
Computational Approach and Cognitive Assistant for Evidence-Based Reasoning in Intelligence Analysis

Gheorghe Tecuci*, David A. Schum,
Dorin Marcu, Mihai Boicu

Learning Agents Center, Volgenau School of Engineering,
George Mason University, 4400 University Dr., Fairfax, VA, USA
E-mails: {tecuci, dschum, dmarcu, mboicu}@gmu.edu

* Corresponding author

Abstract: This paper presents a computational approach to intelligence analysis and its current implementation into a cognitive assistant called Disciple-CD. Intelligence analysis is viewed as ceaseless discovery of evidence, hypotheses, and arguments in a non-stationary world, involving cooperative processes of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary tests of hypotheses. Disciple-CD helps intelligence analysts formulate hypotheses, develop arguments that reduce complex hypotheses to simpler and simpler ones, collect evidence to evaluate the simplest hypotheses, assess the relevance, believability, and inferential force of evidence, and finally the likeliness of the hypotheses.

Keywords: Science of evidence, intelligence analysis, artificial intelligence, cognitive assistant, discovery, abduction, deduction, induction, abstraction, substance-blind classification of evidence, relevance, believability, inferential force or weight, assumption

Biographical notes: Gheorghe Tecuci (PhDs, University of Paris-South and Polytechnic Institute of Bucharest) is Professor of Computer Science and Director of the Learning Agents Center in the Volgenau School of Engineering of George Mason University. He has followed a career-long interest in the development of a computational theory and technology allowing non-computer scientists to develop cognitive agents that incorporate their problem solving expertise, and can act as cognitive assistants to experts, as expert consultants to non-experts, and as intelligent tutors to students.

David A. Schum (PhD, Ohio State University) is Professor of Systems Engineering and Operations Research, Professor of Law, and Chief Scientist of the Learning Agents Center at George Mason University, and Honorary Professor of Evidence Science at University College London. He has followed a career-long interest in the study of the properties, uses, discovery and marshaling of evidence in probabilistic reasoning.

Dorin Marcu (PhD, George Mason University) is Senior Software and Knowledge Engineer and Research Assistant Professor in the Learning Agents Center, Volgenau School of Engineering, George Mason University. His main research interests are in intelligent user interfaces, mixed-initiative interaction, abstraction of reasoning, and cognitive assistants.

Mihai Boicu (PhD, George Mason University) is Associate Professor of Applied Information Technology and Associate Director of the Learning Agents Center in the Volgenau School of Engineering of George Mason University. He has expertise in problem solving and multistrategy learning in dynamic and evolving representation spaces, mixed-initiative interaction, multi-agent systems architecture, collaboration and coordination, abstraction-based reasoning, knowledge representation and knowledge acquisition.

1 Introduction

Intelligence analysts and others face the difficult task of drawing defensible and persuasive conclusions from masses of information of all kinds that come from a variety of different sources. Many books and papers have been written on the obvious complexity of such tasks (Schum, 1987, 2001a; Cooper, 2008; Tecuci et al., 2011a). The mass of evidence upon which conclusions eventually rest has five major characteristics that make conclusions drawn from evidence necessarily probabilistic in nature. Our evidence is always *incomplete* no matter how much we have, and is commonly *inconclusive* in the sense that it is consistent with the truth of more than one hypothesis or possible explanation. Further, the evidence is frequently *ambiguous*; we cannot always determine exactly what the evidence is telling us. A mass of evidence is in most situations *dissonant* to some degree; some of it favors one hypothesis or possible explanation but other evidence favors other hypotheses. Finally, all of our intelligence evidence comes from sources having any possible gradation of *believability or credibility* shy of perfection. Arguments, often stunningly complex, are necessary in order to establish and defend the three major credentials of evidence: its *relevance, believability, and inferential force or weight*. These arguments rest upon both *imaginative* and *critical reasoning* on the part of intelligence analysts.

But these assorted evidential characteristics are not the only elements of the complexity of intelligence analysis tasks. For example, in many situations the analysts have to rapidly perform their analyses without having much time for collecting relevant evidence, and for performing a detailed analysis.

Identifying the complexities of intelligence analysis is actually the easy part. What is not so easy are efforts to assist analysts in coping with the complexities of the evidential reasoning tasks they routinely face.

