Computing in Science and Engineering, 20(6):9-21, November/December 2018

Evidence-based Reasoning in Intelligence Analysis: Structured Methodology and System

Gheorghe Tecuci
Louis Kaiser
Dorin Marcu
Chirag Uttamsingh
Mihai Boicu
Learning Agents Center, Volgenau School of Engineering, George Mason University

This paper presents a scientific method-based approach and system that helps intelligence analysts and others to reason better when addressing issues involving incomplete, contradictory, ambiguous, and missing information, by synergistically integrating the analyst's imagination and expertise with the computer's knowledge and critical reasoning.

Intelligence analysts face the astonishingly complex task of drawing defensible and persuasive conclusions from masses of imperfect information from a variety of different sources.[1-5]

This paper presents research aimed at assisting analysts in coping with the complexities of the evidential reasoning tasks they routinely face. It introduces a structured methodology based on the scientific method for the process of hypotheses generation, evidence collection, and hypotheses analysis.

It then presents Cogent (cognitive assistant for cogent analysis), an intelligent system that incorporates this structured methodology and facilitates a synergistic integration of analyst's imaginative reasoning and expertise with agent's knowledge and critical reasoning, to develop Wigmorean argumentations[2,4,6,7] for answering intelligence questions. This system is the latest in a sequence of increasingly more practical cognitive assistants for intelligence-analysis: Disciple-LTA[8-10], TIACRITIS[11], Disciple-CD[4], and a previous version of Cogent.[12]

The use of Cogent in answering an intelligence question is illustrated through an example.

ANALYSIS IN THE FRAMEWORK OF THE SCIENTIFIC METHOD

The analysis implemented in Cogent follows the scientific method of generating and testing hypotheses based on evidence, as illustrated in Figure 1. An intelligence analyst collaborates with Cogent to answer an intelligence question. During the first phase (question in search of answers) of this three-phase process, the analyst has to imagine possible answers to the question. Each answer is a hypothesis whose probability has to be assessed based on evidence. In the second phase (hypothesis in search of evidence), each hypothesis is put to work to guide the discovery of evidence that is relevant to that hypothesis. The hypothesis is decomposed into simpler and simpler hypotheses, with the simplest hypotheses suggesting the evidence that may be used to assess the hypotheses. During the third phase (evidentiary assessment of hypotheses), the probability of each hypothesis is assessed based on the discovered evidence.

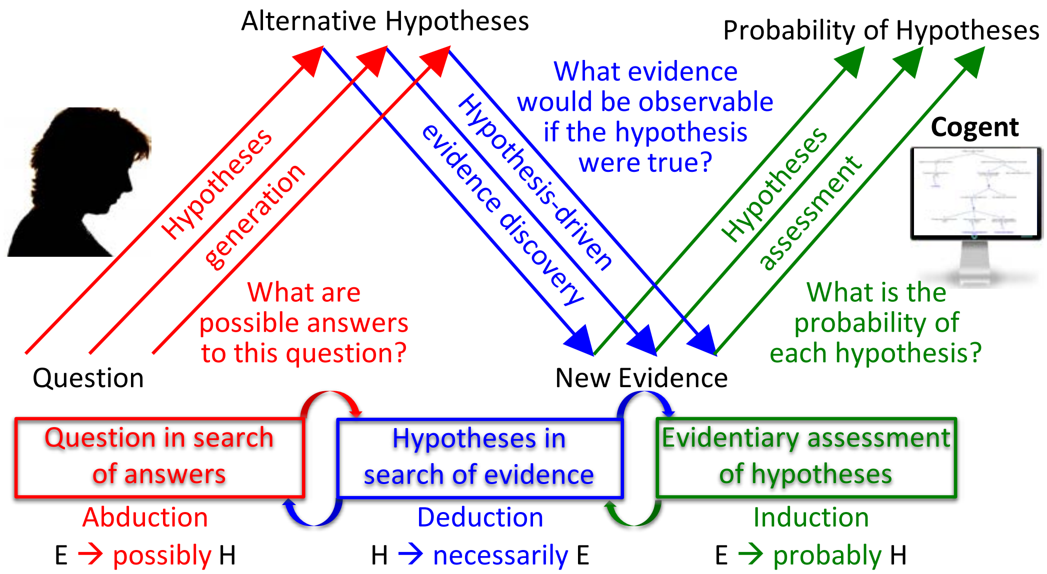

Figure 1. Intelligence analysis in the framework of the scientific method.

Notice, however, that this is not a linear process. The discovery of new evidence may lead to new answers to the question that, in turn, may lead to the discovery of new evidence and revised assessments of probability.

The process in Figure 1 integrates all types of reasoning: abductive, deductive, and inductive. Hypotheses generation involves *abductive* reasoning that shows that something is *possibly* true. The hypothesis-driven evidence discovery involves *deductive* reasoning which shows that something is *necessarily* true, and the hypotheses assessment involves *inductive* reasoning that shows that something is *probably* true.

In this paper we will focus on the third phase from Figure 1, evidentiary assessment of hypotheses.

BACONIAN PROBABILITIES WITH FUZZY QUALIFIERS

Hypothesis assessment is necessarily probabilistic in nature because our evidence is always *incomplete*, usually *inconclusive* (it is consistent with the truth of more than one hypothesis), frequently *ambiguous*, and commonly *dissonant* (some of it favors one hypothesis but other evidence favors other hypotheses), and has various degrees of *credibility*.[2,4]

Unfortunately, none of the probability views can cope with all the mentioned characteristics of evidence. For example, the conventional Subjective Bayesian view cannot cope with ambiguities or imprecision in evidence. On the other hand, the Fuzzy view, which uses symbolic probabilities, such as likely, can naturally cope with such imprecisions. But none of them can account for the incompleteness of the coverage of evidence. For this there is the Baconian view where the probability of a hypothesis depends on how complete the evidence is or how many questions recognized as being relevant remain unanswered by the evidence we do have. While on the Bayesian probability scale “0” means *disproof*, on the Baconian scale, “0” simply means *lack of proof*. A hypothesis now having “0” Baconian probability can be revised upward in probability as soon as we have some evidence for it. But we cannot revise upward in probability any hypothesis disproved, or having “0” conventional probability.[4]

A probabilistic system based on ideas from both the Baconian view and the Fuzzy view may potentially cope with all the five characteristics of evidence. Moreover, the use of similar min/max probability combination rules by the Baconian and the Fuzzy views facilitates the development of such an integrated system.[2,4,13-15] These rules are much simpler than the Bayesian probability combination rules, which is important for the understandability of the analysis.

