DISCIPLE-LTA: ANALYST'S COGNITIVE ASSISTANT

Gheorghe Tecuci, Mihai Boicu, Cindy Ayers, Dorin Marcu, Vu Le, Cristina Boicu, Marcel Barbulescu

Disciple-LTA is a personal cognitive assistant to an intelligence analyst that addresses several systemic problems associated with the current analytic process related to effective analytic assistance, capturing of analytic expertise, and tutoring new analysts. Disciple-LTA can rapidly acquire and maintain analytic expertise which currently takes years to establish, is lost when analysts separate from service, and is costly to replace. It can help an analyst to rapidly find solutions to complex analytical problems, share intelligence, collaborate with complementary experts and their agents, and report results of research. It can also help new intelligence analysts learn the reasoning processes involved in making intelligence judgments and solving intelligence analysis problems. While presenting an overview of all these capabilities of Disciple-LTA, this paper concentrates on the analytic assistance provided to an end-user.

GEORGE MASON UNIVERSITY

RESEARCH REPORT 1

VOLUME 2007

Learning Agents Center MASON

DISCIPLE-LTA: ANALYST'S COGNITIVE ASSISTANT

Gheorghe Tecuci[1, 2], Mihai Boicu[1], Cindy Ayers[2], Dorin Marcu[1], Vu Le[1], Cristina Boicu[1], Marcel Barbulescu[1]

[1] *Learning Agents Center, George Mason University, Fairfax, VA, 22030, USA*

[2] *Center for Strategic Leadership, US Army War College, Carlisle, PA, 17013, USA*

tecuci@gmu.edu, http://lac.gmu.edu

Keywords: structured argumentation, hypothesis analysis, assumption-based reasoning and what-if scenarios, capturing analytic expertise, tutoring new analysts.

***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 the U.S. Army War College, the Department of the Army, the Department of Defense, or any other 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 product 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.*

10 March, 2007

1. INTRODUCTION

The Disciple-LTA personal cognitive assistant has the potential to alleviate several systemic problems associated with the current analytic process related to effective analytic assistance, capturing of analytic expertise, and tutoring new analysts (Lowenthal, 1999; National Commission on Terrorist Attacks Upon the United States, 2004; Wheaton, 2001). As illustrated in Figure 1, it is envisioned that an analyst will be assisted by a personal Disciple-LTA cognitive assistant. Once the analyst specifies the intelligence problem or task to solve (e.g. "Assess whether location-A is a training base for terrorist operations.") Disciple-LTA will rapidly develop an initial analysis which can be further refined by its human counterpart. The analyst, through his/her Disciple-LTA assistant, may collaborate with other experts and their agents to solve problems requiring expertise in multiple areas, or develop alternative solutions using different points of view (i.e. biases, assumptions, prior knowledge), to reduce the effect of a narrowly focused analytic mindset. Disciple-LTA agents will thus provide a more rapid, more rigorous, and more comprehensive solution to complex analytic problems by facilitating the process of evaluating and sharing intelligence data while acting in collaboration with other analysts and their agents.

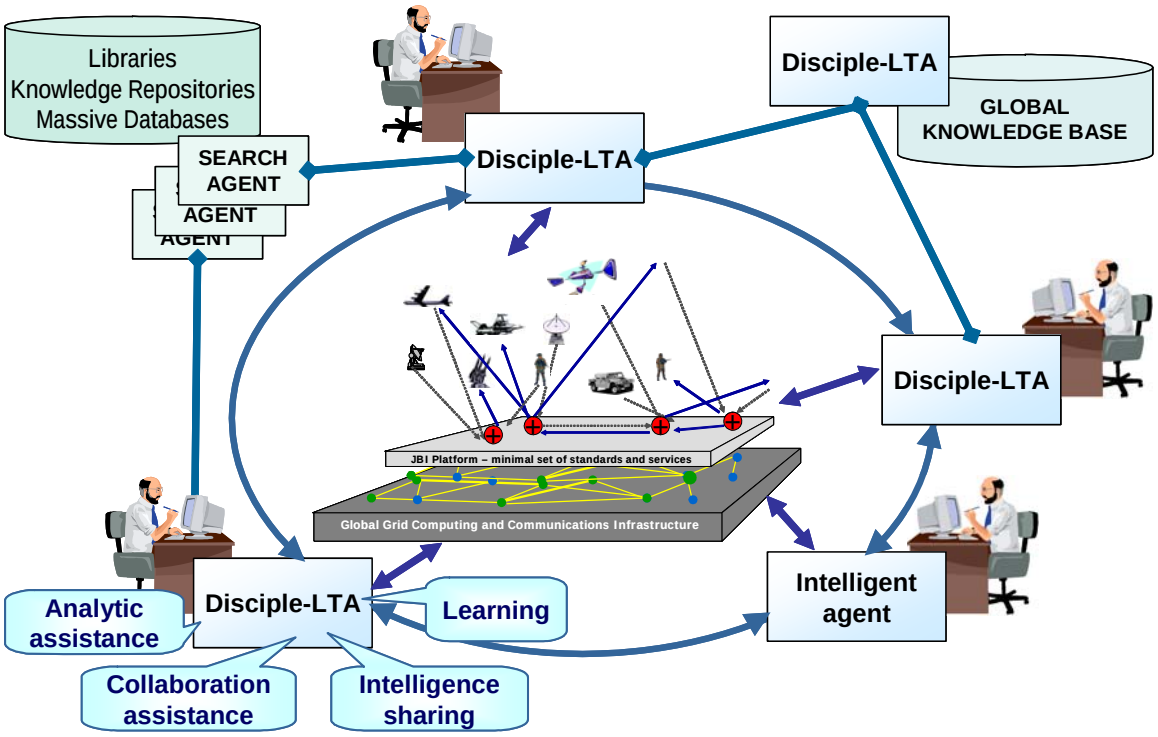

Figure 1: Use of Disciple-LTA in an operational environment.