This paper presents a systematic computational approach to hypotheses analysis which is based on the emerging science of evidence (Schum, 2009), artificial intelligence (Tecuci, 1998), logic and probability (Schum, 2001a). This approach is implemented in the Disciple-CD cognitive assistant which helps automate significant portions of the hypotheses analysis process, helping intelligence analysts in coping with many analytic complexities. The name Disciple-CD is an abbreviation of Disciple cognitive assistant for Connecting the Dots..

In the next section we present a view of intelligence analysis as a process of ceaseless discovery in a non-stationary world, process involving evidence in search of hypotheses (primarily through abductive reasoning), hypotheses in search of evidence (primarily through deductive reasoning), and evidential tests of hypotheses (primarily through inductive reasoning), all going on at the same time. In Section 3 we show how Disciple-CD helps performing this process by employing a general divide and conquer reasoning strategy called problem reduction and solution synthesis. Then Section 4 discusses the abstraction of reasoning used in Disciple-CD to facilitate the browsing and understanding of complex arguments. Section 5 discusses the Disciple-CD implementation of the hypotheses in search of evidence process. Section 6 defines the major credentials of evidence (relevance, believability and inferential force or weight) and how they are used to assess hypotheses. After that, Section 7 presents a substance-blind classification of evidence and how it is used in assessing the believability of evidence. Section 8 presents how an analyst can perform analyses at different levels of details, and use assumptions to deal with lack of evidence or analysis time. Section 9 presents the learning and reuse of analytic expertise, and Section 10 concludes the paper with a summary of the main features of developed computational approach and Disciple-CD.

2 Discovery: Generating Hypotheses, Evidence, and Arguments

All intelligence analyses, in common with analytic activities in any other context, begin with the asking of questions about matters of interest. These questions can arise from the analysts themselves or from other persons, such as the policy or decision makers, who are being served by intelligence analysts. These questions can concern possible explanations for events or situations in the past or possible predictions about events or situations in the future. In many cases these questions are bound together. In order to predict possible events in the future we need accurate explanations for related events in the past. The field of intelligence analysis has many inherent difficulties, but none seem more difficult than the fact that analysts must provide their explanations or predictions in a non-stationary world. In short, the world keeps changing as analysts are trying their best to understand it well enough to provide explanations or to make predictions. One consequence is that we have continuing streams of new information, some items of which we will assess as being relevant evidence regarding our explanations or predictions. An explanation for some

pattern of past events analysts have previously regarded as correct may now seem incorrect in light of new evidence just discovered today. A prediction regarded as very likely today may be overtaken by events we will learn about tomorrow. In fact, the very questions we have asked yesterday may need to be revised or may even seem unimportant in light of what we learn today. One consequence of all of this is that the process of discovery or investigation in intelligence analysis is a ceaseless activity. It would be a drastic mistake to view discovery in intelligence analysis as being a stationary activity in a non-stationary world.


What exactly does discovery involve, or what needs to be discovered in intelligence analysis? The answer is: *hypotheses, evidence, and arguments linking hypotheses and evidence*. From observations we make, or questions we ask, we generate alternative hypotheses or propositions offered in explanation for past events or possible predictions about future events. In the continual streams of data or information provided to intelligence analysts only a minute fraction of these data are justified as being termed evidence. *Data or items of information only become evidence when their relevance to hypotheses being considered is established by defensible and persuasive arguments*. What is true is that establishing these three ingredients of all intelligence analyses is a very complex activity involving *imaginative* as well as *critical reasoning*. Discovery in intelligence analysis involves mixtures of all three forms of reasoning that have been identified: *abduction, deduction, and induction*. 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, 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 (Peirce, 1898; 1901).

