

***Disclaimer:** The views and opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of any agency of the U.S. government. Examples of analysis performed within this article are only examples. They should not be utilized in real-world analytic products as they are based only on very limited and dated open source information. Assumptions made within the analysis are not reflective of the position of any U.S. government entity.*

## **Disciple-LTA: Learning, Tutoring and Analytic Assistance<sup>1</sup>**

**Abstract:** This paper presents an overview of Disciple-LTA, a new type of analytic tool that synergistically integrates three complex capabilities. It can rapidly *learn* the analytic expertise which currently takes years to establish, is lost when analysts separate from service, and is costly to replace. It can *tutor* new intelligence analysts how to systematically analyze complex hypotheses. Finally, it can *assist* the analysts to analyze complex hypotheses, collaborate, and share information.

**Keywords:** structured argumentation, assumptions-based analysis, learning analytic expertise, tutoring new analysts, mixed-initiative analytic assistance, terrorism.

### **1. Introduction**

Disciple-LTA is a unique and complex analytic tool that integrates powerful capabilities for analytic assistance, learning and tutoring. At the basis of Disciple-LTA is an approach to software agent development where a subject matter expert, such as an intelligence analyst, can teach a Disciple software agent how to solve problems in a way that resembles how the expert would teach a student or a new analyst when solving problems in collaboration (Tecuci, 1998). The expert analyst will formulate a specific problem, such as “Assess whether Iran is pursuing nuclear power for peaceful purposes”, and will explain the agent how she or he performs this analysis. From this specific analysis Disciple-LTA will learn general analysis rules that will allow it to solve similar problems, such as “Assess whether Venezuela is pursuing nuclear power for peaceful purposes.” As Disciple-LTA learns from the expert analyst, it becomes an increasingly useful collaborator, solving the problems formulated by the analyst similarly to how the analyst would solve them himself or herself, only much faster. This allows the human analyst to act as the orchestrator of the analytic process, guiding the high level exploration of the reasoning space, while Disciple-LTA implements this guidance, together forming a very powerful team.

Critical to the success of Disciple-LTA as an analytic tool is the use of a systematic approach to hypothesis analysis which is both natural for a human analyst and appropriate for an automatic software agent. This approach facilitates the teaching of a Disciple-LTA agent by an expert analyst. It also allows a trained Disciple-LTA agent to act as a natural extension of the analytical reasoning capabilities of a human analyst, as well as teach new analysts in a way that is similar to how it was taught.

This paper presents a very brief overview of Disciple-LTA’s capabilities, from a user’s point of view. The next section introduces Disciple-LTA’s approach to hypothesis analysis. Then section 3 shows how Disciple-LTA helps an analyst to rapidly analyze complex hypotheses and perform assumptions-based analysis. Section 4 illustrates the teaching of Disciple-LTA by an expert analyst, and section 5 discusses how Disciple-LTA can teach new analysts. Finally, section 6 presents an overview of other capabilities of Disciple-LTA.

---

<sup>1</sup> Authors: G.Tecuci, M. Boicu, D. Marcu, V. Le, C. Boicu, George Mason University, Fairfax, VA, 22030, email: tecuci@gmu.edu

## 2. Hypothesis analysis through problem reduction and solution synthesis

Disciple-LTA is a knowledge-based software agent that can analyze complex hypotheses through a divide and conquer approach. For instance, let us consider that Disciple-LTA is given the complex hypothesis analysis problem from the top of Figure 1:

Assess whether **Al Qaeda** has **nuclear weapons**.

Disciple-LTA successively reduces this problem to simpler and simpler hypothesis analysis problems, guided by questions and answers:

What factors should I consider to determine whether **Al Qaeda** has **nuclear weapons**?

Characteristics associated with possession of **nuclear weapons** and current evidence that is has **nuclear weapons**.

Therefore I have to

Assess whether **Al Qaeda** has **nuclear weapons** based on the characteristics associated with the possession of **nuclear weapons**.

Assess whether there is current evidence that **Al Qaeda** has **nuclear weapons**.

Each of these two hypothesis analysis problems is reduced in a similar way, guided by questions and answers, as illustrated in Figure 1. This problem reduction process continues until the resulting problems are simple enough to find their solutions.

