In: Phillips-Wren, G., Jain, L.C., Nakamatsu, K., Howlett, R.J. (eds.)
Advances in Intelligent Decision Technologies, SIST 4, pp. 1-10.
Springer-Verlag Berlin Heidelberg, 2010.

Intelligence Analysis as Agent-Assisted Discovery of Evidence, Hypotheses and Arguments

Gheorghe Tecuci[1], David Schum[1], Mihai Boicu[1], Dorin Marcu[1], Benjamin Hamilton[2]

[1] Learning Agents Center, George Mason University, Fairfax, VA, USA
{tecuci, dschum, mboicu, dmarcu}@gmu.edu

[2] Department of Defense, USA

Abstract. This paper presents a computational approach to intelligence analysis which is viewed as mixed-initiative discovery of evidence, hypotheses and arguments by an intelligence analyst and a cognitive assistant. The approach is illustrated with the analysis of wide area motion imagery of fixed geographic locations where the goal is to discover threat events such as an ambush or a rocket launch. This example is used to show how the Disciple cognitive assistants developed in the Learning Agents Center can help the analysts in coping with the astonishing complexity of intelligence analysis.

Keywords: intelligence analysis, science of evidence, wide-area motion imagery, discovery, cognitive assistants, learning, evidence-based reasoning, mixed-initiative reasoning

1 Introduction

Problem-solving and decision-making depends critically on accurate intelligence that needs to be discovered in an overwhelming amount of mostly irrelevant, always incomplete, usually inconclusive, frequently ambiguous, and commonly dissonant information with various degrees of believability about a highly complex and dynamic world. This is an astonishingly complex process where each analytic task is unique and always requires mixtures of imaginative and critical reasoning. Indeed, hypotheses about situations of interest must be generated by imaginative thought and then subjected to critical evidence-based analysis.

We are researching a computational theory of intelligence analysis which forms the basis for the development of cognitive assistants that help intelligence analysts in coping with this complexity. Part of this theory is a view of intelligence analysis as mixed-initiative discovery of evidence, hypotheses and arguments by intelligence analysts (who are capable of imaginative reasoning, have broad subject matter expertise, and have access to evidence from a wide variety of sources) and

their cognitive assistants (that are capable of critical reasoning and have both domain-specific knowledge and general knowledge from the Science of Evidence).
In the next section we present a sample intelligence analysis problem (analysis of wide-area motion imagery) that will be used to illustrate the developed approach. After that, the following five sections present the processes of discovery of hypotheses, evidence and arguments. Then, section 8 concludes the paper with a discussion on how the Disciple cognitive assistants developed in the Learning Agents Center, which are capable of analytic assistance, learning, and tutoring, can help in coping with the astonishing complexity of intelligence analysis [14, 15].

2 Sample Problem: Analysis of Wide-Area Motion Imagery

Capabilities exit today to persistently monitor fixed geographic locations (such as conflict areas) as wide as 100 km^2, for long periods of time, using electro-optic sensors (see Fig. 1). This leads to the collection of huge amounts of data to be used either in real-time analysis or in forensic analysis. During real-time analysis, analysts attempt to discover impeding threat events (e.g., ambush, kidnapping, rocket launch, false check-point, suicide bomber, IED) in time to react. During forensic analysis, the analysts backtrack from such an event (e.g., an ambush) in order to discover the participants, possible related locations and events, and the specific movement patterns [2]. The problem however is that the manual analysis of these huge amounts of data would require thousands of analysts.
We will use this sample analysis problem to illustrate our approach.

Fig. 1. Wide area motion imagery (WAMI).

3 Discovery of Hypotheses

Let us consider an analyst who, while reviewing wide area motion imagery (WAMI) of a region of Iraq, notices evidence of road work at 1:17am, an unusual time for such an activity. The question is: *What possible threat does this evidence suggest?* Through a flash of insight, the analyst may abductively leap to the hypothesis H_k that there is an ambush threat at that location [6]. Attempting to justify the hypothesis, the analyst may generate the following abductive inference steps shown also in the left hand side of Fig. 2 (as we know, abductive inference indicates that something is possibly true):

E^*_i: *There is evidence of road work at 1:17am at location L_1.*

→ E_i: *It is possible that there is indeed road work at location L_1.*

→ H_a: *It is possible that the road work is for blocking the road*

→ H_c: *It is possible that there is an ambush preparation at location L_1.*

→ H_k: *It is possible that there is an ambush threat at location L_1.*

So here we have evidence in search of hypotheses where a newly discovered item of evidence searches for hypotheses that would explain it.