But now we must return to intelligence analysis being a ceaseless discovery-related activity performed in a non-stationary world. On at least some accounts it may appear that the generation of a productive hypothesis occurs as a result of a single glorious episode of abductive or imaginative reasoning on the part of a particular intelligence analyst (Schum, 2001b). Barring clairvoyance or divine intervention, this seems quite unlikely. Tying discovery to just abductive reasoning overlooks the true complexity of discovery in intelligence analysis and in many other contexts. Remember that we have three things to be discovered in intelligence analysis: hypotheses, evidence, and arguments linking evidence to hypotheses. The fact that the world is changing all the while we are trying to understand it means that we have evidence in search of hypotheses, hypotheses in search of evidence, and evidential tests of hypotheses *all going on at the same time*. What this means is that discovery in intelligence analysis involves mixtures of abductive, deductive, and inductive reasoning. By means of abductive reasoning we generate hypotheses from evidence we gather; by deductive reasoning, we make use of our hypotheses to generate new lines of inquiry and evidence; and by inductive reasoning

we test hypotheses on the basis of the evidence we are discovering. Such testing depends on the relevance and the believability of our evidence. These factors combine in further complex ways to allow us to assess the inferential force or weight of the evidence we are considering.

We have developed a computational theory of how an intelligent agent may support an analyst in this process of ceaseless discovery of evidence, hypotheses, and arguments linking them. This is briefly described by using the diagram in Figure 1.

Figure 1 Discovery of evidence, hypotheses and arguments



Consider that you are an intelligence analyst and receive a tip from an informant that a company owned by Aum Shinrikyo, a Japanese apocalyptic sect, has recently acquired a significant amount of sarin. This is evidence E^* at the bottom left of Figure 1. This tip arouses your curiosity because of your concern about terrorist acts in the Tokyo area. The question is: *What hypothesis would explain this observation?* You experience a flash of insight that Aum Shinrikyo might use chemical weapons in the Tokyo subway. This is hypothesis H at the top-left of Figure 1.

However, no matter how imaginative or important the hypothesis H is, no one will take it seriously unless you are able to justify it. So you develop the chain of abductive

inferences shown in the left side of Figure 1 and in Table 1.

Table 1 Abductive reasoning steps justifying a hypothesis

There is evidence that Aum has acquired a significant amount of sarin (E^*).
Therefore it is possible that Aum has acquired a significant amount of sarin (E).
Therefore it is possible that Aum is pursuing sarin-based chemical weapons (F).
Therefore it is possible that Aum will use chemical weapons in a crowded area (G).
Therefore it is possible that Aum will use chemical weapons in the Tokyo subway (H).

The chain of inferences from the left side of Figure 1 and from Table 1 show the possibility that Aum will use chemical weapons in the Tokyo subway. Can you then conclude that this will actually happen? No because there are many other hypotheses that may explain evidence E^* , as shown in the left side of Figure 1 and discussed below.

First, just because we have evidence E^* about an event E does not mean that the event E has actually happened. That is, just because our informant says that Aum has acquired a significant amount of sarin does not mean that it has indeed done that. At issue here is the believability of this informant. What if the informant is mistaken or deceptive? Thus an alternative hypothesis is E^c : Aum has not acquired a significant amount of sarin.

But let us assume that Aum has indeed acquired a significant amount of sarin, and therefore it is possible that it will use it to build chemical weapons (hypothesis F). But Aum's chemical products company may have bought sarin to manufacture its commercial products (hypothesis F'), not chemical weapons, and so on.

This is the process of *evidence in search of hypotheses*, shown in the left side of Figure 1. We cannot conclude that Aum will use chemical weapons in the Tokyo subway (i.e. hypothesis H) until we consider all the competing hypotheses and show that those on the chain from E^* to H are actually more likely than their competing hypotheses. That is, we need to show that each of E , F , G , and H is more likely than its competing hypotheses. But to analyze all these competing hypotheses and make such an assessment, we need additional items of evidence. How can we get them? As represented in the right side of Figure 1 and discussed below, we put each hypothesis at work to guide us in the collection of additional evidence.

Let us suppose that we have shown that E and F are more likely than their competing hypotheses. For example, F is *very likely* while F' is only *likely*. Now we have to assess G , G' , To obtain evidence relevant to G , we successively decompose G into simpler and simpler hypotheses, as shown in the right side of Figure 1. G would be true if K and M would be true. Then M would be true if the events N , P , and Q would be true. Thus we look for evidence that these events have happened and find N^* , P_1^* , P_2^* , and Q^* .

This is the process of *hypotheses in search of evidence* that guides the evidence collection task.

Now some of the newly discovered items of evidence (e.g. N^*) may trigger new hypotheses (or the refinement of the current hypotheses). Therefore, as indicated at the bottom 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.

