Cognitive Assistants for Analysts

Dr. Gheorghe Tecuci, Dr. Mihai Boicu, Mr. Dorin Marcu, Dr. Cristina Boicu, Mr. Marcel Barbulescu
Learning Agents Center, George Mason University

Professor Cindy Ayers
Center for Strategic Leadership, U.S. Army War College

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 U.S. Army War College aimed at developing a new type of analytic tool that will help alleviate several of the above problems (Boicu et al., 2005; Tecuci et al., 2005a). This tool, called Disciple-LTA (learner, tutor, and assistant), is a personal cognitive assistant to an intelligence analyst that rapidly acquires expertise in intelligence analysis directly from the analyst, including the analyst's prior and tacit knowledge, biases and assumptions. It helps the analyst solve complex problems through a mixed-initiative process that makes possible the synergistic integration of a human's experience and creativity with an automated agent's knowledge and speed. It facilitates collaborative problem solving and intelligence sharing between analysts (and their agents) in a network-centric environment. It can rapidly train new analysts and explain its knowledge in a pedagogical way. 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. Finally, it will function as a repository of the analyst's knowledge-base to provide continuity and historical archive for departing analysts, thus codifying individual and personal analytical expertise, otherwise lost forever to the intelligence community.

Research performed in the creation of the 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 U.S. 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 trained 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 (Tecuci et al., 2005b). Disciple-LTA extends the previous approach in several important directions related to its application to intelligence analysis; its capabilities to represent and reason with written evidence; its ability to act as a mixed-initiative problem solving system; and its use as a tutoring system.

The rest of this paper is organized as follows. The next section presents our vision on the integration of intelligence analysis education and operations that guides the development of Disciple-LTA. Then, section 3 presents a systematic approach to hypotheses analysis which is both natural for a human analyst and appropriate for an automated agent, and is the basis of analyst-agent collaboration. Sections 4, 5, and 6 present the three main uses of Disciple-LTA, as a problem solving assistant, as a learning system, and as a tutoring system. Section 7 discusses the development and experimentation approach of Disciple-LTA as well as the evaluation results obtained from users during and after the experimentation processes. Section 8 discusses how this research addresses some of the main challenges of intelligence analysis, results that have already been achieved, and directions of future research.

2. Integration of Intelligence Analysis Education and Operations

Figure 1 presents a long-term vision of 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. The Disciple-LTA shell has the required knowledge representation, learning and problem solving capabilities, but no specific domain knowledge in its knowledge base.

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 goal of this phase is to create an agent that can act both as a tutor and as a competent collaborator of the human analysts. In our experiments, the subject matter experts are U.S. Army War College (USAWC) professors, as discussed in section 5.

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

The USAWC 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

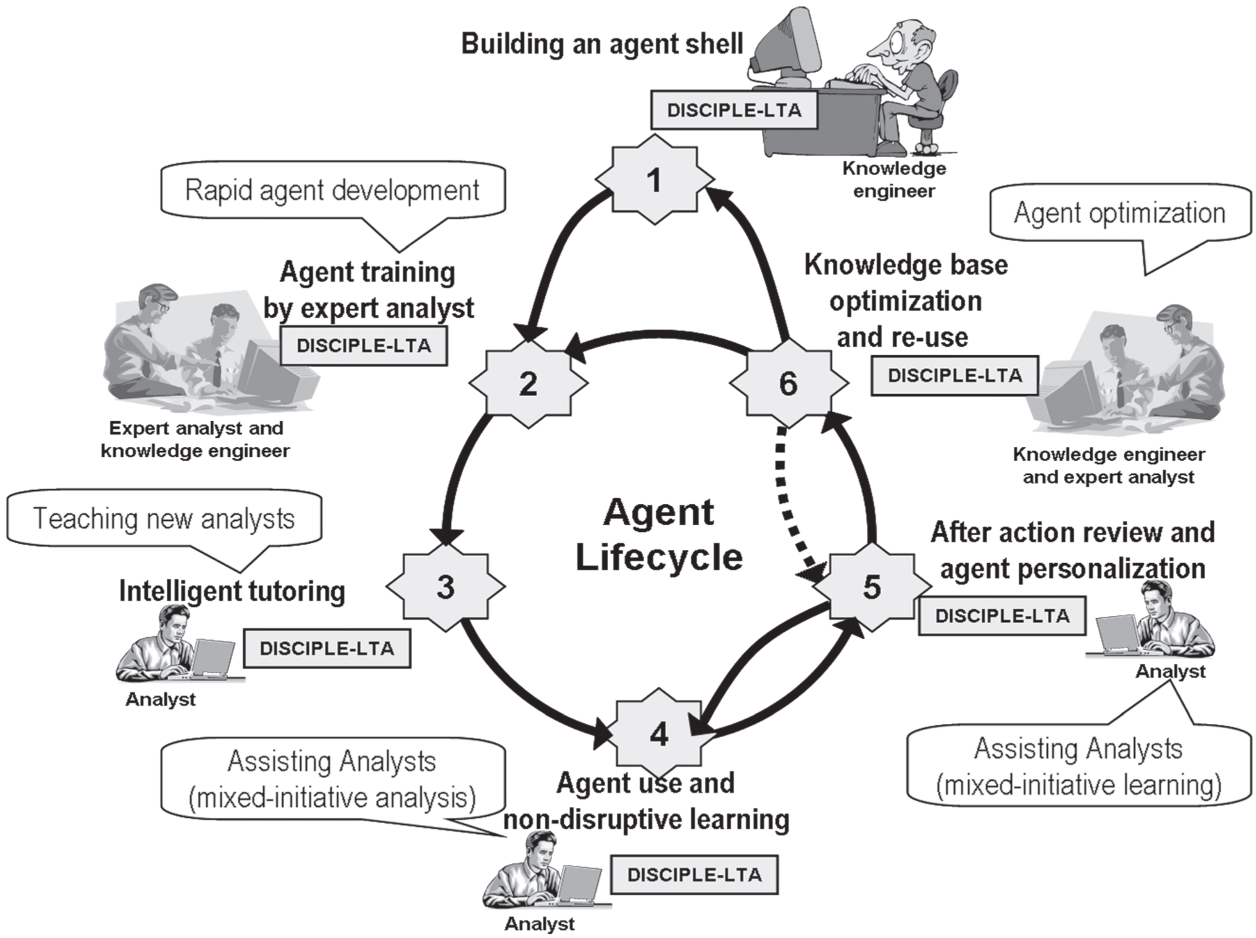

Figure 1. Integration of intelligence analysis education and operations.

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, an 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 line 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).

Figure 2 shows the envisioned use of the Disciple-LTA cognitive assistants in an operational environment. An analyst will be assisted by a personal Disciple-LTA agent. The agent will be able to solve problems either through mixed-initiative reasoning with its analyst, autonomously, or in collaboration with other experts and their agents in a network centric environment.

The Disciple-LTA agents will share and maintain a global knowledge base to facilitate the collaboration between the analysts and their agents. For instance, an analyst, through his/her agent, may collaborate with other experts and their agents to solve problems requiring expertise in multiple areas (e.g. expertise in chemical warfare weapons development, in weapons of mass destruction, or in politics). Similarly, the analyst, through his/her Disciple-LTA agent, may ask the agents of other analysts to provide a solution to the current analysis task from their points of view (i.e. biases, assumptions, prior knowledge), thus reducing the effect of a narrowly focused analytic mindset.

The Disciple-LTA agents will not only assist the human analysts to solve problems and to collaborate, but they will also continuously improve their knowledge by learning from their problem solving experience.

In the rest of the paper we will present current results in realizing this vision and the next steps of our research.

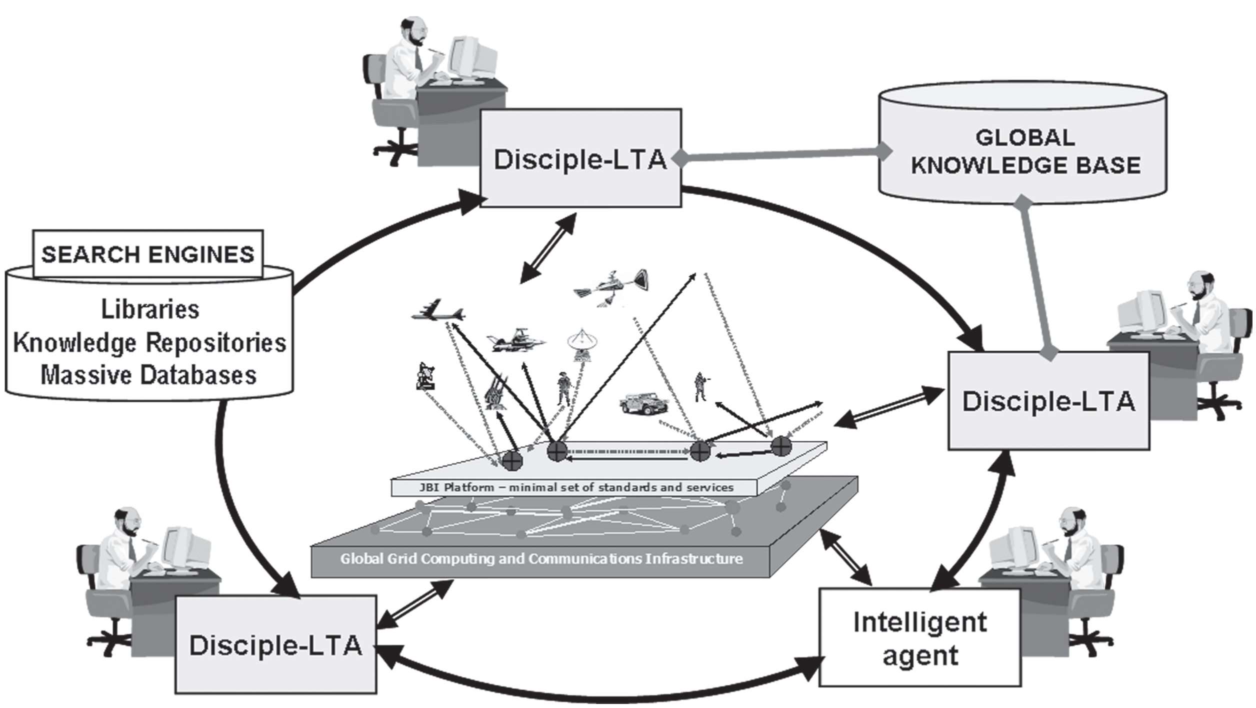

Figure 2. Envisioned use of Disciple-LTA in an operational environment.