Disciple-LTA can rapidly acquire and maintain analytic expertise which currently takes years to establish, is lost when analysts separate from service, and is costly to replace. The Disciple agent is taught by an expert analyst in a way that is similar to how the analyst would teach any student--by showing it examples of how to analyze specific hypotheses, by helping it to understand the analysis process, and by supervising and correcting the agent's analysis of new hypotheses. As a result of this process, Disciple-LTA will acquire the analyst's expertise in hypothesis analysis, including his/her prior and tacit knowledge, biases and assumptions. This will allow the Disciple agent to emulate the reasoning processes of an expert analyst and provide efficient analytic assistance to its end-users. Disciple-LTA will also function as a repository of the analyst's knowledge-base to provide continuity and a historical archive for departing analysts, thus codifying individual and personal analytical expertise, otherwise lost forever to the intelligence community.

The Disciple agents can also teach new analysts in a way that is similar to how Disciple-LTA was itself taught by an expert analyst. The roles are now reversed, with the agent being the expert and the human the learner. Disciple-LTA will, as a teacher, consider typical hypothesis analysis tasks, explaining to the student how to solve them in a systematic way. As a consequence, Disciple-LTA can be used as a very effective classroom tool, helping new intelligence analysts learn the critical thinking and reasoning processes involved in making intelligence judgments and solving intelligence problems.

Disciple-LTA is a unique and complex analytic tool that integrates powerful capabilities for analytic assistance, learning and tutoring. At the basis of these capabilities is the use of a systematic approach to hypothesis analysis which is both natural for a human analyst and appropriate for an automatic agent. This approach is presented in the following section. After that we will present in more detail each of the three main capabilities of Disciple-LTA.

2. HYPOTHESES ANALYSIS THROUGH TASK-REDUCTION AND SOLUTION-SYNTHESIS

One of the main results of this research is the development of a systematic approach to hypothesis analysis which is based on the general task-reduction/solution-synthesis paradigm of problem solving (Durham, 2000; Lowrance et al., 2001; Powel and Schmidt, 1988; Tecuci, 1998). Our approach is illustrated by the reasoning tree in Figure 2. Such a tree is generated by Disciple-LTA and is intended to be a natural and explicit representation of an expert analyst's thread of logic--as if the analyst would be thinking aloud, as discussed in the following example.

