

TEACHING INTELLIGENCE ANALYSIS WITH TIACRITIS

**Tecuci G., Schum D.
Boicu M., Marcu D.**
Learning Agents Center
George Mason University

Hamilton B.
Department
of Defense

Wible B.
Joint Forces
Staff College

ABSTRACT: This paper 1) discusses the astonishing complexity of intelligence analysis by using the popular metaphor of “connecting the dots,” 2) outlines a systematic computational approach, grounded in the Science of Evidence, that allows coping with this complexity, and 3) introduces an innovative intelligent software agent, called TIACRITIS, for teaching intelligence analysts how to perform evidence-based reasoning. TIACRITIS is a web-based system with case studies and knowledge bases incorporating a significant amount of knowledge about evidence, its properties, uses, and discovery. It is a personalizable agent that helps analysts acquire the knowledge, skills, and abilities involved in discovering and processing of evidence and in drawing defensible and persuasive conclusions from it, by employing an effective learning-by-doing approach. It allows analysts to practice and learn how to link evidence to hypotheses through abductive, deductive, and inductive reasoning that establish the basic credentials of evidence: its relevance, believability, and inferential force or weight. Analysts can also experiment with what-if scenarios and study the influence of various assumptions on the final result of analysis.

1. Introduction

One of our most important bulwarks against terrorists and other enemies is the imaginative and critical reasoning abilities of our intelligence analysts who face the highly complex task of drawing defensible and persuasive conclusions from masses of evidence of all kinds from a variety of different sources (Schum, 1987). These conclusions are necessarily probabilistic in nature because the evidence is always incomplete (the analysts can look for more, if they have time), usually inconclusive (it is consistent with the truth of more than one hypothesis or possible explanation), frequently ambiguous (the analysts cannot always determine exactly what the evidence is telling them), commonly dissonant (some of it favors one hypothesis or possible explanation but other evidence favors other hypotheses), and with various degrees of believability (Schum, 2001). Arguments, often stunningly complex, requiring both imaginative and critical reasoning, are necessary in order to establish and defend the three major credentials of evidence: its relevance, believability, and inferential force or weight.

But these assorted evidential characteristics are not the only elements of the complexity of intelligence analysis tasks. A major objective of intelligence analysis is to help insure that the policies and decisions reached by our governmental and military leaders, at all levels, are well informed. Analysts face different requirements in their efforts to serve these policy and decision-making "customers". In some cases they are required to answer questions that are of immediate interest and that do not allow time for extensive research and deliberation on available evidence. In other cases, teams of analysts participate in more lengthy finished intelligence that combines evidence from every

available source. In some cases finished intelligence can refer to long-term assessments on matters of current and abiding interest.

Based on our years of experience in training intelligence analysts, law students and high ranking military officers we have arrived at the following conclusions related to the training of intelligence analysts. The first is that training in the evidential reasoning tasks required in these and other important contexts cannot be learned 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.

Our second conclusion is 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 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. As Douglas MacEachin (1999), former DDI, CIA, remarked, *“Information and expertise are a necessary but not sufficient means of making intelligence analysis the special product that it needs to be. A comparable effort has to be devoted to the science of analysis.”*

Our third conclusion is based on the strong emphasis currently placed in the Intelligence Community on the development of computer-based tools to assist analysts. Analysis tools based on solid theoretical foundations should be integrated into intelligence analysis courses to teach the process of analysis (including the properties, uses, discovery, and marshaling of the evidence) through a learning-by-doing approach.

Recently we have developed Disciple-LTA, a unique and complex cognitive assistant for evidence-based hypothesis analysis which was successfully used at the US Army War College and in several experiments with intelligence analysts (Schum, 2007; Schum et al., 2009a; Tecuci et al., 2005, 2007, 2008). Disciple-LTA has several very distinct differences from other knowledge-based "expert systems" developed in the field of Artificial Intelligence over the years (Jackson, 1999). The name Disciple, by itself, suggests that it learns about intelligence analysis through its interaction with experienced intelligence analysts. Indeed, instead of being programmed by a knowledge engineer (which is a very long, difficult and error-prone process), Disciple-LTA learns (L) its expertise directly from expert analysts who can teach it in a way that is similar to how they would teach a person. Disciple-LTA can also serve as a tutor (T) for novice and experienced analysts, and it can assist (A) in the performance of analytic tasks. Most importantly, it has a stock of established knowledge about evidence, its properties, and its uses for evidence-based assessment of hypotheses, which is employed during all these L, T, A processes. However, Disciple-LTA was developed as a standalone software system that has to be installed on the analyst's workstation. This impedes considerably its use by analysts in their environment.

Disciple-LTA is at the basis of an innovative software agent for teaching intelligence analysts that addresses the three conclusions we have just mentioned. This agent, called TIACRITIS (Teaching Intelligence Analysts Critical Thinking Skills), is a web-based system with case studies and knowledge bases incorporating a significant amount of knowledge about evidence, its properties, uses, and discovery. It is a personalizable agent that will help analysts acquire the knowledge, skills, and abilities involved in discovering and processing of evidence and in drawing defensible and persuasive conclusions from it, by employing an effective learning-by-doing approach. The analysts will practice and learn how to link evidence to hypotheses through abductive, deductive, and inductive reasoning that establish the basic credentials of evidence: its relevance, believability or credibility, and inferential force or weight. As we know, deduction shows that something is necessarily true, induction shows that something is probably true, and abduction shows that something is possibly true. The identification of abductive reasoning was first made by the American philosopher Charles S. Peirce (1998, 1901), who argued that we will not generate any new ideas, in the form of hypotheses, by deductive or inductive reasoning. He identified abductive reasoning as being associated with imaginative, creative, or insightful reasoning.

They will experiment with what-if scenarios and study the influence of various assumptions on the final result of analysis. The use of TIACRITIS in intelligence analysis courses is supported by a dedicated textbook which includes a wide array of examples of the use of TIACRITIS 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.

The rest of this paper is organized as follows. The next section discusses the complexity of intelligence analysis by using the popular metaphor of "connecting the dots." Then Section 3 outlines a systematic computational approach, grounded in the Science of Evidence (Anderson et al., 2005), that allows coping with the astonishing complexity of intelligence analysis. Section 4 presents some of the basic capabilities of the web-based TIACRITIS system for teaching intelligence analysts how to perform a rigorous analysis like the one outlined in Section 3. Section 5 discusses the personalized teaching and learning supported by TIACRITIS and its associated textbook. The paper is concluded with a summary of future work.

2. Intelligence Analysis as Connecting the Dots

One way to describe the process of intelligence analysis is to employ a metaphor that is so often used, and misused, in discussions of intelligence analyses. How many times have you seen accounts in printed sources or heard "talking heads" on television say that our intelligence analysts have failed to "connect the dots" in some important matter concerning our national security? On some accounts, it is made to sound as if the process of connecting intelligence-related dots is as simple as the tasks most of us faced as children when we were asked to use a pencil to connect some collection of numbered dots on a printed page. If the dots were connected in an appropriate order, a representation of some familiar object or person, such as Santa Claus, a witch, a valentine, or a puppy would be revealed. Our belief is that critics employing this metaphor in criticizing intelligence analysts have very little awareness of how astonishingly difficult the process of connecting the dots can be in so many contexts, especially in

intelligence analysis. The main goal of TIACRITIS and the associated textbook is to help the intelligence analysts understand this complexity and be able to cope with it.

Here comes an account of a variety of problems that may so easily be overlooked in discussions of the task of connecting the dots in situations as complex as intelligence analysis.

2.1 Problem 1: There is More than One Kind of Dot to Be Connected

It is so easy to assume that the only kind of dot to be connected concerns details in the collected observable information or data that may eventually be considered as evidence in some analysis. We might refer to these dots as being evidential dots. The second form of dot concerns ideas an analyst has about how some evidential dot, or a collection of evidential dots, is connected to matters that the analyst is trying to prove or disprove. We commonly refer to the matters to be proved or disproved as hypotheses. Hypotheses represent possible alternative conclusions an analyst could entertain about matters of interest in an analysis. These other dots, that we call idea dots, come in the form of intermediary hypotheses in chains of reasoning or arguments constructed by the analyst to link evidential dots to the top level hypotheses. Each of these idea dots refer to sources of uncertainty or doubt the analyst believes to be interposed between his/her evidence and his/her hypotheses. This is precisely where imaginative reasoning is involved. The essential task is for the analyst to imagine what evidential dots mean as far as hypotheses or possible conclusions are concerned. Careful critical reasoning is then required to check on the logical coherence of sequences of idea dots in the analyst's arguments or chains of reasoning. In other words, does the meaning that the analyst has attached to sequences of idea dots make logical sense?

The TIACRITIS system knows about the two forms of dots: evidential and idea dots. It will assist the analysts in forming chains of reasoning or sequences of idea dots that will allow them to form defensible and persuasive arguments from a mass of evidential dots.