There is also the issue of using a numerical probability scale, which is required by a Bayesian view, as opposed to a symbolic scale required by a Fuzzy view.[4] While a numerical probability is much more precise, it is not at all clear how an analyst would be able to defend a subjective assessment that, for instance, might assess the probability of a hypothesis H_k as exactly 77%. Analysts would arrive at different probability assessments, which would impede their collaboration. Because words are less precise than numbers, there will often be less disagreement about a verbal or fuzzy probability.

Starting from such considerations, we have defined an intuitive and easy to use system of *Baconian probabilities*[13,14] with *Fuzzy qualifiers*[15]. This system uses the following positive probability scale that is a refinement of the scale provided in the Intelligence Community Directive 203[16]: lacking support (0-50%) < barely likely (50-55%) < likely (55-70%) < more than likely (70-80%) < very likely (80-95%) < almost certain (95-99%) < certain (100%).

If the evidence does not support the truthfulness of the hypothesis **H** (i.e., **H** is lacking support), then it may support the truthfulness of its negation, **not H**. In such a case, the probability of **H** may be expressed using the following negative probability scale: no chance (0%) < almost no chance (1-5%) < very unlikely (5-20%) < more than unlikely (20-30%) < unlikely (30-45%) < barely unlikely (45-50%).

SAMPLE PROBLEM: SURFACE-TO-AIR MISSILE (SAM) SALE

To illustrate the evidence-based assessment of hypotheses with Cogent, we will use the problem from Table 1. The answers to the intelligence question “Which SAM system is Manada selling Sindia?” are constrained by the available information which suggest three possibilities or hypotheses: Destructor, Devastator, and Demolisher. We have to assess the probability of each of these hypotheses.

Table 1. Problem description: Surface-To-Air Missile Sale

Situation: The country Manada produces three different surface-to-air missile (SAM) systems (a combination of missiles and radars). At a press conference on 1 May, Manada announced that it had signed a contract to sell the country Sindia the Destructor SAM system. About one-half of the time Manada lies about which SAM system countries are buying (based on numerous observations over several years).

Question: Which SAM system is Manada selling Sindia?

Available Information:

According to an intercepted communication, on 1 January 2016, Manada provided specifications on its three SAM systems to the country Zupzicia, which was interested in buying SAMs to protect its major hydroelectric power plant. The Destructor SAM had a maximum range of 500 km and could destroy targets flying at altitudes between 500 to 10,000 meters. The Devastator SAM had a maximum range of 900 km and could destroy targets flying at altitudes between 200 and 12,000 meters. The Demolisher SAM had a maximum range of 850 km and could destroy targets flying at altitudes between 1,300 and 10,000 meters. The official from Manada told officials from Zupzicia that Manada was 90 percent confident in the reliability of these specifications.

In a training exercise in 2016 the Devastator engaged a target at a range of 880 km, according to data from a highly accurate technical collection radar.

According to a new source, who worked at the Kelman Institute for a short period of time in 2010, the original planned range of the Devastator was to be 650 km. Other reporting from this source has been corroborated. The Kelman Institute completed development of the Devastator SAM in 2012.

A reliable source at the Ferris Institute, the design and production facility for the Destructor SAM, reported in June 2016 that the Destructor SAM was being modified to extend its range to 700 km. Development of this longer version Destructor SAM was to be completed in November 2017. Redesign of this SAM began in April 2016.

According to intercepted communications, a delegation from Sindia confirmed on 30 May 2017 that it would visit the Tantrum SAM Test Facility on 24 June 2017 to observe a test of the SAM it was buying.

A new source at the Ferris Institute reported in early April 2017 that the budget of the Ferris Institute was not increased in January 2017. The new source has provided three other reports; information in the source's earlier reports has been corroborated. Until vetting of the source is complete, the source will not be considered a reliable source. The official managing this source has stipulated that at this point the credibility of the source should be considered in the range of 70-80 percent.

At 0800 on 24 June 2017, a missile-tracking radar station based in the country Javia, which borders on Manada, detected a SAM launching from Tantrum. The SAM flew about 680 km and destroyed a target aircraft flying at 1,000 meters. The station's radars had been recently calibrated. Only one SAM test was conducted at the Tantrum facility on 24 June.

ASSESING SIMPLE HYPOTHESES

One can directly assess a simple hypothesis based on an item of evidence by assessing the credentials of the evidence: credibility, relevance, and inferential force. This is illustrated in the left hand side of Figure 2.

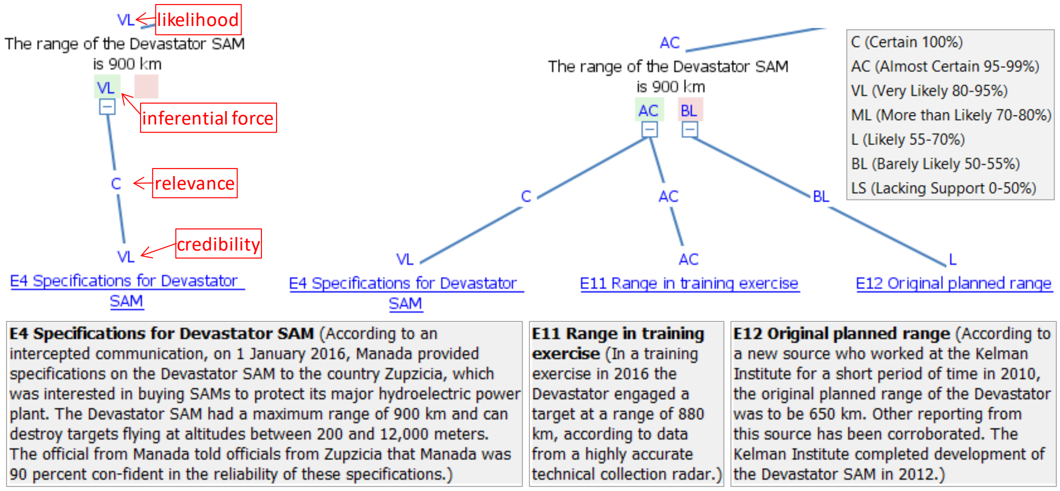

Figure 2. Hypothesis assessments based on the credentials of evidence.