4 Discovery of Evidence

A great challenge in any intelligence analysis task is the massive amount of data that needs to be searched quickly, especially during the real-time use of the system. The diagram in the middle of Fig. 2 illustrates the deductive process involved in putting the generated hypothesis at work to guide the search for new relevant evidence in the WAMI data. The question is: *Assuming that the threat is real, what other events or entities should be observable?* The deductive reasoning process for answering this question successively reduces the assessment of the top-level hypothesis H_k to the assessment of simpler hypotheses, ultimately resulting in precise queries to be answered from the WAMI data (as we know, deductive inference indicates that something is necessarily true):

Let us assume H_k, *that there is an ambush threat at location L_1 after 1:17am.*

→ H_b: *L_1 should be an ambush location.*

H_c: *There should be ambush preparation at L_1 around 1:17am.*

H_q: *There should be ambush execution at L_1 (if forensic analysis).*

If this is real-time analysis occurring soon after 1:17am, then the ambush has not yet been executed and the third sub-hypothesis (H_q) will not be considered. However, if this is forensic analysis, then H_q should also be considered.

Let us now assume H_b, *that L_1 is indeed an ambush location.*

→ H_d: *L_1 should be on a route of the blue forces after 1:17am.*

H_e: *There should be cover at location L_1.*

This guides the analyst to search for the following evidence:

- *Search for evidence that L_1 is on a planned blue route after 1:17am.*

- *Search for evidence in the WAMI data that there is cover at location L_1.*

A similar analysis of the hypothesis H_c (*There is an ambush preparation at L_1 around 1:17am*) leads to the following queries for specific events and entities in the WAMI data and from other sources (shown as shaded circles in Fig. 2):

- *Search for evidence in the WAMI data that there is departure of vehicle V_1 from facility F_1 before 1:17am.*
- *Search for evidence that F_1 is a suspected terrorist facility.*
- *Search for evidence in the WAMI data that there is arrival of vehicle V_1 at location L_1 short before 1:17am.*
- *Search for evidence in the WAMI data that personnel P_1 descends from vehicle V_1 at location L_1 short before 1:17am.*

Notice that these are precise queries that can be answered very fast. Being based on evidence, the answers will be probabilistic, such as:

It is almost certain that there is arrival of vehicle V_1 at location L_1 at 1:09am.

It is very likely that personnel P_1 descends from vehicle V_1 at L_1 at 1:10am.

These probabilistic solutions and other discovered evidence will be used to assess the likelihood of the top level hypothesis H_k, as discussed in Section 5.

The above has illustrated the deductive process of <u>hypotheses in search of evidence</u> that leads to the discovery of new evidence that may favor or disfavor them. Some of the newly discovered items of evidence may trigger new hypotheses or the refinement of the current hypothesis. Therefore, as indicated at the bottom of Fig. 2, the processes of evidence in search of hypotheses and hypotheses in search of evidence take place at the same time, and in response to one another.

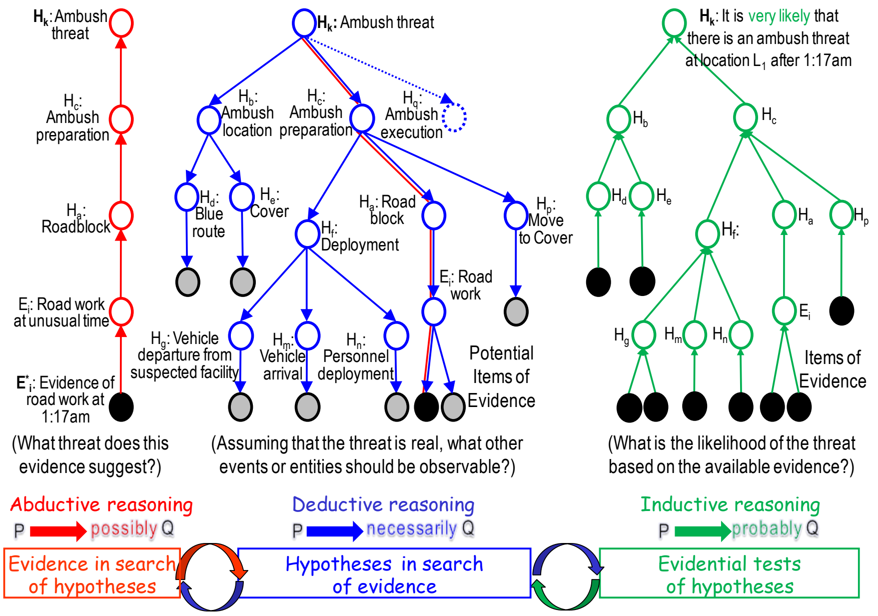

Fig. 2. Discovery of evidence, hypotheses and arguments.