2.2 Problem 2: Which Evidential Dots Should be Connected?

Here is where the astonishing complexity of intelligence analysis begins to arise, even in the simplest of cases in any form of analysis based on evidence. Considering just the evidential dots or details mentioned in the previous section, any connections among them would require analysts to consider possible combinations of these individual details, leading to an exponential explosion. Suppose an analyst considers having some number **N** of evidential dots. The question is: How many combinations **C** of two or more evidential dots are there for **N** evidential dots? The answer is given by the following expression: $C = 2^N - [N + 1]$. Here are a few examples showing how quickly **C** mounts up with increases in **N**: for **N** = 10, **C** = 1013; for **N** = 50, **C** = $1.13(10)^{15}$; for **N** = 100, **C** = $1.27(10)^{30}$. Given the array of sensing devices and human observers available to our intelligence services, the number **N** of potential evidential dots is as large as one wishes to make it. In most analyses **N** would certainly be greater than 100 and would increase as time passes. Remember that we live in a non-stationary world in which things change and the analysts find out about new things all the time. So, in most cases, even if the analysts had access to the world's fastest computer, they could not possibly examine all possible evidential dot

combinations even when **N** is quite small. However, if the analysts examined these dots separately or independently they would not perceive these new possibilities, particularly the evidential synergism in which two or more evidence items mean something quite different when examined jointly than when examined separately or independently. This is one of the most interesting and crucial evidence subtleties or complexities that have, quite frankly, led to intelligence challenges in the past: inability to identify and exploit evidential synergisms.

TIACRITIS has been designed to teach analysts identify which evidential dot combinations to look for in which there possibly are the evidential synergisms mentioned above.

2.3 Problem 3: Which Evidential Dots Can Be Believed?

This is one of the most important, challenging, and interesting problems in any area of intelligence analysis. From some source, a sensor of some sort or a person, an analyst obtains an evidential dot E^*_i , saying that a certain event E has occurred. But just because this source says that this event occurred does not entail that it did occur. So, a basic inference the analyst encounters is whether or not E did occur based on the evidence E^*_i . Clearly, this inference rests upon the believability or credibility of source i . There are some real challenges here in assessing the believability of source i . We might easily have said in Problem 1 above that there are even distinctions to be made in what we have termed evidential dots. Some of these dots arise from objects the analysts obtain or from sensors that supply them with records or images of various sorts. So one major kind of evidential dot involves what can be termed tangible evidence that analysts can observe for themselves to see what events it may reveal. In many other cases analysts have no such tangible evidence but must rely upon the reports of human sources who allegedly have made observations of events of interest to the analysts. Their reports to the analysts come in the form of testimonial evidence or assertions about what they have observed. The origin of one of the greatest challenges in assessing the believability of evidence is that analysts must ask different questions about the sources of tangible evidence from those they ask about the sources of testimonial evidence. The crucial importance of evidence believability assessments results from the fact that such assessments form the very foundation for all arguments the analysts make from evidence to possible conclusions. The TIACRITIS system knows very much about this important task.

2.4 Problem 4: What Are the Connections between Evidential Dots and Hypotheses?

From years of experience teaching law students to construct defensible and persuasive arguments from evidence, we have found that most of them often experience difficulty in constructing arguments from single items of evidence; they quickly become overwhelmed when they are confronted with argument construction involving masses of evidence. But they gain much assistance in such tasks by learning about argument construction methods devised nearly a hundred years ago by a world-class evidence scholar named John H. Wigmore (1937). Wigmore was the very first person to study carefully what today we call inference networks which provide the connections between evidential dots and hypotheses. TIACRITIS will help analysts learn how to develop Wigmorean probabilistic inference

networks that link evidence to hypotheses by establishing the believability, the relevance and the inferential force of evidence.

2.5 Problem 5: What Do Our Arguments Mean?

This section considers the direction and force of arguments based on the combined evidence considered. By direction we refer to the hypothesis analysts believe their evidence favors most. By force we mean how strongly analysts believe the evidence favors this hypothesis over alternative hypotheses they have considered. There are two uncontroversial statements we can make about the force or weight of evidence. The first is that the force or weight of evidence has vector-like properties. What this means is that evidence points analysts in the direction of certain hypotheses or possible conclusions with varying degrees of strength. The second is that the force or weight of evidence is always graded in probabilistic terms indicating the analysts' uncertainties or doubts about what the evidence means in terms of its inferential direction and force. But beyond these two statements controversies begin to arise.

So, the question remains: How should one assess and combine the assorted uncertainties in complex arguments in intelligence analysis, and in any other context where the task is to make sense out of masses of evidence? The problem is that there are several quite different views among probabilists about what the force or weight of evidence means and how it should be assessed and combined across evidence in either simple or complex arguments. Each of these views has something interesting to say, but no one view says it all. There is a further difficulty as far as judgments of the weight or force of evidence is concerned. Analysts, or teams of analysts, may agree about the construction of an argument but disagree, often vigorously, about the extent and direction of the force or weight this argument reveals. There may be strong disagreement about the credibility of sources of evidence or disagreements about the strength of relevance linkages. These disagreements can only be resolved when arguments are made carefully and are openly revealed so that they can be tested by colleagues.

There is one final matter of interest in making sense out of masses of evidence and complex arguments. Careful and detained argument construction might seem a very laborious task, however necessary it is. However, let us now consider the task of revealing the conclusions resulting from an analysis to some policy-making "customer" who has decisions to make that rest in no small part on the results of an intelligence analysis. What this "customer" will probably not wish to see is a detailed inference network analysis that displays all of the dots that have been connected and the uncertainties that have been assessed and combined in the process. A fair guess is that this "customer" will wish to have a narrative account or a story about what the analysis predicts or explains. In some cases, at least, "customers" will require only short and not extensive narratives. This person may say: "Just tell me the conclusions you have reached and briefly why you have reached them". So the question may be asked: Why go to all the trouble to construct defensible and persuasive arguments when the "customers" may not wish to see their details?

There is a very good answer to the question just raised. The narrative account of an analysis must be appropriately anchored on the evidence available to the analyst. What analysts wish to be able to tell is a story that they believe contains some truth; i.e. it is not just a good story. The virtue of careful and critical argument construction is that it will allow analysts to anchor their narrative not only on their imagination, but also on the care they have taken to subject their analyses to critical examination. There is no telling what questions analysts might be asked about their analyses. Care in constructing their arguments from their evidence is the best protection they have in dealing with "customers" and other critics who might have entirely different views regarding the conclusions the analysts have reached. TIACRITIS is designed to allow analysts to critically evaluate the arguments they have constructed.

2.6 Problem 6: Whose Dots Should be Connected?

One obvious answer is that all the potential evidential dots collected by any intelligence service that bear upon a problem involving our nation's security should be shared or brought together. Since 9/11/2001 there are so many examples of potential relevant evidence, gathered by different intelligence services, that was never shared across agencies and offices. The basic problem this creates is that the extremely important evidential synergisms discussed in Problem 2 can never be detected and exploited in reaching analytic conclusions. In some cases this has resulted in the failure to reach any conclusion at all in some important matter. This forms the basis for one of the major criticisms of our intelligence services in their failure to "connect the dots".

As will be shown in the next sections, careful argument construction will help reveal the incompleteness of available evidence. An analyst might easily observe that not all questions that should be asked about the problem at hand have in fact been answered. This raises questions such as:

- Have any other agencies or offices attempted to answer these questions that the analyst believes have gone unanswered?
- If these other agencies have gathered such evidence, how can the analyst best justify or be able to have ready access to it?
- What collection efforts should be mounted to gather evidence necessary in order to provide more complete assessments of evidence necessary to form more productive conclusions? In many cases such evidence may never have been collected. In such cases analysts can play important roles in directing effective and productive evidence collection efforts. In so many instances it seems that we try to collect everything with the hope of finding something.

3. A Computational Approach to Intelligence Analysis

Coping with the astonishing complexity of intelligence analysis requires a systematic approach, grounded in the Science of Evidence, which integrates imaginative and critical reasoning to "connect the dots", in a continuously changing world, through ceaseless discovery and testing of evidence, hypotheses, and arguments.

3.1 Discovery of Evidence, Hypotheses and Arguments

Figure 1 illustrates such a systematic approach which involves evidence in search of hypotheses (through abductive or imaginative reasoning), hypotheses in search of evidence (through deductive reasoning), and evidential tests of hypotheses (through inductive reasoning), all going on at the same time. Mavis, a counterterrorist analyst, reads in today's Washington Post an article where a person named Willard reports that a canister of cesium-137 has gone missing from a company XYZ in MD. This item of evidence (E_i^*) leads Mavis to abductively leap to the hypothesis H_k that a dirty bomb will be set off in the Washington DC area. Asked to justify it, Mavis provides the following abductive reasoning shown in the left hand side of Figure 2: From this article (evidence E_i) I infer that it is possible that the cesium-137 canister is indeed missing (fact E_i), that it is possible that it was stolen (H_a), that it was stolen by someone associated with a terrorist organization (H_c), that the organization will use it to construct a dirty bomb (H_e), and that the bomb will be set off in the Washington DC area (H_k). So, in this case, we have evidence in search of hypotheses where an item of evidence “searches” for hypotheses that explain it. Notice here the connections between an evidence dot (E_i^*), idea dots (H_a, H_c, H_d), and hypothesis dot (H_k).