3. A Systematic Approach to Hypotheses Analysis

3.1 Hypotheses Analysis through Task Reduction and Solution Synthesis

One of the main results of this research is the development 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 and solution synthesis 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, 1988). Our approach is illustrated by the reasoning tree in figure 3. Such a tree is jointly developed by the analyst and his/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/she would be thinking aloud, as discussed in the following sample problem.

We need to:

> Assess whether Al Qaeda has nuclear weapons.

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 factors should we consider to determine whether Al Qaeda has nuclear weapons?

The answer is:

> Reasons to obtain nuclear weapons, desire to obtain nuclear weapons, ability to obtain nuclear weapons, and current evidence that it has nuclear weapons.

This answer leads to the reduction of the above top level task to four 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 first task is reduced to five simpler tasks. 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. Such tasks are the five tasks mentioned above, the second of which is:

> Assess whether Al Qaeda considers self defense as a reason to obtain nuclear weapons.

The next step is to search for and analyze pieces of evidence that are relevant to each of the tasks T_i. To illustrate this, let us consider the analysis of the task "Assess whether Al Qaeda considers self defense as a reason to obtain nuclear weapons," as illustrated in figure 3:

> Is there any potentially relevant piece of evidence?

> Dawn-Mir01-a which mentions that Al Qaeda would use nuclear weapons to defend itself.

Dawn-Mir01-a represents the following fragment from an interview given by Osama bin Laden to Hamid Mir entitled "Osama claims he has nukes: If U.S. uses N-arms it will get same response:" *"We ourselves are the target of killings, destruction, and atrocities. We are only defending ourselves. This is defensive jihad. We want to defend our people and our land. That is why I say that if we don't get security, the Americans, too would not get security."* (Mir, 2001).

We therefore need to:

> Assess to what extent the piece of evidence Dawn-Mir01-a supports or undermines the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.
>
> What factors determine how a piece of evidence supports or undermines a hypothesis?
>
> The information provided by the piece of evidence and the extent to which it can be trusted.

We therefore can reduce the above task to two simpler tasks:

> Assess to what extent Dawn-Mir01-a supports or undermines the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons, assuming that we believe that the information provided by Dawn-Mir01-a is true.
>
> Assess the credibility of Dawn-Mir01-a.

The first of these tasks can be immediately solved (as shown in figure 3):

> Taking into account only the information provided by Dawn-Mir01-a what is the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons?
>
> Very high because Dawn-Mir01-a states that Al Qaeda would use nuclear weapons to defend itself.

Therefore we can conclude that:

> The information from Dawn-Mir01-a supports to a very high degree the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.

The second task "Assess the credibility of Dawn-Mir01-a." is solved through a similar process of task reduction / solution synthesis, as will be discussed in the next section. However, for now, let us assume that we have found its solution:

> The credibility of Dawn-Mir01-a is medium.

The next step is to combine these two solutions into the solution of the task:

> Assess to what extent the piece of evidence Dawn-Mir01-a supports or undermines the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.

Solution synthesis is also guided by questions and answers (see figure 3), as illustrated below.

We have determined that:

> The information from Dawn-Mir01-a supports to a very high degree the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.
>
> The credibility of Dawn-Mir01-a is medium.

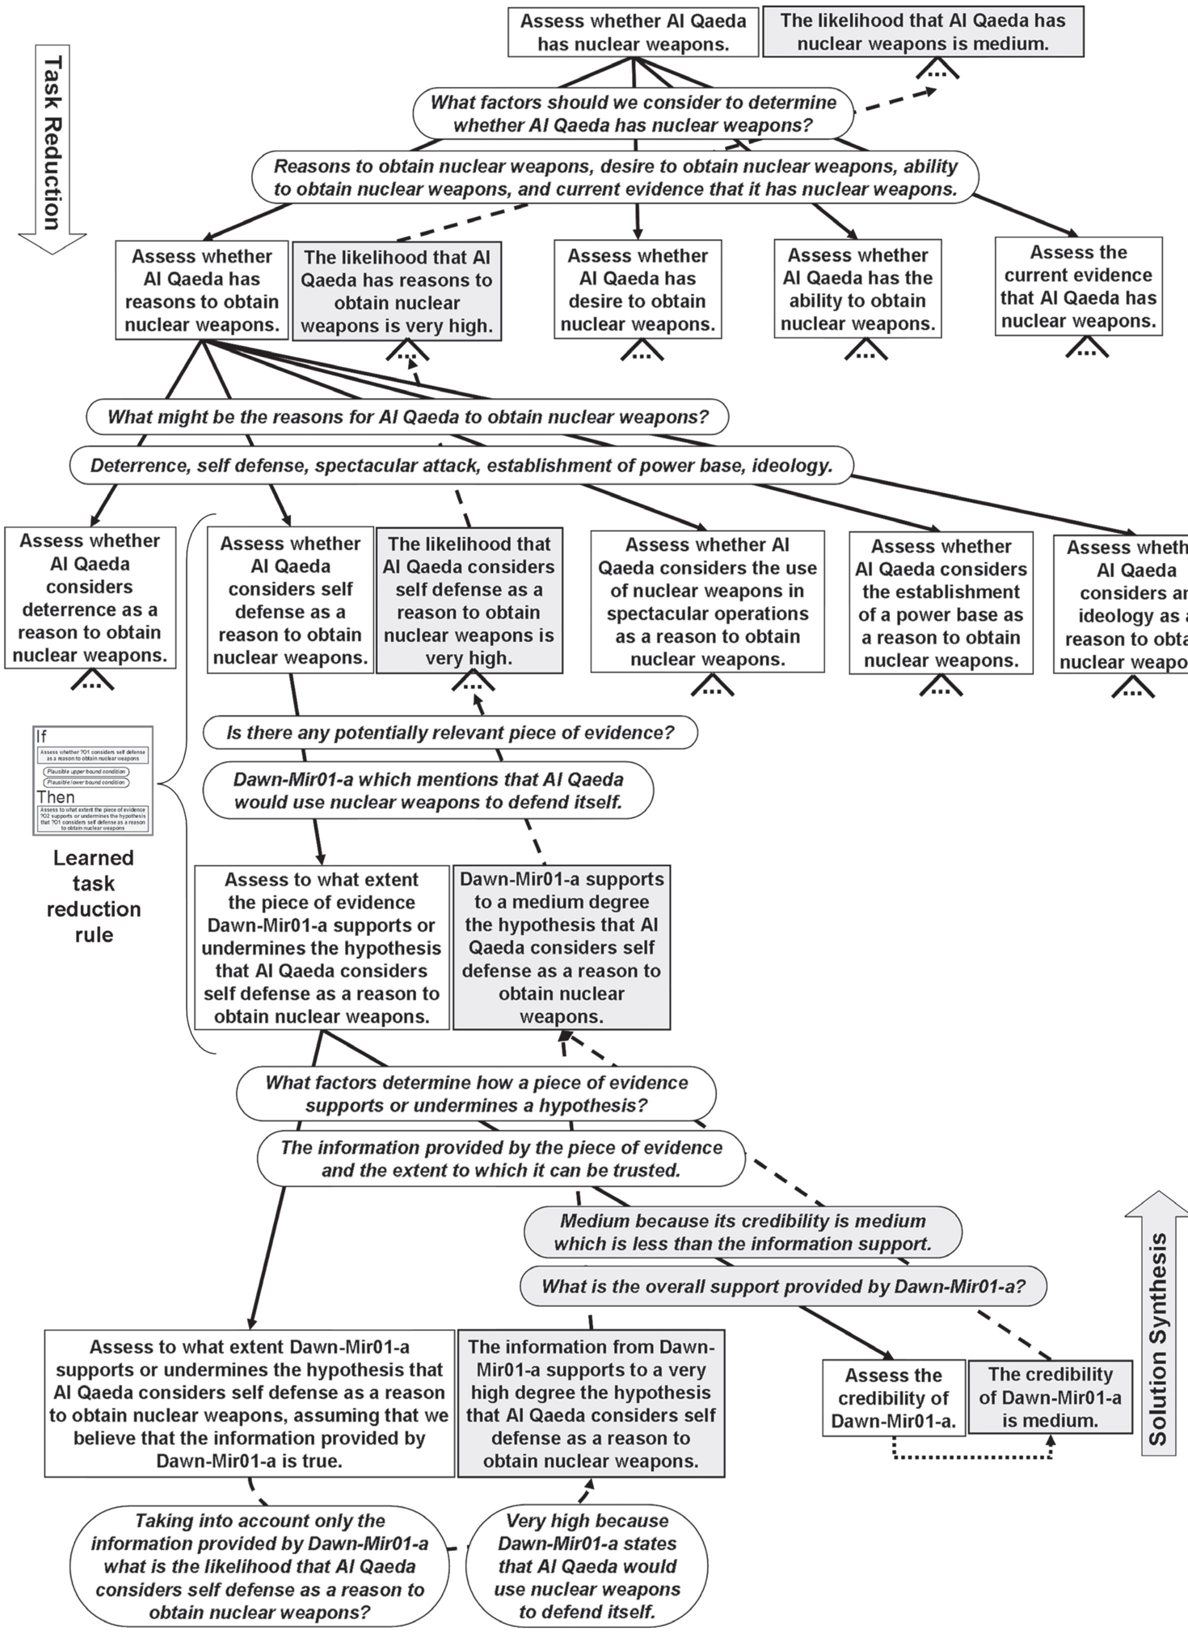

Figure 3. Hypothesis analysis through task reduction and solution synthesis.

In this case:

> What is the overall support provided by Dawn-Mir01-a?
>
> Medium because its credibility is medium and is less than the information support.

Therefore we can conclude that:

> Dawn-Mir01-a supports to a medium degree the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.

To summarize these last steps, a task T was reduced to two simpler tasks, the solutions of these simpler tasks have been found, and these solutions have been combined into the solution of T.