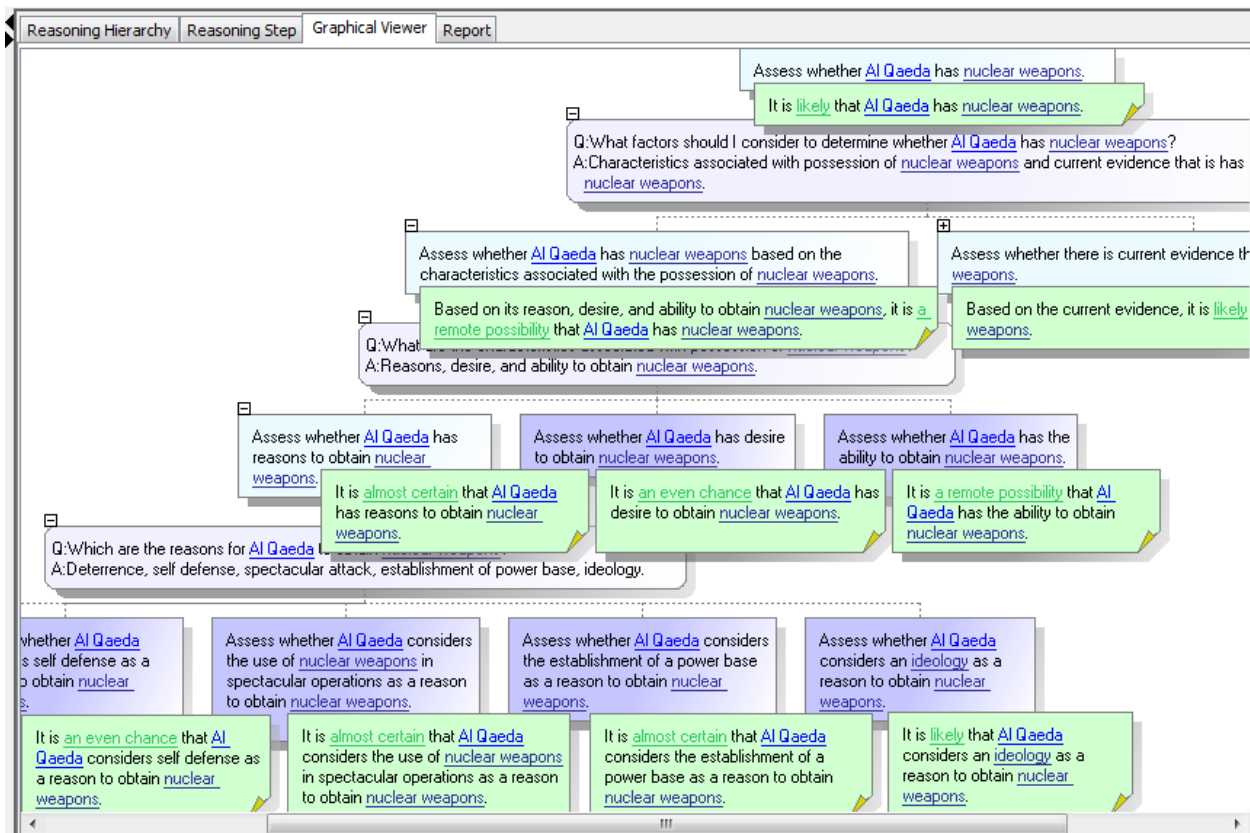


Figure 1: Hypothesis analysis through problem reduction and solution synthesis.

Let us consider the leaves of the tree from Figure 1. One of them represents the simpler hypothesis analysis problem:

Assess whether **Al Qaeda** considers the use of **nuclear weapons** in spectacular operations as a reason to obtain **nuclear weapons**.

The solution of this problem appears in a green background:

It is **almost certain** that **Al Qaeda** considers the use of **nuclear weapons** in spectacular operations as a reason to obtain **nuclear weapons**.

This and the other solutions of the problems from the bottom of Figure 1 are successively combined, from bottom-up, to synthesize the solutions of the upper-level problems, until the solution of the top level problem is obtained:

It is **likely** that **Al Qaeda** has **nuclear weapons**.

Disciple-LTA uses five symbolic probabilities (such as **likely** or **almost certain**) to express the solutions of the problems. They correspond to the National Intelligence Council's standard estimative language. However, this could easily be changed to consider more or fewer symbolic probabilities and to associate specific probability intervals with each of them (Kent 1994; Heuer, 1999). Notice also that some words appear in blue. They correspond to entities (such as **Al Qaeda** or **nuclear weapons**) that are represented in the knowledge base of Disciple-LTA (Tecuci et al, 2007).

Figure 1 shows only the top part of the problem reduction and solution synthesis tree. The problems from the bottom of Figure 1 are themselves reduced to simpler problems in order to solve them. Overall, the entire reasoning tree has over 1,700 nodes. To help browse and understand such a complex analysis, Disciple-LTA displays a simplified version of it, as illustrated in the left-hand side of Figure 2.