First we assess the *credibility* of the item of evidence E4 by answering the question, *"What is the probability that this item of evidence is true?"* The credibility of E4 is VL (very likely, 80-95%) because the source indicates that the confidence level in the reported numbers is 90%.

Next we assess the *relevance* of E4 to the hypothesis by answering the question, *"Assuming that the evidence is true, what is the probability of the hypothesis?"* This relevance is C (certain, 100%) because if the evidence is true then the hypothesis above it is also true.

Finally, Cogent determines the *inferential force* of E4 on the hypothesis. The inferential force answers the question, *"What is the probability that the hypothesis is true, based only on this item of evidence?"* The relevance of E4 is certain, but its credibility is only very likely. Therefore the probability that "The range of the Devastator SAM is 900 km" is only very likely.

In general, the inferential force of an evidence item is determined as the smaller between its credibility and its relevance. Indeed, an evidence item that is not credible would not convince us that the hypothesis is true, no matter how relevant the evidence item is. Therefore the inferential force in this circumstance would be low. Similarly, it is not enough for the evidence item to be credible, if the information provided is not relevant to the hypothesis. The inferential force will be high only if the evidence item is both highly relevant and credible.

In this case, because we have considered only one item of evidence, the probability of the hypothesis "The range of the Devastator SAM is 900 km" is given by the inferential force of this evidence item. However, as shown in the right hand side of Figure 2, there are two items of evidence favoring this hypothesis (E4 and E11, shown under the left, green square) and one item of evidence disfavoring it (E12, shown under the right, red square).

The inferential force of E11 is AC (almost certain, 95-99%), and the combined inferential force of E4 and E11 is the maximum of the two inferential forces, that is AC. The inferential force of the disfavoring evidence item E12 is limited by its relevance, which was assessed as BL (barely likely, 50-55%) because the source was not working at the institute when development of the Devastator SAM was completed, raising the possibility that the planned range of this SAM was later increased.

The Baconian probability system requires considering either **H** or **not H** as probably true, but not both at the same time.[13,14] To assess the likelihood of a hypothesis that has both favoring and disfavoring evidence, Cogent uses an "on balance" function that balances the inferential force of the favoring evidence with that of the disfavoring evidence. If the inferential force of the disfavoring evidence is higher than or equal to that of the favoring evidence, then Cogent concludes that, based on all the available evidence, **H** is lacking support. If, however, the inferential force of the favoring evidence is greater than that of the disfavoring evidence, then the probability of **H** is determined by lowering the inferential force of the favoring evidence, based on the inferential force of the disfavoring evidence.

ASSESSING COMPLEX HYPOTHESES

The previous section presented a simple way of directly assessing a hypothesis based on evidence. This works well when it is obvious that the evidence favors or disfavors the hypothesis, as was illustrated in Figure 2. But it does not work well for complex hypotheses. Indeed, how confident would you be in assessing the relevance of the evidence from Figure 2 to the hypothesis "Manada is selling Devastator"?

And there is an additional challenge "How can you find evidence to assess such complex hypotheses?"

Answers to both these challenges are provided by Cogent, which guides the analyst in developing an argumentation structure that successively reduces a complex hypothesis to simpler and simpler hypotheses, down to the level of very simple hypotheses such as the one presented in Figure 2. Then the analyst has to search for evidence that is relevant to these very simple hypotheses, and to assess their credibility and relevance, as was illustrated in the previous section. Once this is done, Cogent automatically composes these assessments, from bottom up, based on the logic embedded in the argumentation, finally assessing the probability of the top-level hypothesis.

In general, a hypothesis is assessed by considering both favoring arguments and disfavoring arguments. There may be one argument of each type, several, or none. Each argument reduces the top-hypothesis to simpler hypotheses for which favoring and disfavoring arguments are again considered.

Let's consider the Demolisher hypothesis in the context of the information from Table 1. There is both favoring and disfavoring evidence for this hypothesis.

- The favoring evidence is that the range of the missile tested on 24 June (when Sindian officials were to observe a test of the SAM they were buying) is consistent with the range of the Demolisher.
- The disfavoring evidence is that the engagement-altitude capability of the tested SAM is not consistent with the capability of the Demolisher.

The favoring argument for the Demolisher hypothesis is shown in Figure 3. This is an AND argument because all of the sub-hypotheses below the "&" have to have some credibility to support the hypothesis above. The relevance of this argument is C (certain, 100%) because if all of the sub-hypotheses below the "&" are true, the hypothesis above is true. The inferential force of this argument is VL (very likely, 80-95%), the smallest among the probabilities of the sub-hypotheses and the relevance of the argument.