We have just illustrated how the task "Assess whether Al Qaeda considers self defense as a reason to obtain nuclear weapons" was performed based on a relevant piece of evidence (i.e. Dawn-Mir01-a). In general, however, we may have several relevant pieces of evidence and all of them have to be used in solving the above tasks. Thus, when we ask the question "Is there any potentially relevant piece of evidence?" we have to consider all the available pieces of evidence. Then we have to analyze the support or the undermining of the current hypothesis provided by each piece of evidence found. After that, we have to combine all the found solutions into a global assessment of the extent to which all the available pieces of evidence support or undermine the analyzed hypothesis.

Essentially, the more complex solution synthesis process for the initial task proceeds from the bottom up, successively combining the solutions of sub-tasks into the solutions of tasks, until the solution of the initial task "Assess whether Al Qaeda has nuclear weapons" is found, by taking into account all the available evidence and knowledge.

The next section discusses a systematic approach to the analysis of a piece of evidence.

3.2 Analysis of a Piece of Evidence

We are also developing a systematic approach to evidence analysis which identifies different types of evidence and defines analytic 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 the task "Assess the credibility of Dawn-Mir01-a." As indicated in the previous section, Dawn-Mir01-a represents a fragment from a newspaper article where the author, Hamid Mir, cites Osama bin Laden. Considering Schum's types of evidence, Dawn-Mir01-a is second hand testimonial evidence. Therefore, to analyze its credibility, one would need to consider both the credibility of Osama bin Laden (the original source), and the credibility of Hamid Mir (who may have potentially altered the information). Figure 4 illustrates this analysis in the task reduction / solution synthesis framework.

We need to:

> Assess the credibility of Dawn-Mir01-a.

Was Dawn-Mir01-a obtained as a second-hand evidence and, if yes, what is the reporting chain?

Yes, Dawn-Mir01-a is testimonial evidence of Osama bin Laden cited in Dawn-Mir01 by Hamid Mir.

Therefore we need to:

Assess the credibility of Hamid Mir as the reporter of Dawn-Mir01-a.

Assess the credibility of Osama bin Laden as the source of Dawn-Mir01-a.

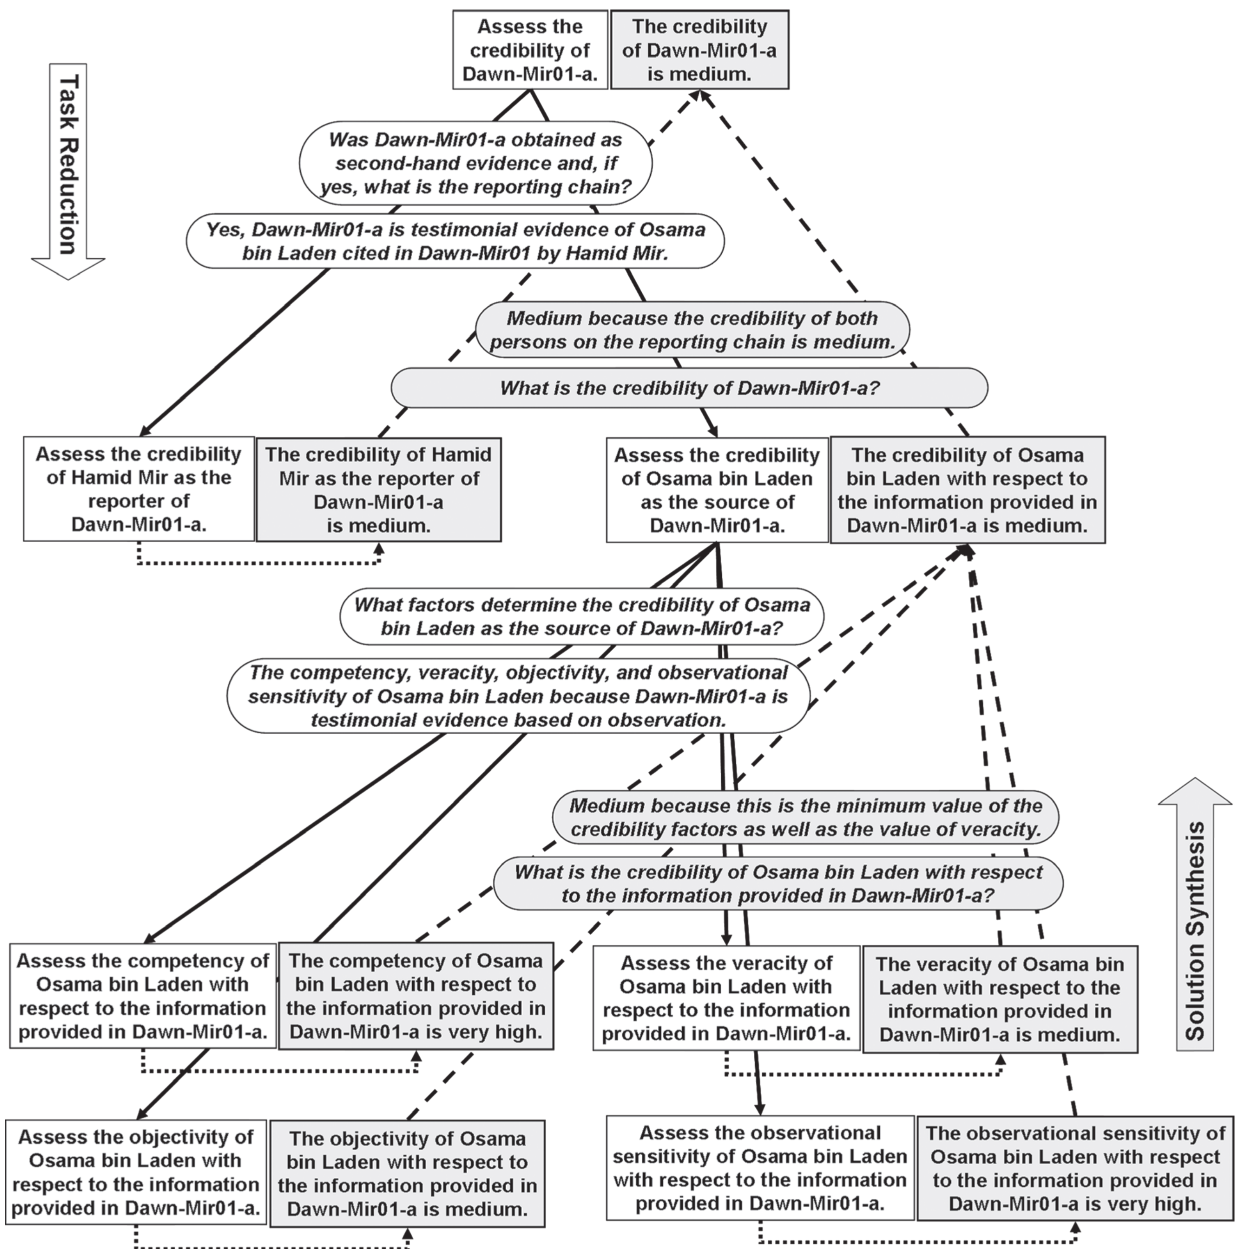

Figure 4. Credibility analysis through task reduction and solution synthesis.

Let us consider the credibility of Osama bin Laden:

> What factors determine the credibility of Osama bin Laden as the source of Dawn-Mir01-a?
>
> The competency, veracity, objectivity, and observational sensitivity of Osama bin Laden because Dawn-Mir01-a is testimonial evidence based on observation.

Therefore we need to solve the following four tasks, one for each credibility factor:

> Assess the competency of Osama bin Laden with respect to the information provided in Dawn-Mir01-a.
>
> Assess the veracity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a.
>
> Assess the objectivity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a.
>
> Assess the observational sensitivity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a.

The *veracity* of an observer refers to the degree to which that observer believes that the event actually occurred (i.e. is Osama bin Laden 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/her senses and does not let personal motivations or expectations determine what he/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).

Estimating the above components of credibility may themselves be complex tasks, solved also through task reduction and solution synthesis. For instance, to estimate the veracity of Osama bin Laden the analyst may need to consider previous statements made by him and the extent to which they proved to be true or false.

After the components of Osama bin Laden's credibility are evaluated, they need to be combined into an overall estimation of his credibility, as illustrated at the bottom of figure 4 and in the following.

We have determined that:

> The competency of Osama bin Laden with respect to the information provided in Dawn-Mir01-a is very high.
>
> The veracity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a is medium.
>
> The objectivity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a is medium.
>
> The observational sensitivity of Osama bin Laden with respect to the information provided in Dawn-Mir01-a is very high.

In this case:

> What is the credibility of Osama bin Laden with respect to the information provided in Dawn-Mir01-a?
>
> Medium because this is the minimum value of the credibility factors as well as the value of veracity.

Therefore we can conclude that:

> The credibility of Osama bin Laden with respect to the information provided in Dawn-Mir01-a is medium.

The credibility of Hamid Mir is determined in a similar manner as being medium. Then both these credibility estimates are combined into an overall credibility for the piece of evidence Dawn-Mir01-a (see the top of figure 4).

The above is an illustration of how a piece of evidence can be rigorously analyzed, based on its type, in the task reduction / solution synthesis framework. Our current research considers how to analyze other types of evidence in the presented framework.

3.3 Summary of the Hypothesis Analysis Process

The analysis from figure 3 and figure 4 illustrates the following general hypothesis analysis process:

1) A complex intelligence analysis task T is successively reduced to simpler tasks T_i that either have known solutions or can be solved through evidence analysis. As a result, the analyst determines what kind of pieces of evidence to look for.
2) Potentially relevant pieces of evidence for each task T_i are identified.
3) The identified pieces of evidence are analyzed through task reduction and solution synthesis and a solution for each task T_i is obtained.
4) The solutions of the tasks T_i are successively combined, from bottom-up, to obtain the solution for the initial task T.

The next sections present a brief overview of the three roles of Disciple-LTA: as a problem-solving assistant, as a learning system, and as a tutor.

4. Disciple-LTA as a Problem Solving Assistant

The main problem-solving engine of Disciple-LTA is based on the general task reduction / solution synthesis paradigm illustrated in the previous section. To be able to generate reasoning trees like the ones from figure 3 and figure 4, the knowledge base of a Disciple agent is structured into an object ontology and a set of general problem solving rules.