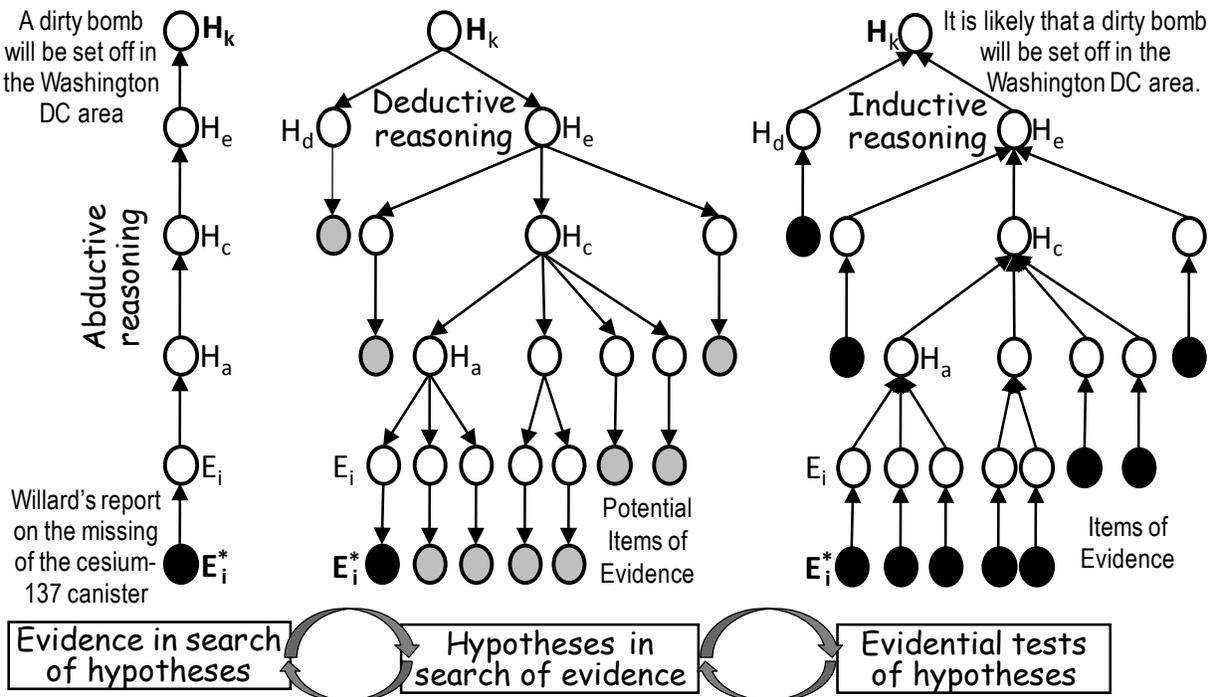


Figure 1. Connecting the dots through the discovery of evidence, hypotheses and arguments.

The diagram in the middle of Figure 1 illustrates the process of hypotheses in search of evidence. Once the new hypothesis H_k has been generated, the analyst has to assess it, through deductive reasoning. The reasoning might start as follows. If H_k were true, there are sub-hypotheses, listed as H_d and H_e , that would be necessary and sufficient to make H_k true. If H_d were true then one would need to observe E_p . Similarly, H_e and its sub-hypotheses allow the analyst to deduce potential items of evidence (shown as shaded circles) that bear upon them. So here we have hypotheses in search of evidence that

guides the process of collecting new evidence. Notice that this helps answering the question: Which evidential dots should be connected? as well as Whose dots should be connected?

Now, some of the newly discovered items of evidence may trigger new hypotheses (or the refinement of the current hypothesis). So, as indicated at the bottom left of Figure 1, 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.

The combination of evidence in search of hypotheses and hypotheses in search of evidence results in hypotheses which have to be tested, through inductive reasoning, based on the discovered items of evidence, as illustrated in the right-hand side of Figure 1. This is a probabilistic inference network that shows how the evidence at the leaves of the network (shown as black circles) is linked to hypotheses (e.g., H_a , H_c , H_e , H_d , H_k) through arguments that establish the relevance, believability and inferential force of evidence with respect to those hypotheses. The result of the testing process is the likelihood of the considered hypothesis (e.g., H_k : It is likely that a dirty bomb will be set off in the Washington DC area). Building such a probabilistic inference network helps answer the questions: Which evidential dots should be connected? Which evidential dots can be believed? What are the connections between evidential dots and hypotheses? and What does our argument mean?

3.2 Analysis of Competing Hypotheses

For each abducted hypothesis from the left hand side of Figure 1 (such as “the cesium-137 canister

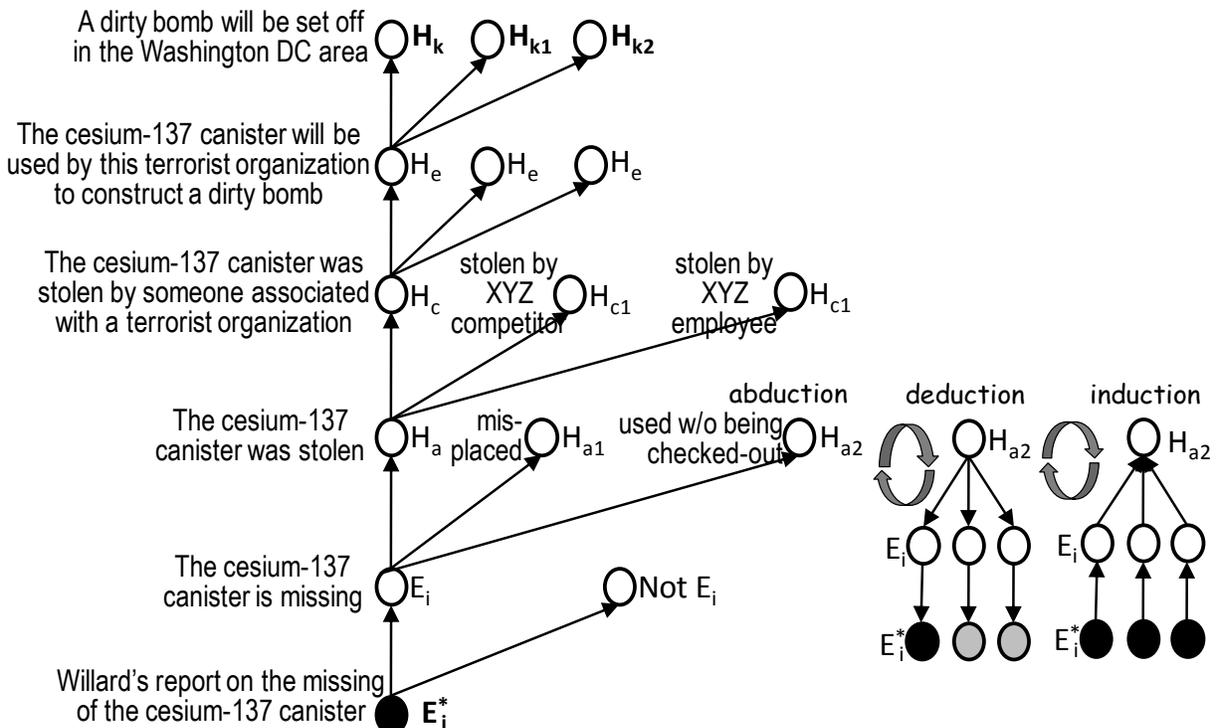


Figure 2. Analysis of competing hypotheses.

was stolen”) one would need to consider competing hypotheses (e.g., “the cesium-137 canister was misplaced” or “the cesium-137 canister was used in a project without being checked-out”), as illustrated in Figure 2. Moreover, for each such competing hypothesis one has to search for relevant evidence and use the evidence to test it, as discussed above and illustrated in Figure 2.

4. TIACRITIS: A Web-based System for Teaching Intelligence Analysis

The following sections present some of the basic capabilities of TIACRITIS, a web-based system for teaching intelligence analysts how to perform a rigorous analysis like the one outlined above.

4.1 Wigmorean Networks for Hypothesis Analysis

Let us consider again evidence E^*_i (Willard’s report on the missing of the cesium-137 canister from the company XYZ), shown at the bottom left of Figure 1. Just because an analyst has evidence of an event does not mean that the event actually happened. Thus, as indicated at the bottom of Figure 2, two hypotheses need to be tested: E_i (The cesium-137 canister is missing) and Not E_i (The cesium-137 canister is NOT missing).

As indicated by the tree from the middle of Figure 1, one needs to put the hypothesis of interest at work to guide in the collection of additional items of evidence that will be relevant to proving or disproving the hypothesis. Then the collected evidence will be used in assessing the likelihood of the hypothesis, as indicated in the right hand side of Figure 1.