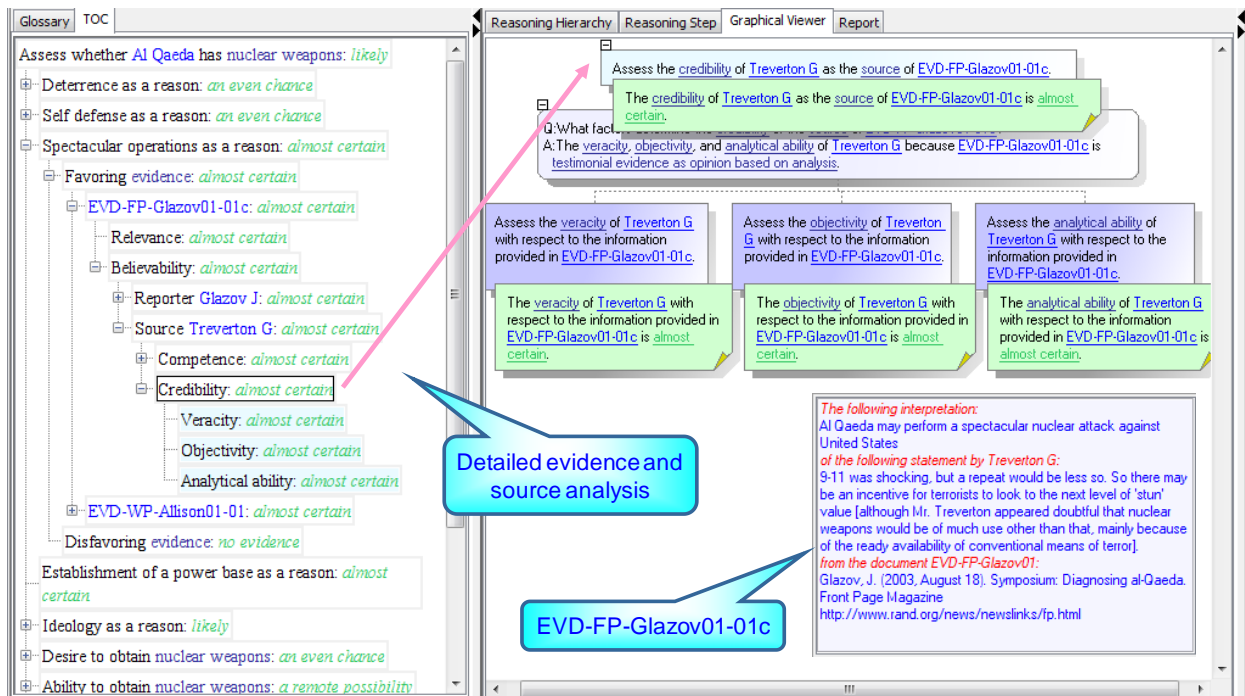


Figure 2: Evidence and source analysis.

Under the initial problem and its solution (Assess whether [Al Qaeda](#) has [nuclear weapons](#): [likely](#)) Disciple-LTA lists the short names of the bottom problems from Figure 1 and their solutions:

Deterrence as reason: [an even chance](#)

Self defense as reason: [an even chance](#)

Spectacular operations as reasons: [almost certain](#)

...

Each of these simpler problems is solved by using the same reduction and synthesis approach. Let us consider “Spectacular operation as reason” which is a short name for “Assess whether [Al Qaeda](#) considers the use of [nuclear weapons](#) in spectacular operations as a reason to obtain [nuclear weapons](#).” As indicated in the left-hand side of Figure 2, to solve this hypothesis analysis problem Disciple-LTA considered both favoring evidence and disfavoring evidence. Disciple-LTA has found two pieces of favoring evidence, EVD-FP-Glazov01-01c and EVD-WP-Allison01-01, and it has analyzed to what extent each of them favors the hypothesis that Al Qaeda considers the use of nuclear weapons in spectacular operations as a reason to obtain nuclear weapons. EVD-FP-Glazov01-01c is shown in the bottom right of Figure 2. It is a fragment from a magazine article published in the Front Page Magazine by Glazov J. where he cites Treverton G. who stated that Al Qaeda may perform a spectacular nuclear attack against United States (Glazov, 2003). To analyze EVD-FP-Glazov01-01c, Disciple-LTA considered both its relevance and its believability (Schum, 2001; 2007). The believability of EVD-FP-Glazov01-01c depends both on the believability of Glazov J. (the reporter of this piece of information) and the believability of Treverton G. (the source). The believability of the source depends on his competence and his credibility. The credibility of Treverton G. depends on his veracity, objectivity, and analytical ability.

