# Computational Theory and Cognitive Assistant for Intelligence Analysis

Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, David Schum, Katherine Russell
Learning Agents Center, Volgenau School of Engineering, George Mason University, Fairfax, VA 22032

***Abstract*— This paper presents elements of a computational theory of intelligence analysis and its implementation in a cognitive assistant. Following the framework of the scientific method, this theory provides computational models for essential analysis tasks: evidence marshaling for hypotheses generation, hypotheses-driven evidence collection, and hypotheses testing through multi-INT fusion. Many of these models have been implemented in a web-based cognitive assistant that not only assists an analyst in coping with the astonishing complexity of intelligence analysis, but it also learns from their joint analysis experience.**

***Intelligence analysis, scientific method, cognitive assitant, evidence-based reasoning, mixed-initiative reasoning, discovery, ontology, rules, learning, evidence collection, hypotheses testing***

## I. INTRODUCTION

The purpose of Intelligence Analysis is to answer questions arising in the decision-making process. Often stunningly complex arguments, involving both *imaginative and critical reasoning*, are necessary in order to establish and defend the *relevance*, the *believability*, and the *inferential force* of evidence with respect to the questions asked. The answers are necessarily probabilistic in nature because evidence is always *incomplete* (we can look for more, if we have time), usually *inconclusive* (it is consistent with the truth of more than one answer), frequently *ambiguous* (we cannot always determine exactly what the evidence is telling us), commonly *dissonant* (some of it favors one answer but other evidence favors other answers), and has various degrees of *believability* shy of perfection [1, 2]. Not only is this process highly complex, but it often needs to be performed in a very short period of time.

Given these characteristics of intelligence analysis, we believe that it can be best performed through the mixed-initiative integration of human imagination and computer knowledge-based reasoning [3]. To this purpose we are developing a *Computational Theory of Intelligence Analysis* which is grounded in the science of evidence [4], artificial intelligence, logic, and probability. This theory provides computational models for essential analysis tasks: evidence marshaling for hypotheses generation, hypotheses-driven evidence collection, and hypotheses testing through multi-INT fusion. Many of these models have already been implemented in the TIACRITIS web-based cognitive assistant. The first version of TIACRITIS was developed to help intelligence analysts learn critical thinking skills for evidence-based reasoning, through a hands-on approach, based on predefined analysis cases [2, 5]. That version has now been significantly extended with new capabilities that allow intelligence analysts to formulate and analyze their own hypotheses, and also to learn from the performed analyses.

This paper provides an overview of the current status of the computational theory of intelligence analysis, and its implementation in the extended version of TIACRITIS.

## II. INTELLIGENCE ANALYSIS AS CEASELESS DISCOVERY OF EVIDENCE, HYPOTHESES, AND ARGUMENTS

Within the framework of the scientific method, we view intelligence analysis as ceaseless discovery of evidence, hypotheses, and arguments in a non-stationary world. It involves a collaborative process of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary testing of hypotheses (see Fig. 1). Through *abductive reasoning* (which shows that something is *possibly* true) we generate hypotheses from our observations; through *deductive reasoning* (which shows that something is *necessarily* true) we use our hypotheses to generate new lines of inquiry and discover new evidence; and through *inductive reasoning* (which shows that something is *probably* true) we test our hypotheses with the discovered evidence. Therefore, in this paper we will illustrate the discovery of evidence, hypotheses, and arguments with an analysis example, and then we will show how the same analysis is performed with TIACRITIS.

In our analysis example, Mavis, a counterterrorism analyst, reads in today's Washington Post that a canister containing cesium-137 is missing from the warehouse of the Company XYZ in MD (see evidence E at the bottom-left of Fig. 2). The

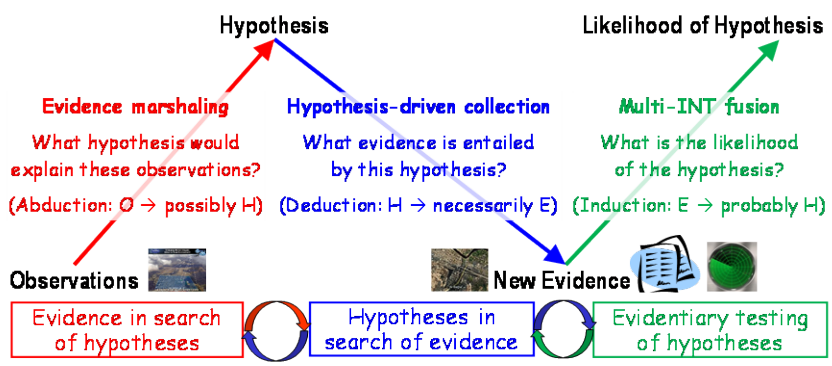


Figure 1. Framework of the Computational Theory of Intelligence Analysis.

This research was partially supported by the National Geospatial-Intelligence Agency, by the Department of Defense, and by George Mason University. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Defense or the U.S. Government.

question is: *What hypothesis would explain this observation?*

Through *imaginative reasoning*, Mavis *abductively* infers that a dirty bomb will be set off in the Washington, DC area. However, no matter how imaginative or important this hypothesis is, no one will take it seriously unless Mavis and her cognitive assistant, TIACRITIS, are able to justify it. So they develop the chain of abductive inferences shown in the left hand side of Fig. 2. We have evidence that the cesium-137 canister is missing (E). Therefore it is possible that it is indeed missing ($H_1$). It is possible that it was stolen ($H_2$). It is possible that it was stolen by someone associated with a terrorist organization ($H_3$). It is possible that the terrorist organization will use the cesium-137 canister to build a dirty bomb ($H_4$). It is possible that the dirty bomb will be set off in the Washington, DC area ($H_5$).

But these are not the only hypotheses that explain E. Just because there is evidence that the cesium-137 canister is missing does not mean that it is indeed missing. At issue here is the believability of the source of this information. Thus an alternative hypothesis is that the cesium-137 canister is not missing ($H'_1$). But let us assume that it is missing. Then it is possible that it was stolen ($H_2$). But it is also possible that it was misplaced ($H'_2$), or maybe it was used in a project at the XYZ Company ($H''_2$). But let us suppose that it was stolen ($H_2$). Then it is possible that it was stolen by someone associated with a terrorist organization ($H_3$). But it is also possible that it was stolen by a competitor ($H'_3$), or maybe it was stolen by an employee ($H''_3$), and so on. This is the process of *evidence in search of hypotheses* that would explain it.