The *object ontology* is a hierarchical representation of the objects and type of objects from the application domain, together with their properties and relationships (Fensel 2000; Stanescu et al., 2003). The objects to be represented include different types of intelligence sources, such as HUMINT

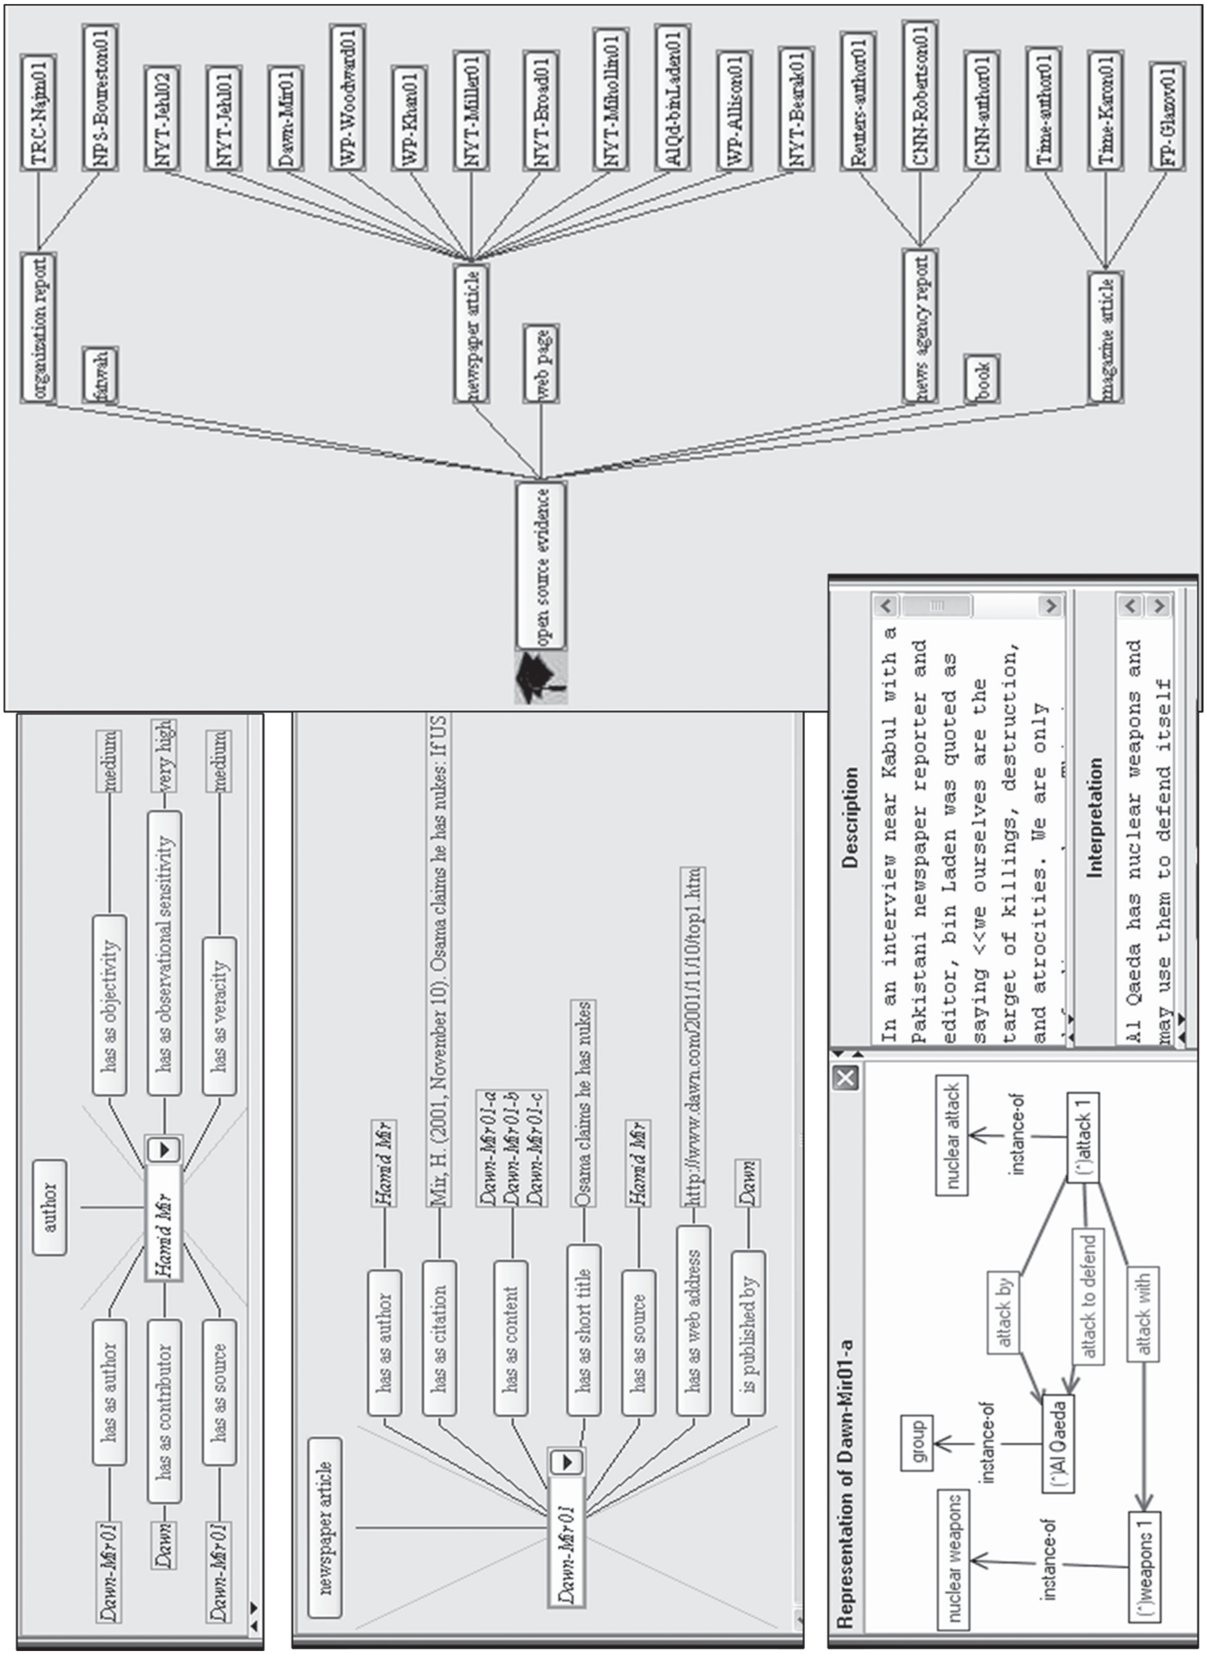

Figure 5. Fragments from the object ontology of Disciple-LTA.

(e.g. intelligence obtained from agents, informers, observants) or OSINT (e.g. intelligence obtained from books, newspaper articles, magazine articles, and research reports available to the public), as well as descriptions of entities (such as states, terrorist groups, individuals) and different types of weapons (e.g. nuclear, chemical, biological, radiological), etc. depending on the types of tasks to be solved by Disciple-LTA.

Figure 5 represents several fragments from the current object ontology of Disciple-LTA, shown in the interface of some of the system's viewers. The right hand side pane illustrates the interface of the Graphical Tree Browser. It shows a (rotated) hierarchy fragment the top of which is open-source evidence. Under it there are several types of open-source evidence such as, organization report, fatwah, newspaper article, web page, news agency report, etc. Under the category of newspaper article are the actual articles represented in the ontology, one of which is Dawn-Mir01, an article from the Dawn newspaper, authored by Mir.

The two panes from the upper left side of figure 5 illustrate the interface of the Association Browser. The top pane shows the description of Hamid Mir, the author of Dawn-Mir01. The middle pane shows the description of Dawn-Mir01. The descriptions of these entities are in terms of their properties and relationships. For instance, Dawn-Mir01 has as content Dawn-Mir01-a which was used as a piece of evidence in the tree from figure 3.

Part of the representation of Dawn-Mir01-a is shown in the interface of the Evidence Editor, located in the bottom-left side of figure 5. The upper right side of this pane shows the natural language description of Dawn-Mir01-a. Under it there is the analyst's interpretation of this description: "Al Qaeda has nuclear weapons and may use them to defend itself." At the bottom left side of figure 5 there is the formal representation of this interpretation: (*)attack1 is a nuclear attack. This is an attack by (*)Al Qaeda, an attack to defend (*)Al Qaeda, and an attack with (*)weapons1 which are nuclear weapons. Notice that some objects are preceded by *. This is to indicate that these are not the objects themselves, but views of these objects from Dawn-Mir01-a. Different pieces of evidence will generally have different and even contradictory views on some objects. For instance, in Dawn-Mir01-a, (*)Al Qaeda is viewed as a group (because this is how presumably Osama bin Laden is viewing it), while in other parts of the ontology (which represent analyst's knowledge), Al Qaeda is a terrorist group.

The reasoning trees like the ones from figure 3 and figure 4 are generated through the application of task reduction and solution synthesis rules. An example of a task reduction rule is presented in figure 6. This is an IF-THEN rule that indicates the condition under which one could solve the IF task by solving the

IF: Assess the credibility of *?O1*.

Q: Was *?O1* obtained as a second-hand evidence and, if yes, what is the reporting chain?

A: Yes, *?O1* is a testimonial evidence of *?O3* cited in *?O4* by *?O2*.

MAIN CONDITION

VAR	BOUND
?O1	(testimonial evidence)
?O2	(source)
?O3	(source)
?O4	(evidence)

VAR	RELATIONSHIP	VAR
?O4	has as content	*?O1*
?O1	has as source	*?O3*
?O4	has as source	*?O2*

THEN: Assess the credibility of *?O2* as the reporter of the *?O1*.
Assess the credibility of *?O3* as the source of the *?O1*.

Figure 6. Task reduction rule.