Then we use the identified evidence to assess the likeliness of G . In the developed computational theory of intelligence analysis (Tecuci et al., 2011b), hypotheses assessment is based on a combination of ideas from the Baconian probability system (Cohen, 1977; Schum, 2001a) and from the Fuzzy probability system (Zadeh, 1983; Schum, 2001a), and uses a symbolic probability scale. In particular, in Disciple-CD, the likeliness of a hypothesis (e.g., G) may have one of the following ordered values:

no support < likely < very likely < almost certain < certain

In this scale, **no support** means that our evidence does not support the conclusion that the hypothesis is true. This may change if new evidence favoring it is later found.

The likeliness of an upper-level hypothesis (e.g., G) is obtained from the likeliness of its sub-hypotheses (i.e., K and M) by using min or max Baconian and Fuzzy combination functions, depending on whether the sub-hypotheses K and M represent necessary and sufficient conditions for the hypothesis G , sufficient conditions, or just indicators.

Competing hypotheses to G (e.g., G') are assessed in a similar way and the most likely one is selected. But if no hypothesis is more likely than all of its competitors, then the processes of hypotheses in search of evidence, and evidence in search of hypotheses have to be resumed, as shown by the arrows at the bottom of Figure 1.

As will be discussed in the next sections, Disciple-CD supports the analyst in performing such an analysis. First, it has a significant amount of knowledge about evidence from the science of evidence. Second, it is a learning system that can learn general analysis rules from an expert analyst and uses them in future analyses.

3 Computational Approach to Hypothesis Analysis: Problem Reduction and Solution Synthesis

Disciple-CD partially implements the approach discussed in the previous section. It allows its user to formulate hypotheses, develop arguments that reduce complex hypotheses to simpler and simpler ones, collect evidence relevant to the simplest hypotheses, and finally assess the relevance, the believability, and the inferential force of evidence, and the likeliness of the hypotheses. We will present these capabilities in the following.

First the analyst has to specify the hypothesis to analyze. If this is a completely new hypothesis, like the following one, then the analyst will formulate it in natural language,

selecting the specific entities used in it (i.e., Aum Shinrikyo and sarin-based weapons):

Assess whether [Aum Shinrikyo](#) is pursuing [sarin-based weapons](#).

This guides Disciple-CD to learn a general hypothesis analysis pattern: "Assess whether ?O1 develops ?O2." In the future, when the analyst needs to assess a similar hypothesis, such as "Assess whether [Z-group](#) is pursuing [botulinum-based weapons](#)", he/she can simply instantiate this pattern.

Next the analyst and Disciple-CD need to develop an evidence-based argumentation for assessing the likeliness of the hypothesis. They do this by employing a general divide-and-conquer approach to problem solving, called *problem reduction and solution synthesis*, which has a grounding in the problem reduction representations developed in Artificial Intelligence (Nilsson, 1977; Powell and Schmidt, 1988; Tecuci, 1988; 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; 2001a).

A guiding principle in the development of Disciple-CD was to employ a problem solving paradigm that is both natural enough for the analyst and formal enough to be automatically executed by the Disciple-CD agent. This resulted in an inquiry-driven approach where a complex hypothesis is assessed by:

- Successively reducing it, from top-down, to simpler and simpler hypotheses, guided by introspective questions and answers;
- Assessing the simplest hypotheses based on evidence; and
- Successively combining, from bottom-up, the assessments of the simpler hypotheses, until the assessment of the top-level hypothesis is obtained.

This approach is illustrated in Figure 2 and discussed in the following.

Table 2 shows how a question and its answer guide the reduction of the top-level hypothesis analysis problem from the top of Figure 2 into four sub-problems.

Table 2 Problem reduction guided by a question and answer

We have to:

Assess whether [Aum Shinrikyo](#) is pursuing [sarin-based weapons](#).

Which is a sufficient condition?

[Aum Shinrikyo](#) has reasons and desire to pursue [sarin-based weapons](#), and is developing capabilities to secretly acquire and use them.