The analyst and TIACRITIS need to assess each of these hypotheses before they can conclude that a dirty bomb will be set off in the Washington, DC area. During this process, they would also need to discover who will set off the dirty bomb, and where and when it would be set off.

Starting with $H_1$, each hypothesis is deductively put to work to guide the collection of additional evidence (see the blue tree in the middle of Fig. 2). Assuming that the cesium-137 canister is indeed missing ($H_1$), what other things should be observable? Which are the necessary conditions for an object to be reported as missing from a warehouse? It was in the warehouse ($H_{11}$), it is no longer there ($H_{12}$), and no one has checked it out ($H_{13}$). This leads Mavis 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 appears to have been forced (see bottom right of Fig. 2). Ralph's testimony provides several items of evidence which are relevant for the hypotheses $H_{11}$, $H_{12}$, and $H_{13}$. This is *hypothesis in search of evidence* that guides the analyst in collecting new evidence.

Mavis and TIACRITIS have now collected more relevant evidence, and the question is: What is the likelihood that the cesium-137 canister is missing, based on the available evidence? To answer this question, they build a Wigmorean probabilistic inference network that shows how the evidence is *fused* through an argument that establishes its relevance, its believability, and its inferential force on the intermediate hypotheses $H_{11}$, $H_{12}$, and $H_{13}$ and on the top-level hypothesis H. They conclude that it is very likely the cesium-137 canister

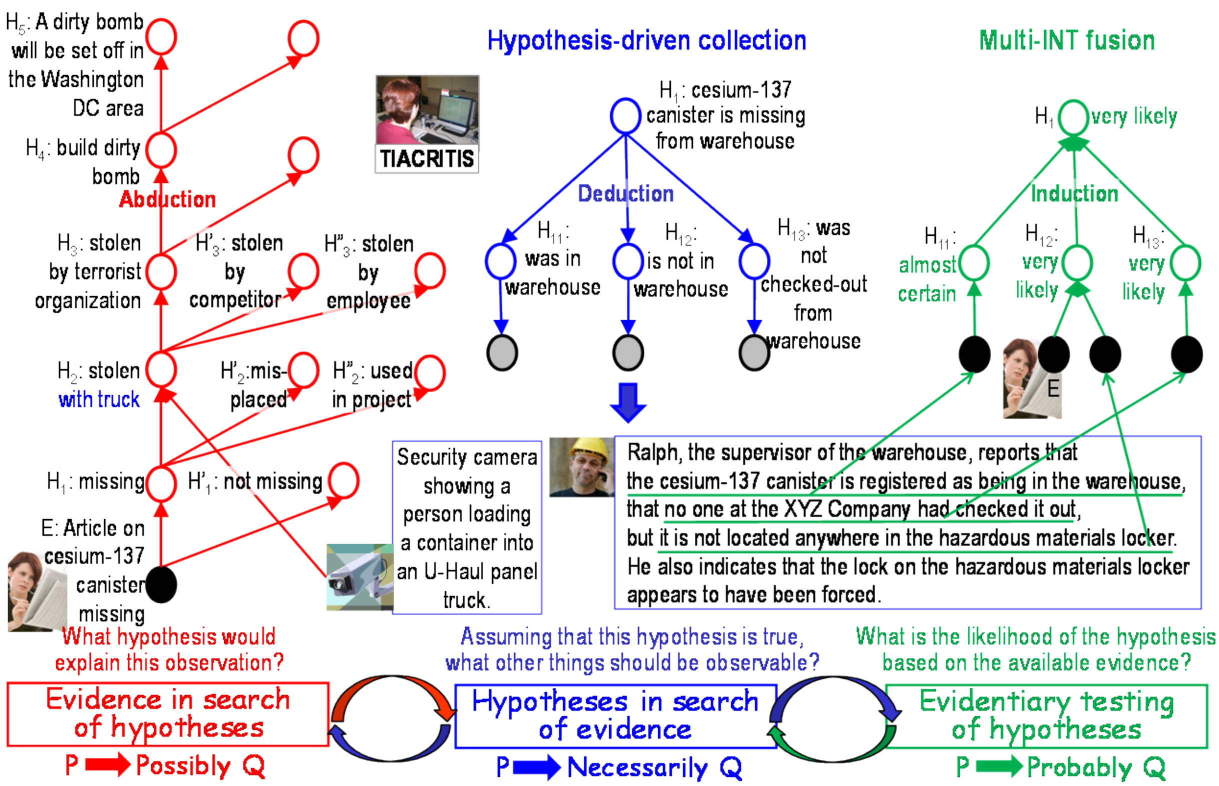


Figure 2. Discovery of evidence, hypotheses, and arguments.

is missing (see the green tree in the right hand side of Fig. 2).

Now, some of the newly discovered items of evidence may trigger new hypotheses, or the refinement of the current hypotheses. Therefore these processes of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary testing of hypotheses, take place at the same time, and in response to one another, as indicated by the arrows at the bottom of Fig. 2. For example, during her investigation of the security camera of the XYZ warehouse, Mavis discovers a video segment showing a person loading a container into a U-Haul panel truck. Therefore the hypothesis $H_2$ is refined to "the cesium-137 canister was stolen with the U-Haul panel truck" (see the left part of Fig. 2).

Having concluded that the cesium-137 canister is missing, Mavis and TIACRITIS now have to establish whether the cesium-137 canister was stolen with a truck ($H_1$), misplaced ($H'_1$), or used in some project ($H''_1$). Each of these hypotheses is put to work to guide the collection of relevant evidence which is then used to assess it, as illustrated in Fig. 3.

Assuming that the cesium-137 canister was stolen with a truck ($H_2$), what other things should be observable? The current evidence suggests the following scenario of how the cesium-137 might been stolen: The truck entered the company, the canister was stolen from the locker, the canister was loaded into the truck, and the truck left with the canister (see the blue tree in the right side of Fig. 3). Such scenarios have enormous heuristic value in advancing the investigation because they consist of mixtures of what is taken to be factual and what is conjectural. Conjecture is necessary in order to fill in natural gaps left by the absence of evidence. Each such conjecture opens up a new avenue of investigation, and the discovery of additional evidence, if the scenario turns out to be true. In this case, for instance, Mavis is led to check whether the truck entered the XYZ parking area. She investigates the record of the security guard and discovers that a panel truck bearing Maryland license plate number MDC-578 was in the XYZ parking area the day before it was discovered that the cesium-137 canister was missing (see the bottom of Fig. 3).