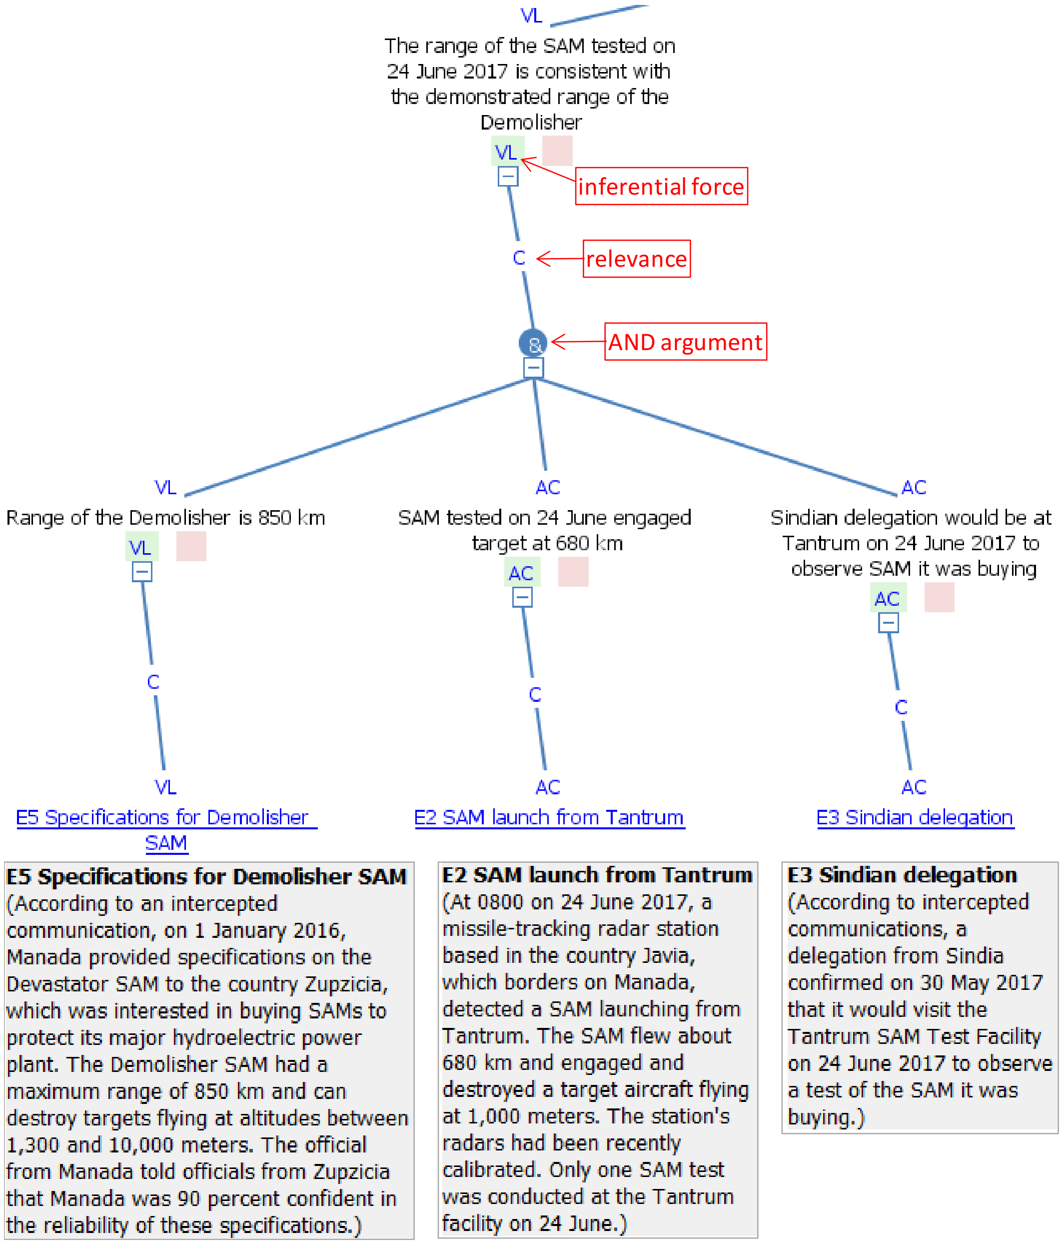

Figure 3. The favoring argument for the “Manada is selling the Demolisher” hypothesis.

The relevance of each evidence item from Figure 3 to the sub-hypothesis above it is C (certain, 100%) because, if the evidence item is true, the sub-hypothesis above it is also true. The credibility of E5 is VL (very likely, 80-95%) because the source indicated a confidence level of 90% and we have no reason to doubt its truthfulness. The credibility of both E2 and E3 is AC (almost certain, 95-99%) because the corresponding information is from an intercepted communication. Once all these assessments are made by the analyst, Cogent determines the probability of the top hypothesis in Figure 3 as VL (very likely, 80-95%).

The disfavoring argument for the Demolisher hypothesis is shown in bottom part of Figure 4. This is again an AND argument. We judge its relevance as VL (very likely, 80-95%) because the target was engaged at 1,000 meters which is significantly outside the known target-engagement range of Demolisher (1,300 – 10,000 meters). The inferential force of this argument is VL (very likely, 80-95%), the smallest among the probabilities of the sub-hypotheses and the relevance of the argument.

The top part of Figure 4 shows the assessment of the Demolisher hypothesis. The relevance of its favoring argument (or reason) is only BL (barely likely, 50-55%) because another missile, the Devastator, also has a range that is compatible with the SAM tested on 24 June. So, based only on the range, the probability that Manada is selling Demolisher is about 50%.