Figure 3 and Figure 4 show two views of the corresponding reasoning tree generated by TIACRITIS for “Assessing whether the cesium-137 canister is missing from the XYZ warehouse.” TIACRITIS generated this tree by employing a general divide-and-conquer approach to problem solving, called problem-reduction/solution-synthesis, which has a grounding in the problem reduction representations developed in Artificial Intelligence (Nilsson, 1971; Tecuci, 1988; Tecuci 1998), and in the argument construction methods provided by the noted jurist John H. Wigmore (1937), the philosopher of science Stephen Toulmin (1963), and the evidence professor David Schum (1987; 2001). This approach uses expert knowledge to successively reduce a complex 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 tree from Figure 3 shows the step-by-step, top-down, reduction of the top level problem into simpler problems. This reduction tree is also shown in Figure 4, this time together with the corresponding synthesis tree. The synthesis tree shows the step-by-step composition of the solutions of the problems from the reduction tree, as will be explained in the following.

First, TIACRITIS decomposes the top problem from Figure 3 into three simpler problems, guided by a question and its answer, as also shown in Table 1. In this particular case, the three sub-problems are simple enough to be directly solved through evidence analysis. Thus one would need to look for evidence that is directly relevant to them. Let us focus on the second sub-problem (problem [P3]). As indicated in Figure 3, one would need to consider both the favoring evidence (e.g., problem [P5]) and

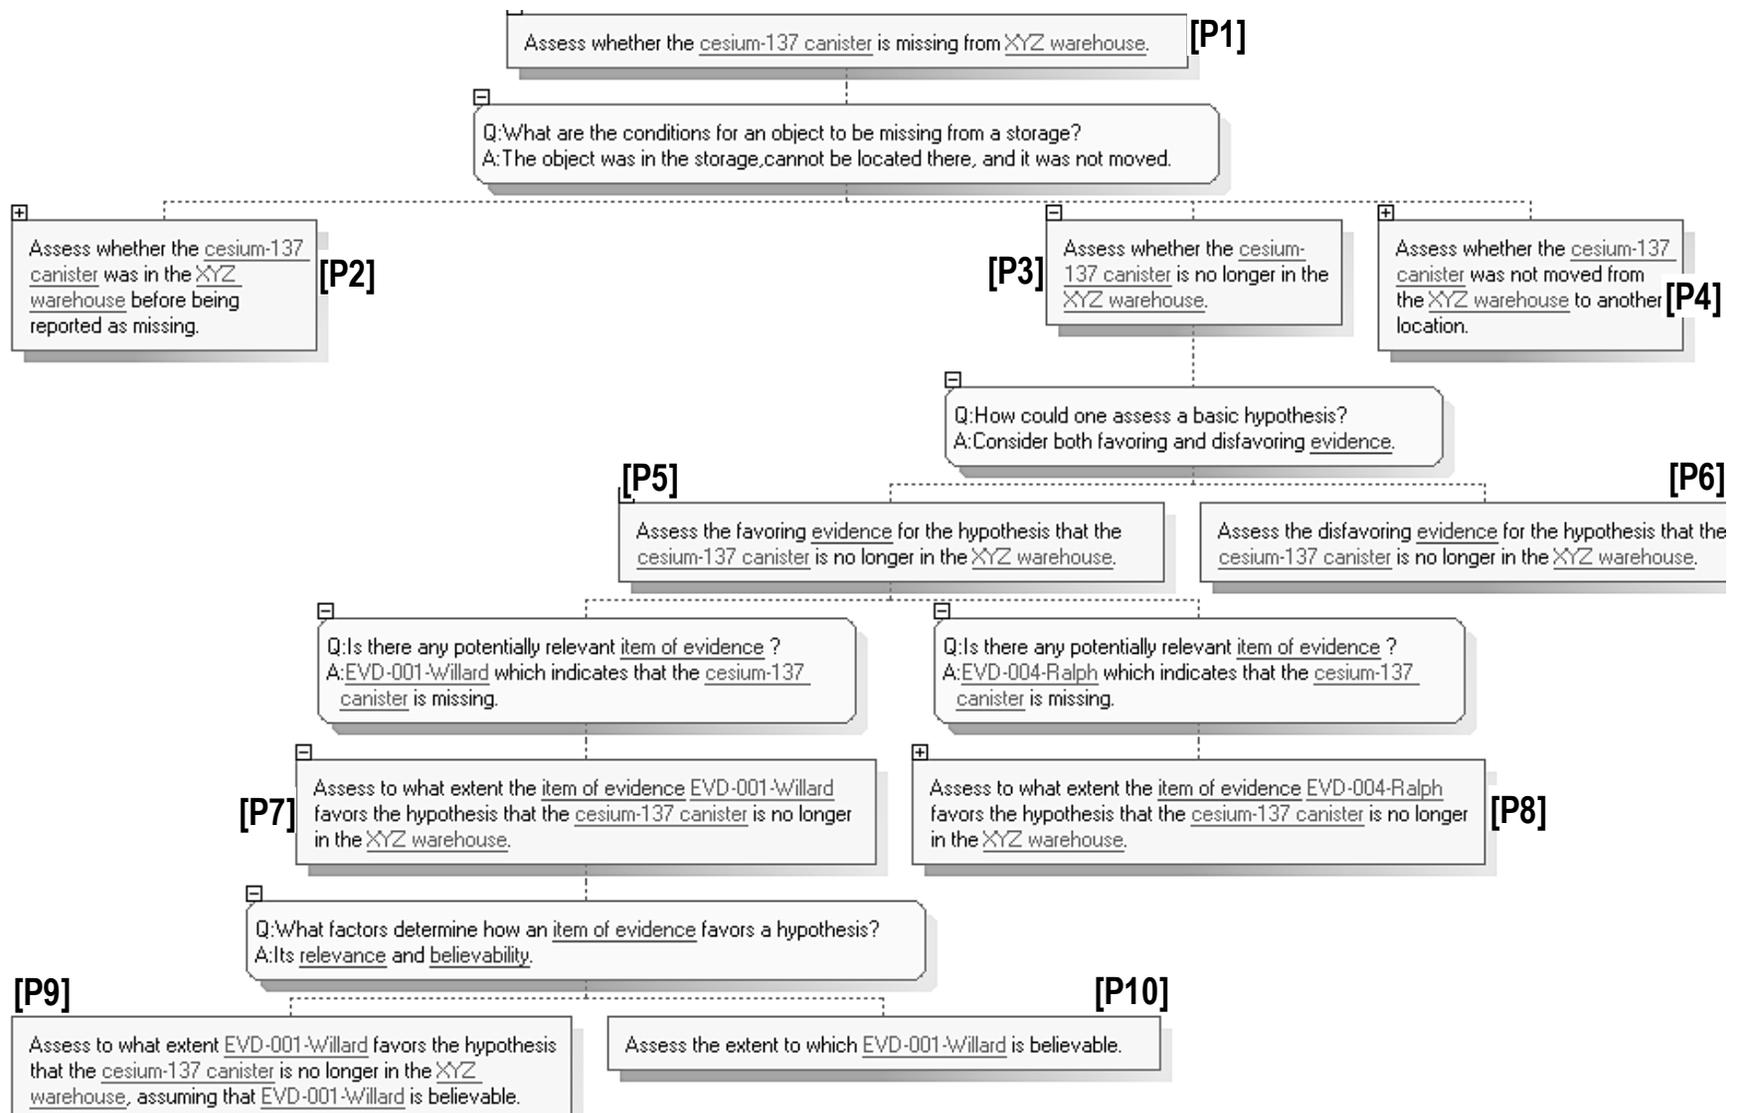


Figure 3. Reduction of a hypothesis analysis problem to simpler problems.