Fusing all the discovered evidence, Mavis and TIACRITIS conclude that it is very likely that the cesium-137 canister was stolen with the MDC-678 truck. However, they now need to also assess $H'_2$ and $H''_2$. They do not find any relevant evidence for $H'_2$. In searching for evidence relevant to $H''_2$, Mavis contacts Grace, the Vice President for Operations at XYZ. Grace tells Mavis that no one at the XYZ Company had checked the canister out for work on any project. She says that the XYZ Company has other projects involving hazardous materials but none that involves the use of cesium-137. As a result, it is concluded to be very unlikely that the cesium-137 canister was used in a project at the XYZ Company.

Through such *spiral hybrid reasoning*, where abductions, deductions, and inductions feed on each other in recursive calls, Mavis and TIACRITIS continuously generate and update intermediate alternative hypotheses, use these hypotheses to guide the collection of relevant evidence, and use the evidence

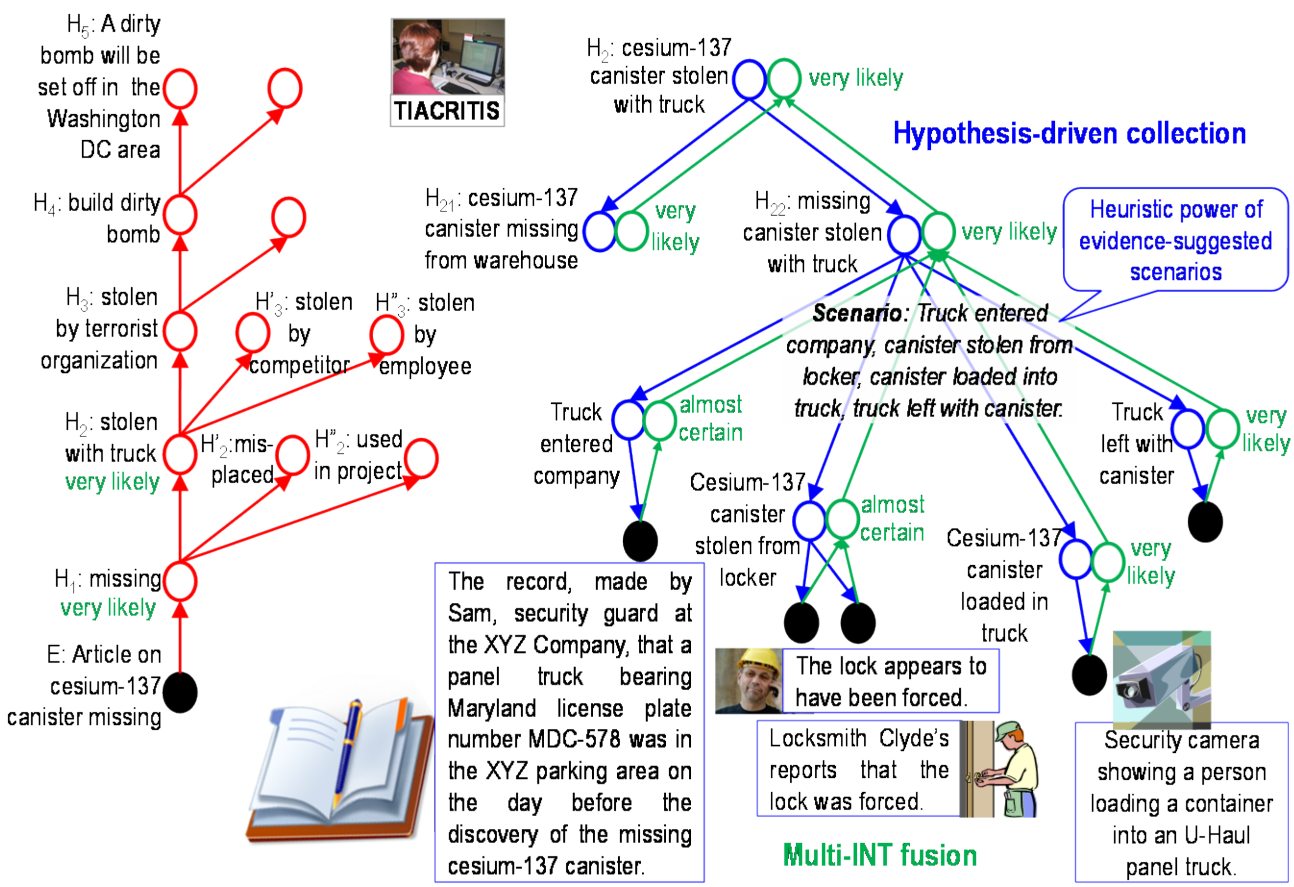


Figure 3. Spiral hybrid reasoning involving synergistic abductive, deductive, and inductive steps.

to test these hypotheses, until the likelihood of the top-level hypothesis is assessed. At the same time, TIACRITIS learns reasoning patterns from the analyst, and becomes increasingly more knowledgeable, as will be illustrated in Section IV.

## III. Hypothesis Analysis Through Problem Reduction and Solution Synthesis

The analyst and TIACRITIS analyze hypotheses by employing a general divide and conquer approach, called *problem reduction and solution synthesis*, which combines the deductive and inductive reasoning trees, as shown in the right hand side of Fig. 3. This approach is grounded in the problem reduction representations developed in artificial intelligence [6-8], and in the argument construction methods provided by the noted jurist John H. Wigmore [9], the philosopher of science Stephen Toulmin [10], and the evidence professor David Schum [1]. In this approach, which is illustrated in Fig. 4, the problem of assessing a complex hypothesis H is successively reduced to the assessment of simpler and simpler hypotheses, down to the level of elementary hypotheses. Then these elementary hypotheses (e.g., $H_2$) are assessed based on the available evidence. Finally, the solutions of these assessments are successively combined, from bottom-up, to obtain the solution of the top level hypothesis assessment.

In Fig. 4 the assessment of the hypothesis H is reduced to the assessment of three simpler hypotheses, $H_1$, $H_2$, and $H_3$. The middle hypothesis $H_2$ is assessed based on the available evidence. As indicated in Fig. 4, one has to consider both *favoring evidence* and *disfavoring evidence*. In this example there are two items of favoring evidence, $E_1$ and $E_2$. Therefore one has to assess to what extent each of them favors the hypothesis $H_2$. This requires the assessment of the *relevance* and *believability* of $E_1$, and of its *inferential force* on $H_2$.

