

TIACRITIS System and Textbook: Learning Intelligence Analysis through Practice

Gheorghe Tecuci¹, Mihai Boicu¹, Dorin Marcu¹, David Schum¹
Benjamin Hamilton²

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

² Department of Defense, USA

Abstract. This paper presents the TIACRITIS web agent and textbook for teaching intelligence analysts the critical thinking skills needed to perform evidence-based reasoning. They are based on a computational theory which views Intelligence Analysis as ceaseless discovery of evidence, hypotheses, and arguments, in a complex world that is changing all the time. TIACRITIS helps students learn about the properties, uses, and marshaling of evidence upon which all analyses rest, through regular practice involving analyses of evidence in both hypothetical and real situations.

Keywords: intelligence analysis, cognitive assistants, education and training

1 Introduction

The purpose of Intelligence Analysis is to answer complex questions arising in the decision-making process. Complex arguments, requiring both *imaginative* and *critical reasoning*, are necessary in order to establish and defend the *relevance*, *believability*, and *inferential force* of evidence with respect to the questions asked. The answers are necessarily probabilistic in nature because our evidence is always *incomplete*, usually *inconclusive*, frequently *ambiguous*, commonly *dissonant*, and with various degrees of *believability* [1, 2]. Moreover, the analysts are often required to answer questions very quickly, with insufficient time for extensive research of the available evidence.

How should the analysts be trained for such astonishingly complex tasks?

First, we think that learning to perform such complex evidential reasoning tasks cannot be done effectively just by listening to someone discuss his/her own analyses, or just by giving students lectures and assigned readings on the topics. What is absolutely necessary is *regular practice involving analyses of evidence* using either hypothetical situations or examples drawn from actual situations. In short, evidential analysis is mastered best by performing analyses contrived to illustrate the wide variety of subtleties or complexities so often encountered in actual evidential analyses. Second, based on our inspection of the materials offered in several courses for training intelligence analysts, it appears that analysts are so often trained in the

production of intelligence analyses (i.e., how to write analysis reports) rather than upon the actual *process of analysis* itself. Very little training is offered regarding the properties, uses, discovery, and marshaling of the evidence upon which all analyses rest. Our third conclusion is based on the strong emphasis currently placed in the Intelligence Community on the development of structured analytic methods and computer-based tools to assist analysts. These tools, however, to be really useful, need to have solid theoretical foundations grounded in the *Science of Evidence*.

To address the above issues, we are developing a *Computational Theory of Intelligence Analysis* which is briefly introduced in Section 2. This theory is at the basis of a textbook and web-based system, called TIACRITIS, for teaching intelligence analysts the critical thinking skills required to perform evidence-based reasoning, through a hands-on, learning by doing approach. The textbook is briefly introduced in Sections 3. Section 4 illustrates the use of the TIACRITIS system.

2 Computational Theory of Intelligence Analysis

We are developing a *Computational Theory of Intelligence Analysis*, grounded in the Science of Evidence, Artificial Intelligence, Logic, and Probabilities, to be used as a basis for building advanced cognitive assistants that:

- Support intelligence analysts in coping with the astonishing complexity of providing accurate explanations and predictions in a non-stationary world;
- Help intelligence analysts learn critical thinking skills for evidence-based reasoning through an effective hands-on approach.

As illustrated in Fig. 1, we view intelligence analysis as a process of ceaseless discovery in a non-stationary world involving mixtures of *abductive*, *deductive*, and *inductive* reasoning for *evidence in search of hypotheses*, *hypotheses in search of evidence*, and *evidential tests of hypotheses*, all going on at the same time. By means of abductive reasoning we generate hypotheses from evidence we gather; by deductive reasoning we make use of our hypotheses to generate new lines of inquiry and evidence; and by inductive reasoning we test hypotheses on the basis of the evidence we are discovering. Such testing depends on the relevance and believability of our evidence. These factors combine in further complex ways to allow us to assess the inferential force or weight of the evidence we are considering [3]. Based on this computational theory we have developed cognitive assistants for intelligence analysts that synergistically integrate three complex capabilities. They can rapidly *learn* the analytic expertise which currently takes years to establish, is lost when analysts sepa-




Fig. 1. Intelligence Analysis as discovery of evidence, hypotheses, and arguments.

rate 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 evaluating the likelihood of hypotheses by developing Wigmorean probabilistic inference networks [4] that link evidence to hypotheses in argumentation structures that establish the relevance, believability and inferential force or weight of evidence. A first prototype of such an agent is Disciple-LTA [5, 6] which is at the basis of the TIACRITIS system introduced in this paper.