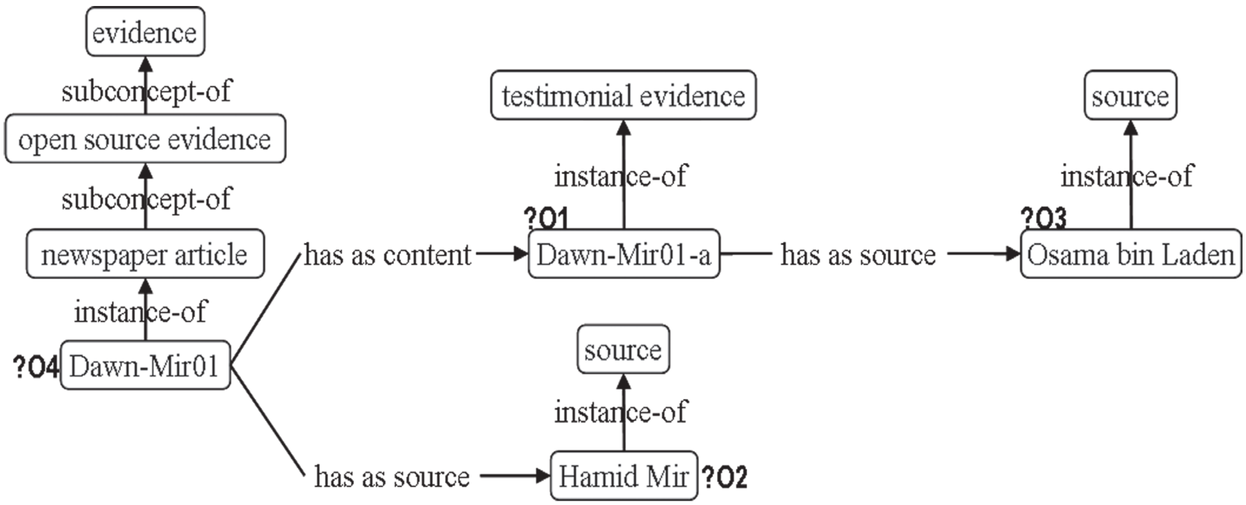

Figure 7. Fragment from the object ontology of Disciple-LTA.

THEN tasks. There are two alternative forms of the condition, one for the analyst and the other for the agent.

The condition for the analyst consists of a question (Q) and its answer (A). That is, if the question Q has the answer A, then one can reduce the IF task to the THEN tasks. Notice that the rule in figure 6 was used to generate the task reduction step from the top of figure 4 by instantiating ?O1 to Dawn-Mir01-a, ?O2 to Hamid Mir, ?O3 to Osama bin Laden, and ?O4 to Dawn-Mir01. Thus the rule's tasks, question and answer are used to generate reasoning trees that can be naturally followed by an analyst.

The condition for the agent (called MAIN CONDITION in figure 6) consists of relationships that have to exist between the rule's variables in order to reduce the IF task to the THEN tasks. More specifically, in order to determine whether a rule like the one in figure 6 can be applied to reduce a task, such as "Assess the credibility of Dawn-Mir01-a" from the top of figure 4, the agent has to do two things. First, it has to check that the IF task of the rule (i.e. "Assess the credibility of ?O1") matches the current task (which is true for ?O1 instantiated to Dawn-Mir01-a). Second, it has to determine that the MAIN CONDITION of the rule is satisfied. This means that the agent has to find objects in its ontology that match all the rule's variables and satisfy the condition. In our case the instances of the variables indicated in the previous paragraph have all the features specified in the MAIN CONDITION, as shown in figure 7. Indeed, Dawn-Mir01 (which matched ?O4) is a newspaper article, which is an "open source evidence" (see also the right hand side of figure 5), which is, in turn, contained within the broader category of "evidence." Thus the condition "?O4 is an evidence" is satisfied. The conditions "?O4 has as source ?O2" and "?O4 has as content ?O1" are also satisfied because "Dawn-Mir01 has as source Hamid Mir" and "Dawn-Mir01 has as content Dawn-Mir01-a" (see these features also in the middle left pane from figure 5), and so on.

5. Disciple-LTA as a Learner

The Disciple-LTA learning agent shell 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 and a tutor for new analysts. The process of developing

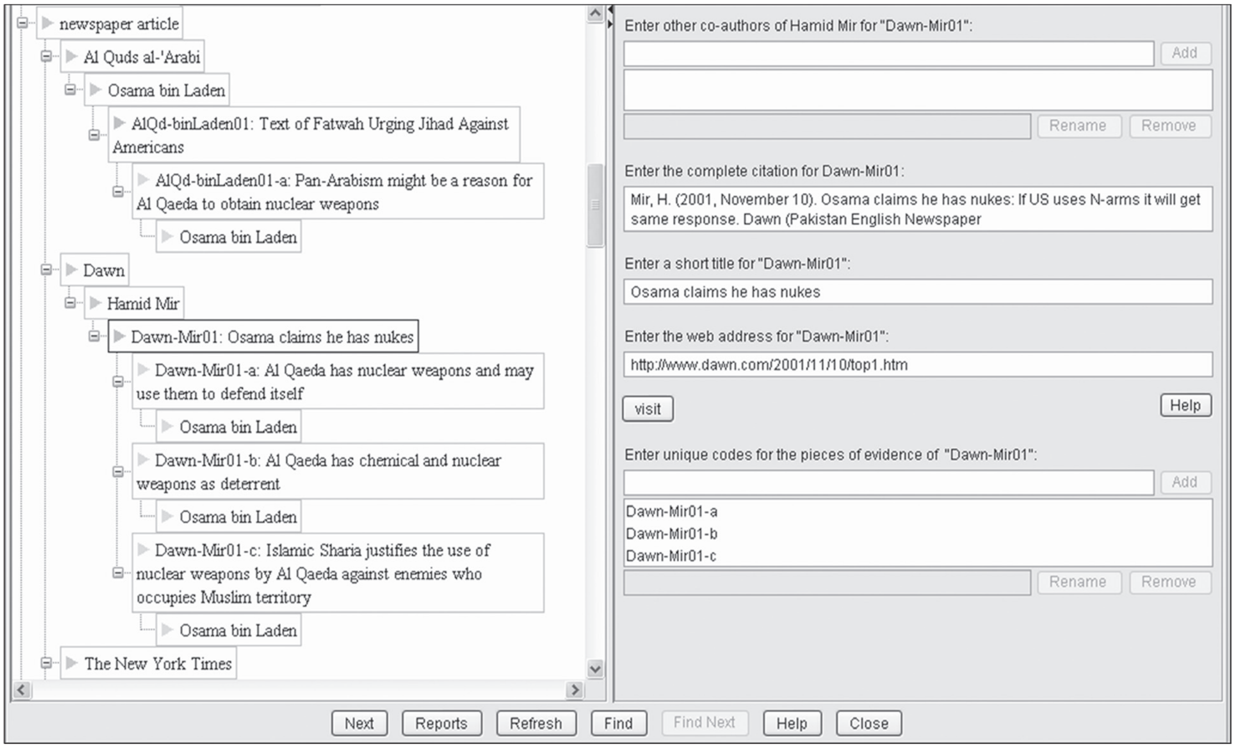

Figure 8. The interface of the Evidence Elicitation module.

a specific Disciple-LTA agent for a class of intelligence analysis tasks consists of two main phases, which are briefly discussed in the following:

1) The development of an initial object ontology, 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.

5.1 Ontology Development

The object ontology is developed by importing ontological knowledge from existing knowledge repositories and by using the ontology building tools of Disciple. While, in general, the ontology is developed by a knowledge engineer and an expert analyst, several of the ontology modules have been developed to be easily used by the analyst, with no assistance from a knowledge engineer. Two of such modules are the Graphical Tree Browser and the Association Browser mentioned in section 4. Another one is the Evidence Elicitation module, the interface of which is shown in figure 8. This module guides the analyst to describe an open-source piece of evidence, such as a newspaper article (e.g. Dawn-Mir01), by requesting specific pieces of information (e.g. "Enter the complete citation for Dawn-Mir01" or "Enter a short title for Dawn-Mir01"). Based on the information provided by the analyst, Disciple-LTA generates an internal representation of the piece of evidence. For instance, based on the dialog illustrated in figure 8, Disciple-LTA generated the internal representation of Hamid Mir and Dawn-Mir01 shown in the two upper-left panes from figure 5.

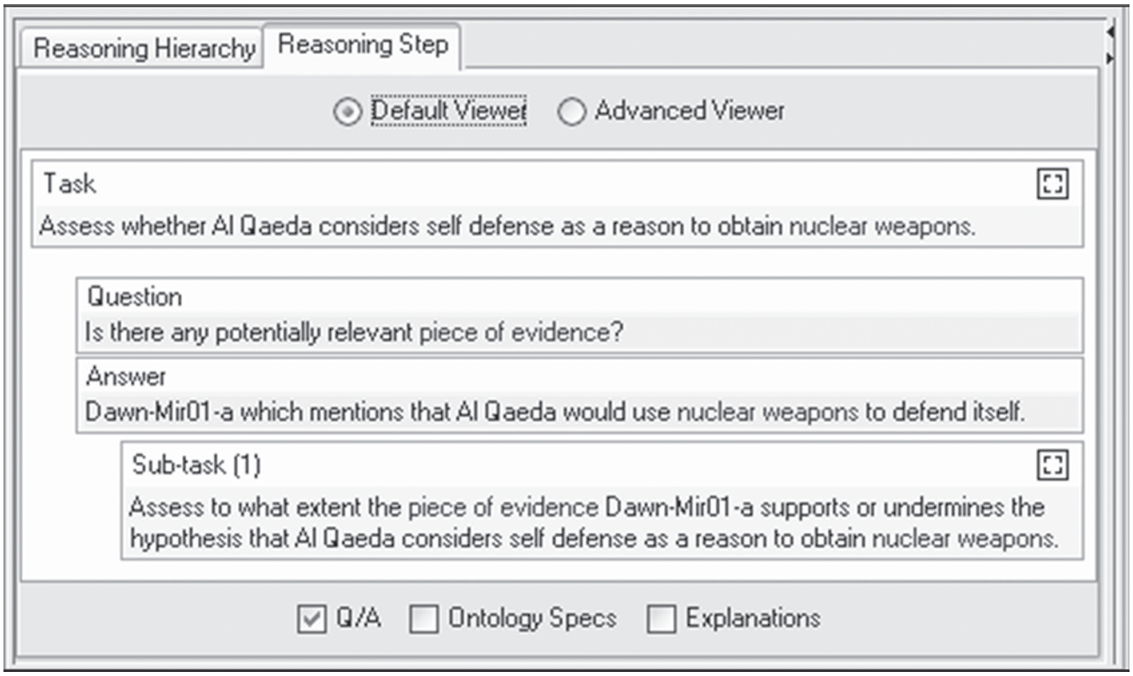

Figure 9. Task reduction step from figure 3 shown in the Reasoning Step Viewer.