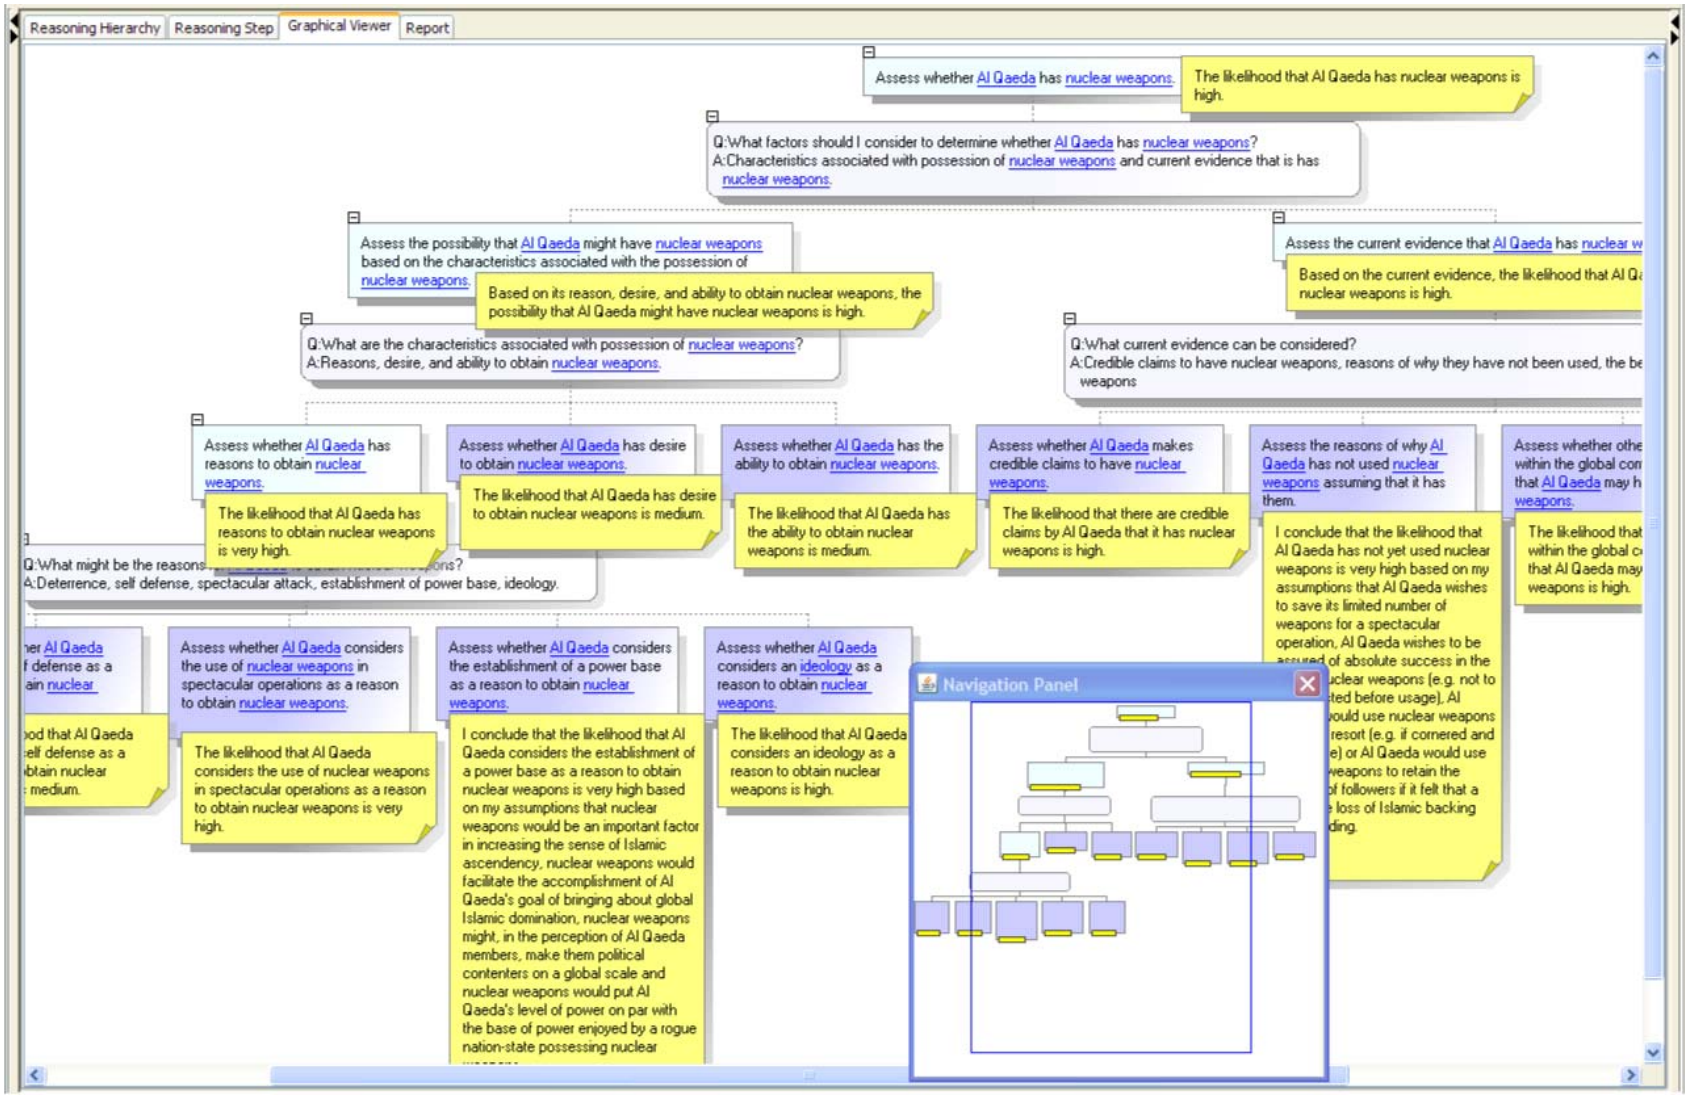

Figure 2: Hypothesis analysis through task-reduction and solution-synthesis.

I need to: Assess whether Al Qaeda has nuclear weapons. [1]

In order to perform this assessment task, Disciple-LTA will ask itself a series of questions. The answer to each question will lead to the reduction of the current assessment task to simpler tasks. For instance, the first question asked (followed by its answer) is:

Q: What factors should I consider to determine whether Al Qaeda has nuclear weapons?
A: Characteristics associated with possession of nuclear weapons and current evidence that it has nuclear weapons.

This leads to the reduction of task [1] to two simpler tasks:

Assess the possibility that Al Qaeda might have nuclear weapons based on the characteristics associated with the possession of nuclear weapons. [2]
Assess the current evidence that Al Qaeda has nuclear weapons. [3]

Each of these tasks is further reduced, guided by corresponding questions and answers (e.g. [2]):

What are the characteristics associated with possession of nuclear weapons?
Reasons, desire, and ability to obtain nuclear weapons.

Therefore I need to:

Assess whether Al Qaeda has reasons to obtain nuclear weapons.
Assess whether Al Qaeda has desire to obtain nuclear weapons.
Assess whether Al Qaeda has the ability to obtain nuclear weapons.

The purpose of these successive task-reduction steps is to reduce the initial intelligence analysis task to a set of simpler tasks that either have known solutions or can be solved through evidence analysis. An example of a task that should be solved through evidence analysis is:

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

The system identifies pieces of evidence that are relevant to the task, and analyzes them through task-reduction and solution-synthesis (as discussed in the next section), to obtain the task's solution:

The likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is medium.