The *relevance* answers the question: So what? How does this item of evidence bears on what we are trying to prove or disprove? The *believability* answers the question: Can we believe what this item of evidence is telling it? The *inferential force or weight* answers the question: How strong is this item of relevant evidence in favoring or disfavoring various alternative hypotheses we are entertaining?

As indicated before, all these assessments are probabilistic and, in our research, we have considered symbolic probabilities with names that are similar to those from the US National Intelligence Council's standard estimative language. For example, as shown in the table from the left side of Fig. 4, indicating that a hypothesis is "likely" is equivalent to saying that its probability of being true is between 0.55 and 0.75. Of course, the actual symbolic probabilities and the associated intervals from Fig. 4 are just examples. A user may decide to use other names for symbolic probabilities, as well as other associated intervals, as discussed by Kent [11] and Weiss [12].

In this example let us assume the following solutions for the relevance and the believability of $E_1$: *"If we believe $E_1$ then $H_2$ is almost certain"* and *"It is likely that $E_1$ is true."* These assessments need to be composed to assess the inferential force of $E_1$ on $H_2$. TIACRITIS uses the *"minimum"* composition function, because an item of evidence needs to be both very relevant and very believable to convince us that the hypothesis is true. As a result, the assessed the inferential force of $E_1$ on $H_2$ is: *"Based on $E_1$ it is likely that $H_2$ is true."* The inferential force of $E_2$ on $H_2$ is similarly assessed by TIACRITIS as *almost certain*. Then TIACRITIS composes the inferential force of $E_1$ on $H_2$ with the inferential force of $E_2$ on $H_2$, by using the *"maximum"* function because it is enough to be convinced by one item of evidence that the hypothesis is true. As a result, TIACRITIS assesses the following inferential force of the favoring evidence (i.e. both $E_1$ and $E_2$) on $H_2$: *"Based on the favoring evidence it is almost certain that $H_2$ is true."* Through a similar process TIACRITIS assesses the inferential force of the disfavoring evidence on $H_2$, and then the likelihood of $H_2$ based on both the favoring and the disfavoring evidence. $H_1$ and $H_3$ are assessed in a similar way as *very likely* and *likely*, respectively. Then the assessments of $H_1$, $H_2$, and $H_3$ are combined by TIACRITIS through a function selected by the analyst, such as *minimum* (all three hypotheses required to be true), *maximum* (one hypothesis required to be true), *average*, or *weighted sum*, into the assessment of the top level hypothesis H.

TIACRITIS is able to significantly help the analyst because it has a lot of knowledge about evidence. This includes an *ontology of evidence*, a fragment of which is shown in the bottom-right part of Fig. 4. This ontology distinguishes between different types of *tangible* and *testimonial evidence*. For each such type, TIACRITIS automatically employs a specific believability assessment procedure. For instance, in the case of an item of *demonstrative tangible evidence* which is a representation or image of a tangible thing (e.g., the record of the security camera in Fig. 2), its believability depends on its *authenticity*, *accuracy*, and *reliability*. Also, the believability of *unequivocal testimonial evidence based upon direct observation* (such as Ralph's testimony in Fig. 2) depends on source's *competence* and *credibility*. Competence depends on *access* and *understandability*, while credibility depends on *veracity*, *objectivity*, and *observational sensitivity* [1, 2].

This knowledge allows TIACRITIS to automatically reduce the assessment of complex hypotheses to the assessment of the relevance and believability credentials of evidence, as well as to automatically compose these assessments, once they are made by the analyst.

## IV. Illustration of the Use of TIACRITIS

TIACRITIS allows its users to formulate hypotheses, develop argumentation structures to assess them, collect evidence, associate evidence to elementary hypotheses, assess and justify the relevance and the believability of evidence, make assumptions with respect to certain sub-hypotheses, select the composition functions for determining the inferential force of evidence, and assess the hypotheses. We will illustrate these capabilities with the example of assessing the hypothesis $H_2$ and its argumentation structure from the right side of Fig. 3.

Using TIACRITIS, the analyst formulates the hypothesis analysis problem in English and selects its instances, as shown in the top part of Fig. 5. Selecting the instances allows TIACRITIS to learn the following general hypothesis analysis pattern: "Assess whether a ?O1 was stolen from the ?O2 with the ?O3."