3 Intelligence Analysis Textbook

The TIACRITIS textbook has been written for persons throughout the Intelligence Community and those it serves, including collectors of intelligence information, evaluators of incoming intelligence information at various levels and in different offices, and even policy-making "customers" of intelligence analysts. The matters discussed are applicable regardless of the subject of an intelligence analysis and the kinds of intelligence information required, such as HUMINT, IMINT, SIGINT, MASINT, and Open Source information.

The textbook teaches basic knowledge about the properties, uses, and marshaling of evidence to show students what is involved in assessing the relevance, believability, and inferential force credentials of evidence. It includes a wide array of examples of the use of the TIACRITIS system and hands on exercises involving both real and hypothetical cases chosen to help students recognize and evaluate many of the complex elements of the analyses they are learning to perform. Each chapter starts with a presentation of some important matter, such as assessing the believability of evidence. Then the students are asked to use TIACRITIS and experiment with what they have just been taught. Both the textbook and TIACRITIS are easily customizable by selecting the chapters and the case studies to be used.

Also discussed in the textbook is how the intelligence analysis concepts and methods embedded into TIACRITIS (e.g., the systematic approach to the development of argumentation structures, the ontology of evidence and the associated procedures for assessing the believability of evidence, the drill-down analysis and assumptions-based reasoning) help analysts perform better analyses, no matter what analysis methods they use. In particular, it shows how the very popular Richards J. Heuer's Analysis of Competing Hypothesis (ACH) method [7] can be improved by employing the concepts and methods embedded into TIACRITIS.

4 The TIACRITIS System

TIACRITIS is a web-based system with knowledge bases and case studies incorporating a significant amount of knowledge about evidence, its properties, uses, and discovery. Each knowledge base includes an ontology that defines both general concepts for evidence-based reasoning [8], and domain-specific concepts from an application domain. It also includes learned problem reduction rules and solution synthesis rules which are represented with the concepts from the ontology. These

knowledge bases allow TIACRITIS to automatically generate argumentation structures for hypotheses testing, as illustrated in Fig. 2.

The case studies are designed to learn about and practice with one new important matter at a time, such as analyzing hypotheses through reduction and synthesis, making assessments and assumptions in arguments, assessing the believability of evidence, analyzing competing hypotheses, etc. To provide an intuitive understanding of the use of TIACRITIS we present the case study “Hypothesis analysis and evidence search” which helps students learn how to search for relevant evidence on the Internet. This case study also guides the student to practice with many of the matters introduced in the previous ones. As shown at the bottom of Fig. 2, the student is instructed to select a hypothesis analysis problem and browse its analysis tree to see how it is reduced to simpler hypotheses that have to be assessed by searching for evidence on the Internet. The student will then define search criteria for the elementary hypotheses, will invoke specific search engines with those criteria, copy relevant information into TIACRITIS, define evidence from this information, associate evidence with the corresponding hypotheses, and evaluate its relevance and believability, with the goal of assessing the likelihood of the top level hypothesis.

The student is first instructed to select the **Hypothesis** menu at the top of the window. As a result, TIACRITIS will display a list of hypotheses to select from, including the option to define a new hypothesis. Next the student is instructed to select the hypothesis analysis problem “*Assess whether United States will be the glo-*

The screenshot shows the Reasoner module interface with the following components:

- Navigation Bar:** Case Study, Hypothesis, Reasoner, Evidence, Assumption, Description, Test, Preferences, Gheorghe Tecu
- Abstract Analysis:**
 - United States will be a global leader in wind power within the next decade
 - reasons
 - United States imports huge quantities of oil
 - favoring evidence
 - disfavoring evidence
 - huge oil import is security threat
 - huge oil import is big burden on economy
 - huge oil consumption is damaging the environment
 - desire
 - people of United States
 - major political parties
 - energy industries of United States
 - capability
 - homegrown scientific knowledge
 - technical knowledge
 - economic resources
 - significant natural resources

- Q&A Dialog:**
- Q: Which is a potential reason?
- A: Significant internal production of wind power...
- Assessment Boxes:**
- Assess whether United States imports huge quantities of oil.
- Assess whether quantities of oil security threat to United States.
- Instructions:**
- Select **Hypothesis**
- Select the hypothesis analysis problem: *Assess whether United States will be a global leader in wind power within the next decade*
- Reasoner** is automatically selected
- Browse the analysis of the above hypothesis to see how it is reduced to simpler hypotheses that you have to assess by searching for evidence on the Internet.
- Select **Evidence**

Fig. 2. The interface of the Reasoner module.

bal leader in wind power within the next decade.” As a result, the **Reasoner** module is automatically invoked, generating and displaying the analysis tree from Fig. 2. Notice that the left panel displays an abstraction of the decomposition tree where the top level hypothesis is successively decomposed into simpler and simpler hypotheses, with the simplest one (such as “United States imports huge quantities of oil”) to be assessed based on favoring and disfavoring evidence.