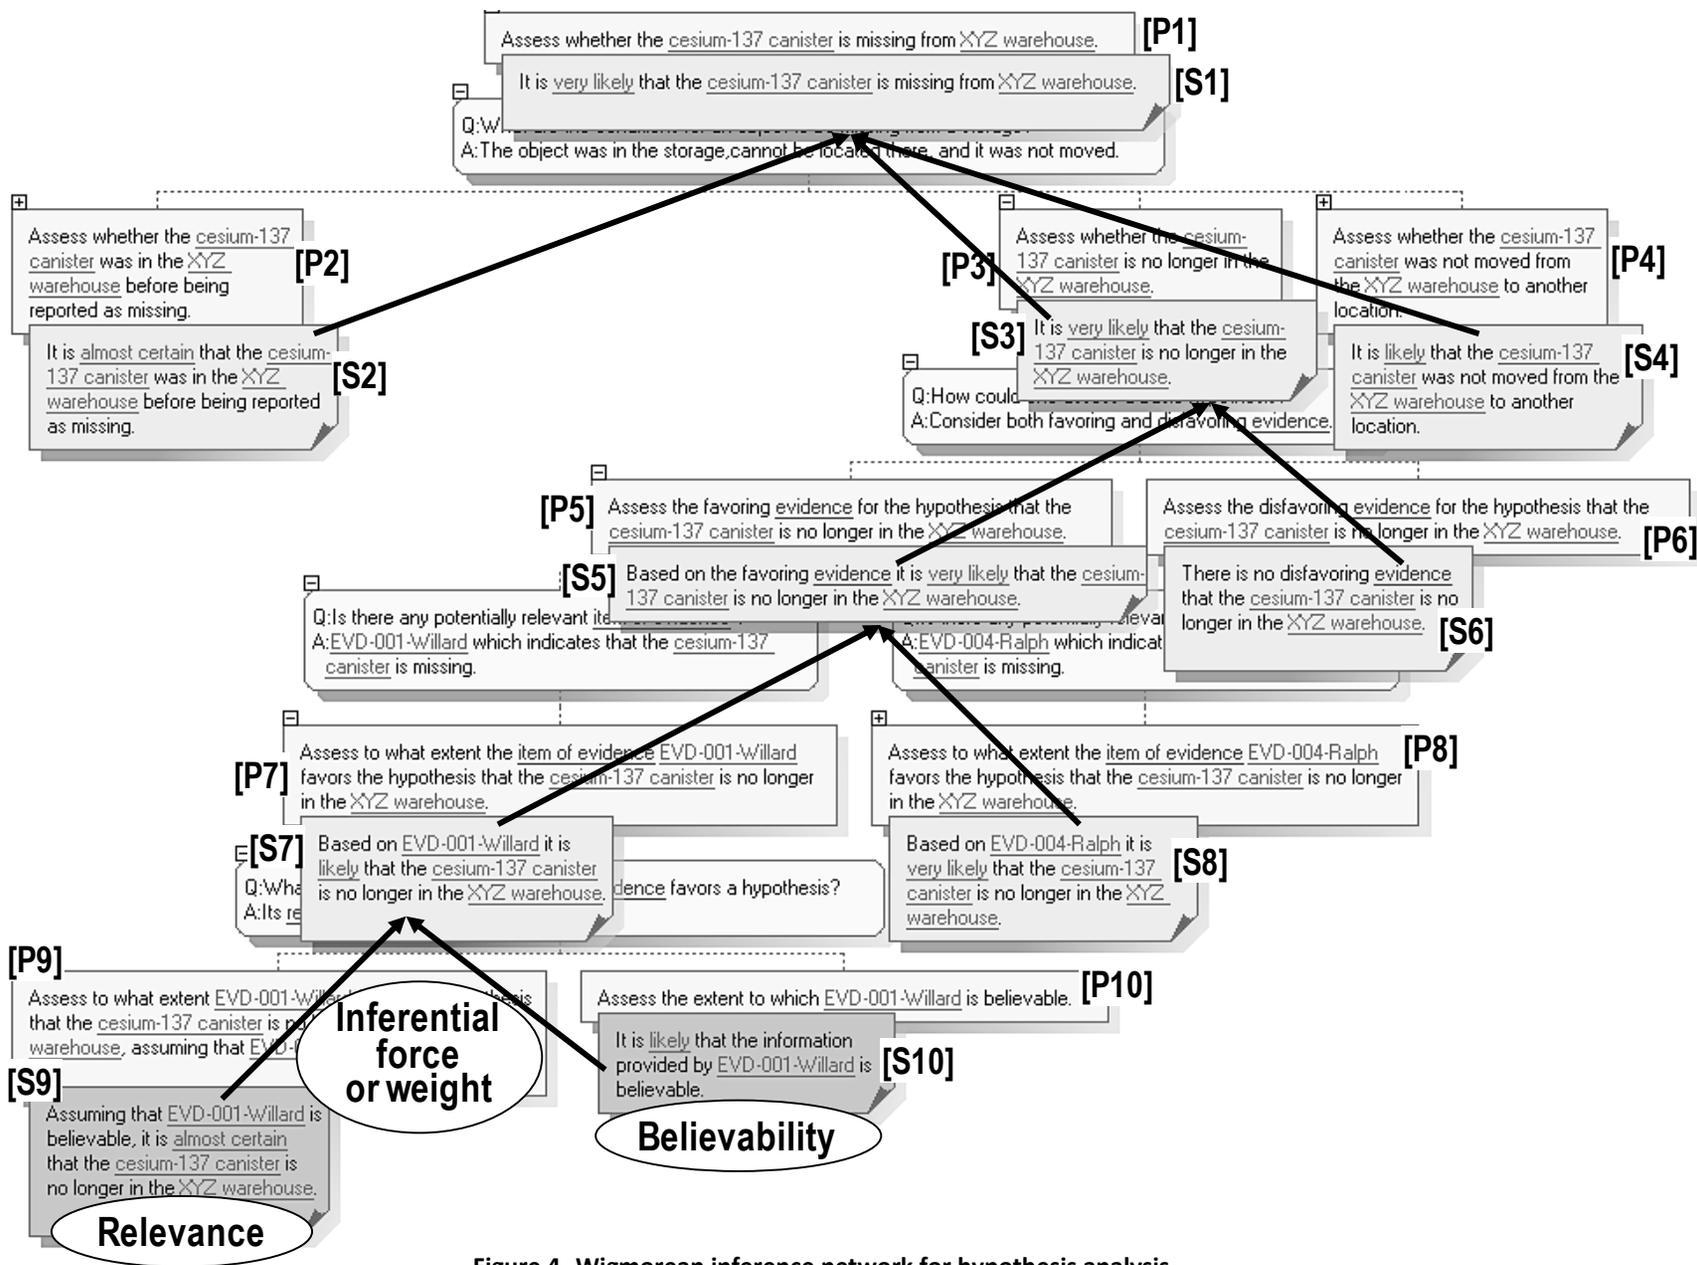


Figure 4. Wigmorean inference network for hypothesis analysis.

the disfavoring evidence (problem [P6]). That is, one would need to look for evidence that the cesium-137 canister was (or was not) in the warehouse, evidence that currently it is (or it is not) in the warehouse, and evidence that the cesium-137 canister was (or was not) moved.

Table 1. The top problem reduction step from Figure 3.

<i>I have to</i>	
Assess whether the cesium-137 canister is missing from XYZ warehouse.	[P1]
What are the conditions for an object to be missing from a storage?	
The object was in the storage, cannot be located there, and it was not moved.	
<i>Therefore I have to</i>	
Assess whether the cesium-137 canister was in the XYZ warehouse before being reported as missing.	[P2]
Assess whether the cesium-137 canister is no longer in the XYZ warehouse.	[P3]
Assess whether the cesium-137 canister was not moved from the XYZ warehouse to another location.	[P4]

This directs the analyst to contact Ralph, the supervisor of the warehouse, who reports that the cesium-137 canister is registered as being in the warehouse, that no one at the XYZ Company had checked it out, but it is not located anywhere in the hazardous materials locker. He also indicates that the lock on the hazardous materials locker appear to have been forced. Ralph’s testimony provides the analyst with several items of evidence, one of them being directly relevant to [P5]: “EVD-004-Ralph: Ralph’s report that the canister is not located anywhere in the hazardous materials locker.” This corroborates the previous evidence item E*, referred as EVD-001-Willard” in Figure 3. As shown in Figure 3, one needs to assess the extent to which each of these two items of evidence favors the hypothesis that the cesium-137 canister is no longer in the XYZ warehouse (i.e., problems [P7] and [P8]). The bottom part of Figure 3 and Table 2 show how problem [P7] is reduced to two problems, one for assessing the relevance of EVD-001-Willard, and the other for assessing its believability.

Table 2. The problem reduction step from the bottom of Figure 3.

<i>I have to</i>	
Assess to what extent the item of evidence EVD-001-Willard favors the hypothesis that the cesium-137 canister is no longer in the XYZ warehouse.	[P7]
What factors determine how an item of evidence favors a hypothesis?	
Its relevance and believability.	
<i>Therefore I have to</i>	
Assess to what extent EVD-001-Willard favors the hypothesis that the cesium-137 canister is no longer in the XYZ warehouse, assuming that EVD-001-Willard is believable.	[P9]
Assess the extent to which EVD-001-Willard is believable.	[P10]

The relevance answers the question: So what? How does this datum or item of information, whatever it is, bear on what an analyst is trying to prove or disprove? The believability answers the question: Can we believe what this item of intelligence information is telling us?

Assessing the relevance and the believability of an item of evidence may be done by employing the same problem reduction/solution synthesis approach illustrated here. However, in this case, the solutions of these assessments were provided by the analyst as assumptions, as indicated by the boxes with darker background at the bottom of Figure 4:

Assuming that EVD-001-Willard is believable, it is almost certain that
the cesium-137 canister is no longer in the XYZ warehouse. [S9]
It is likely that the information provided by EVD-001-Willard is believable. [S10]

By compositing the solutions [S9] and [S10] (e.g., through a “min” function) the agent assesses the inferential force or weight of EVD-001-Willard on the considered hypothesis:

Based on EVD-001-Willard it is likely that the cesium-137 canister is no longer
in the XYZ warehouse. [S7]

As illustrated above, the inferential force or weight answers the question: How strong is this item or body of relevant evidence in favoring or disfavoring various alternative hypotheses or possible conclusions being entertained?