Therefore we have to:

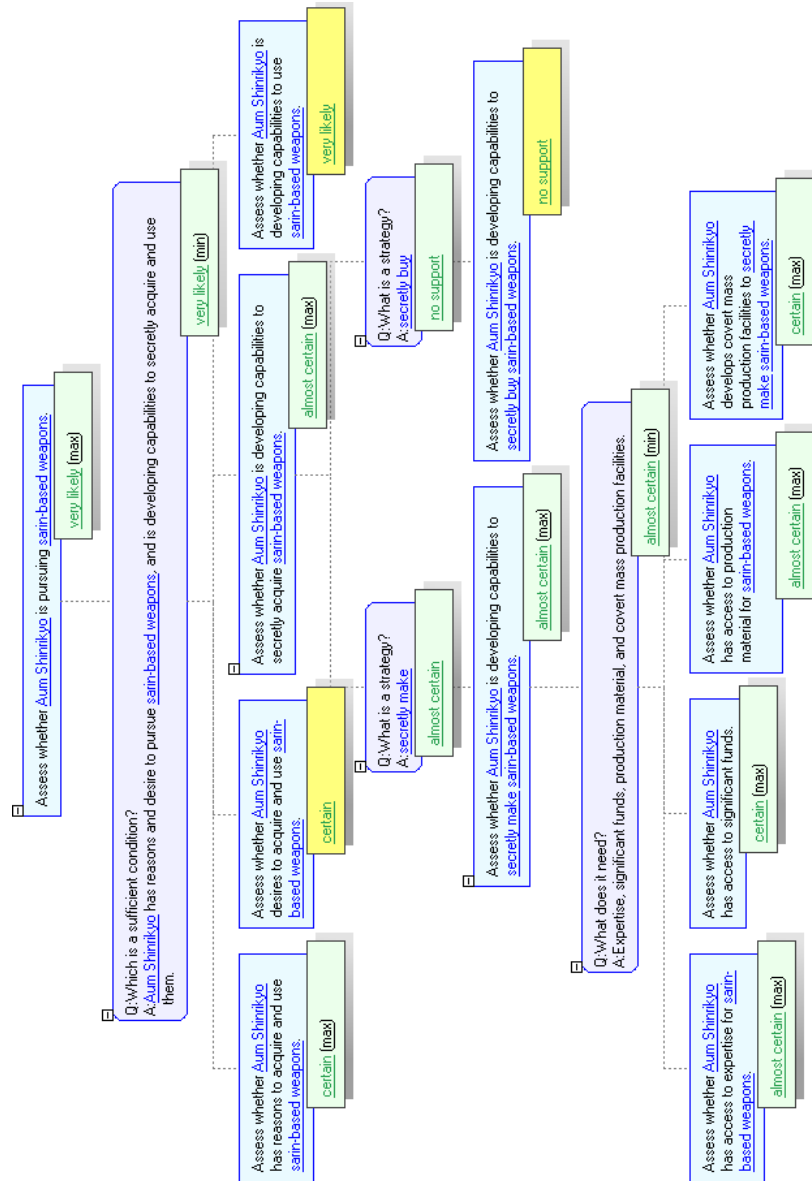
Assess whether [Aum Shinrikyo](#) has reasons to acquire and use [sarin-based weapons](#).

Assess whether [Aum Shinrikyo](#) desires to acquire and use [sarin-based weapons](#).

Assess whether [Aum Shinrikyo](#) is developing capabilities to secretly acquire [sarin-based weapons](#).

Assess whether [Aum Shinrikyo](#) is developing capabilities to use [sarin-based weapons](#).

Figure 2 Problem reduction and solution synthesis



Let us now assume that we have obtained the following assessments or solutions of the four sub-problems in Table 2: **certain**, **certain**, **almost certain**, and **very likely**, respectively (as shown Figure 2). Then the solution of the top-level problem, corresponding to this sufficient condition, is the minimum of the four solutions (i.e., **very likely**), because each of the four sub-hypotheses would need to be true to assure that the top-level hypothesis is true. This value and the min function that produced it are associated with the question/answer pair and displayed as “**very likely (min)**”, as shown in Figure 2.

In general, as will be illustrated below, there may be more than one strategy to reduce a hypothesis to simpler hypotheses, each resulting in a possibly different solution. In such a case, the solution of the hypothesis should be taken as the maximum of all these possible solutions. In this particular example, since we only have one strategy, the solution of the top-level problem is $\max(\text{very likely}) = \text{very likely}$.