When all the solutions of the subtasks of a given task have been obtained, they are combined into the solution of the task. Thus, the solutions of the leaf tasks from Figure 2 are successively combined, from the bottom-up, to obtain the solution for the initial task [1]:

The likelihood that Al Qaeda has nuclear weapons is high.

3. ANALYSIS OF A PIECE OF EVIDENCE

We are also developing a systematic approach to evidence analysis which distinguishes between different types of evidence and defines systematic 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).

As indicated in the previous section, a task that should be solved through evidence analysis is:

Assess whether Al Qaeda considers self defense as a reason to obtain nuclear weapons. [4]

This task is first reduced by Disciple-LTA to two simpler tasks:

Assess the evidence that supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons. [5]

Assess the evidence that rejects the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons. [6]

Then, potentially relevant pieces of evidence for each task are identified. Let us consider the reduction of task [5]:

I need to: [7]

Assess the evidence that supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.

Q: Is there any potentially relevant piece of evidence?

A: EVD-Dawn-Mir01-01c which mentions that Al Qaeda would use nuclear weapons to defend itself.

Therefore I need to

Assess to what extent the piece of evidence EVD-Dawn-Mir01-01c supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons. [8]

EVD-Dawn-Mir01-01c represents the statement:

Al Qaeda has nuclear weapons and may use them to defend itself.

This statement is the analyst's interpretation of the following paragraph attributed to Osama bin Laden in an interview conducted by Hamid Mir and published in the Dawn newspaper:

"We [Al Qaeda] have chemical and nuclear weapons [...] 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).

Task [8] is further reduced to two simpler tasks:

Assess to what extent EVD-Dawn-Mir01-01c supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons, assuming that we believe the information provided by EVD-Dawn-Mir01-01c. [9]

Assess the extent to which the information provided by EVD-Dawn-Mir01-01c can be trusted. [10]

Task [9] can be immediately solved by Disciple-LTA:

Q: Taking into account only the information provided by EVD-Dawn-Mir01-01c what is the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons?

A: very high because EVD-Dawn-Mir01-01c states that Al Qaeda would use nuclear weapons to defend itself.

Therefore I conclude:

Based only on the information from EVD-Dawn-Mir01-01c, the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is very high. [11]

To solve task [10], Disciple-LTA has to consider how the piece of evidence EVD-Dawn-Mir01-01c was obtained. It was obtained as testimonial evidence of Osama bin Laden cited in the EVD-Dawn-Mir01 article of Hamid Mir. Considering Schum's types of evidence (Schum, 2001), EVD-Dawn-Mir01-01c is second hand testimonial evidence. Therefore, one has to solve the following tasks:

Assess the extent to which one can trust Hamid Mir as the reporter of EVD-Dawn-Mir01-01c. [12]

Assess the extent to which one can trust Osama bin Laden as the source of EVD-Dawn-Mir01-01c. [13]

The second of these tasks, task [13], is further reduced by Disciple-LTA to two simpler tasks:

Assess the competency of Osama bin Laden as the source of EVD-Dawn-Mir01-01c. [14]

Assess the credibility of Osama bin Laden as the source of EVD-Dawn-Mir01-01c. [15]

Task [15] is reduced to three simpler tasks, each assessing Osama bin Laden's veracity, objectivity, and observational sensitivity, respectively, as indicated in Figure 3. 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 (i.e. veracity, objectivity, and observational sensitivity) may themselves be complex tasks that can be solved through task-reduction and solution-synthesis (Schum and Morris, 2007).

After the components of Osama bin Laden's credibility are evaluated, they need to be combined into an overall estimation of his credibility, which is the solution of task [15], as illustrated in Figure 3:

The credibility of Osama bin Laden as the source of EVD-Dawn-Mir01-01c is medium. [16]

Then Osama bin Laden's credibility [16] has to be combined with his competence (the solution of task [14]), to produce an estimate of how much we could trust him as the source of EVD-Dawn-Mir01-01c. This represents the solution of task [13]:

One can trust to a medium degree Osama bin Laden as the source of EVD-Dawn-Mir01-01c. [17]

Disciple-LTA determines in a similar manner how much one could trust Hamid Mir as the reporter of this information (i.e. the solution of task [13]), and then combines this with the trust in Osama bin Laden ([17]), to estimate how much one could trust the piece of evidence:

One can trust to a medium degree the information provided by EVD-Dawn-Mir01-01c. [18]

Solution [18] (which is the solution of task [10]) is combined with solution [11] of task [9] to obtain the following solution of task [8]:

Based on EVD-Dawn-Mir01-01c, the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is medium. [19]

In general, however, there will be several pieces of evidence that are relevant to a given task. For