5 Discovery of Arguments

The discovered evidence (shown as black circles at the right hand side of Fig. 2) can now be used to discover an argument that assesses, through inductive inference, the likelihood of the hypothesis H_k (e.g.,*"It is very likely that there is an ambush threat at location L_1 after 1:17am"*). As we know, inductive inference indicates that something is probably true.

Fig. 3 shows a Wigmorean probabilistic inference network that combines the deductive reasoning tree and the inductive reasoning tree from Fig. 2. This network has a well-defined structure, which has a grounding in the problem reduction representations developed in Artificial Intelligence [4, 10], and in the argument construction methods provided by the noted jurist John H. Wigmore [17], the philosopher of science Stephen Toulmin [16], and the evidence professor David Schum [5, 7]. This approach uses expert knowledge and evidence to successively reduce a complex hypothesis analysis problem to simpler and simpler problems, to find the solutions of the simplest problems, and to compose these solutions, from bottom-up, to obtain the solution of the initial problem. The Wigmorean network shows how evidence is linked to hypotheses through arguments that establish the relevance, believability and inferential force or weight of evidence [5, 9].

As shown in Fig. 3, the assessment of hypothesis H_k is reduced to the assessment of two simpler hypotheses: H_b and H_c. Then H_b is reduced to H_d and H_e. Each of these two hypotheses is assessed by considering both favoring evidence and disfavoring evidence. Let us assume that there are two items of favoring evidence for H_d: E_{d1} and E_{d2}. For each of them one would need to assess the extent to which it favors the hypothesis H_d. This requires assessing the relevance, believability, and inferential force or weight of evidence.

Relevance answers the question: *So what? How does this item of information bear on what the analyst is trying to prove or disprove?*

Believability (or credibility) answers the question: *Can we believe what this item of intelligence information is telling us?*

Inferential force or weight answers the question: *How strong is this item of relevant evidence in favoring or disfavoring various alternative hypotheses or possible conclusions being entertained?*

Let us assume the following solutions for the relevance and the believability of E_{d1}: *"If we believe E_{d1} then H_d is almost certain"* and *"It is likely that E_{d1} is true."*

In this example, almost certain and likely are symbolic probabilities for likelihood similar to those from the DNI's standard estimative language, but other scales for uncertainty can easily be used [18].

The relevance of E_{d1} (almost certain) is combined with its believability (likely), for example through a "min" function, to determine E_{d1}'s inferential force or weight on H_d: *"Based on E_{d1} it is likely that H_d is true."*

Similarly one assesses the inferential force of E_{d2} on H_d: *"Based on E_{d2} it is almost certain that H_d is true."*