Ideally, a hypothesis would be reduced to several sub-hypotheses that would represent a necessary and sufficient condition. However, many times it will not be easy to identify a necessary and sufficient condition to reduce a hypothesis. In such a case, a second best reduction would be a sufficient condition, like the one from the top of Figure 2. A common example of a sufficient condition is a scenario in which the considered hypothesis would be true. Another common example is a strategy or plan to perform an action which would result in a state where the hypothesis would be true. The middle part of Figure 2, for instance, shows two possible strategies for Aum Shinrikyo to secretly acquire sarin-based weapons: it can secretly make them or it can secretly buy them. Each strategy has to be assessed and their maximum represents the assessment of the hypothesis that Aum Shinrikyo is developing capabilities to secretly acquire sarin-based weapons.

Most of the time when we are assessing a hypothesis we only have indicators. An indicator is weaker than a sufficient condition. If we determine that a sufficient condition is satisfied (e.g., a scenario has actually happened), we may conclude that the hypothesis is true. But we cannot draw such a conclusion just because we have discovered an indicator. However, we may be more or less inclined to conclude that the hypothesis is true, based on the strength of the indicator. Therefore, Disciple-CD distinguishes between three types of indicators, of different strengths: “**likely** indicator”, “**very likely** indicator”, and “**almost certain** indicator”. A “**likely** indicator” is one which, if discovered to be true, would lead Disciple-CD to conclude that the considered hypothesis is **likely**. Similarly, a “**very likely** indicator” would lead to the conclusion that the hypothesis is **very likely**, and an “**almost certain** indicator” would lead to the conclusion that the hypothesis is **almost certain**. Now if **I** is a **likely** indicator of **H** and the likeliness of **I** is assessed to be **almost certain**, then Disciple-CD concludes that the likeliness of **H** is **likely**, the minimum between **likely** and **almost certain**. In general, *the likeliness of a hypothesis **H** based on an indicator **I** is the minimum between the likeliness of the indicator and the strength of the indicator* (which could be **likely**, **very likely**, or **almost certain**). Disciple-CD does not

consider the type “**certain** indicator” because this would be a sufficient condition. It also does not consider the type “**no support** indicator”. Indeed, if it does not provide any support to the hypothesis, then it is not an indicator of it.

4 Abstraction of Reasoning

Analyses of complex hypotheses from masses of evidence result in the generation of very large reasoning trees, some with thousands of nodes. To help develop, browse, understand, and update such a complex analysis, Disciple-CD enables the analyst to develop an abstraction of it which only shows the abstracted names of the main sub-problems considered in the analysis. The analyst only needs to click on the hypothesis to be abstracted and specify its abstraction. The use of abstractions changes the view of the analysis tree as illustrated in Figure 3 and explained in the following.

The left panel in Figure 3 shows the abstraction of the entire analysis tree. The nodes in this tree are abstractions of the main hypothesis analysis problems from the detailed analysis, and their solutions. In particular, the top node of the abstract tree corresponds to the top node of the detailed tree in Figure 2. This node has four children: reasons, desire, capabilities to acquire, and capabilities to use. These are the abstractions of the four sub-problems of the top problem in Figure 2. The analyst can expand an abstract hypothesis (to show its abstract sub-hypotheses) by clicking on its “+” label. Clicking on a “-“ label collapses the tree.

When the analyst clicks on an abstract hypothesis, such as “production material: **almost certain**” in the left panel of Figure 3, Disciple-CD displays its detailed reduction in the right panel. The detailed reduction shows the full descriptions of the problems, the question/answer pair, and the synthesis functions. The abstract reduction and the corresponding detailed reduction are also shown in Table 3.