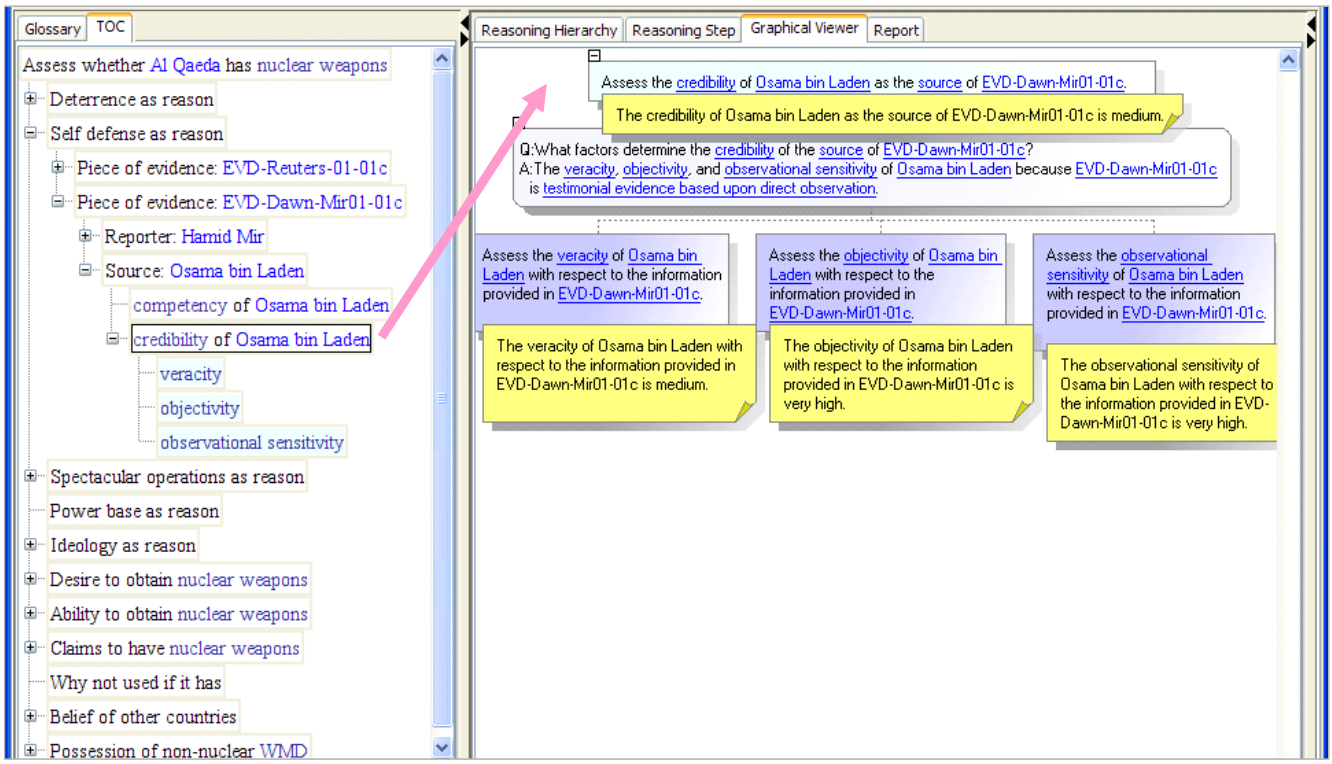

Figure 3 Abstract analysis tree and concrete analysis of source credibility.

instance, EVD-Reuters-01-01c (*Al Qaeda has nuclear weapons that would be used inside the United States or to defend against a US nuclear or chemical attack)* is another piece of evidence that is relevant to task [5]. Therefore, another reduction of task [5] is the following one:

I need to: [20]

Assess the evidence that supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons.

Q: Is there any potentially relevant piece of evidence?

A: EVD-Reuters-01-01c which mentions that Al Qaeda would use nuclear weapons to defend itself.

Therefore I need to

Assess to what extent the piece of evidence EVD-Reuters-01-01c supports the hypothesis that Al Qaeda considers self defense as a reason to obtain nuclear weapons. [21]

Following the reasoning discussed above for task [8], Disciple obtains a solution of task [21]:

Based on EVD-Reuters-01-01c, the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is low. [22]

The solutions [19] and [22] are then combined into the following solution of task [5]:

Based on the supportive evidence, the likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is medium. [23]

Solution [23] is combined with the solution of task [6] to obtain the solution of task [4]:

The likelihood that Al Qaeda considers self defense as a reason to obtain nuclear weapons is medium. [24]

This is indeed a long and complex process, but it is instantaneously performed by Disciple-LTA, as discussed in the next section.

4. ANALYTIC ASSISTANCE

4.1 RAPID HYPOTHESIS ANALYSIS

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. The analyst specifies the intelligence task to solve (e.g. "Assess whether Al Qaeda has nuclear weapons" or "Assess whether Hezbollah has nuclear weapons") and Disciple-LTA immediately generates a systematic and comprehensive analysis (the Al Qaeda task, for example, has over 1700 nodes). In order to facilitate its understanding and navigation by the analyst, Disciple-LTA also generates an abstraction of the reasoning tree that shows simplified descriptions of the main tasks, as indicated in the left-hand side of Figures 3 and 5. This abstract tree can be regarded as a table of contents for the navigation of the detailed tree. Its top node represents the initial task. Under

it are the main abstract subtasks of the initial task. Thus, to determine whether Al Qaeda has nuclear weapons, one has to assess whether it uses deterrence as a reason to acquire them, as well as self-defense, spectacular operation, power base and ideology, whether Al Qaeda has desire to obtain nuclear weapons, ability to obtain nuclear weapons, whether there are credible claims that it has nuclear weapons, why it has not used them assuming that it has the weapons, what other countries believe, and whether it possesses non-nuclear WMD. Each of these abstract tasks corresponds to one of the leaf tasks in Figure 2. Thus the tree in Figure 2 shows the reduction of the top level task into its main subtasks.

