A Learning and Reasoning System for Intelligence Analysis

Mihai Boicu, Gheorghe Tecuci, Cindy Ayers[1], Dorin Marcu, Cristina Boicu, Marcel Barbulescu, Bogdan Stanescu, William Wagner, Vu Le, Denitsa Apostolova, Adrian Ciubotariu

Learning Agents Center, George Mason University, MSN 4A5, 4400 University Drive, Fairfax, VA 22032, USA
[1] Center for Strategic Leadership, US Army War College, Carlisle Barracks, Carlisle, PA, 17013, USA
{mboicu, tecuci, dmarcu, ccascava, mbarbule, bstanesc, wwagner, vle3, dapostol, aciubota}@gmu.edu, cindy.ayers@us.army.mil

Abstract

This paper presents a personal cognitive assistant, called Disciple-LTA, that can 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.

Intelligence Analysis through Task Reduction and Solution Composition

Disciple-LTA builds on Disciple-RKF (Tecuci et al., 2002) and advances the Disciple approach to the development of knowledge-based agents by subject matter experts (Tecuci, 1998) with respect to the application to intelligence analysis, and the tutoring, problem solving and learning capabilities, as discussed below.

One of the most important contributions of Disciple-LTA is the developing and implementation of 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, and consists of the following steps:

1) A complex intelligence analysis task T is successively reduced to simpler tasks that either have known solutions, or can be solved through evidence analysis.

2) Potentially relevant pieces of evidence for each unsolved task are identified.

3) The identified pieces of evidence are analyzed using the task reduction paradigm and a solution for each unsolved task is obtained.

4) The solutions of the simplest tasks are successively combined to obtain the solution of the initial task T.

The reductions and the compositions are guided by questions and answers, as if the analyst or the agent would be thinking aloud, asking themselves how to reduce the current task or to compose the current solutions.

Evidence analysis (steps 2 and 3), which is inspired by the theory of evidence developed by Schum (2001), identifies different types of evidence (tangible, unequivocal testimonial, equivocal testimonial, missing tangible or testimonial, and authoritative records) and defines analyses procedures that are specific to each type. 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 piece of evidence is potentially relevant to the tasks "Assess whether there are explosive experts in the vicinity of Location-A." In Schum's terminology, this is 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, as illustrated in Figure 1. 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 in 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."

Disciple-LTA

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 hypotheses evaluation, information collection, and report generation. The main problem-solving engine is based on the task-reduction paradigm discussed in the previous section. To be able to generate a

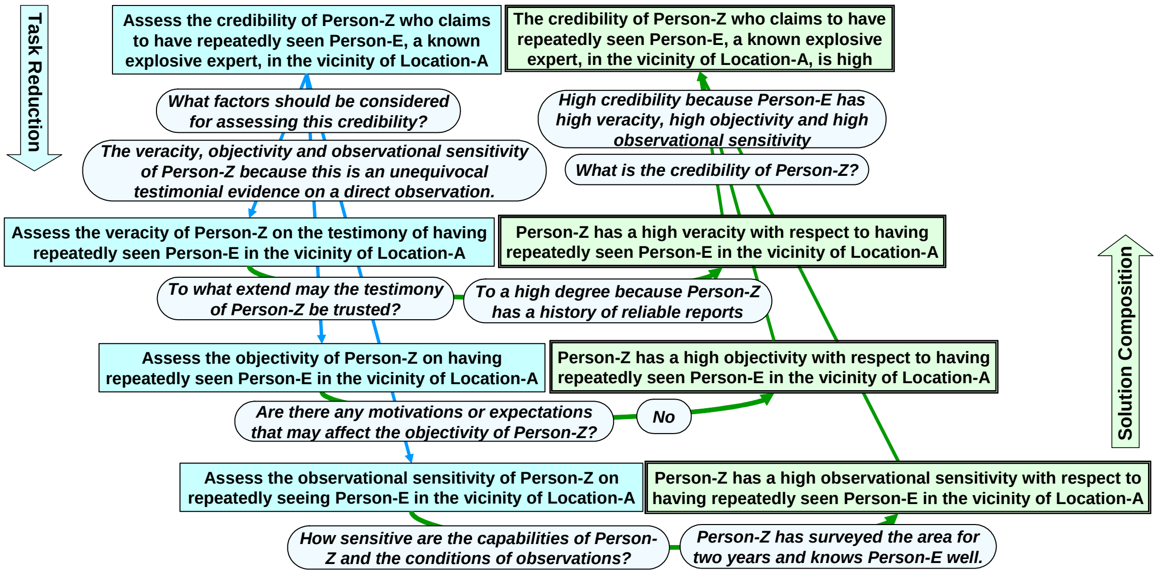

Figure 1: Credibility analysis through task reduction and solution composition

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 intelligence analysis domain, together with their properties and relationships. 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 the solutions of the simpler tasks can be combined into the solution of the complex task.

The ***learning agents*** of Disciple-LTA facilitate the rapid development of the knowledge base by capturing the problem solving expertise of experienced analysts. Many of these learning agents are developments of the corresponding learning agents of the Disciple-RKF system for center of gravity analysis (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 editor and a modeling assistant to help the analyst express her reasoning using the task reduction paradigm, a learning agent to refine task reduction rules, a learning agent to learn and refine 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 the specific domain, which is performed jointly by a knowledge engineer and an expert intelligence analyst, and 2) The teaching of the Disciple-LTA agent, 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-left 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. As Disciple-LTA is trained by the analyst, it starts to act more as an assistant than a student, contributing to the analysis process, and learning from it.

The ***tutoring agents*** of Disciple-LTA enable it to teach new analysts how to perform intelligence analysis. 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 and will explain to the student analyst how to solve them.

We are experimenting with Disciple-LTA at the US Army War College where 7 military experts, who have no prior knowledge engineering experience, have been introduced to Disciple-LTA, using it as a tutoring system, problem solving assistant, and learner, over ten 3 hours long sessions course, as part of the Military Applications of Artificial Intelligence course, taught in Spring 2005.

References

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.; 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.

Tecuci G., Boicu M., Ayers C., Cammons D., Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving, in *Proc. of the First Int. Conf. on Intelligence Analysis,* McLean, VA, 2-6 May 2005.