Similarly the agent assesses the inferential force or weight of EVD-004-Ralph:

Based on EVD-004-Ralph it is very likely that the cesium-137 canister
is no longer in the XYZ warehouse. [S8]

By composing the solutions [S7] and [S8] (e.g., through a “max” function) the agent assesses the inferential force/weight of the favoring evidence (i.e., of EVD-001-Willard and EVD-004-Ralph):

Based on the favoring evidence it is very likely that the cesium-137 canister
is no longer in the XYZ warehouse. [S5]

Through a similar process one assesses the disfavoring evidence for the same hypothesis:

There is no disfavoring evidence that the cesium-137 canister is no longer
in the XYZ warehouse. [S6]

Composing [S5] and [S6] one obtains:

It is very likely that the cesium-137 canister is no longer in the XYZ warehouse. [S3]

The problems [P2] and [P4] are solved through a similar process:

It is almost certain that the cesium-137 canister was in the XYZ warehouse
before being reported as missing. [S2]

It is likely that the cesium-137 canister was not moved from the XYZ warehouse
to another location. [S4]

Then [S2], [S3], and [S4] are composed (e.g., through “average”) into the solution [S1] of [P1]:

It is very likely that the cesium-137 canister is missing from the XYZ warehouse. [S1]

4.2 Abstraction of Analysis

Analyses of complex hypotheses from masses of evidence result in the generation of very large reasoning trees with thousands of nodes. To help browse and understand such a complex analysis, TIACRITIS will display an abstraction of it which only shows abstractions of the main sub-problems considered in the analysis, as illustrated in the left-hand side of Figure 5.

The top line in the left-hand side of Figure 5 is the analyzed problem and its solution:

Assess whether the cesium-137 canister is missing from the XYZ warehouse: very likely.

Under it are the abstractions of its three main sub-problems and their solutions:

cesium-137 canister was in the XYZ warehouse: almost certain

cesium-137 canister is no longer in the XYZ warehouse: very likely

cesium-137 canister was not moved from XYZ warehouse: likely

These abstract problems can be expanded to browse their abstract sub-problems. Notice that this abstract view shows very clearly the main lines of reasoning. Thus, for “cesium-137 canister is no longer in the XYZ warehouse” TIACRITIS considers both favoring evidence and disfavoring evidence. It identified two items of favoring evidence, EVD-001-Willard and EVD-004-Ralph. The inferential force of EVD-004-Ralph (very likely) was obtained based on its relevance (almost certain) and its believability (very likely). Because EVD-004-Ralph is testimonial evidence based upon direct observation, its believability is given by the believability of its source which is Ralph. The believability of Ralph depends on his competence and credibility (Schum et al., 2009). Competence involves access (Has Ralph observed the event E himself or had access to the information he reports?) and understandability (Has Ralph understood what was being observed well enough to provide us with an intelligible account?). Credibility involves veracity (Does Ralph believe that event E occurred?), objectivity (Was Ralph’s belief based on his sensory evidence, or was it based on what Ralph expected or wished to observe?), and observational sensitivity (How good was the sensory evidence Ralph received under the conditions in which this observation was made?). TIACRITIS knows how to assess the believability of various types of items of evidence (Tecuci et al., 2009).

When the analyst clicks on a problem in the abstract view (e.g., “Credibility: very likely”), TIACRITIS displays the detailed reasoning for that problem, as illustrated in the right hand side of Figure 5. In particular, the analyst provided the solutions for Ralph’s veracity, objectivity, and observational sensitivity as assumptions. Then these solutions were automatically composed by TIACRITIS (through a min function) into an estimation of Ralph’s credibility.

⊖ Assess whether the cesium-137 canister is missing from XYZ warehouse: very likely

⊕ cesium-137 canister was in XYZ warehouse: almost certain

⊖ cesium-137 canister is no longer in XYZ warehouse: very likely

⊖ favoring evidence: very likely

⊕ EVD-001-Willard: likely

⊖ EVD-004-Ralph: very likely

..... relevance: *almost certain*

⊖ believability: very likely

⊖ Source Ralph: very likely

⊖ Competence: almost certain

..... Access: *almost certain*

..... Understandability: *almost certain*

⊖ **Credibility: very likely**

⊕ Veracity: *almost certain*

..... Objectivity: *very likely*

..... Observational sensitivity: *almost certain*

..... disfavoring evidence: no evidence

⊕ cesium-137 canister was not moved from XYZ warehouse: likely

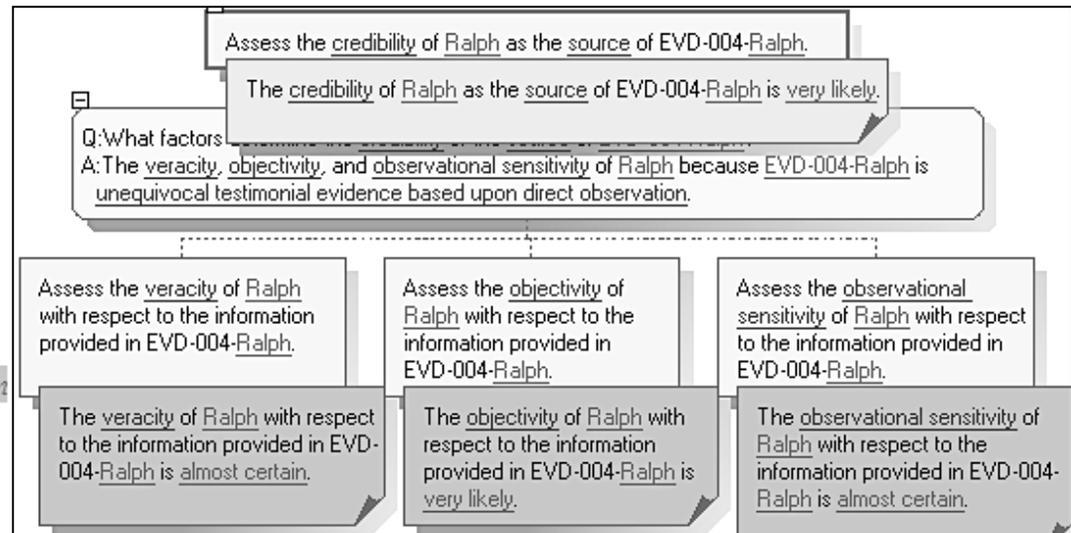


Figure 5. Abstract reasoning and detailed reasoning.

4.3 Drill-Down Analysis and Assumption-based Reasoning

There is an important virtue of TIACRITIS that bears directly on the difficulties analysts so frequently face. This system allows the user to decompose an analysis to different levels of refinement. Stated another way, this system allows the user to "drill down" to various levels in order to reach conclusions about necessary analytic ingredients. In some cases, as in current analyses, analysts will not have the time to deal with all of the complexities their own experience and TIACRITIS makes evident. In other situations, analysts will not have access to the kinds of information necessary to answer all questions regarding elements of an analysis that seem necessary. In such instances TIACRITIS provides mechanisms necessary to define assumptions of various sorts and to show the extent to which conclusions rest upon them.

Consider, for example, the problem [P10], "Assess the extent to which EVD-001-Willard is believable", at the bottom of Figure 4. In this case the analyst provided its solution (i.e., [S10]) as an assumption: "It is likely that the information provided by EVD-001-Willard is believable." In the previous section, however, we have shown how one could actually assess the believability of an item of evidence (e.g., EVD-004-Ralph) in terms of lower-level believability credentials. In fact, one could go even deeper than illustrated in Figure 5, by further reducing the assessment of Ralph's veracity, objectivity and observational sensitivity (Schum and Morris, 2007; Schum et al., 2009b). It will be up to the analyst to decide how deep he or she would like to drill-down in such an analysis.

The analyst could also experiment with alternative assumptions. For example, he or she may also assume that Ralph's veracity has a different value (e.g., unlikely), and determine how this new assumptions changes the analysis which is automatically updated by TIACRITIS.

5. Teaching and Learning with TIACRITIS

TIACRITIS and the associated textbook have been designed to be helpful in intelligence analysis courses that are necessary, or frequently advisable, for persons throughout the Intelligence Community and those it serves. This includes collectors of intelligence information, evaluators of incoming intelligence information at various levels and in different offices, and even the policy-making "customers" of intelligence analysts. Everyone in the business of evaluating and acting upon intelligence information should be knowledgeable regarding the many subtleties and complexities of reasoning based on evidence. Moreover, what is being taught is 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.

TIACRITIS may be used both as a teaching tool by the instructor, and as a learning tool by the students. The instructor can use it to explain and illustrate, with actual examples generated by TIACRITIS, all the topics discussed in the intelligence analysis course. The student can use it to practice intelligence analysis, as discussed in the next section.