One can browse the abstract tree by expanding and collapsing the abstract tasks. The left hand side of Figure 3 shows the main abstract subtasks of the task "Self-defense as reason." Thus, there are two relevant pieces of evidence: EVD-Reuters-01-01c and EVD-Dawn-Mir-01-01c. The analysis of EVD-Dawn-Mir-01-01c involves the analysis of the reporter Hamid Mir, and the analysis of the primary source Osama bin Laden. The analysis of Osama bin Laden involves the analysis of its competency and its credibility. When the analyst clicks on one of these abstract tasks (e.g. "credibility of Osama bin Laden"), the right-hand side of the solution viewer shows the detailed reasoning for that task, as illustrated in the right hand side of Figure 3.

4.2 ASSUMPTION-BASED REASONING AND WHAT-IF SCENARIOS

The analyst can select any task from the reasoning tree and provide an assumption for that task as well as a justification of the assumption. For instance, the analyst can select the task (see Figure 2)

> Assess whether Al Qaeda makes credible claims to have nuclear weapons.

and make the assumption that its solution is

> The likelihood that there are credible claims by Al Qaeda that it has nuclear weapons is very high.

In this case Disciple-LTA automatically updates its reasoning tree and generates the resulting solution:

> Based on its reason, desire, and ability to obtain nuclear weapons, the possibility that Al Qaeda might have nuclear weapons is high. However, according to the current evidence, the likelihood that it actually has nuclear weapons is very high.

Using the assumptions capability of Disciple-LTA, the analyst can investigate various what-if scenarios by considering different solutions to certain tasks and determining how they change the final result of the analysis.

We are developing a capability that will allow Disciple-LTA to constructively challenge an assumption made by the analyst when it has enough information to infer a different result.

4.3 MIXED-INITIATIVE EVIDENCE RETRIEVAL AND REPRESENTATION

A capability that is currently under development is to associate various search criteria with the tasks that need to be solved though evidence analysis. Consider, for instance, reduction [7], discussed in section 3. The answer from this reduction indicates why EVD-Dawn-Mir01-01c is a relevant piece of evidence (Al Qaeda would use nuclear weapons to defend itself). This allows Disciple-LTA to suggest criteria for searching pieces of evidence that are potentially relevant to certain analysis tasks, such as [5] or [6]. Then specialized search engines can be invoked to retrieve relevant pieces of evidence to be used in the analysis. In particular, the above search criteria identifies EVD-Reuters-01-01c as another piece of evidence that is relevant to task [5].

We are currently developing a capability for mixed-initiative representation of the retrieved pieces of evidence, allowing Disciple-LTA to automatically analyze them.

Disciple-LTA will be able to automatically share the represented pieces of evidence and their analyses with other analysts and their cognitive assistants.

4.4 REPORT GENERATION

A report generation capability is under development. Once implemented, Disciple-LTA will be able to automatically generate analytic reports based on a completed (or partially completed) hypothesis analysis tree. Reports of differing levels of abstraction can be created, showing logic, specifying the assumptions made, citing sources of data used, and indicating the manner in which the analysis was performed. The analyst can then finalize the report using Microsoft Word or an XML editor.

4.5 COLLABORATION

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 2. This task is reduced to several simpler assessment tasks and therefore, to perform it, one has to perform these simpler tasks. Each of these simpler assessments could, in principle, be completed by a different expert. Then their solutions would be combined into the solution of the initial task.

Disciple-LTA will enable multiple analysts to work collaboratively. It will allow tracking of the contribution of each of the participants, the development of alternative analyses in solving the subtasks, and the publication of reports to facilitate side bar discussions to debate differences.

5. CAPTURING OF ANALYTIC EXPERTISE

5.1 RAPID ACQUISITION OF ANALYTIC EXPERTISE

Disciple-LTA can rapidly acquire, use, and preserve the prior and tacit knowledge of an expert analyst who can directly teach it how to analyze hypotheses in a natural way--similar to the manner in which the analyst would teach a person. The main feature of Disciple-LTA is its ability to rapidly learn analysis rules 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 (Buchanan and Wilkins, 1993).

Let us consider the task from the top of Figure 4 (Assess whether there are states with nuclear weapons that may be willing to sell nuclear weapons to Al Qaeda). Because Disciple-LTA does not know how to perform this task, the analyst has to provide the analysis from Figure 4. However, from each task reduction step (consisting of a task, a question, an answer and one or several subtasks), Disciple-LTA will learn a general task reduction rule which will help it to solve similar tasks in a similar way. Thus, the rules learned from the North Korea case in Figure 4 will allow Disciple-LTA to infer whether other states (such as Russia or Pakistan) may be willing to sell nuclear weapons to Al Qaeda, as illustrated in Figure 5.

During interactive problem solving with the analyst, Disciple-LTA will learn not only from the contributions of the analyst, but also from its own successful or unsuccessful problem solving attempts, which will lead to the refinement of the learned rules.