5.2 Agent Teaching and Rule Learning

Once an initial object ontology has been developed, the analyst may start teaching the agent his/her problem solving expertise, including his/her problem solving strategies, prior and tacit knowledge. The analyst considers typical intelligence analysis tasks, such as the one from the top of figure 3, 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, which will allow it to solve similar problem solving tasks (Boicu, 2002).

Consider, for instance, the reasoning step from the middle of figure 3, this time shown in the interface of the Reasoning Step Viewer from figure 9. From this reasoning step, Disciple-LTA has learned the IF-THEN task reduction rule from figure 10. This rule indicates under what condition that IF task can be reduced to the simpler THEN task. Notice that the tasks, question, and answer are generalizations of the corresponding expressions from the example, where each specific object (i.e. Al Qaeda and Dawn-Mir01-a) was generalized to a variable (i.e. ?O1 and ?O2, respectively). When this rule is used in problem solving, the variables ?O1 and ?O2 are instantiated to specific objects, as discussed in section 4. The role of the rule's condition is to restrict the possible values of these variables so that any instantiation of the rule makes sense. Thus the main goal of the rule learning process is to determine this condition, which is done as briefly explained below.

First of all, let us notice that the analyst's explanation of why the task reduction step from figure 9 makes sense is given by the question and answer following the top task:

Is there any potentially relevant piece of evidence? [1]

Dawn-Mir01-a which mentions that Al Qaeda would use nuclear weapons to defend itself.

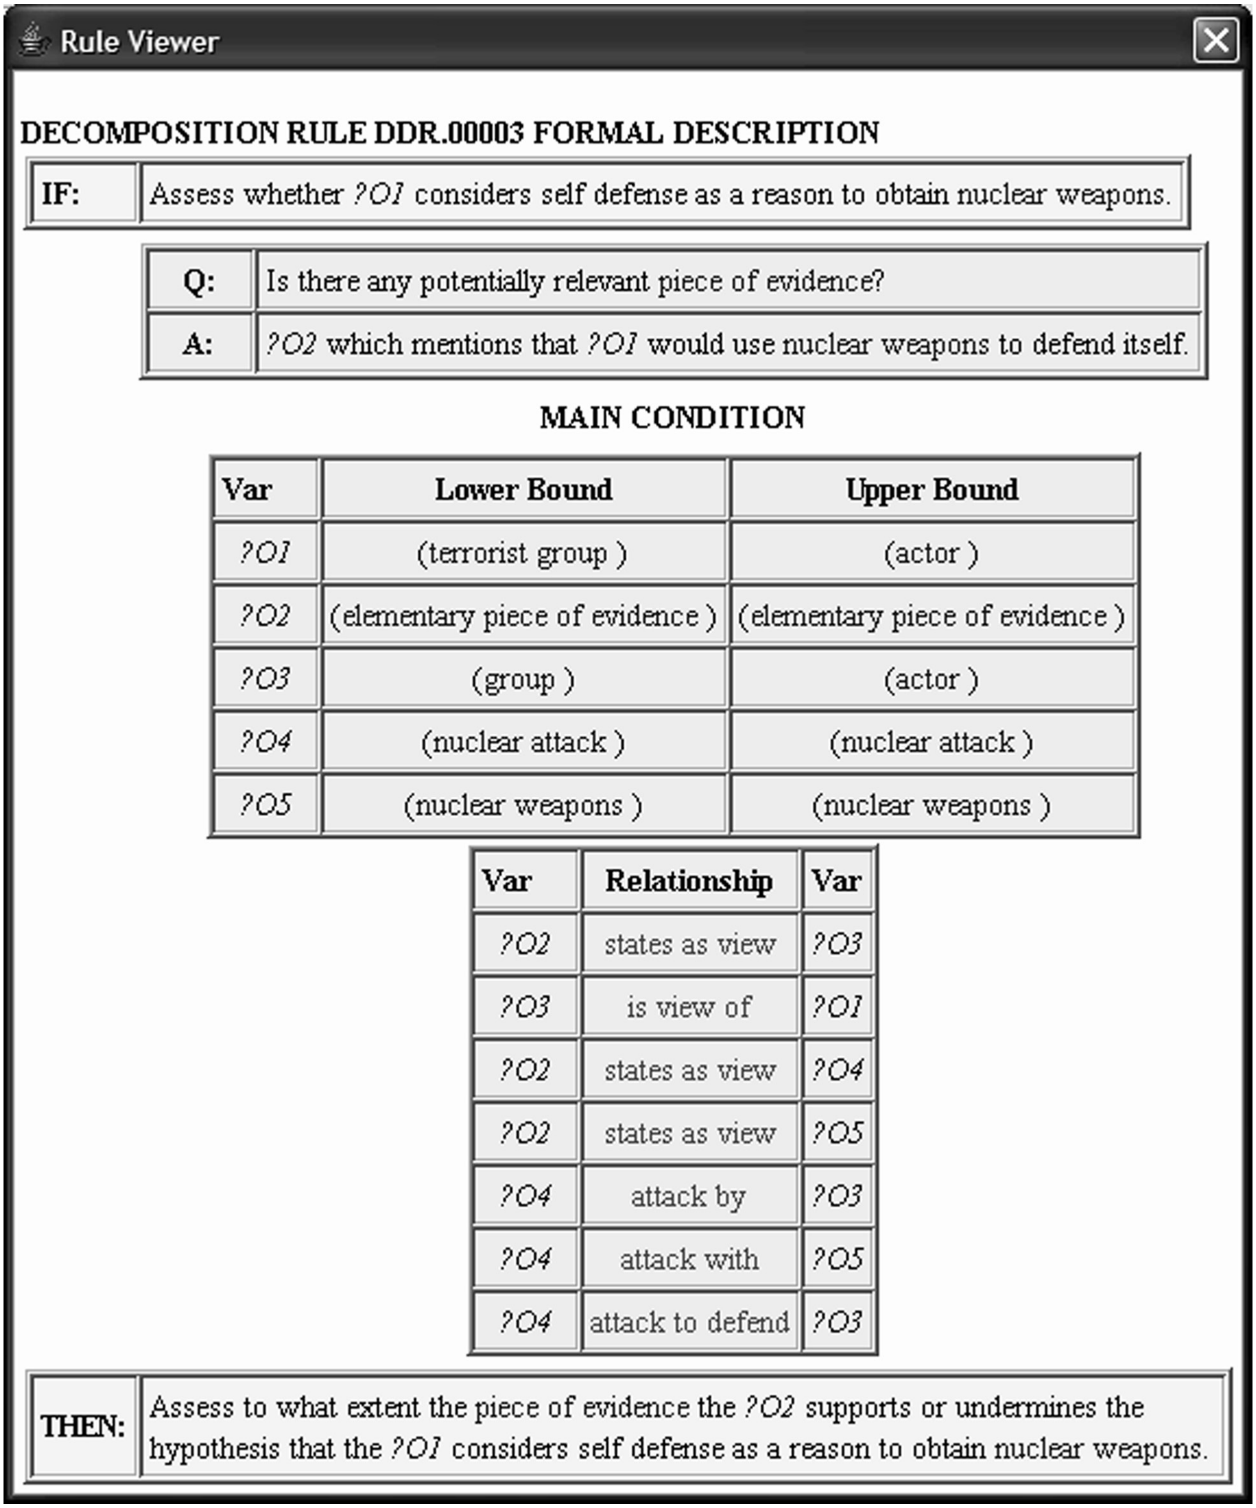

Rule Viewer

DECOMPOSITION RULE DDR.00003 FORMAL DESCRIPTION

IF:	Assess whether *?O1* considers self defense as a reason to obtain nuclear weapons.

Q:	Is there any potentially relevant piece of evidence?
A:	*?O2* which mentions that *?O1* would use nuclear weapons to defend itself.

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(elementary piece of evidence)	(elementary piece of evidence)
?O3	(group)	(actor)
?O4	(nuclear attack)	(nuclear attack)
?O5	(nuclear weapons)	(nuclear weapons)

Var	Relationship	Var
?O2	states as view	*?O3*
?O3	is view of	*?O1*
?O2	states as view	*?O4*
?O2	states as view	*?O5*
?O4	attack by	*?O3*
?O4	attack with	*?O5*
?O4	attack to defend	*?O3*

THEN:	Assess to what extent the piece of evidence the *?O2* supports or undermines the hypothesis that the *?O1* considers self defense as a reason to obtain nuclear weapons.

Figure 10. A task reduction rule learned from the example in figure 9.

However the question-answer pair is in natural language and the analyst has to help Disciple-LTA to "understand" its meaning, that is, to determine its representation in terms of the objects and features from the agent's object ontology.

First let us notice that Dawn-Mir01-a is a piece of evidence that includes some statements about Al Qaeda. However these statements may or may not be true. Therefore they are not represented in the agent's knowledge base as facts about Al Qaeda, but as facts about the view of Al Qaeda expressed by Dawn-Mir01-a, view denoted by (*)Al Qaeda. Thus the names of all the objects mentioned in Dawn-Mir01-a (e.g. Al Qaeda, weapons, attack) are preceded by (*). This aspect was also mentioned in section 4, where we have presented the representation of Dawn-Mir01-a inside the Evidence Editor (see the lower left pane from figure 5).

In the object ontology of Disciple-LTA, Dawn-Mir01-a is represented as a piece of evidence that mentions an attack by Al Qaeda, an attack with nuclear weapons, attack intended to defend Al Qaeda, which is in essence the meaning of the above question-answer pair. Thus the following expressions (called explanation pieces) express the meaning of the question-answer pair [1] in the agent's ontology:

Dawn-Mir01-a —is—> piece of evidence [2]

Dawn-Mir01-a —states as view—> (*)Al Qaeda —is view of—> Al Qaeda

Dawn-Mir01-a —states as view—> (*)weapons 1 —is—> nuclear weapons

Dawn-Mir01-a —states as view—> (*)attack 1 —is—> nuclear attack

(*)attack 1 —attack by—> (*)Al Qaeda

(*)attack 1 —attack with—> (*)weapons 1

(*)attack 1 —attack to defend—> (*)Al Qaeda

Disciple-LTA will use analogical reasoning with previously learned rules, as well as general heuristics, to hypothesize the meaning of the question-answer pair [1]. It will generate plausible explanation fragments, like the above ones [2], ordered by their plausibility. Then the analyst will select those explanation fragments that best express the meaning of the question-answer pair [1]. As illustrated above, each explanation fragment proposed by the agent is a relationship (or a relationship chain) involving instances, concepts and constants from the task reduction step and from the knowledge base. The analyst may also help the agent to propose the right explanation pieces by proving hints, such as pointing to a relevant object that should be part of them.