When the user of Disciple-LTA clicks on a problem, such as “Credibility” from the left-hand side of Figure 2, Disciple-LTA displays the details on how it solved that problem, as shown in the right-hand side of Figure 2. For example, to “Assess the [credibility](#) of [Treverton G](#) as the [source](#) of [EVD-FP-Glazov01-01c](#)” Disciple-LTA has assessed his veracity, objectivity, and analytical ability. Then the results of these assessments ([almost certain](#), [almost certain](#) and [almost certain](#)) have been combined into an assessment of the credibility ([almost certain](#)). Disciple-LTA may use different synthesis functions for the solutions (such as, minimum, maximum, average, etc.), depending on the types of the problems. A simplified representation of the synthesis process is displayed in the left hand side of Figure 2, where the solutions appear in green, attached to the corresponding problems.

### 3. Analytic assistance and assumption-based analysis

When a human analyst uses Disciple-LTA, she or he selects the hypothesis analysis problem to solve and Disciple-LTA “instantly” generates the analysis, as illustrated in Figures 1 and 2. After that the analyst may browse, modify and/or extend the analysis. For example, the analyst can select any problem from the reasoning tree and provide a solution in the form of an assumption. This capability is illustrated in Figure 3. In this case the analyst selected the problem:

Assess whether [Al Qaeda](#) had reasons not to use [nuclear weapons](#), assuming that it has them.

Then the analyst has made the assumption that the solution of this problem is (see the bottom right-hand side of Figure 3):

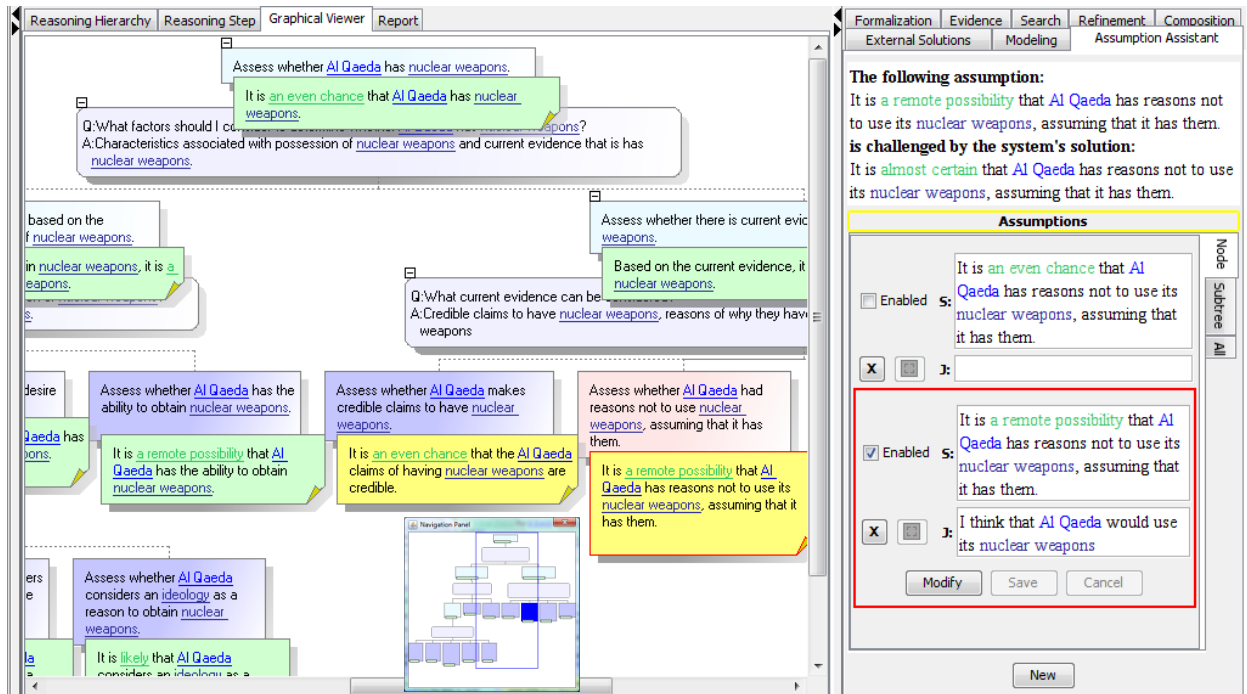


Figure 3: Assumptions-based analysis.

It is a remote possibility that Al Qaeda has reasons not to use its nuclear weapons, assuming that it has them.

The analyst has also provided the following justification for this solution:

I think that Al Qaeda would use its nuclear weapons.

This assumption has immediately been used to update the analysis and change the result to

It is an even chance that Al Qaeda has nuclear weapons.

The assumptions made by the human analyst appear with a yellow background, to be distinguished from the other solutions. In addition, the above assumption also has a red contour because it is challenged by Disciple-LTA, as indicated in the upper right-part of Figure 3.