5.2 LEARNING 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).

6. TUTORING NEW ANALYSTS

The main concept behind the tutoring function of Disciple-LTA is to take advantage of the expertise of a senior analyst—the expert analyst responsible for teaching a Disciple-LTA initially—by utilizing "trained" Disciple agents as tutors for new analysts. These agents would teach in a way that is similar to how Disciple-LTA was itself taught. Thus the roles would be reversed, with the agent being the expert—remembering that the human expert's knowledge could be used exponentially in the agent-based instruction of many human students.

6.1 LESSON DESIGN AND GENERATION

An instructor can use Disciple-LTA to define a sequence of lessons for teaching an analyst how to systematically analyze hypotheses by using the task-reduction and solution synthesis approach described in sections 2 and 3. The instructor has various options to design the lessons but, in principle, each lesson will introduce general strategies on how to solve certain tasks, such as, how to assess the credibility of a human source or how to assess the credibility of tangible evidence. The lesson will also

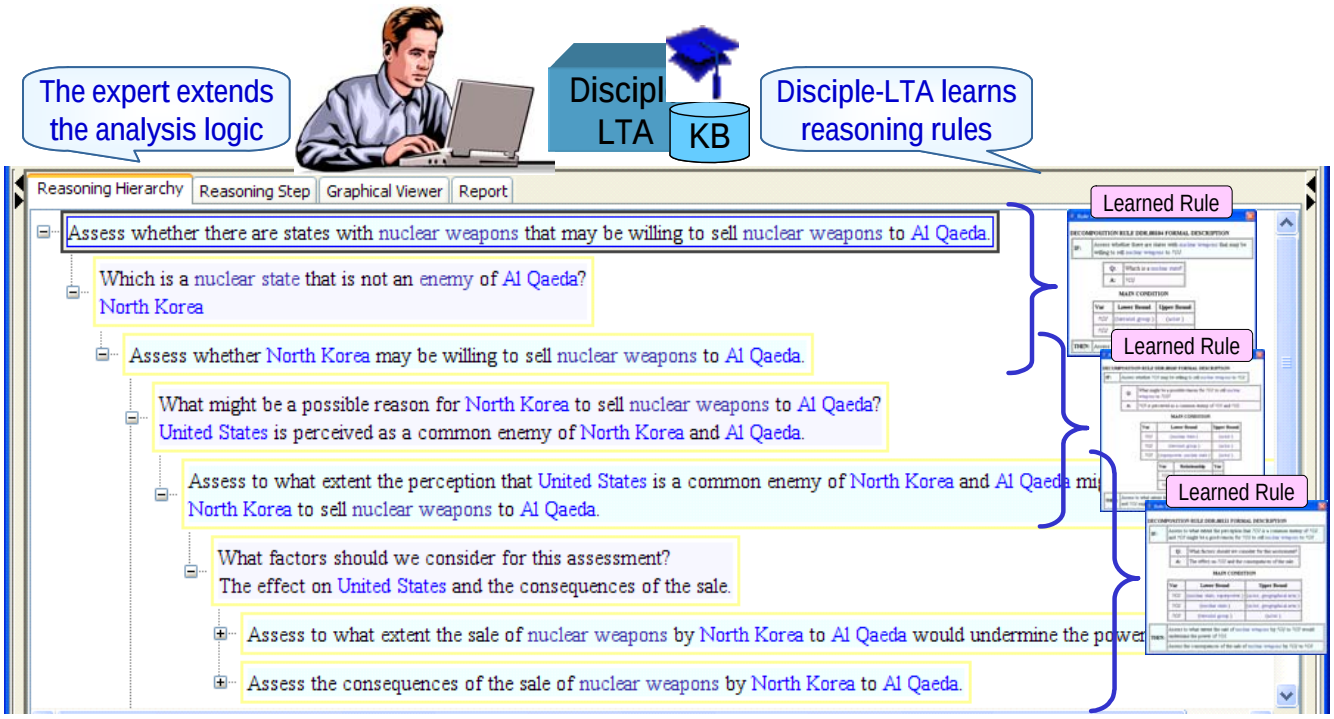

Figure 4: Disciple-LTA teaching by the expert analyst.

