitive assistant, called Disciple-LTA,

that helps an intelligence analyst to systematically solve complex intelligence analysis tasks faster and better—an assistant that learns and uses analyst's preferred problem solving strategies, biases and assumptions, but can also constructively challenge them and consider alternative what-if scenarios. Disciple-LTA facilitates the retention of the expertise and the training of new analysts. Its ability to rapidly acquire subject matter expertise (Tecuci et al., 2002) allows also the development of agents that reason consistently with the culture of the data source—agents that can further improve the analysis process.

Acknowledgments

The research on the Disciple approach was sponsored by several US Government agencies including Defense Advanced Research Projects Agency, Air Force Research Laboratory, and Air Force Office of Scientific Research.

References

- Boicu, M. 2002. *Modeling and Learning with Incomplete Knowledge*. PhD Thesis in Information Technology, Learning Agents Laboratory, George Mason University.
- DeBellis, M. and Haapala, C. 1995. User-Centric Software Engineering. *IEEE Expert*, 10, 34-41.
- Durham S. 2000. Product-Centered Approach to Information Fusion, *AFOSR Forum on Information Fusion*, Arlington, VA, 18-20 October, 2000.
- Lowenthal, M. M. 1999. *Intelligence: From Secrets to Policy*. Washington, DC: Congressional Quarterly Press.
- Lowrance J.D.; Harrison I.W. and Rodrigues A.C. 2001. Capturing Analytic Thought. *First International Conference on Knowledge Capture*, October, 2001, pp. 84-91.
- National Commission on Terrorist Attacks Upon the United States. 2004. *The 9/11 Commission Report*. New York: WW.Norton & Company.
- Powel G.M. and Schmidt C.F. 1988. A First-order Computational Model of Human Operational Planning, *CECOM-TR-01-8*, US Army CECOM, Fort Monmouth, New Jersey.
- Schum D.A. 2001. *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.
- Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London, England: Academic Press.
- Tecuci G.; Boicu M.; Bowman M.; 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*, 22(2): 43-61, AAAI Press.
- Tecuci G.; Boicu M.; Marcu D.; Stanescu B.; Boicu C. and Comello J. 2002. Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain, in *AI Magazine*, 24(4): 51-68, AAAI Press.
- Wheaton, K. J. 2001. *The Warning Solution: Intelligent Analysis in the Age of Information Overload*. Fairfax, VA: AFCEA International Press.