The analyst may experiment with different assumptions and determine their influence on the global result of the analysis. For instance, the right hand side of Figure 3 shows two different assumptions for the same problem. Each of them could be enabled, by simply clicking on the corresponding button, and the analysis is instantly updated. The analyst can also view all the assumptions corresponding to the entire reasoning tree, or a part of it, and can enable, disable, delete and define new assumptions. Thus, the assumptions allow the analyst to hypothesize a solution for a problem that cannot be solved by Disciple-LTA, to change a solution generated by the system, or to experiment with different what-if scenarios.

As shown above, the analysis tree generated by Disciple-LTA makes very clear the analysis logic, what evidence was used and how, what assumptions have been made, and what is not known. This illustrates an ability of a decision-support system to generate a solution to a complex problem in a very transparent way. The generated analysis is very natural for a human analyst, it is very precise and very detailed, and it is generated very fast. This allows the human

analyst to critically evaluate the reasoning process of Disciple-LTA, accept parts of it, modify other parts, and produce an analysis which she or he would consider to be her or his own.

#### 4. Teaching Disciple-LTA

An important capability of Disciple-LTA is that it can rapidly capture the analytic expertise directly from an expert analyst, who can teach it in a way that is similar to how the analyst would teach another person. Disciple-LTA represents this analytic expertise in a knowledge base consisting of an ontology and a set of reasoning rules (Tecuci et al., 2007). The ontology is a description of important concepts used in analysis (such as “piece of evidence”, “credibility”, “source”). The rules are general if-then structures that express the conditions under which a general problem can be reduced to simpler ones or partial solutions combined into more complete ones. These rules could be quite complex, but they are learned by Disciple-LTA, rather than being manually defined by a knowledge engineer, as illustrated in the following.

A sub-problem of “Assess whether *Al Qaeda* has nuclear weapons” is “Assess whether there are states with nuclear weapons willing to sell nuclear weapons to *Al Qaeda*.”

Because Disciple-LTA does not know how to solve this sub-problem, its solution has to be provided by the human analyst. The analyst may provide the solution in the form of an assumption, as discussed in the previous section. Alternatively, the human analyst may perform an analysis to obtain the solution, and teach Disciple-LTA at the same time, as illustrated in Figure 4.