Based on the example from figure 9 and its explanation [2], and using the object ontology as the basis for generalization, Disciple-LTA learned the rule from figure 10. Notice that this rule is only partially learned, and will need to be further improved. Indeed, instead of a single applicability condition (MAIN CONDITION), it has two bounds for this condition, a Lower Bound and an Upper Bound. The Lower Bound is obtained from the minimally general generalization of the example and its explanation, in the context of the current ontology, while the Upper Bound is obtained from the maximally general generalization. For instance, Al Qaeda (?O1) has been generalized to a terrorist group in the Lower Bound and to an actor in the Upper Bound. Based on additional examples, Disciple will determine whether the exact set of values for ?O1 should be terrorist group, or actor, or something in between, such as group.

From each of the reasoning steps in figures 3 and 4, Disciple-LTA learned a reasoning rule like the one in figure 10. As Disciple-LTA is trained by the analyst and learns new rules, the analyst-agent interaction evolves from a teacher-student interaction toward an interaction where both collaborate in performing an intelligence analysis task. For instance, based on the rules learned from the tree in

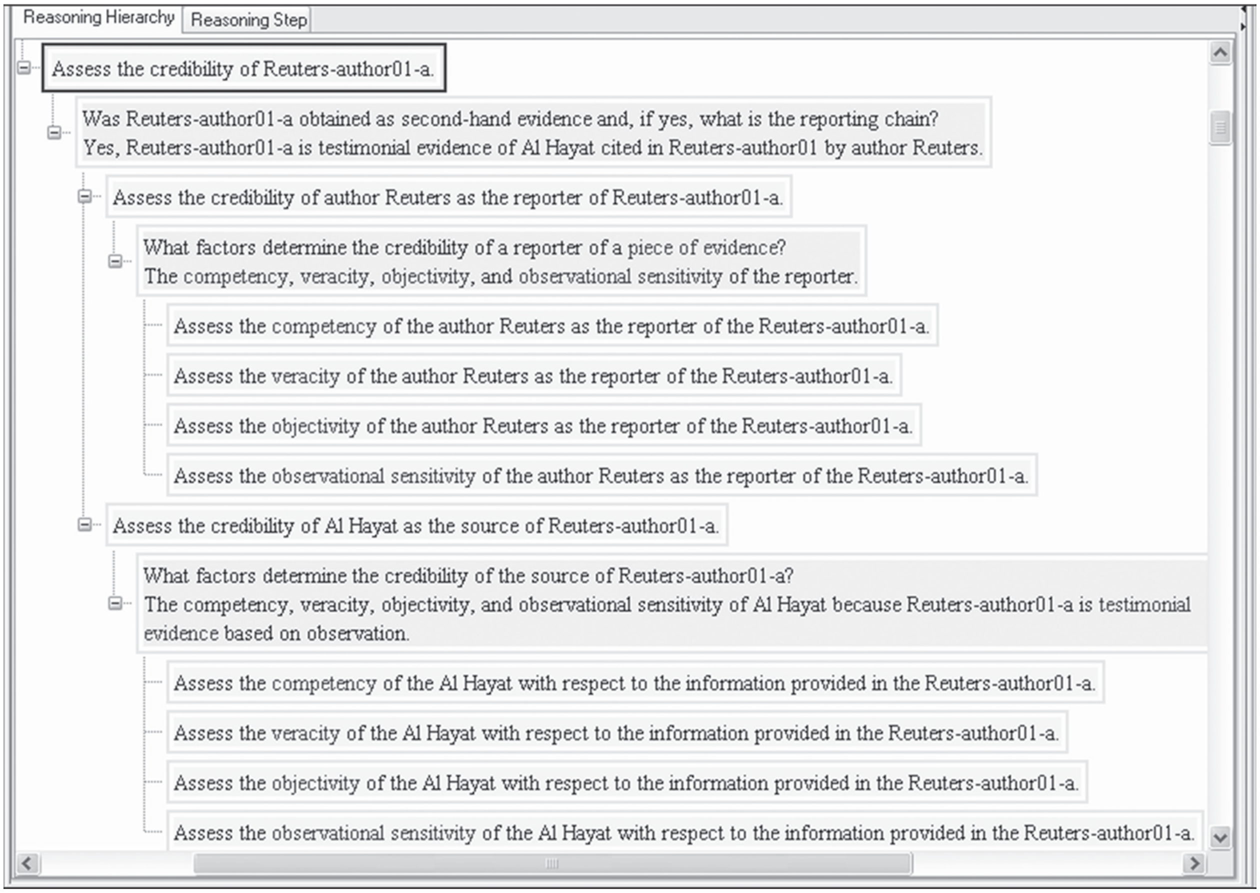

Figure 11. The interface of the Reasoning Hierarchy Viewer.

figure 4, Disciple-LTA automatically generated the tree from figure 11 (shown in the interface of the Hierarchical Viewer).

During interactive problem solving with the analyst, 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. For instance, Disciple-LTA will generalize the lower bound condition of the rule in figure 10 when proposing correct reductions in order to include a generalization of these new cases. It will also specialize the upper bound condition when proposing incorrect reductions, to avoid making similar mistakes in the future. Alternatively, it may add except-when conditions that should not be satisfied in order for the rule to be applicable. Thus, in time, the lower and upper bounds of the learned rules will converge toward one-another, and new except-when conditions may be added to the rules, leading to complex, but accurate reasoning rules that represent analyst's expertise.

6. Disciple-LTA as a Tutor

The main idea of using Disciple-LTA as a tutor 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 3, and will explain to the student analyst how to solve them in a systematic way.

Our current research concentrates on developing the tutoring capabilities of Disciple-LTA by developing modules for learning and refining tutoring knowledge, for learning a student's model, for test generation, and for student evaluation. Experimentation and evaluation of these modules was performed at the U.S. Army War College during the Spring 2006 offering of the course "Military Applications of Artificial Intelligence: Intelligence Analysis."

7. Experimentation Environment and Experimental Results

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 takes 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 use an experimentation environment similar to one which was very successfully used for Disciple-RKF, as part of the DARPA's Rapid Knowledge Formation Program (Tecuci et al., 2002). Experimentation is conducted during the Military Applications of Artificial Intelligence (MAAI) course, taught every Spring at the U.S. Army War College. MAAI is an elective attended by intelligence analysts and individuals with an interest in artificial intelligence and intelligence analysis; but generally, there are no students with knowledge engineering experience. The first evaluation took place in Spring 2005.

Throughout the conduct of the course, we have collected useful feedback on the Disciple approach and the current implementation of the tools used by the students. The feedback was collected in the form of questionnaires where the students were requested to express their agreement or disagreement with various statements, taking one of the following 5 positions: strongly agree, agree, neutral, disagree, and strongly disagree. Sample evaluation results are presented in table 1.

The top part of table 1 contains the evaluation results for the task reduction paradigm discussed in section 3. The second part of table 1 contains the evaluation results for the Evidence Elicitation Tool, the interface of which is shown in figure 8. The third part of table 1 contains the evaluation of the Rule Learning Tool which was described in section 5.2. Finally, the bottom part of table 1 contains some general assessments of Disciple. These and other results provided both valuable feedback on the current development of Disciple-LTA, and a significant support of its usefulness to the Intelligence Community.

In the next section we review some of the challenges of intelligence analysis and how the current or future versions of Disciple-LTA can address them.

The Task-Reduction Paradigm Used in Disciple-LTA	SA	A	N	D	SD
The task-reduction paradigm is adequate to express the logic for analyzing an issue	3	4			
The task-reduction logic is easy to understand	4	3			
I can use the task-reduction paradigm with an acceptable amount of effort	1	6			

Evidence Elicitation Tool	SA	A	N	D	SD
The Evidence Elicitation Tool provides a good way to describe a piece of evidence	1	5			
It is easy to learn how to use the Evidence Elicitation Tool	3	2	1		
It is easy to use the Evidence Elicitation Tool to describe a piece of evidence	3	3			

Rule Learning Tool	SA	A	N	D	SD
It is easy to learn how to use the Rule Learning Tool	1	2	4		
It is easy to use the Rule Learning Tool	2	4	1		
It is easy to search for explanations	1	6			
It is easy to direct Disciple to generate explanations related to certain instances	2	4	1		
The proposed explanations are easy to understand	2	5			
The most relevant explanations are proposed first	3	4			

General Assessments	SA	A	N	D	SD
The use of Disciple is well suited to learn basic artificial intelligence concepts	2	2			
The use of Disciple is well suited to learn basic intelligence analysis concepts		4			
The use of Disciple is a useful learning experience	2	2			
Disciple should be used in future sessions of this course	1	3			
It is easy to learn how to use Disciple	1	3			
It is easy to use Disciple	1	3			
The learning experience offered by Disciple justifies the effort required for using it	2	2			

Table 1. Sample evaluation results.

8. Addressing Intelligence Analysis Challenges: Current Results and Future Research

8.1 Difficult to acquire and retain analytic expertise

Intelligence analysis expertise takes years to establish, is lost when analysts separate from service, and is costly to replace. As discussed in section 5, Disciple-LTA agents can rapidly acquire, use, and preserve the prior and tacit knowledge of an expert analyst who can directly teach an agent how to analyze hypotheses, in a natural way, similarly to how the analyst would teach a person. The main feature of the Disciple-LTA is its capability to rapidly learn rules like the one from figure 10, based only on specific examples and explanations. This is in contrast with other approaches to knowledge bases and agents development, such as CYC (Siegel et al., 2005), where the rules have to be manually defined, verified and corrected by knowledge engineers and subject matter experts, through a process that is long, difficult, and error-prone.