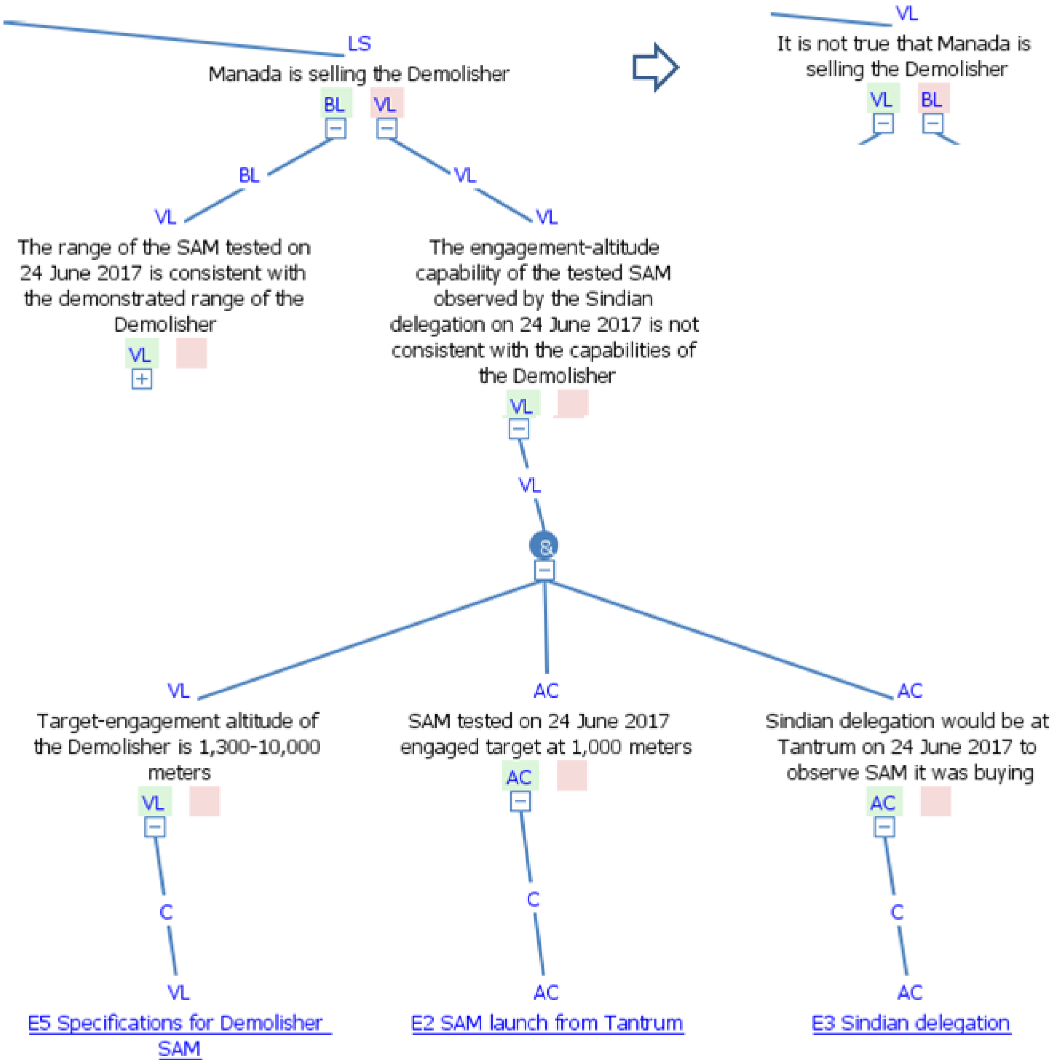

Figure 4. The assessment of the Demolisher hypothesis.

The relevance of the disfavoring argument (reason) of the Demolisher hypothesis is assessed as VL (very likely, 80-95%) because this information has a much higher probability of precluding the Demolisher as the SAM being sold.

Because the inferential force of the disfavoring argument is higher than the inferential force of the favoring argument, Cogent concludes that the Demolisher hypothesis lacks support.

A top-level hypothesis that lacks support may be changed to negation to determine the likelihood of its negation, as shown in Figure 5:

It is very likely (80-95%) not true that Manada is selling the Demolisher.

This can be more naturally expressed by using the negative probability scale:

It is very unlikely (5-20%) that Manada is selling the Demolisher.

Note that, in some cases, the opposite hypothesis also may lack support, in which case you cannot conclude whether that SAM is being sold or not.

Let us now consider the assessment of Destructor hypothesis, shown in Figure 5.

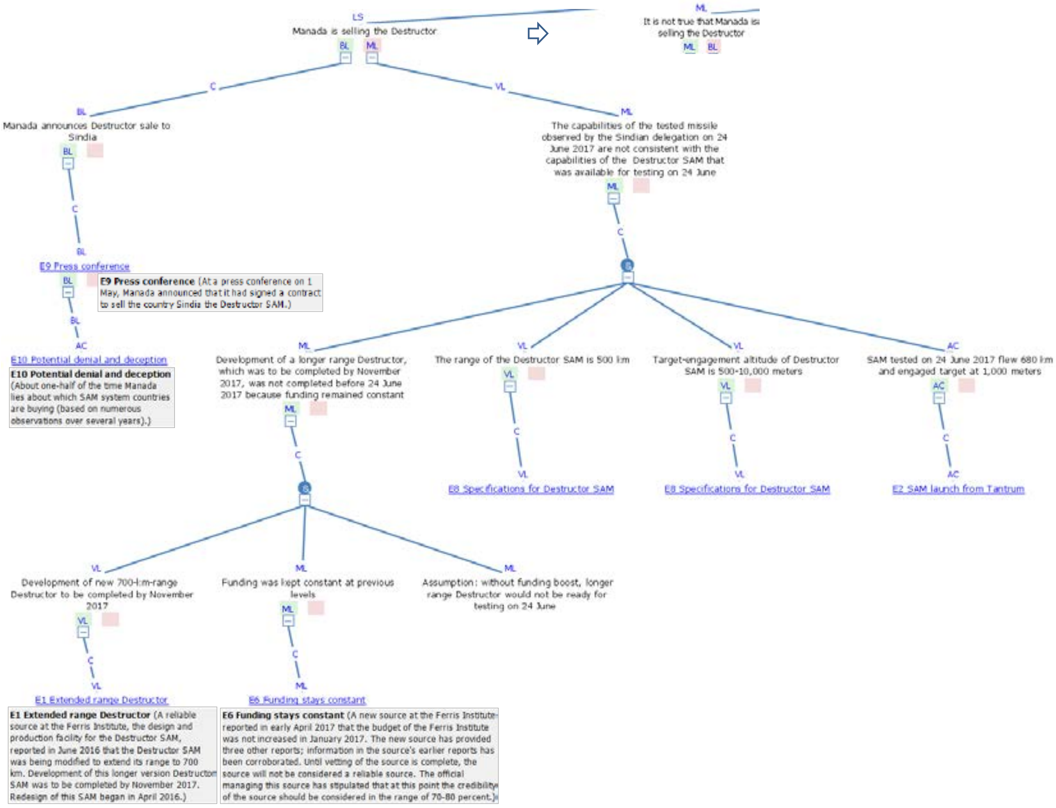

Figure 5. The assessment of the “Manada is selling the Destructor” hypothesis.

The favoring argument is that Manada announced that it was selling the Destructor SAM. It is based on the **E9 Press conference** evidence (see Figure 5).

The relevance of E9 to the sub-hypothesis above it (Manada announces Destructor sale to Sindia) is obviously C (certain, 100%). But what is its credibility? We have evidence **E10 Potential denial and deception** that raises questions about the credibility of the government announcement.

We therefore assess the credibility of E9 as being only BL (barely likely, 50-55%). As a result, the inferential force of the favoring argument is only BL.

The disfavoring argument for the Destructor hypothesis is that the range of the SAM tested on 24 June greatly exceeds the range of the Destructor. The relevance of this disfavoring argument is assessed as VL (very likely, 80-95%) because this information has a much higher probability of precluding the Destructor as the SAM being sold.

Notice the AND argument favoring this disfavoring argument. A component of this argument is that the development of a longer range Destructor, which was to be completed by November 2017, was not completed before 24 June 2017 because funding remained constant. This is itself an AND argument with three components. The first is the hypothesis that the development of new 700-km-range Destructor was to be completed by August 2017. This is judged as VL (very likely, 80-95%) based on evidence **E1 Extended range Destructor** from a reliable source.