The student is next instructed to browse this analysis tree. As she clicks on an abstract hypothesis in the left panel (e.g. “reasons”), the right panel displays the detailed decomposition of the corresponding hypothesis analysis problem, as illustrated in the right side of Fig. 2 and in Fig. 3.

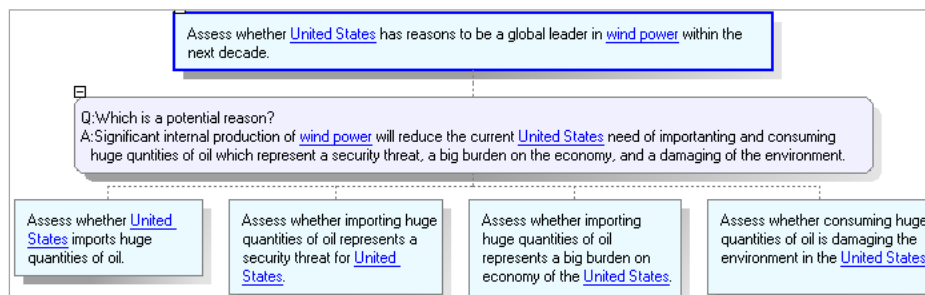


Fig. 3. Detailed reasoning step.

Next the student is instructed to select the **Evidence** menu (see the top of Fig. 2) and is explained the operation modes shown in the upper part of the left panel in Fig. 4. Since [COLLECTION GUIDANCE] is selected, the left panel shows the elementary hypotheses and their evidential support. When the student clicks on one such hypothesis, the right panel provides more details about it. It also allows the student to associate search criteria with the selected hypothesis. For example, in the situation illustrated in Fig. 4, the student has selected the hypothesis “United States imports huge quantities of oil,” and has associated two search criteria with it, the second one being currently selected. Clicking on one of the available search engines (i.e., BING, GOOGLE, YAHOO) will invoke it with the selected search criterion. The student will use these capabilities to associate search criteria with elementary hypotheses and to search for relevant evidence on the Internet. In this example the student has identified a relevant article by Daniel Workman. She is instructed to copy it into TIACRITIS, extract items of evidence from it, specify the types of these items of evidence, and associate them with the corresponding hypotheses.

The right panel in Fig. 5 shows the defined characteristics of such an item of evidence, *EVD-001-US-top-oil-importer*. Notice its description and the item of information from which it was extracted, *INFO-001-US-oil-import*, which is the entire article. The student was then instructed to select its type from a comprehensive list of possible types. Since she selected “unequivocal testimonial evidence obtained at second hand”, she was prompted to specify both the name of the source of the testimony (i.e. *Daniel Workman*), and that of the primary source (*US Energy Information Administration*).




Fig. 4. Defining search criteria for a given hypothesis.

The bottom part of the right panel displays the list of all the elementary hypotheses from the analysis tree, under the label “Irrelevant to,” each followed by four commands: [FAVORS], [DISFAVORS], [REASONING], [COLLECTION]. The student may select one of the first two commands to indicate that the current item of evidence (i.e., *EVD-001-US-top-oil-importer*) favors or disfavors that hypothesis. In the illustration from Fig. 5, the student indicated that this item of evidence favors the hypothesis that the “*United States imports huge quantities of oil.*” As a result, the “*Favors*” label was created and this hypothesis was moved under it.

The student is next instructed to select the [REASONING] command following the above hypothesis. As a result, the **Reasoner** module is invoked with “favoring evidence” for that hypothesis selected, as shown in the left hand side of Fig. 6. Notice that TIACRITIS has automatically generated the reasoning tree for the assessment of the relevance and believability of *EVD-001-US-top-oil-importer* with respect to the