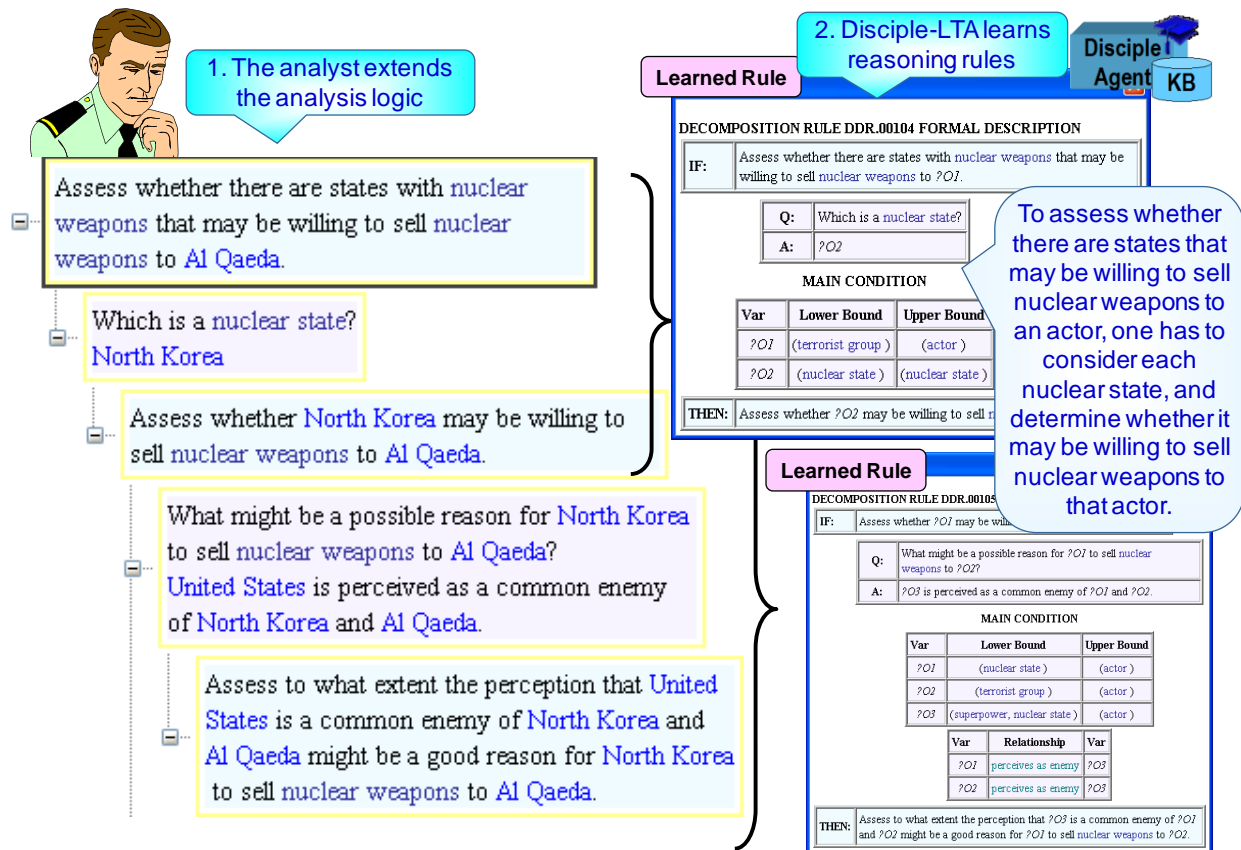


Figure 4: Reasoning rules learned from the analyst’s solution.

First the analyst considers a specific nuclear state (North Korea, in this illustration) which leads him or her to “Assess whether North Korea is willing to sell nuclear weapons to Al Qaeda.” Then the analyst considers each possible reason for North Korea to sell nuclear weapons to Al Qaeda (e.g. “United States is perceived as a common enemy of North Korea and Al Qaeda”, or “Desperate need of revenues because the financial reserves of North Korea are at a very low level”) and analyses each of them in terms of both positive consequences and negative consequences, in order to obtain a global evaluation of whether it makes sense for North Korea to sell nuclear weapons to Al Qaeda.

From each reasoning step, consisting of a problem, question, answer and sub-problem(s), Disciple-LTA learns a general reasoning rule, as illustrated in the right-hand side of Figure 4. For instance, from the first reasoning step Disciple-LTA has learned the following rule: “To assess whether there are states that may be willing to sell nuclear weapons to an actor, one has to consider each nuclear state, and determine whether it may be willing to sell nuclear weapons to that actor.” Disciple-LTA has immediately applied the learned rule in problem solving and has considered the other nuclear states, in order to assess whether any of them might be willing to sell nuclear weapons to Al Qaeda, as illustrated in Figure 5. However, the analyst has pointed to the United States, indicating that this reasoning step is incorrect because the United States perceives Al Qaeda as an enemy. As a result, Disciple-LTA has refined its rule as follows: “To assess whether there are states that may be willing to sell nuclear weapons to an actor, one has to consider each nuclear state which is not an enemy of that actor, and determine whether that state may be willing to sell nuclear weapons to that actor.” With this new rule, Disciple-LTA will no longer consider, not only the United States, but also United Kingdom, France, and any other nuclear state which perceives Al Qaeda as an enemy.

1. Disciple-LTA applies the learned rule

2. The analyst critiques the reasoning

3. Disciple-LTA refines the rule with an except-when condition

Disciple Agent KB

Assess whether there are states with nuclear weapons willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?  
North Korea

Assess whether North Korea is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?  
United States

Assess whether United States is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?  
United Kingdom

Assess whether United Kingdom is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?  
Russia

Assess whether Russia is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?  
Pakistan

Assess whether Pakistan is willing to sell nuclear weapons to Al Qaeda.

Which is a nuclear state?

This is wrong!  
United States will not sell nuclear weapons to Al Qaeda because it perceives it as an enemy.

To assess whether there are states that may be willing to sell nuclear weapons to an actor, one has to consider each nuclear state which is not an enemy of that actor, and determine whether that state may be willing to sell nuclear weapons to that actor.

Rule Viewer

Refined Rule LE DDR.00205 FORMAL

IF: Assess whether there are states with nuclear weapons willing to sell nuclear weapons to ?O1.

Q: Which is a nuclear state?  
A: ?O2

MAIN CONDITION

Var	Lower Bound	Upper Bound
?O1	(terrorist group)	(actor)
?O2	(nuclear state)	(nuclear state)

EXCEPT WHEN CONDITION 1

Var	Lower Bound	Upper Bound
?O2	(nuclear state)	(actor)
?O1	(terrorist group)	(actor)

Var	Relationship	Var
?O2	perceives as enemy	?O1

THEN: Assess whether ?O2 is willing to sell nuclear weapons to ?O1.

Figure 5: Rule refined based on the analyst’s critique.

In this way Disciple-LTA learns complex reasoning rules through a natural and simple interaction with the expert analyst, capturing his/her analytic expertise.