The second component is the hypothesis that the funding was kept constant at previous levels. This is judged as ML (more than likely, 70-80%) based on evidence **E6 Funding stays constant** from a source whose credibility is considered in this range by the official managing this source.

The third component is the assumption that, without a funding boost, the longer range Destructor would not be ready for testing on 24 June. We assess the probability that this assumption is true as ML (more than likely, 70-80%).

As a result of these assessments, the inferential force of the disfavoring argument is ML (more than likely, 70-80%), which is higher than the BL (barely likely, 50-55%) inferential force of the favoring argument. As a result, the hypothesis that Manada is selling the Destructor is lacking support. Negating this hypothesis, we obtain that it is more than unlikely (20-30%) that Manada is selling the Destructor.

Figure 6 presents the three top-level hypotheses and their main arguments. Let us consider the "Manada is selling the Devastator" hypothesis and its favoring argument. We lack direct evidence that states Manada is selling the Devastator. Given that is possible that more than one SAM has the same capabilities of the tested missile, the relevance of the favoring argument depends somewhat on our assessment of the likelihood that the capabilities of the other SAMs also meet the capabilities of the tested missile, and thus are potential candidates for being sold. We assessed as being more than unlikely (20-30%) that Manada is selling the Destructor, and very unlikely (5-20%) that Manada is selling the Demolisher. Thus, a reasonable probability assessment for the relevance in question here is L (likely, 55-70%) because the sum of the probabilities of the three hypotheses should be 100%. We therefore conclude that Manada is likely (55-70%) selling the Devastator SAM to Sindia.

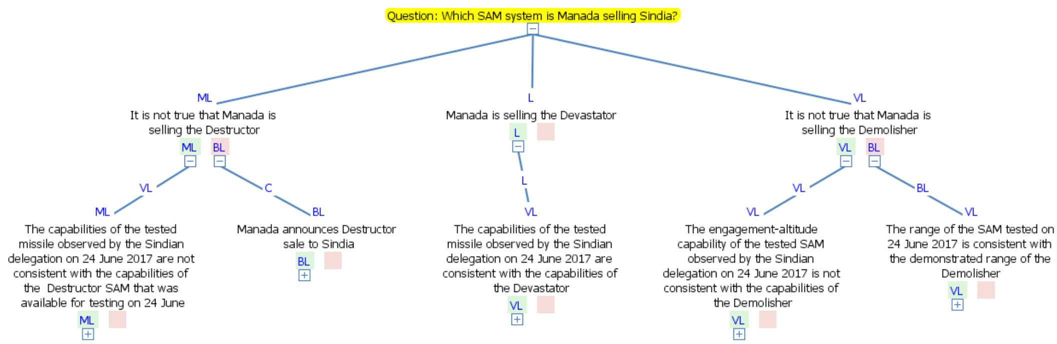

Figure 6. The assessment of the three competing hypotheses.

ANALYSIS IN THE CONTEXT OF EVOLVING INFORMATION AND AN EVER CHANGING WORLD

Intelligence analysis has many difficulties, but none seems more difficult than the fact that analysts must provide their answers while the world keeps changing. An answer that appeared correct in the past may now seem incorrect in light of new evidence just discovered today. A prediction regarded as very likely today may be overtaken by events we will learn about tomorrow. In fact, the very questions we have asked may need to be revised or may even seem unimportant in light of new information.

To illustrate, let us assume that, after performing the above analysis, we have just received the evidence **E14 Funding was increased** that directly contradicts the previous item of evidence **E6 Funding stays constant**.

Not only is E14 more recent than E6, but it also comes from a reliable source, and therefore its credibility is VL (very likely, 80-95%). We therefore have to conclude that funding was increased to accelerate the development of a longer range Destructor. As a result, the argumentation for the Destructor hypothesis is updated as shown in Figure 7 and explained below.

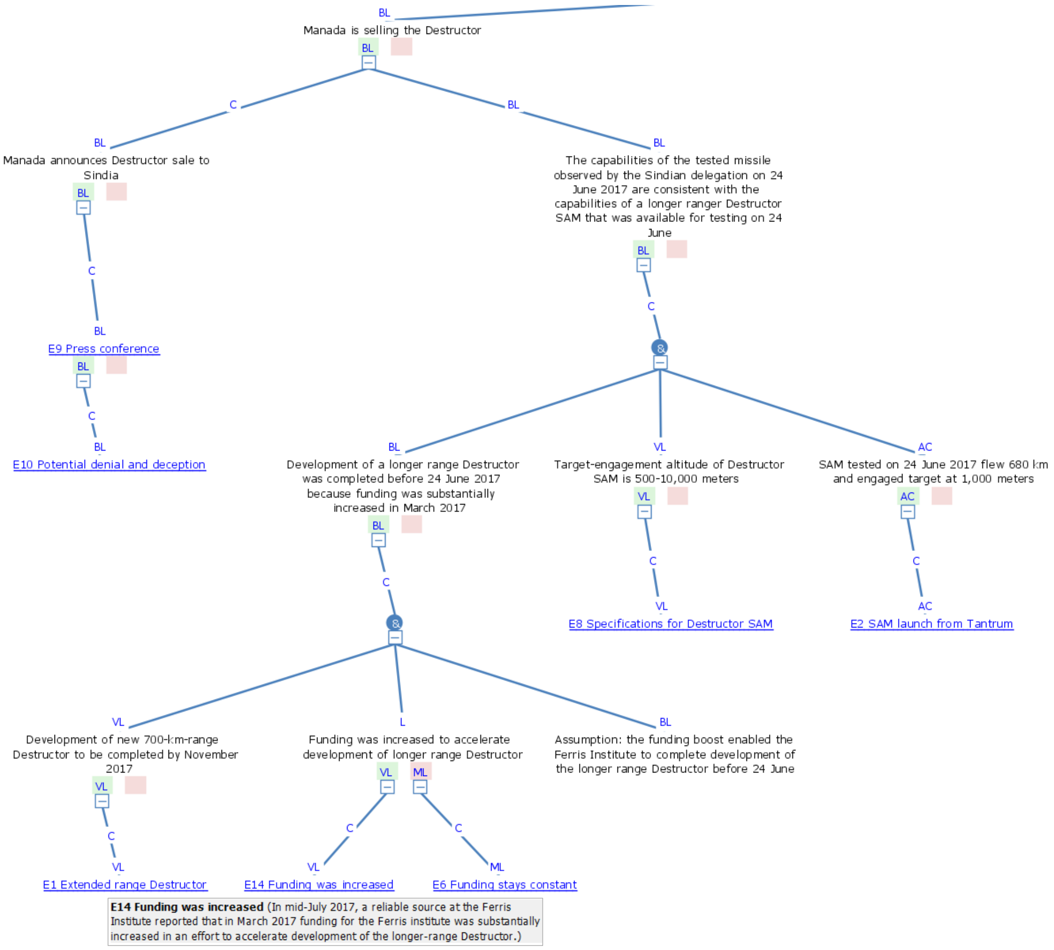

Figure 7. Revised assessment of the Destructor hypotheses.