5.1 Learning-by-Doing

One of the main goals of TIACRITIS is to enable a learning-by-doing approach through regular practice involving analyses of evidence using either hypothetical situations or examples drawn from actual situations. To illustrate some of the exercises for this practice let us continue with our example from the previous sections where we have already established that the cesium-137 canister is missing (i.e., E_i in Figure 2). The next step is to consider the competing hypotheses H_a , H_{a1} , and H_{a2} .

We can provide the students with additional information, represented in Table 3, which may help them assess these hypotheses.

Table 3. Additional information on the mission of the cesium-137 canister.

Grace, the Vice President for Operations at XYZ, tells us that no one at the XYZ Company had checked out the canister for work on any project. She says that the XYZ Company had other projects involving hazardous materials but none that involved the use of cesium-137.

We have asked a professional locksmith named Clyde who said that the lock had been forced, but it was a clumsy job.

There is a security perimeter around the XYZ warehouse and employee parking area having just one gate that is controlled by a guard. On the day before the missing canister was observed, the security guard Sam recorded that a panel truck having Maryland license MDC-578 was granted entry at 4:45PM just before the XYZ closing hour at 5:00PM. The driver of this vehicle showed the guard a manifest containing items being delivered to the XYZ warehouse. This manifest contained a list of packing materials allegedly ordered by the XYZ Company. The vehicle was allowed to enter the parking area. But at 8:30PM this same vehicle was allowed to exit the parking area. A different guard was on duty in the evenings and noticed that his records showed that this vehicle had been permitted entry and so he allowed the vehicle to exit the parking area.

The panel truck carrying the license plate # MD-578 is registered in the name of a truck-rental company called TRUXINC, that is located in Silver Spring MD. The manager of this agency showed records indicating that this truck was rented to a person who gave his name as Omer Riley, having as his listed address: 6176 Williams Ave. in Silver Spring. The truck was rented on the day before Willard's discovery of the missing cesium-137, and it was returned the day after he made his discovery.

There is no residence at 6176 Williams Ave. in Silver Spring, MD.

An examination of the panel truck rented by Omer Riley revealed minute traces of cesium-137.

In a simpler exercise the students would have to identify the “dots” in the text from Table 3 which are fragments that represent relevant items of evidence for the considered hypotheses. Some of these dots are the following ones:

EVD-006-Grace: Grace’s report that there is no current project that uses cesium-137.

EVD-007-Clyde: Locksmith Clyde’s report that the lock was forced.

EVD-008-GardReport: XYZ Company security guard record that a panel truck bearing Maryland license plate # MDC-578 left the XYZ parking area at 8:30PM on the day before Willard's discover of the missing cesium-137 canister.

After identifying the dots, the students would need to associate them to the various nodes in the Wigmorean networks corresponding to the analyses of the considered hypotheses. For example, Figure 6 shows the Wigmorean network for the hypothesis analysis problem “Assess whether cesium-137 canister was stolen with the MDC 578 truck”, as well as some of the new evidence associated with sub-

hypotheses in this network (e.g., EVD-007-Clyde associated as favoring evidence to the hypothesis that the hazardous materials locker was forced).

After the association of all the relevant items of evidence and the definition of all the necessary assumptions (e.g., the relevance of EVD-007-Clyde is very likely, the believability of EVD-007-Clyde is likely), the students can inspect the analyses of the various hypotheses. TIACRITIS may then compare these analyses in terms of their evidential support, pointing to areas where additional evidence would be desirable. TIACRITIS also allows a comparison between the analysis of a hypothesis performed by one student and the analysis of the same hypothesis performed by a different student, to uncover differences in evidence used and assumptions made.

A more complex version of this exercise will also include the evaluation of the believability of various items of evidence (e.g., demonstrative tangible evidence, testimonial evidence based upon direct observation, testimonial evidence obtained at second hand, etc.), based on their characteristic credentials (e.g., veracity, objectivity, accuracy, etc.), as discussed in (Tecuci et al., 2009).

Yet another version of the exercise would ask the students to experiment with what-if scenarios, by changing the assumptions and studying how this changes the evaluation performed by TIACRITIS.

To simulate the evolution of the world, students may then be given additional items of evidence that would need to be integrated into the reasoning trees.

In the above exercises, TIACRITIS reduced the considered hypothesis to a set of “elementary” hypotheses with which the students had to associate the items of evidence. In more complex versions of these exercises, TIACRITIS does not reduce completely the initial hypothesis (or does not reduce it at all). Thus, the students need to reduce the hypothesis or the sub-hypotheses provided by TIACRITIS, before associating the items of evidence. An even more complex exercise does not even provide the initial hypothesis, asking the students to abduce it by studying the dots (the items of evidence). More complex versions may ask the students to abduce several hypotheses.

Yet an even more complex exercise would formulate a real world problem, such as “Assess whether Algeria supports Madani Bin Massie”, and ask the students to look for evidential dots on the Internet, solve the problem and compare their solutions with the solutions provided by their colleagues and by TIACRITIS.

TIACRITIS may also generate test questions for assessing students’ knowledge of evidence-based reasoning. For example, TIACRITIS may generate a small fragment of a Wigmorean network and ask the student whether the reasoning is complete (e.g., includes all the necessary sub-problems of a reduced problem), incomplete (misses some sub-problems but the present ones are correct), or incorrect (includes incorrect sub-problems). Alternatively, TIACRITIS may show a problem and a set of potential sub-problems, asking the student to select the correct sub-problems.

One important thing is that, with TIACRITIS, one can define any combination of these exercises. No traditional book can match this capability.

⊖ Assess whether the cesium-137 canister was stolen with the MDC-578 truck: very likely

⊖ hazardous materials locker forced: likely

⊖ favoring evidence: likely

⊖ EVD-005-Ralph: likely

relevance: *likely*

believability: *almost certain*

⊖ EVD-007-Clyde: likely

relevance: *very likely*

believability: *likely*

disfavoring evidence: no evidence

⊖ MDC-578 truck was inside XYZ company: almost certain

⊖ favoring evidence: almost certain

⊖ EVD-008-GuardRecord: almost certain

relevance: *almost certain*

believability: *almost certain*

disfavoring evidence: no evidence

⊖ MDC-578 truck transported cesium-137 canister: likely

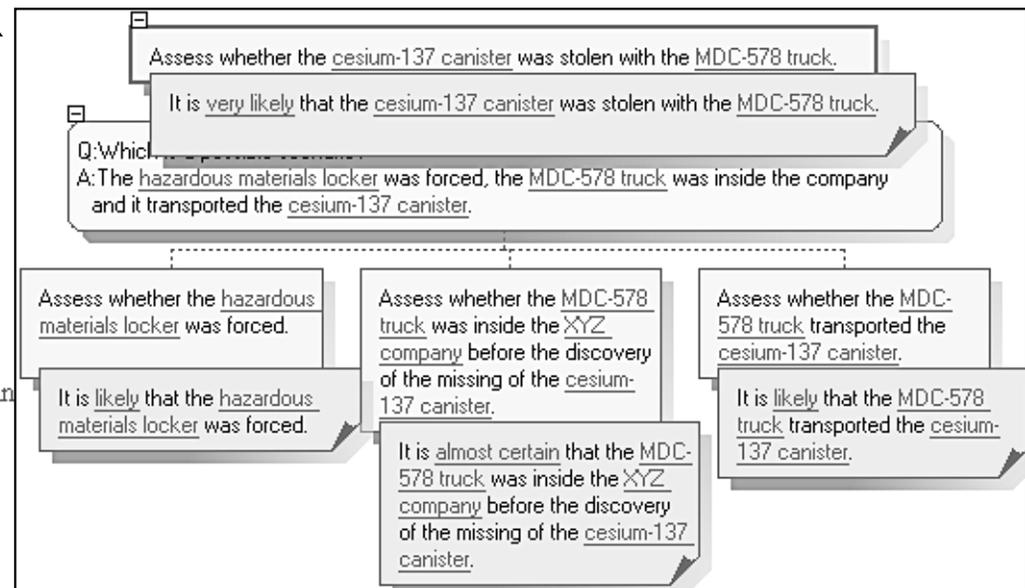


Figure 6. Associating evidence to hypotheses.

5.2 Customized Instruction with TIACRITIS

Another important capability of TIACRITIS and the associated textbook is that they allow personalization to a specific organization and student population by selecting the topics to be taught in the course, as well as the case studies and knowledge bases with problems that correspond to the type of problems that will be faced by these future analysts or users of intelligence.

To illustrate this capability, we will assume that TIACRITIS has a knowledge repository with a wide variety of knowledge bases and case studies. Let us further consider an instructor from a specific organization, such as Joint Forces Staff College or US Army Intelligence Center, who needs to teach a specific intelligence analysis course to a specific population of students. First the instructor will interact with TIACRITIS to select the topics to be taught in the course. The instructor will also select a set of application domains that will be of interest to the students, and a set of case studies in those domains. The instructor can also select types of exercises and the desired levels of difficulties for these exercises.