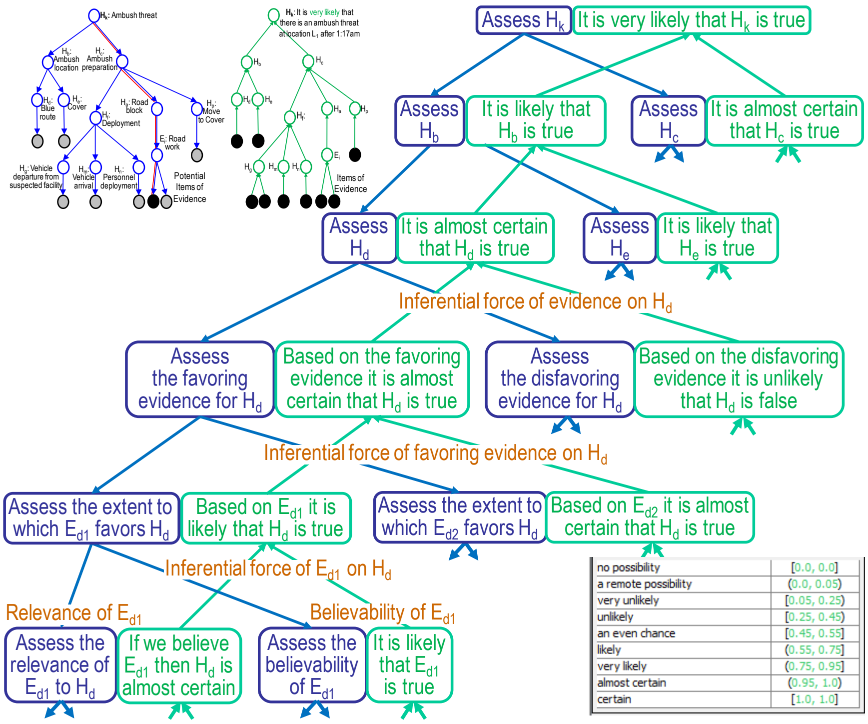

Fig. 3. Wigmorean probabilistic inference network for hypothesis assessment.

By composing the above solutions (e.g., through "max") one assesses the inferential force of the favoring evidence (i.e., E_{d1} and E_{d2}) on H_d: *"Based on the favoring evidence it is almost certain that H_d is true."*
Similarly one assesses the inferential force of the disfavoring evidence on H_d: *"Based on the disfavoring evidence it is unlikely that H_d is false."*
Now because there is very strong evidence favoring H_d and there is weak evidence disfavoring H_d, one concludes: *"It is almost certain that H_d is true."*
H_e is assessed in a similar way: *"It is likely that H_e is true."* Then the assessments of H_d and H_e are composed (through "min") into the assessment of H_b: *"It is likely that H_b is true."* Finally, this assessment is composed with the assessment of H_c (*"It is almost certain that H_c is true."*), through "average", to obtain the assessment of H_k (*"It is very likely that* H_k *is true."*)

6 Believability of Evidence

Above we have discussed the process of evidence-based hypothesis assessment down to the level where one has to assess the relevance and the believability of an item of evidence. In this section we will show how a Disciple agent helps in assessing the believability of evidence. This is based on its stock of established knowledge about evidence, its properties, uses, and discovery from the emerging Science of Evidence [1, 5, 7, 8], which is itself based upon 700 years of experience in the Anglo-American system of law. For example, the right-hand side of Fig. 4 shows a substance-blind classification of recurrent forms and combinations of evidence based, not on substance or content, but on the inferential properties of evidence [9].

This classification is important because each type of evidence has specific believability credentials, as well as a well-defined procedure for assessing its believability, as shown in the left hand side of Fig. 4.

In this classification, wide area motion imagery is demonstrative tangible evidence (i.e., a representation or image of a tangible thing), which has three believability attributes: authenticity, reliability, and accuracy.

Authenticity addresses the question: Is this object what it is represented as being or is claimed to be?

Reliability is especially relevant to various forms of sensors that provide us with many forms of demonstrative tangible evidence. A system, sensor, or test of any kind is reliable to the extent that the results it provides are repeatable or consistent. For example, a sensing device is reliable if it provides the same image or report on successive occasions on which this device is used.

Finally, the accuracy concerns the extent to which the device that produced the representation of the real tangible item had a degree of sensitivity (resolving power or accuracy) that allows us to tell what events were observed.

For testimonial evidence we have two basic sources of uncertainty: the competence and the credibility of the source (see bottom left-side of Fig. 4). Competence involves access and understandability. Credibility involves veracity (or truthfulness), objectivity, and observational sensitivity under the conditions of observation [9].

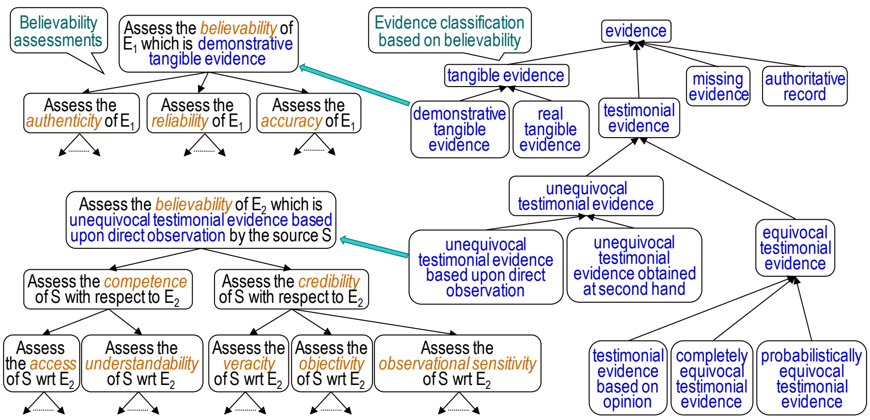

Fig. 4. Types of evidence and their believability assessments.