Figure 3 Abstract and detailed views of an analysis tree

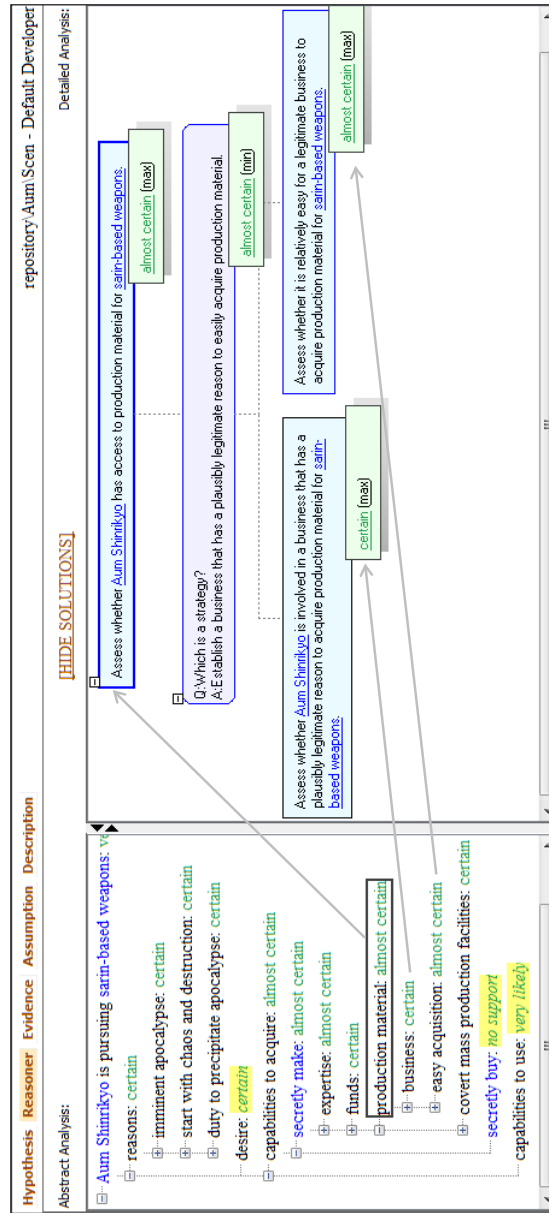


Table 3 An abstract reduction and the corresponding detailed reduction from Figure 3

production material: almost certain	
business: certain	
easy acquisition: almost certain	
is the abstraction of	
<i>We have to:</i>	
Assess whether Aum Shinrikyo has access to production material for sarin-based weapons .	almost certain (max)
Which is a strategy?	
Establish a business that has a plausibly legitimate reason to easily acquire production material.	almost certain (min)
<i>Therefore we have to:</i>	
Assess whether Aum Shinrikyo is involved in a business that has a plausibly legitimate reason to acquire production material for sarin-based weapons .	certain (max)
Assess whether it is relatively easy for a legitimate business to acquire production material for sarin-based weapons .	almost certain (max)

5 Hypotheses in Search of Evidence

Up to this point we have discussed how a hypothesis is reduced to simpler and simpler hypotheses, down to the level of hypotheses that are to be directly assessed based on relevant evidence. But how can the analyst find relevant evidence?

With an analytic tool based on the well-known *Analysis of Competing Hypotheses* method (Heuer, 1999, pp. 77-88; Heuer, 2008, pp. 251 - 265), where no argumentation structure is built, the analyst receives no guidance on how to search for evidence that is relevant to the initial hypothesis that [Aum Shinrikyo](#) is pursuing [sarin-based weapons](#). Therefore evidence collection is a very difficult task whose result has a direct influence on the quality of the performed analysis.

With Disciple-CD, we put this hypothesis at work to guide us in the collection of relevant evidence, as was discussed in Section 2 and illustrated in the right side of Figure 1. More specifically, as discussed in Section 3, we successively decompose our hypothesis into simpler and simpler hypotheses, and then search for evidence that is relevant to the leaf hypotheses. But this is a much simpler task because the leaf hypotheses are much more precise and limited in scope. Indeed, it is one thing to look for evidence that [Aum Shinrikyo](#) is pursuing [sarin-based weapons](#) (the top-level hypothesis in Figure 2), and quite another thing to look for evidence that, for instance, [Aum Shinrikyo](#) has access to significant funds (a leaf hypothesis in Figure 2). What, in fact, we have done, was

to reduce the complex task of searching for evidence for the top hypothesis, to simpler tasks of searching for evidence for the leaf hypotheses.

Now, to search for evidence that is relevant to the leaf hypotheses, Disciple-CD guides us to associate search criteria with them, and to invoke various search engines on the Internet, as illustrated in Figure 4. The top part of this figure shows the leaf hypothesis “Aum Shinrikyo preaches that there is an imminent apocalypse” for which there is currently no evidence. The bottom part shows two search criteria defined by the analyst: “Aum Shinrikyo preaches imminent apocalypse” and “Aum Shinrikyo predicts apocalypse”. The analyst can invoke BING, GOOGLE, or YAHOO with any one of these criteria to search for relevant evidence on the Internet. Other search engines could be integrated with Disciple-CD to search in other repositories.