Cogent can help analysts to properly evaluate new information by guarding against over-assessing or under-assessing its significance. New information, which is inserted into an updated Cogent argumentation, is evaluated within the construct of all the previous information on the subject in question, not in isolation. The prevailing analytic conclusion is not easily discarded without a strong basis for doing so. Similarly, Cogent makes it more difficult to minimize or discount the impact of new information because it requires that the prevailing analytic judgment be reassessed and take into account the new information.

Notice that we now are making the assumption that the funding boost enabled the Ferris Institute to complete development of the longer range Destructor before 24 June, but the probability of this assumption is only BL (barely likely, 50-55%). A consequence of the new information is that the disfavoring argument is changed into a favoring one because now the capabilities of the tested missile observed by the Sindian delegation on 24 June 2017 are consistent with the capabilities of a longer ranger Destructor SAM that was available for testing on 24 June. So now we have two missiles whose capabilities are consistent with the tested one: Destructor and Devastator. Therefore the probability that Manada is selling Destructor, based only on this feature is about 50%, and the relevance of this new favoring argument is only BL (barely likely, 50-55%).

Similarly, the relevance of the favoring argument of the Devastator hypothesis is now BL (see the middle of Figure 6), and the probability Manada is selling Devastator is also BL.

The conclusion is that, at this point, based on the available information, it is equally likely that Manada is selling Destructor and Devastator, and there is almost no chance (1-5%) that it is selling the Demolisher (because the sum of the probabilities should be 100%).

COGENT INTERFACE

Cogent has numerous capabilities enabling an analyst develop a rigorous argumentation and a corresponding production report that is both defensible and persuasive. Figure 8 illustrates its interface. The upper left side is the whiteboard area where the argumentation is developed using simple right-click and drag-and-drop operations. The upper right side is the assistants' area, each assistant helping in performing a group of related operations. The expanded one is the Evidence Assistant that enables the definition of the evidence items and dragging-and-dropping them in the argumentation. For example, the analyst may select a paragraph from an Internet document and then drag-and-drop it under a hypothesis as favoring or disfavoring evidence. As a side-effect, a corresponding evidence item is automatically created in the Evidence Assistant.

The Argument Assistant helps with developing the argumentation, as was illustrated in the previous sections. The Analytics Assistant checks the argumentation for potential errors and guides the analyst in addressing them. For example, it detects when the sum of the probabilities of disjoint hypotheses is either under or above 100%. It also warns of potential confirmation, satisficing or absence of evidence bias. The Report Assistant generates a structured report corresponding to the developed argumentation and helps the analyst in transforming it into a production report which is the final result of the analysis.

Online help is accessible through the tabs shown at the bottom of Figure 8. Under the Instructions tab there is a sequence of instructions that guide the user to solve the current problem. When the user clicks on a node in the argumentation, the Operations Help tab lists all the operations that can be performed at that node, and illustrates each of them. The EBR Training tab provides access to the videos and presentations of the evidence-based theory that is at the basis of Cogent. When the user clicks on a report component, the Report Help tab lists all the operations that can be performed at that component, and illustrates each of them. The Report Training tab provides access to the videos and presentations on how to transform a structured report generated by Cogent into a short, clear, and persuasive production report. The General Help tab provides access to videos and presentations on other aspects of Cogent, for example on how to use it within Amazon AppStream environment.

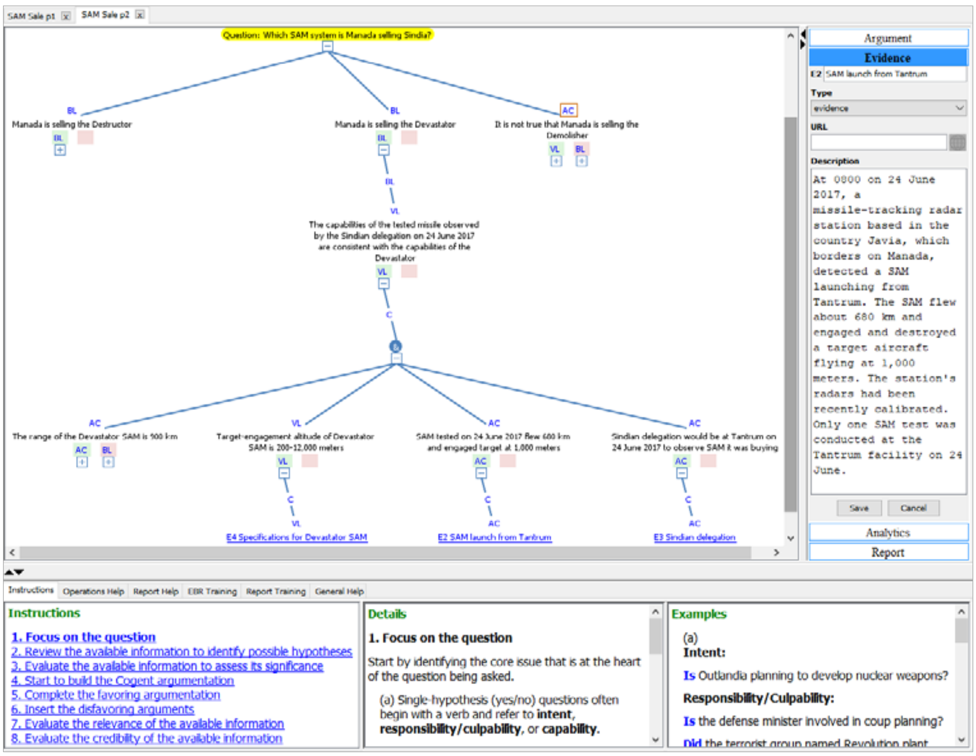

Figure 8. Cogent interface.