7 Analysis of Competing Hypotheses

Just because we have evidence of an event (e.g., E^*_i: evidence of road work at 1:17am) does not mean that the event actually occurred. Thus, as indicated in Fig. 5, we need to test two hypotheses: E_i (There is road work …) and Not E_i (There is no road work …). Similarly, for each abduced hypothesis (e.g., H_a: Roadblock), one would need to consider competing hypotheses (e.g., H_{a1}: Road repair). Moreover, for each such competing hypothesis one has to search for relevant evidence and use this evidence to test it, as discussed in the previous sections.

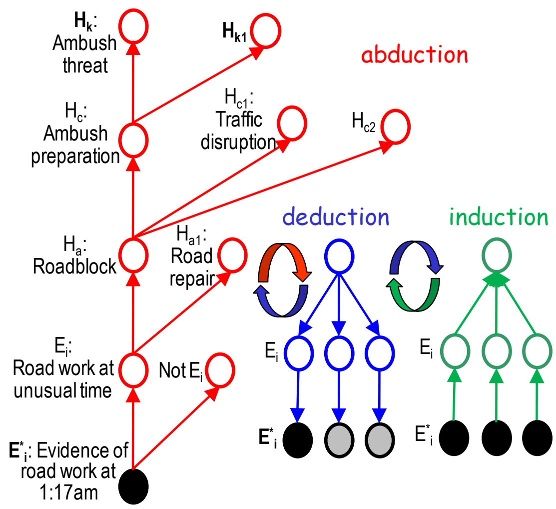

Fig. 5. Analysis of competing hypotheses.

8 Cognitive Assistants for Learning, Teaching, and Analysis

The researched computational theory of intelligence analysis is being implemented in Disciple cognitive assistants that synergistically integrate three complex capabilities. They can rapidly learn the analytic expertise which currently takes years to establish, is lost when analysts separate from service, and is costly to replace. They can tutor new intelligence analysts how to systematically analyze complex hypotheses. Finally, they can assist the analysts in analyzing complex hypotheses, collaborate, and share information [14, 15].

The problem solving engine of a Disciple assistant employs a general divide-and-conquer approach to problem solving, called problem-reduction/solution-synthesis, which was illustrated in Fig. 3. To exhibit this type of problem solving behavior, the knowledge base of the agent contains an ontology which describes both general concepts for evidence-based reasoning (see Fig. 4) and domain-

specific concepts from an application domain. The knowledge base also includes a set of learned problem reduction and solution synthesis rules which are represented with the concepts from the ontology. A problem reduction rule expresses how and under what conditions a generic problem can be reduced to simpler generic problems. Reduction rules are applied to automatically reduce assessment problems to simpler problems, as illustrated in Fig. 3. Similarly, a solution synthesis rule expresses how and under what conditions the solutions of generic sub-problems can be combined into the solution of a generic problem. These rules are applied to automatically perform compositions such as those from Fig. 3.
The cognitive assistant also includes a complex learning engine that uses multi-strategy methods (e.g., learning from examples, from explanation, and by analogy) to allow a subject matter expert to teach it in a way that is similar to how the expert would teach a person [10, 11, 14]. For instance, the expert will show the agent how to perform an analysis, as it was illustrated in Fig. 2, and will help it to understand each inference step. The agent, on the other hand, will attempt to learn a general reduction and synthesis rule from each such step and will extend its ontology. Moreover, the acquired knowledge will be pedagogically tuned [3], the agent solving new problems and explaining its reasoning similarly to how the expert taught it. This makes the agent an effective tool for teaching new intelligence analysts.
A trained Disciple cognitive assistant can help an analyst cope with the astonishing complexity of intelligence analysis through the use of mixed-initiative reasoning, a type of collaboration between humans and automated agents that mirror the flexible collaboration between people. It consists of an efficient, natural interleaving of contributions by the analyst and the agent that is determined by their relative knowledge and skills and the problem-solving context, rather than by fixed roles, enabling each of them to contribute what it does best, at the appropriate time [12, 13]. The analyst will act as the orchestrator of the reasoning process, guiding the high-level exploration, while the agent will implement this guidance by taking into account the analyst's preferred problem solving strategies, assumptions and biases. For example, the agent discovers evidence in the WAMI data of road work at location L_1, at 1:17am, an unusual time for such an activity, and alerts the analyst. As a result, the analyst directs the agent to analyze the hypothesis that there is an ambush threat and the agent develops the reasoning tree from the middle of Fig. 2, which will guide it to search for additional relevant evidence in the WAMI data and from other sources. The identified evidence is then used by the agent to evaluate the likelihood of the considered hypothesis, as was discussed in Section 5 and illustrated in Fig. 3. This reasoning tree makes very clear the analysis logic, what evidence was used and how, what assumptions have been made, and what is not known. This allows the analyst to critically evaluate the reasoning process, to accept parts of it, to modify other parts, and to produce an analysis which s/he would consider her/his own. The emphasis is on enhancing analyst's creativity, relying on the human to take the most critical decisions, and only to critique and correct the more routine ones that are proposed by the agent.