## 5. Teaching new analysts

Once Disciple-LTA has been taught how to systematically analyze complex hypotheses, it can be used to tutor new analysts. Figure 6 shows a fragment the first part of a lesson on how to assess the support provided by a piece of evidence to a hypothesis. First Disciple-LTA presents the reduction strategy at an abstract level, as illustrated in the top part of Figure 6. Then it generates concrete examples of the abstract strategy, as illustrated in the middle part of Figure 6.

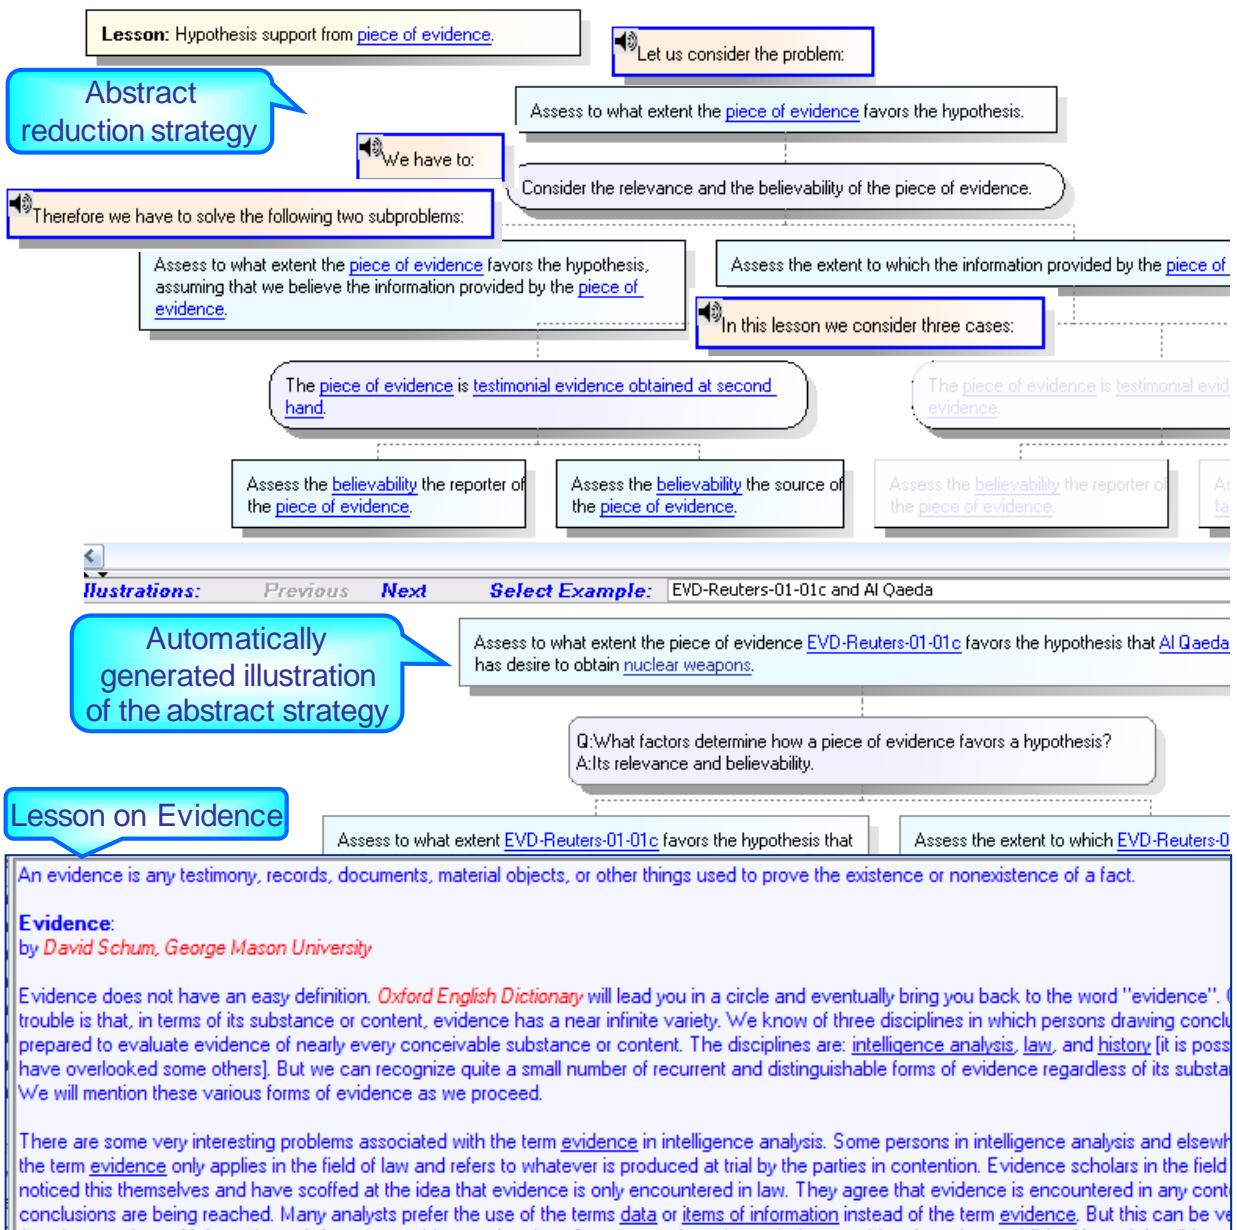


Figure 6: Lesson fragment – Hypothesis support from a piece of evidence.



By clicking on various concepts (such as “evidence” or “testimonial evidence obtained at second hand”) the user will receive tutorials on those concepts, consisting of definitions and intelligence analysis stories (Schum, 2007), as illustrated at the bottom part of Figure 6.

The lesson will continue with the presentation of the synthesis strategy, first at an abstract level, and then illustrated with examples generated by Disciple-LTA from its knowledge base.

Disciple-LTA was also taught to generate test questions for the student analyst. One type of test question consists in showing the analyst an analysis fragment which the analyst has to judge as being correct, incorrect, or incomplete. Another type of test question will give the analyst an analysis problem and a list of potential sub-problems, asking the analyst to select the correct sub-problems. Disciple-LTA could be asked to provide a hint, and will also provide feedback on the answer received from the student.

## **6. Other capabilities of Disciple-LTA**

In addition to its capabilities for learning, analytic assistance and tutoring, the current version of Disciple-LTA has several capabilities for evidence search, collaboration, and report generation, as briefly described in the following.

### ***Evidence search***

Disciple-LTA distinguishes between complex hypotheses that are reduced to simpler hypotheses (such as “Assess whether [Al Qaeda](#) has [nuclear weapons](#)”) and basic hypotheses which are solved through evidence analysis (such as “Assess whether [Al Qaeda](#) considers the use of [nuclear weapons](#) in spectacular operations as a reason to obtain [nuclear weapons](#)”). With each basic hypothesis, Disciple-LTA and the analyst associate search criteria for identifying favorable and unfavorable pieces of evidence. Then Disciple-LTA can invoke external search engines (such as Yahoo or Google) to retrieve potentially relevant pieces of evidence which can be selected by the analyst and included into the analysis.