Let us now consider the actual use of the personalized TIACRITIS system in the planned course. The instructor will discuss a specific topic (which is presented in the textbook) and will guide the students through hand-on practice with TIACRITIS. After that, however, each student may further personalize his or her copy of TIACRITIS by selecting various characteristics, such as application environment (e.g., matters in Iraq as opposed to China) or type of problem (e.g., tactical versus strategic intelligence, or explaining past events versus predicting future ones). This, of course, assumes that TIACRITIS has the corresponding knowledge bases. But such knowledge bases can be rapidly developed with Disciple-LTA.

TIACRITIS is being transitioned to the Joint Forces Staff College. A course that will use it is a 4-week focus study that meets one day a week, for two hours. The current course provides traditional classroom instruction on analytic tools and techniques, critical thinking, assessing risk, decision making and practical exercises, etc. The plan is to use the immersive educational environment furnished by TIACRITIS, based on specially developed knowledge bases with case studies and exercises.

6. Conclusions

As indicated in Section 1, TIACRITIS is based on the stand-alone Disciple-LTA system. As opposed to Disciple-LTA, TIACRITIS is accessible on the web and was designed specifically to be used for teaching intelligence analysts. Special attention has been directed to its ease of use and to the development of those capabilities that are most useful for analysts' teaching and learning. The current version of TIACRITIS does not include modules for the development of knowledge bases and case studies, which have to be done with Disciple-LTA.

We are planning on transitioning TIACRITIS to as many IC and DOD organizations as possible and, in the process, significantly extending its repository of knowledge bases, case studies, and exercises. One of our goals is to continuously increase the training effectiveness of TIACRITIS. Another goal is to extend TIACRITIS with capabilities for agent learning and knowledge base development. They include both the

development of the capabilities currently present in Disciple-LTA, and the development of new capabilities. Yet another goal is to evolve TIACRITIS into a system that can be used in an operational environment because training with TIACRITIS is most effective when this system is also part of the analytic work environment.

7. Acknowledgements

Cristina Boicu has contributed to the development of TIACRITIS. This material is based on research performed in the Learning Agents Center and partially supported by several agencies of the U.S. Government, including the Department of Defense, the National Geospatial-Intelligence Agency, the Air Force Office of Scientific Research (FA9550-07-1-0268) 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.

We are very grateful to Phillip Hwang, Don Kerr, Joan McIntyre, Kelcy Allwein, Keith Anthony, Cindy Ayers, Susan Durham, Sharon Hamilton, Jim Homer, David Luginbuhl, and George Stemler for their suggestions and support.

8. REFERENCES

Anderson T., Schum D., Twining W., *Analysis of Evidence*, Cambridge University Press, 2005.

Jackson, P., *Introduction to Expert Systems*, Addison-Wesley, Essex, England, 1999.

MacEachin D., Forward to "Heuer R.J., *Psychology of Intelligence Analysis*," Center for the Study of Intelligence, Central Intelligence Agency, 1999.
<http://www.au.af.mil/au/awc/awcgate/psych-intel/art2.html>

Nilsson, N.J., *Problem Solving Methods in Artificial Intelligence*, NY: McGraw-Hill, 1971.

Peirce, C. S., *Reasoning and the Logic of Things*, 1898. In Ketner, K. (Ed.), Harvard University Press, Cambridge, MA, 1992.

Peirce, C. S., *Abduction and Induction*, 1901. In *Philosophical Writings of Peirce*. Buchler, J. (Ed.), Dover, New York, NY, 1955, pp 150-156.

Schum D.A., *Evidence and Inference for the Intelligence Analyst* (2 Vols), University Press of America, MD: Lanham, 1987.

Schum D.A., *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press, 2001.

- Schum D.A., Lessons and Stories about Concepts Encountered in Disciple-LTA, *Research Report 2*, Learning Agents Center, George Mason University, August 2007. http://lac.gmu.edu/publications/2007/SchumD_Concepts_in_Disciple_LACRR.pdf
- Schum, D., Morris, J., Assessing the Competence and Credibility of Human Sources of Evidence: Contributions from Law and Probability, *Law, Probability and Risk*, Vol. 6, pp 247-274, 2007.
- Schum D., Tecuci G, Boicu M., Analyzing Evidence and its Chain of Custody: A Mixed-Initiative Computational Approach, *International Journal of Intelligence and Counterintelligence*, Vol. 22, pp 298-319, 2009a. <http://lac.gmu.edu/publications/2009/Schum et al - Chain of Custody.pdf>
- Schum, D., Tecuci, G., Boicu, M., Marcu, D., Substance-Blind Classification of Evidence for Intelligence Analysis, in *Proceedings of the Conference "Ontology for the Intelligence Community,"* George Mason University, Fairfax, VA, 20-22 October 2009b. <http://lac.gmu.edu/publications/2009/Schum-OIC09.pdf>
- Tecuci G., *Disciple: A Theory, Methodology and System for Learning Expert Knowledge*, Thèse de Docteur en Science (in English), University of Paris-South, France, 1988.
- Tecuci, G., *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, Academic Press, 1998. http://lac.gmu.edu/publications/1998/TecuciG_Building_Intelligent_Agents/default.htm
- Tecuci G., Boicu M., Ayers C., Cammons D., Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving, in *Proceedings of the First International Conference on Intelligence Analysis*, McLean, VA, 2-6 May, 2005. <http://lac.gmu.edu/publications/data/2005/Tecuci-Disciple-LTA.pdf>
- Tecuci G., Boicu M., Marcu D., Boicu C., Barbulescu M., Ayers C., Cammons D., Cognitive Assistants for Analysts, *Journal of Intelligence Community Research and Development*, 2007. http://lac.gmu.edu/publications/2007/TecuciG_Cognitive_Assistants.pdf
- Tecuci G., Boicu M., Marcu D., Boicu C., Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development*, 2008. <http://lac.gmu.edu/publications/2008/Disciple-LTA08.pdf>
- Tecuci, G., Boicu, M., Schum, D., Marcu, D., Overcoming Intelligence Analysis Complexity with Cognitive Assistants, *Research Report 7*, 45 pages, Learning Agents Center, August 2009, updated October 2009. http://lac.gmu.edu/publications/2009/Tecuci-Overcoming_IA_Complexity.pdf
- Toulmin, S. E., *The Uses of Argument*, Cambridge University Press, 1963.
- Wigmore, J.H., *The Science of Judicial Proof*, Little, Brown & Co., Boston, MA, 1937.

9. About the Authors

Gheorghe Tecuci is Professor of Computer Science in the Volgenau School of Information Technology and Engineering and Director of the Learning Agents Center at George Mason University (<http://lac.gmu.edu>). He is also Visiting Professor and former Chair of Artificial Intelligence at the US Army War College. Dr. Tecuci has followed a career-long interest in the development of a computational theory and technology that allows non-computer scientists to develop intelligent assistants that incorporate their problem solving expertise. He has published over 175 papers.

David Schum holds the rank of professor in the Systems Engineering and Operations Research Department in the Volgenau School of Information Technology and Engineering, and in the George Mason University School of Law. In 2003 he was made an honorary member of the faculty of University College London where he is Honorary Professor of Evidence Science. He is the author of two books and a co-author of three books on evidence and probabilistic reasoning. He has published many articles and monographs in a variety of journals and has given many plenary session talks at conferences.

Mihai Boicu is Assistant Professor of Information Technology in the Volgenau School of Information Technology and Engineering and Associate Director of the Learning Agents Center. He has expertise in knowledge engineering, software engineering, artificial intelligence, machine learning, multi-agent systems and human-computer interaction. He has published over 70 papers in these areas.

Dorin Marcu is Research Assistant Professor in the Learning Agents Center of George Mason University. Dr. Marcu received his Ph.D. in Computer Science from George Mason University. His research interests include mixed-initiative human computer interaction, assumption-based reasoning and instructable agents. He has published around 40 papers.

Benjamin D. Hamilton is a Senior Program Analyst and Science Engineering and Technical Assistance (SETA) contractor to the U.S. Government. His expertise includes analyzing, designing, evaluating, and managing projects involving interactive instruction, performance improvement, and simulation/game-based training. He received his Masters degree in Instructional Technology from Bloomsburg University and his Doctorate of Education in Instructional Technology and Distance Education from Nova Southeastern University.

Benjamin F. Wible is a Visiting Professor and Chair for the National Security Agency at the Joint Forces Staff College in Norfolk, Virginia. He has over 25 years experience within the Intelligence Community specifically with the domains of intelligence analysis, cognitive science, and systems engineering. He is a graduate of the Joint Advance Warfare School, where he received a MS in Strategy and Planning. He is also a graduate of the National Defense Intelligence College, where he received a MS in Strategic Intelligence. He is currently writing his PHD dissertation on human cognition and decision making.