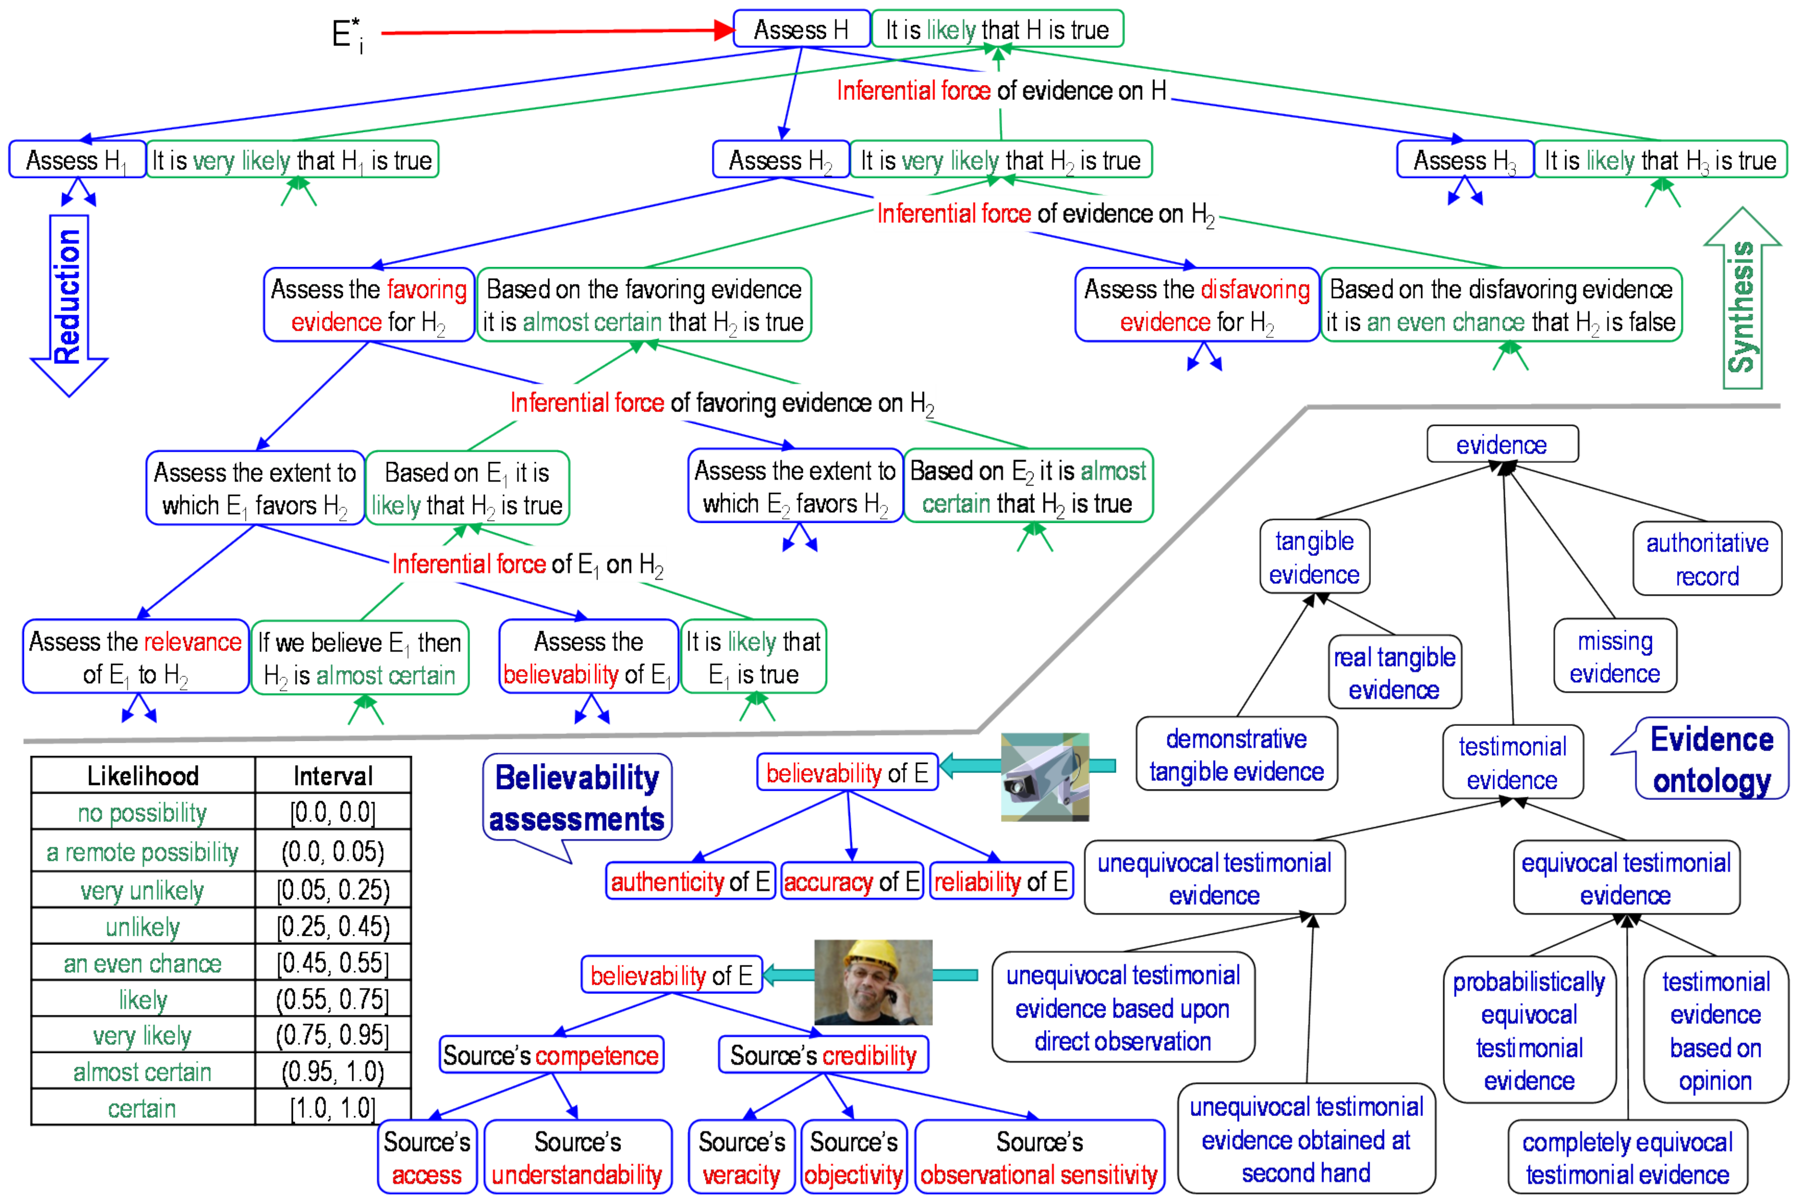


| Likelihood | Interval |
|---|---|
| no possibility | [0.0, 0.0] |
| a remote possibility | (0.0, 0.05) |
| very unlikely | [0.05, 0.25) |
| unlikely | [0.25, 0.45) |
| an even chance | [0.45, 0.55] |
| likely | (0.55, 0.75] |
| very likely | (0.75, 0.95] |
| almost certain | (0.95, 1.0) |
| certain | [1.0, 1.0] |

Figure 4. Evidence-based hypothesis analysis through reduction and synthesis.