8.2 Difficult to train new analysts

As discussed in section 6, Disciple-LTA is currently developed to be used as a tutoring system for new analysts. It is expected that the resulting system will have the following capabilities:

- it can be rapidly trained by a subject matter expert to solve problems in a systematic, pedagogical way, and to explain its knowledge;
- it will allow a subject matter expert and a knowledge engineer to define a curriculum for intelligence analysis, based on a repository of knowledge bases;
- it will allows a teacher to customize a curriculum based on specific course objectives;
- it will be able to dynamically generate lessons based on the content of the knowledge base used in the course, and further customize the curriculum to teach those concepts and problem solving strategies that can be illustrated from the knowledge base.

8.3 Difficult to analyze in reference to the culture of the data source

Disciple-LTA can be trained by a user to solve problems and reason in a manner that is similar to that user. Therefore Disciple-LTA could be trained to reason as an opponent by an expert in the culture of the opponent. For instance, an expert in the reasoning of an Al Qaeda terrorist could train a Disciple agent to reason like a terrorist. Such a trained Disciple agent could then be used to simulate a terrorist and help avoid mirror-imaging in intelligence analysis (i.e. avoid projecting on terrorists the analyst's own values and way of thinking).

8.4 Difficult to collaborate with other analysts and experts

The task reduction paradigm and networking capabilities facilitate collaboration between complementary experts and their agents. Consider, for instance, the task "Assess whether Al Qaeda has nuclear weapons" from the top of figure 3. This task is reduced to four simpler assessment tasks and therefore, to perform it, one has to perform these simpler tasks. Each of these simpler assessments could, in principle, be performed by a different expert. Then their solutions could be combined into the solution of the initial task. This is very important because each task may require a different type of expertise, and no analyst is an expert in all areas. Consider the task "Assess whether Al Qaeda has the ability to obtain nuclear weapons." The performance of this task requires deep knowledge of nuclear weapons because one strategy to obtain them is to acquire and assemble its components.

Disciple-LTA adds several dimensions to the collaboration issue. For instance, when the required experts are not available, their trained Disciple-LTA assistants can act on their behalf, solving the problems in a manner that would closely resemble the way that the experts would have solved them.

Similarly, solutions of certain problems could be stored and reused as solutions of subproblems of other problems. Disciple-LTA facilitates this reuse because it provides not just the solution, but its entire justification which could be checked and updated, in light of new evidence.

8.5 Difficult to share intelligence

The following illustrates this difficulty (Kris, 2005):

> *While I was serving as an intelligence analyst at the US Central Command in Qatar during operations Enduring Freedom and Iraqi Freedom in 2003, my team and I analyzed hundreds of messages and reports each day. We created briefings used by generals Tommy Franks and John Abizaid. A vast amount of information was available to us on Intelink, but there was no simple way to find and use the data efficiently....And while there were hundreds of people throughout the world reading the same materials, there was no easy way to learn what they thought. Somebody had answers to my questions, I knew, but how were we ever to connect? The scary truth is that most of the time analysts are flying half blind.*

Consider the rule from figure 10 learned by Disciple-LTA. This rule allows Disciple to automatically identify pieces of evidence that are potentially relevant to the general task "Assess whether ?O1 considers self defense as a reason to obtain nuclear weapons." Using it, Disciple-LTA automatically identified Reuters-author01-a which represents the following fragment from a Reuters report: *"Al-Qaeda would use the [nuclear] weapons only inside the United States or if the group faced a "crushing blow" which threatened its existence, such as the use of nuclear or chemical weapons against its fighters."* (Reuters, 2004). Moreover, Disciple-LTA also built the analysis tree from figure 11.

As another example, consider Dawn-Mir-01a, which was used in the analysis illustrated in figure 3 and figure 4. An analyst has represented this piece of evidence, as illustrated in the left-hand side of figure 5, and has analyzed its credibility, as illustrated in figure 4, making it possible for the agents of other analysts to automatically retrieve and reason with it. Similarly, stored analyses like the one in figure 4 could be retrieved and reused by the agent of another analyst trying to answer the same questions.

Therefore, Disciple-LTA agents acting as assistants to human analysts, connected on the information grid, and sharing common knowledge bases, as envisioned by the illustration from figure 2, would do much to alleviate the difficulties mentioned in the quotation from the U.S. Central Command analyst (Kris, 2005).

8.6 Difficult to avoid a narrowly focused analytic mindset

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 the analyst's assumptions, preferences and biases. To illustrate this, let us consider again the reasoning tree in figure 3, this time developed by the trained Disciple-LTA agent. The rightmost assessment task from the upper part of figure 3 is "Assess whether Al Qaeda considers an ideology as a reason to obtain nuclear weapons." The agent may reason under the analyst's assumption that Al Qaeda does not consider its ideology as a reason to obtain nuclear weapons, 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, a document is found in which Osama bin Laden states that it is a religious duty for Muslims to acquire nuclear weapons (Time Asia, 2001).

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.

Another way to improve the product of the analytic process is to have the same tasks solved by different analysts and to compare their analyses. While in general this cannot be done by human analysts due to limited resources, it can be done by their agents. Such a comparative analysis may also reveal hard to detect biases and assumptions made by different analysts.

Assumption checking, consideration of *what-if scenarios,* and *comparative analysis* are three important capabilities that we plan to integrate into Disciple-LTA to help an analyst avoid the analytic mindset.

8.7 Difficult to consider multiple hypotheses

The automation of the hypothesis analysis process, the collaboration with other analysts, the reuse of existing analyses and annotated pieces of evidence, not only increase the quality of the analysis, but also reduce the time required to perform it, allowing an analyst to consider more hypotheses, and to compare and contrast them, to arrive at the best possible answers.

8.8 Difficult to rigorously explain the analysis

This difficulty is very clearly stated in the report of *The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction* (CNN, 2005):

> *"The main cause [of the Iraq WMD 'intelligence failure'], the commission said, was the intelligence community's inability to collect good information about Iraq's WMD programs, serious errors in analyzing what information it could gather and a failure to make clear just how much of its analysis was based on assumptions rather than good evidence." "No important intelligence assessment should be accepted without sharp questioning that forces the (intelligence) community to explain exactly how it came to that assessment and what alternatives might also be true."*

As discussed in the previous sections, Disciple-LTA keeps an accessible record of its analysis which *delineates the logic, the evidence and how it was used,* and *the assumptions made*.

9. Conclusions

This paper has presented current research on the development of 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 an 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 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

This research was performed in the Learning Agents Center of George Mason University and in the Center for Strategic Leadership of the Army War College, and benefited from the contribution of several members of these centers, including Vu Le (who developed several tree viewers and modules for tutoring), William Cleckner (who supported the use of Disciple-LTA in the MAAI course), Denitsa Apostolova (who developed the Evidence Editor), Bogdan Stanescu (who developed the Script Editor and components of the Rule Learning Tool), Adrian Ciubotariu (who developed the scripts for the Evidence Elicitation Tool), and William Wagner (who contributed to the development of the Rule and Task Browsers).

The research of the Learning Agents Center is sponsored by several U.S. government agencies including the Air Force Research Laboratory (FA8750-04-1-0527), the Air Force Office of Scientific Research (F9550-07-1-0268), the National Science Foundation (0610743), the Intelligence Community, and the Army War College. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

References

Boicu, M. 2002. *Modeling and Learning with Incomplete Knowledge*. PhD Thesis in Information Technology, Learning Agents Laboratory, George Mason University.

Boicu M.; Tecuci G.; Ayers C.; Marcu D.; Boicu C.; Barbulescu M.; Stanescu B.; Wagner W.; Le V.; Apostolova D.; Ciubotariu A. 2005. A Learning and Reasoning System for Intelligence Analysis, *Proceedings of the Twentieth National Conference on Artificial Intelligence,* AAAI-05, Pittsburgh, Pennsylvania, USA, July 9-13.

CNN.com 2005. On the report of *The Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction* (March 31, 2005).

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.

Fensel D. 2000. Ontologies: A Silver Bullet for Knowledge Management, Springer-Verlag, Berlin.

Kris A. 2005. An Army officer calls for better information gathering. *Wired Magazine.com,* 13[3], http://www.wired.com/wired/archive/13.03/view.html?pg=2?tw=wn_story_top5

Lowenthal, 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.

Mir, H. (2001, November 10). "Osama claims he has nukes: If US uses N-arms it will get same response." *Dawn* (Pakistan English Newspaper).

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*, U.S. Army CECOM, Fort Monmouth, New Jersey.

Reuters (Cairo), 8 Feb 2004, "Arab Newspaper Says Al Qaeda Has Ukrainian Nukes." (Reuters' source: al-Hayat Newspaper.)

Schum D.A. 2001. *The Evidential Foundations of Probabilistic Reasoning,* Northwestern University Press.

Siegel, Nick, B.; Shepard, J. Cabral,; M. Witbrock. 2005. Hypothesis Generation and Evidence Assembly for Intelligence Analysis: Cycorp's Nooscape Application. In Proceedings of the 2005 International Conference on Intelligence Analysis, VA, 2-6 May, 2005.

Stanescu B.; Boicu C.; Balan G.; Barbulescu M.; Boicu M.; Tecuci G. 2003. "Ontologies for Learning Agents: Problems, Solutions and Directions," in *Proceedings of the IJCAI-03 Workshop on Ontologies and Distributed Systems,* Acapulco, Mexico, August, AAAI Press, Menlo Park, CA pp. 75-82.

Tecuci G. 1988. Disciple: A Theory, Methodology and System for Learning Expert Knowledge, 197 pages, Thèse de Docteur en Science, University of Paris-South.

Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London, England: Academic Press.

Tecuci G.; 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.

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

Tecuci G.; Boicu M.; Boicu C.; Marcu D.; Stanescu B.; Barbulescu M. 2005b. The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence,* Vol.21, No.4, 2005, pp. 462-479.

Time Asia (2001, September 14). Exclusive Interview: Conversation with Terror; Osama bin Laden Lashes Out Against the West: TIME's January 1999 Interview. Time Asia. (Text available at: http://www.time.com/time/asia/news/printout/0,9788,174550,00.html)

Wheaton, K. J. 2001. The *Warning Solution: Intelligent Analysis in the Age of Information Overload*. Fairfax, VA: AFCEA International Press.