ANALYST-COGENT SYNERGISTIC INTEGRATION

A main goal of this research is to enable a synergistic integration of the analyst's imagination and expertise with the computer's knowledge and critical reasoning. Consider again the process described in Figure 1. While the analyst has to imagine the questions to ask and to hypothesize their possible answers, Cogent helps with producing a schematic diagram that completely lays out the underlying analytic framework for every analytic conclusion, including the connection between the evidence and various intermediate conclusions in the analysis, the evaluation of the credibility of evidence and its strength in supporting a conclusion, and the role of any assumptions in addressing missing information. Cogent can also detect errors, biases, and potential errors, guiding the analysts in addressing them. It facilitates the analysis of what-if scenarios, where the analyst may make various assumptions and Cogent automatically determines their influence on the analytic conclusion. It also automatically updates the analysis based on new or revised evidence. Once the analysis is finalized, Cogent generates a structured report that the analyst has to transform into a more understandable and persuasive production report. The final report, which includes argumentation fragments and evidence, can be shared with other users, subjected to critical analysis, and correspondingly improved.

CONCLUSION

We have briefly presented a structured methodology and the Cogent cognitive assistant that help analysts better cope with the astonishing complexity of intelligence analysis and perform better reasoning in answering intelligence questions. Additionally Cogent can be used as a teaching and learning tool for critical thinking. The explicit analysis developed, that makes very clear the argumentation logic, what evidence was used and how, what is not known, and what assumptions have been made, greatly facilitates the trainer's job of providing constructive comments on the analysis in question and each employee's analytic skill. Finally, while this paper has focused on intelligence analysis, the presented methodology and system, can be applied, with appropriate adaptations, to almost any investigative research endeavor, including physical sciences, medicine, cybersecurity, criminal investigations, law, and military and business inferences and decisions.[17,18]

ACKNOWLEDGEMENTS

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) under contract number 2017-16112300009, and by George Mason University.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government.

The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

1. Schum, D. A., *Evidence and Inference for the Intelligence Analyst* (2 Volumes). Lanham, MD: University Press of America, 1987.
2. Schum, D. A., *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press, 2001.
3. Schum, D.A., Tecuci, G., Boicu, M., "Analyzing Evidence and Its Chain of Custody: A Mixed-Initiative Computational Approach," *International Journal of Intelligence and CounterIntelligence*, vol. 22, no. 2, pp. 298-319, 2009.
4. Tecuci, G., Schum, D. A., Marcu, D., and Boicu, M., *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*, Cambridge

University Press, 2016.
5. Tecuci, G., Marcu, D., Boicu, M., Schum, D.A., *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, Cambridge University Press, 2016.
6. Wigmore, J. H., "The Problem of Proof," *Illinois Law Review*, vol. 8, no. 2, pp. 77 – 103, 1913.
7. Wigmore, J. H., *The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience and Illustrated in Judicial Trials,* 3rd edition, Little, Brown & Co, Boston, MA, 1937.
8. Tecuci, G., Boicu, M., Ayers, C., Cammons, D., "Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving." *Proceedings of the 1st International Conference on Intelligence Analysis*, McLean, VA, 2005.
9. 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*. Also in Auger J. and Wimbish W. eds. *Proteus Futures Digest*, 303-329, Joint publication of National Intelligence Univ., Office of the Director of National Intelligence, and U.S. Army War College Center for Strategic Leadership, 2007.
10. Tecuci, G., Boicu, M., Marcu, D., Boicu, C., and Barbulescu, M., "Disciple-LTA: Learning, Tutoring and Analytic Assistance," *Journal of Intelligence Community Research and Development,* 2008*.*
11. Tecuci, G., Marcu, D., Boicu, M., Schum, D. A., and Russell, K., "Computational Theory and Cognitive Assistant for Intelligence Analysis," *Proceedings of the Sixth International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS*, pp. 68-75, Fairfax, VA, 16-18 November, 2011.
12. Tecuci G., Marcu D., Boicu M., Schum, D.A., "COGENT: Cognitive Agent for Cogent Analysis," *Proceedings of the 2015 AAAI Fall Symposium "Cognitive Assistance in Government and Public Sector Applications,"* pp.58-65, Arlington, VA, 12-14 November, Technical Report FS-15-02, AAAI Press, Palo Alto, CA, 2015.
13. Cohen, L. J., *The Probable and the Provable,* Clarendon Press, Oxford, 1977.
14. Cohen, L. J., *An Introduction to the Philosophy of Induction and Probability*, Clarendon Press, Oxford, 1989.
15. Zadeh, L., "The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems," *Fuzzy Sets and Systems*, vol. 11, pp. 199 – 227, 1983.
16. ICD 203, Intelligence Community Directive 203, Office of the Director of National Intelligence, 2007.
17. Anderson, T., Schum, D., and Twining, W., *Analysis of Evidence*, Cambridge University Press, 2005.
18. Fox, J., "Arguing About the Evidence: A Logical Approach" in P. Dawid et al. (eds.) Evidence, Inference and Enquiry, Oxford University Press, 2011.

ABOUT THE AUTHORS

Gheorghe Tecuci is Professor of Computer Science and Director of the Learning Agents Center at George Mason University, member of the Romanian Academy, and former Chair of Artificial Intelligence at the U.S. Army War College. Email: tecuci@gmu.edu

Louis Kaiser is Senior Researcher in the Learning Agents Center, former Senior CIA analyst, Course Director and Instructor in the Sherman Kent School's Career Analyst Program. Email: lmk1981pa@yahoo.com

Dorin Marcu is Research Associate Professor in the Learning Agents Center. Email: dmarcu@gmu.edu

Chirag Uttamsingh is Graduate Research Assistant and Ph.D. student in the Learning Agents Center. Email: cuttamsi@masonlive.gmu.edu

Mihai Boicu is Associate Professor of Information Sciences and Technology and Associate Director of the Learning Agents Center at George Mason University. Email: mboicu@gmu.edu