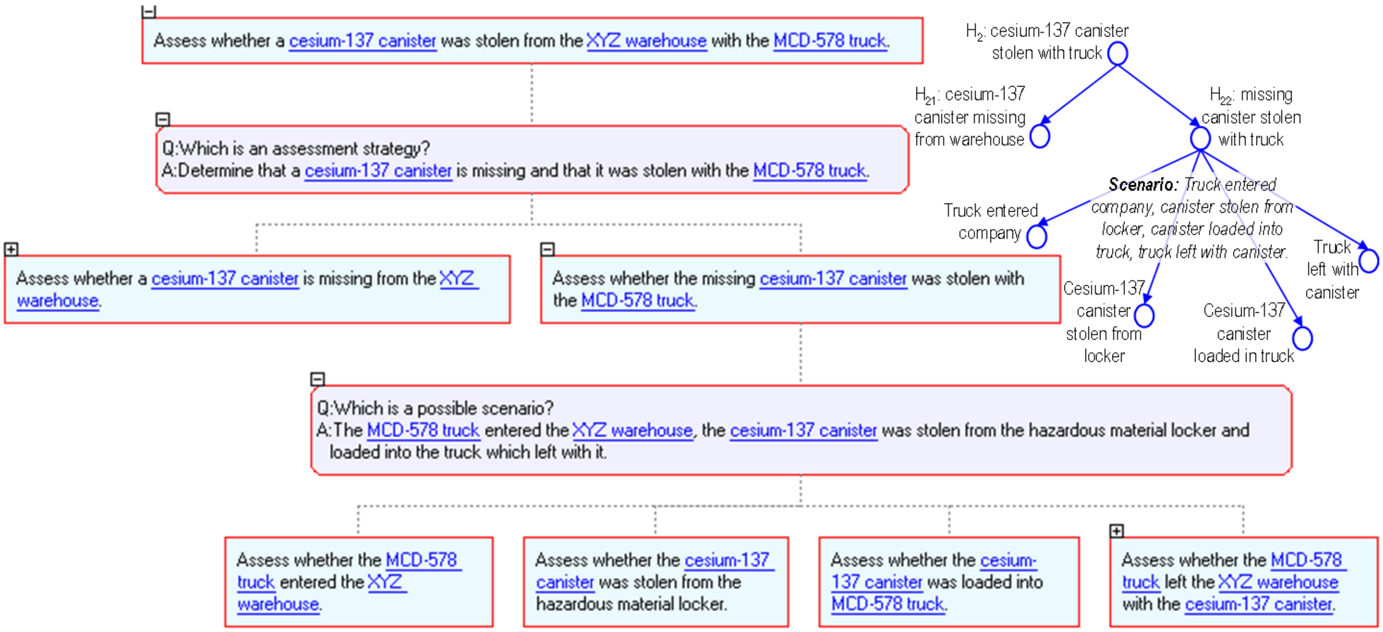


Figure 5. Hypothesis reduction.

As previously described, the analyst and TIACRITIS then reduce this hypothesis analysis problem to simpler and simpler problems, down to the level of elementary hypothesis analysis problems to be solved based on evidence. Notice that each hypothesis analysis problem in Fig. 5 is followed by a question/answer pair which guides its reduction to simpler problems. Thus the top level problem is reduced to two subproblems. The second subproblem is further reduced to four subproblems, based on the scenario discussed in Section II and illustrated in Fig. 3. Some of these reduction steps may be suggested by TIACRITIS, if it has encountered similar steps in past analyses.

Next the analyst will directly assess the elementary hypotheses based on relevant evidence, as discussed below. The analyst may associate any number of search criteria with elementary hypotheses which are then used by TIACRITIS to search for evidence in various repositories, as illustrated in Fig. 6. The top part of this figure shows an elementary hypothesis for which there is no evidence. The bottom part shows a search criteria defined by the analyst, to guide TIACRITIS in searching for relevant evidence on the Internet with BING, GOOGLE, or YAHOO (other search engines and repositories can be added).

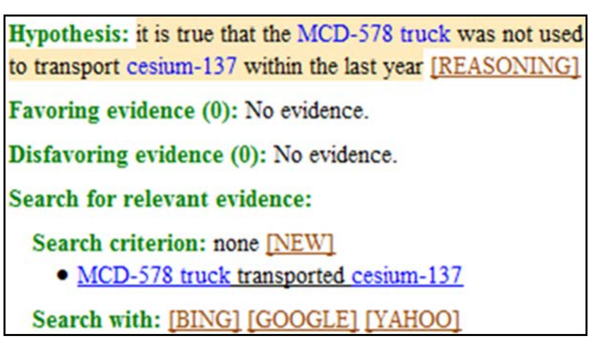


Figure 6. Evidence collection.

The analyst may easily define new items of evidence and may associate them with the hypotheses they favor or disfavor, as illustrated in Fig. 7. The top part of this figure is the description of the evidence item EVD-002-Ralph: Ralph's testimony that the cesium-137 canister is registered as being in the XYZ warehouse. The analyst has selected its type as unequivocal testimonial evidence based upon direct observation. Then the analyst indicated that this item of evidence favors the hypothesis "the cesium-137 canister was in the XYZ warehouse before being reported as missing," as shown in the middle part of Fig. 8.

As a result, TIACRITIS automatically generated the corresponding evidence-based analysis, as shown in Fig. 8. Notice that it considered both favoring and disfavoring

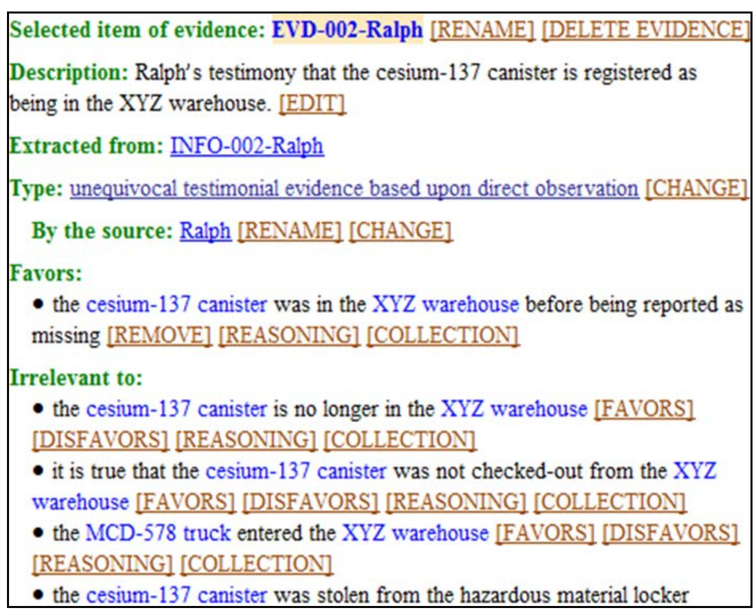


Figure 7. Evidence representation and use.