Acknowledgments We are very grateful to Phil Hwang, Don Kerr, Joan McIntyre, Kelcy Allwein, Keith Anthony, Cindy Ayers, Susan Durham, Sharon Hamilton, Jim Homer, David Luginbuhl, Bill Nolte, George Stemler, and Ben Wible for their suggestions and support. This research was performed in the Learning Agents Center and was partially supported by several agencies of the U.S. Government, including the National Geospatial-Intelligence Agency, the Department of Defense, and the National Science Foundation (0750461). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. 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.

References

1. Anderson T, Schum DA, Twining W: *Analysis of evidence*. Cambridge U. Press (2005)
2. Desai M: *Persistent stare exploitation and analysis system (PerSEAS)*, (2009) http://www.darpa.mil/ipto/solicit/baa/BAA-09-55_PIP.pdf. Accessed 18 April 2010
3. Le V, Tecuci G, Boicu M: Agent shell for the development of tutoring systems for expert problem solving knowledge. *Proc. 9th Int. Conference on Intelligent Tutoring Systems*. LNCS **5091** Springer-Verlag (2008) 228–238, Berlin Heidelberg / Montreal.
4. Nilsson NJ: *Problem solving methods in artificial intelligence*. McGraw-Hill (1971)
5. Schum DA: *Evidence and inference for the intelligence analyst*. University Press of America (1987) Lanham, MD
6. Schum DA: Species of abductive reasoning in fact investigation in law. *Cardozo Law Rev* **22(5-6)** (2001) 1645–1681
7. Schum DA: *The evidential foundations of probabilistic reasoning*. Northwestern University Press (2001)
8. Schum DA: Science of Evidence: Contributions from Law and Probability. *Law Probab Risk* **8** (2009) 197–231
9. Schum DA, Tecuci G, Boicu M, Marcu D: Substance-blind classification of evidence for intelligence analysis. *Proceedings of the Conference "Ontology for the Intelligence Community,"* George Mason University, Fairfax, Virginia, 20-22 October (2009)
10. Tecuci G: *Building intelligent agents: An apprenticeship multistrategy learning theory, methodology, tool and case studies*. Academic Press (1998)
11. Tecuci G, Boicu M, Boicu C, Marcu D, Stanescu B, Barbulescu M: The Disciple-RKF learning and reasoning agent. *Comput Intell* **21(4)** (2005) 462–479
12. Tecuci G, Boicu M, Cox MT: *AI Mag, Special issue on mixed-initiative assistants*. **28(2)** (2007)
13. Tecuci G, Boicu M, Cox MT: Seven aspects of mixed-initiative reasoning: An introduction to the special issue on mixed-initiative assistants. *AI Mag* **28(2)** (2007) 11–18
14. Tecuci G, Boicu M, Marcu D, Boicu C, Barbulescu M: Disciple-LTA: Learning, tutoring and analytic assistance. *J Intell Community Res Dev* (2008) http://lac.gmu.edu/publications/2008/Disciple-LTA08.pdf. Accessed 18 April 2010
15. Tecuci G, Schum DA, Boicu M, Marcu D: Coping with the complexity of intelligence analysis: cognitive assistants for evidence-based reasoning. *Research Report 7* (2010), Learning Agents Center, August 2009, updated April 2010. http://lac.gmu.edu/publications/2009/CompIA.pdf. Accessed 20 April 2010
16. Toulmin SE: *The uses of argument*. Cambridge University Press (1963)
17. Wigmore JH: *The science of judicial proof*. Little, Brown & Co. (1937) Boston, MA
18. Weiss C: Communicating uncertainty in intelligence and other professions. *Int J Intell Count Intell* **21(1)** (2008) 57–85