### ***Collaborative problem solving***

Disciple-LTA is very conducive to collaboration. It solves problems by reducing them to simpler ones which could be solved by different analysts. Once the solutions of these sub-problems are obtained, they are combined by Disciple-LTA into the solution of the initial problem.

### ***Report generation***

Currently, Disciple-LTA has a simple capability for generating a report in Microsoft Word that represents the analysis. This report, the structure of which is defined by the expert analyst, can then be exported into a wiki format.

## **7. Final remarks**

Joan McIntyre has guided this research and development effort.

David Schum has contributed to the incorporation of his theory of credibility analysis into Disciple-LTA and Marcel Barbulescu has contributed to the development of the ontology tools of Disciple-LTA.

The current knowledge base of Disciple-LTA has been developed with the help of Cindy Ayers from the US Army War College, who also assisted in the experimental use of Disciple-LTA in

the “Military Applications of Artificial Intelligence: Intelligence Analysis” course, taught at the US Army War College in Spring 2005, 2006 and 2007.

Disciple-LTA has a stand-alone architecture which requires its installation on the user’s machine. However, work is in progress to develop a web-based architecture of Disciple, where a Disciple client will run in a web browser and will communicate with a Disciple server.

## References

- Glazov J., Symposium: Diagnosing Al Qaeda, *FrontPageMagazine.com*, 18 August 2003.
- Heuer, R.J. Jr., *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington, D.C., 1999.
- Kent S., "Words of Estimated Probability," in Donald P. Steury, ed., *Sherman Kent and the Board of National Estimates: Collected Essays* (CIA, Center for the Study of Intelligence, 1994).
- Schum D.A. 2001. *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.
- Schum D.A., Lessons and Stories about Concepts Encountered in Disciple-LTA, *Research Report 2*, Learning Agents Center, 21 August 2007.
- Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London, England: Academic Press.
- Tecuci G.; Boicu M.; Ayers C.; Cammons D. 2005. 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., Marcu D., Boicu C., Barbulescu M., Ayers C., Cammons D., Cognitive Assistants for Analysts, *Journal of Intelligence Community Research and Development*, 2007