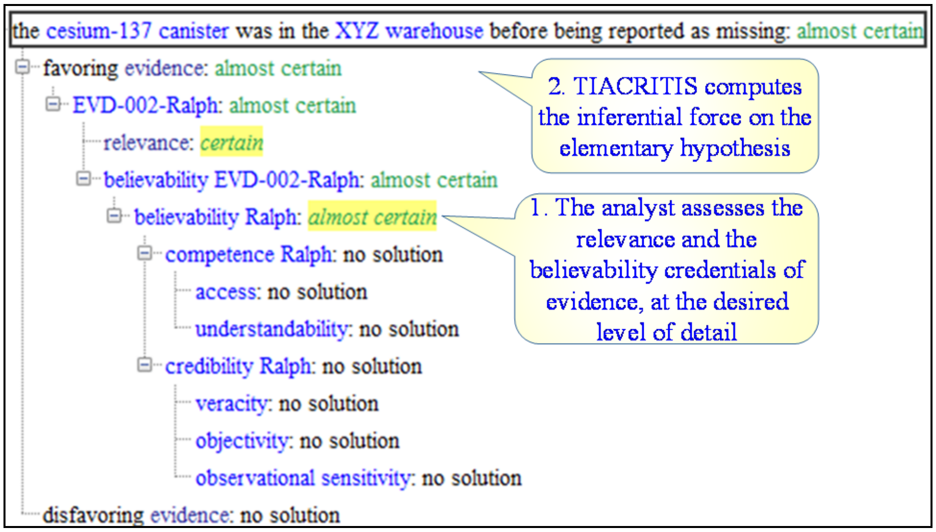


Figure 8. Evidence-based assessment of an elementary hypotheses.

evidence, and included EVD-002-Ralph as favoring evidence for which the analyst needs to assess the relevance and the believability. Because EVD-002-Ralph is unequivocal testimonial evidence based upon direct observation, its believability depends on Ralph's competence and credibility. Competence depends on access and understandability, while credibility depends on veracity, objectivity, and observational sensitivity.

The analyst has assessed the relevance of EVD-002-Ralph as certain and the believability of Ralph as almost certain. Then TIACRITIS has combined these assessments into an inferential force of almost certain, and has computed the likelihood of the corresponding elementary hypothesis.

Notice that although TIACRITIS has provided a detailed believability analysis, the user may drill down into this analysis at the desired level and, in this case, decided to assess directly the believability of Ralph, rather than assessing lower level believability credentials, such as veracity. This is referred to as an *assumption*.

After all the elementary hypotheses have been assessed, either based on evidence or by making assumptions, the user has to select the solution composition functions (e.g., min, max, average, or weighted sum) to be used by TIACRITIS when assessing the likelihoods of the intermediary hypotheses and of the top level hypothesis, as shown in Fig. 9.

TIACRITIS not only supports the analyst in hypotheses analysis, but it also continuously learns to facilitate the analysis of new hypotheses. Consider, for examples, the new hypothesis analyses problem from the top of Fig. 10. TIACRITIS suggests a reduction based on a pattern learned from the analysis in Fig. 5. It also suggests the question for another assessment strategy to be defined by the analyst. Of course, the more TIACRITIS learns, the more useful its suggestions.

## V. Final Remarks

TIACRITIS is an operational web-based system, and is available for education and analysis (see Fig. 11). It includes modules from the Disciple Learning Agent Shell, as well as modules that implement the current version of the computational theory of intelligence analysis. Its use is supported by three textbooks and numerous case studies:

- "Introduction to Intelligence Analysis: A Hands-on

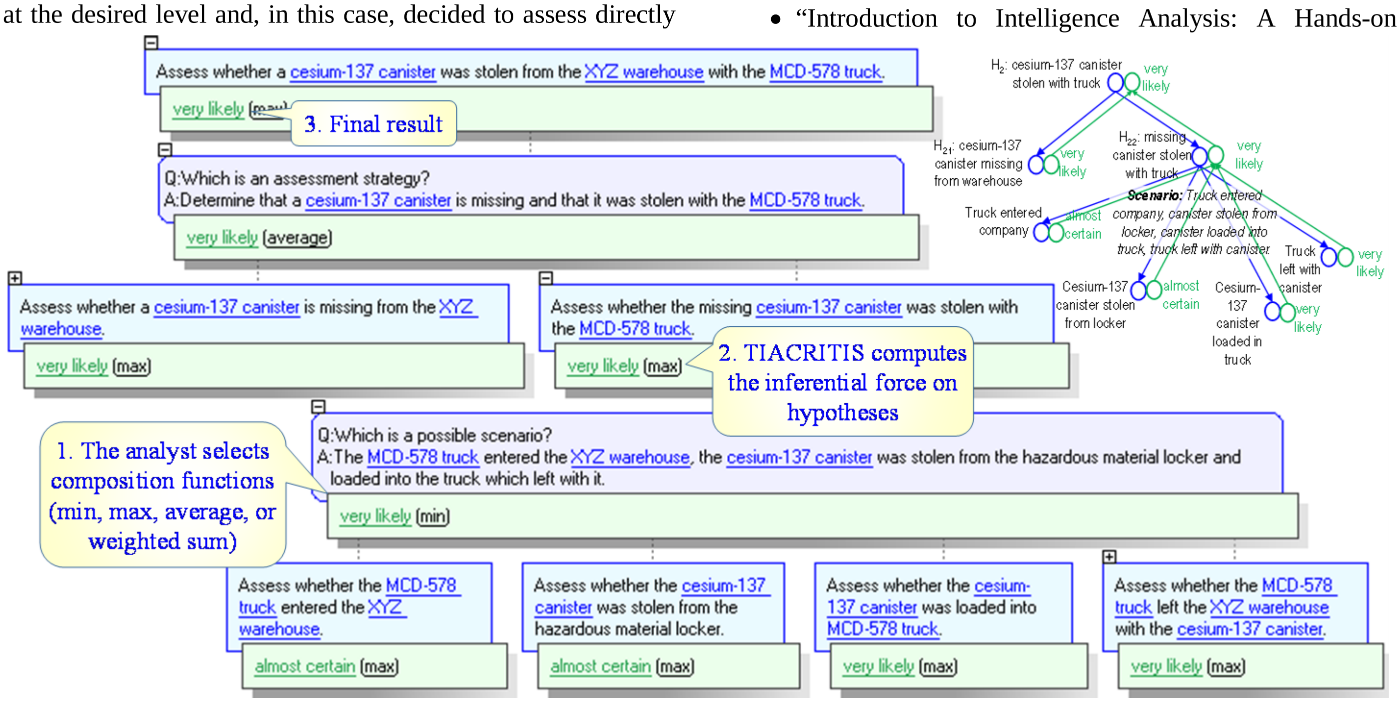


Figure 9. The top part of the hypothesis analysis tree showing the solution composition functions and the likelihoods of the hypotheses.