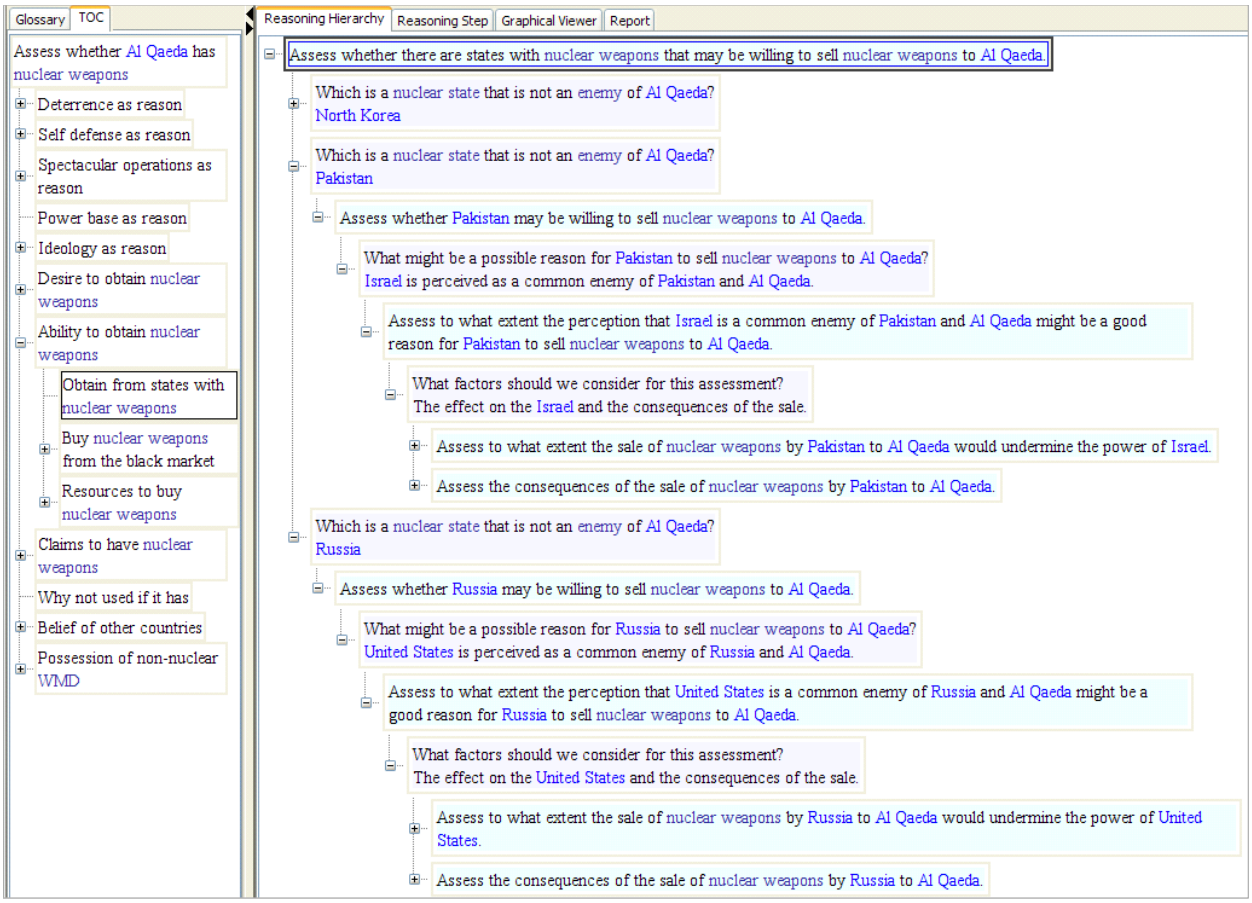

Figure 5: Application of the learned knowledge by Disciple-LTA.

illustrate these strategies with specific examples and will define the various concepts used.

Then a new analyst can interact with Disciple-LTA and learn from it instead of learning from the expert analyst who has taught Disciple. The new analyst also has the option of selecting an application domain that he/she is familiar with because Disciple has the ability to automatically generate lesson examples from that domain, without requiring any change in the lesson structure.

6.2 AUTOMATIC TEST GENERATION

An expert instructor can teach Disciple-LTA how to generate test questions for new analysts, based on specific examples of correct, incomplete, or incorrect reasoning, together with hints and student feedback. Then the agent can automatically generate test questions for students by instantiating the test generation rules in the context of a specific domain knowledge base, and by taking into account the student's model (e.g. the lessons learned). This allows the students to test themselves and to be assessed by Disciple-LTA.

7. EXPERIMENTATION ENVIRONMENT

We are developing Disciple-LTA by 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. 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. Initial experimentation is being conducted in the Military Applications of Artificial Intelligence (MAAI) course, taught every spring at the U.S. Army War College. The first two evaluations took place in the spring of 2005 and 2006, and the third will take place in the spring of 2007. The MAAI elective is attended by senior officers who are experts in military intelligence or have interest in artificial intelligence and intelligence analysis, but generally lack knowledge engineering experience. The elective provides an overview of Intelligence Analysis and Artificial Intelligence, and helps students to link Intelligence Analysis concepts with Artificial Intelligence concepts through the use of Disciple agents. In particular, the students use Disciple-LTA in learning how to systematically analyze hypotheses in an unclassified environment using open-source intelligence, as illustrated in section 2 and 3. The students also teach their agents, incorporating their own knowledge and expertise, while experiencing Disciple-LTA as a cognitive assistant for analysts. U.S. Army War College students assessed Disciple-LTA as being very useful to the analysts, once it has a rich enough knowledge base.

8. CONCLUSIONS

Disciple-LTA is a new type of cognitive assistant developed to help intelligence analysts solve complex problems faster and better—an assistant that learns and employs an analyst's preferred problem solving strategies, biases and assumptions, and can subsequently teach new analysts how to think systematically and critically.

With sufficient knowledge base development, Disciple agents could revolutionize the manner in which intelligence analysis is learned and performed operationally throughout the Intelligence Community. They could provide the proverbial—and greatly needed—"quantum leap" from a field that is bound by human limitations to one that contains infinite possibilities through human-agent interaction.

REFERENCES

Buchanan BG, Wilkins DC, eds. 1993. *Readings in Knowledge Acquisition and Learning: Automating the Construction and Improvement of Expert Systems*. San Mateo, CA: Morgan Kaufmann; 1993.

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.

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.

Lowenthal, M. 1999. *Intelligence: From Secrets to Policy.* Washington, DC: Congressional Quarterly Press.

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,* US 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.

Schum D.A. and Morris J. 2007. Personal Communication.

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.

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.

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