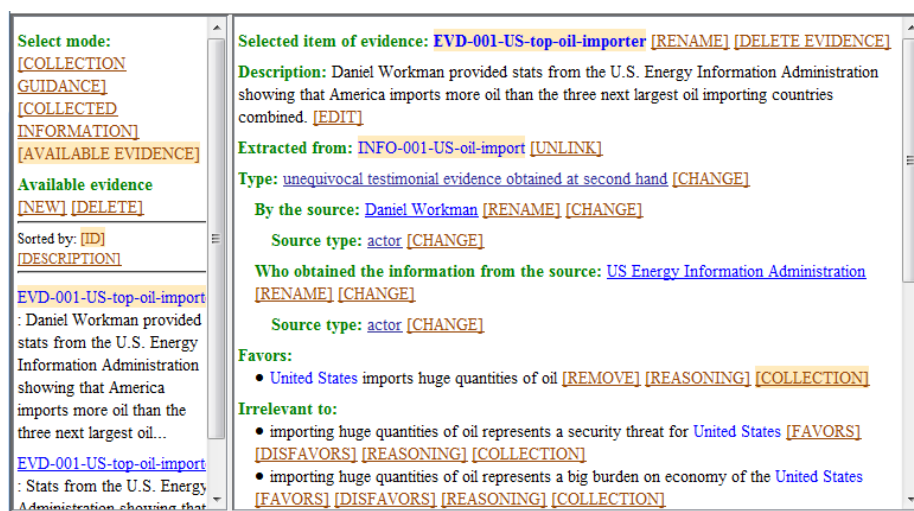


Fig. 5. Defined item of evidence favoring a hypothesis.

hypothesis that the “United States imports huge quantities of oil.” In particular, being *testimonial evidence obtained at second hand*, the *believability* of this item of evidence depends on the believability of its primary and secondary sources which, in turn, depend on their *competence* and *credibility*. Competence involves *access* and *understandability* while *credibility* involves *veracity*, *objectivity* and *observational sensitivity* [2]. The student may either assess these lower level believability credentials, or she may assess upper level ones, or even make a holistic assessment of the believability of the item of evidence. In the illustration of this case study, she clicks on “believability EVD-001-US-top-oil-importer” and then selects the **Assumptions** menu, to specify the believability of this item of evidence as an assumption. As illustrated in the right hand side of Fig. 6, TIACRITIS displays the assessment problem to be solved and a pattern for its solution. What the student needs to do is to select a likelihood, such as “almost certain” or “likely,” from the menu list.




Fig. 6. Evidence assessment.

The relevance is assessed in a similar way. Both assessments are shown in Fig. 7 with a yellow background to indicate that they have been specified as assumptions. Notice also that TIACRITIS has automatically computed the inferential force of this item of evidence on some of the upper level hypotheses.




Fig. 7. Computation of the inferential force of evidence through solution synthesis.

After being guided to perform these operations, the student is instructed to complete the analyses of the top level hypothesis. Then she is instructed to define and solve a new hypothesis analysis problem, such as “*Assess whether China will be the global leader in solar power within the next decade.*”

6 Conclusions

We have presented an intelligent web-based agent and an associated textbook for teaching intelligence analysts through an effective learning by doing approach. Both of them may be used as such or may be easily extended or customized with additional topics and case studies to better serve specific audiences. We plan to continuously increase the training and operational effectiveness of TIACRITIS.

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 and George Stemler for their suggestions and support. Cristina Boicu has contributed to the development of TIACRITIS. Gary Roemmick and Ben Wible have contributed to its transition to the Joint Forces Staff College and to the development of specialized case studies. 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. 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. Schum, D.A.: Evidence and Inference for the Intelligence Analyst. University Press of America, Lanham, MD (1987)
2. Schum, D.A.: The Evidential Foundations of Probabilistic Reasoning. Northwestern University Press (2001)
3. Tecuci, G., Schum, D.A., Boicu, M., Marcu, D., Hamilton, B.: Intelligence Analysis as Agent-Assisted Discovery of Evidence, Hypotheses and Arguments. In: Phillips-Wren, G., Jain, L.C., Nakamatsu, K., Howlett, R.J. (eds.) Advances in Intell. Decision Technologies, SIST 4, pp. 1--10. Springer-Verlag Berlin Heidelberg (2010)
4. Wigmore J.H.: The Science of Judicial Proof. Little, Brown & Co. Boston, MA (1937)
5. Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M., Ayers, C., Cammons, D.: Cognitive Assistants for Analysts. J. Intell. Community Res. Dev. (2007)
6. Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M.: Disciple-LTA: Learning, Tutoring and Analytic Assistance. J. Intell. Community Res. Dev. (2008)
7. Heuer, R.J.: Psychology of Intelligence Analysis, Center for the Study of Intelligence, Central Intelligence Agency, Washington, DC (1999)
8. Schum, D., Tecuci, G., Boicu, M., Marcu, D.: Substance-Blind Classification of Evidence for Intelligence Analysis. In: Ontology for the Intelligence Community, Fairfax, VA (2009)