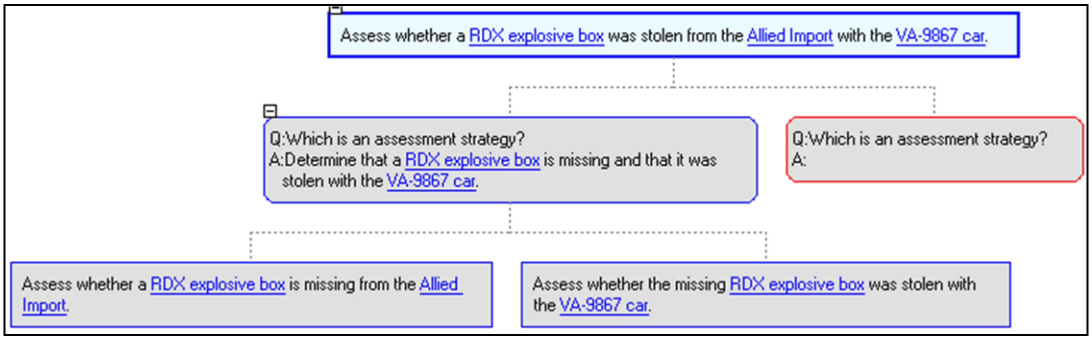


Figure 10. Reductions suggested by TIACRITIS based on learned analysis patterns.

Approach with TIACRITIS” teaches basic knowledge about the properties, uses, and marshaling of evidence to show students how to collect evidence and test hypotheses by assessing the relevance, the believability, and the inferential force of evidence [2].

- “A Practicum in Evidence Marshaling and Argument Construction with TIACRITIS” teaches advanced strategies for organizing and combining analyst’s thoughts and evidence to construct complex arguments from masses of evidence (in preparation).
- “Modeling Violent Extremists with TIACRITIS” teaches an evidence-based methodology for investigating, comprehending, and anticipating the behavior of violent extremists in the war on terror [13].

One main direction of follow-on work is further development of the computational theory and its implementation in TIACRITIS. This includes the development of computational models for evidence marshaling guided by magnets which are powerful heuristics supporting the analysts in hypotheses generation from masses of evidence. Future research also includes the development of more powerful methods for the learning and reuse of analytic expertise, for hypotheses generation through mixed-initiative abduction, for collaborative analysis, for automatic report generation, and for decision-making under uncertainty which integrates the computational theory. Although the focus of the current work was on mixed-initiative analysis involving analysts, TIACRITIS and the theory it is built on can be extended to persistent surveillance and interpretation of dynamic environments by autonomous agents.

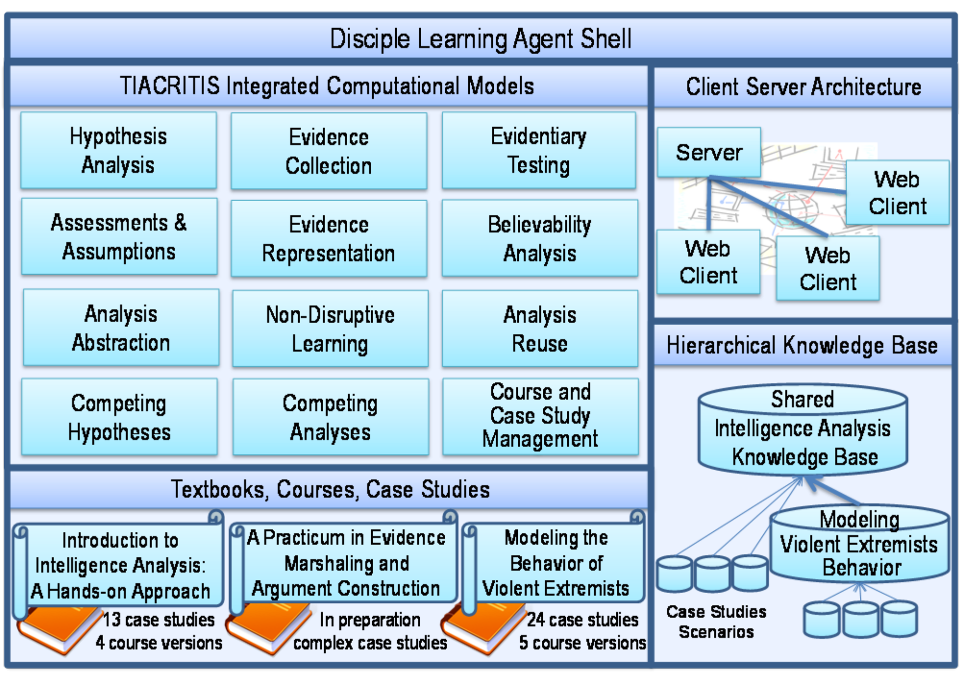


Figure 11. TIACRITIS cognitive assistant and textbooks.

## REFERENCES


[1] Schum D.A. (2001). *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.

[2] Tecuci G., Schum D.A., Boicu M., Marcu D. (2010). *Introduction to Intelligence Analysis: A Hands-on Approach with TIACRITIS*, 220 pages, George Mason University, new edition 2011.

[3] Tecuci G., Boicu M., Cox M.T. (2007). Seven Aspects of Mixed-initiative Reasoning: An Introduction to the Special Issue on Mixed-initiative Assistants. *AI Magazine* 28(2), 11–18.

[4] Schum D.A. (2009). Science of Evidence: Contributions from Law and Probability. *Law Probab Risk* 8, 197–231.

[5] Tecuci G., Boicu M., Marcu D., Schum D., Hamilton B. (2010). TIACRITIS System and Textbook: Learning Intelligence Analysis through Practice, in *Proc of the 5th Int. Conf. on Semantic Technologies for Intelligence, Defense, and Security – STIDS*, 108-115, Fairfax, VA.

[6] Nilsson N.J. (1971). *Problem Solving Methods in Artificial Intelligence*. NY: McGraw-Hill.

[7] Tecuci G. (1988). *DISCIPLE: A Theory, Methodology and System for Learning Expert Knowledge*. Thèse de Docteur en Science, University of Paris-South, France.

[8] Tecuci G. (1998). *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, San Diego, CA: Academic Press.

[9] Wigmore J.H. (1937). *The Science of Judicial Proof*. Boston, MA: Little, Brown & Co.

[10] Toulmin S.E. (1963). *The Uses of Argument*. Cambridge University Press.

[11] Kent S. (1994). Words of Estimated Probability, in Steury D.P., ed., *Sherman Kent and the Board of National Estimates: Collected Essays*, Center for the Study of Intelligence, CIA, Washington, DC.

[12] Weiss C. (2008). Communicating Uncertainty in Intelligence and Other Professions, *International Journal of Intelligence and CounterIntelligence*, 21(1), 57–85.

[13] Boicu M., Tecuci G., Marcu D., Schum D.A., Coughlin S. (2011). *Modeling Violent Extremists with TIACRITIS*, 203 pages, George Mason University, 2011.