Figure 4 Evidence collection

Hypothesis: Aum Shinrikyo preaches that there is an imminent apocalypse [REASONING]

Favoring evidence (0): No evidence.

Disfavoring evidence (0): No evidence.

Search for relevant evidence:

Search criterion: Aum Shinrikyo preaches imminent apocalypse [EDIT] [DELETE] [NEW]

- [Aum Shinrikyo predicts apocalypse](#)
- [Aum Shinrikyo preaches imminent apocalypse](#)

Search with: [BING] [GOOGLE] [YAHOO]

Figure 5 illustrates how the analyst has defined a new item of evidence in Disciple-CD, and has associated it with a hypothesis it favors. The analyst has clicked on the [NEW] label in the upper left pane, and Disciple-CD has opened an editor in the right pane where the analyst has defined the name (e.g. [EVD-002-material-acquisition](#)) and the description of the item of evidence, as shown in the upper right side of Figure 5. The bottom part of the right pane, under the label Irrelevant to, shows the leaf hypotheses, each followed by the labels [FAVORS] and [DISFAVORS]. The analyst has indicated that [EVD-002-material-acquisition](#) favors the hypothesis “it is relatively easy for a legitimate business to acquire production material for sarin based weapons” by clicking on its [FAVORS] label. As a result, Disciple-CD has created the new label Favors and has moved this hypothesis under it, as shown in Figure 5. At the same time, Disciple-CD has introduced the evidence item into the reasoning tree, as will be discussed in the next section.

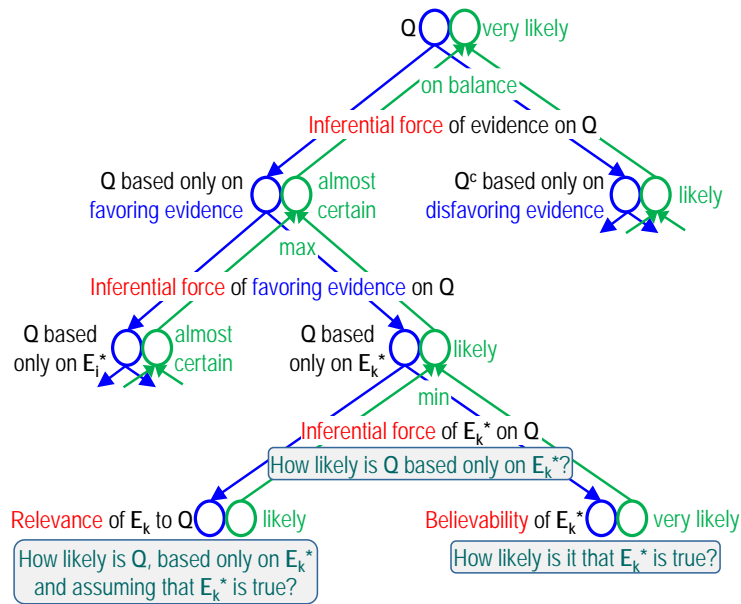
Figure 5 Evidence representation and use

<p>Select mode: [COLLECTION GUIDANCE] [COLLECTED INFORMATION] [AVAILABLE EVIDENCE] [IMPORT EVIDENCE]</p> <p>Available evidence [NEW] [DELETE]</p> <p>Sorted by: [ID] [NAME SUFFIX]</p> <p>EVD-001-chemist: Many have observed that Tsuchiya was impressively skilled as a chemist. Richard Danzig, Marc Sageman, Terrance Leighton, Lloyd Hough, Hidemi Yuki,...</p> <p>EVD-002-material-acquisition: A third factor that made chemical weapons easier to develop was that chemicals of the required purity were readily accessible at low visibility for...</p> <p>EVD-003-chem-exp-transfer: Although he made some improvements and innovations of his own, Tsuchiya mastered the basics of sarin production and a number of other hazardous or...</p> <p>EVD-004-bio-exp-transfer: By contrast, while Endo's efforts to identify, harvest and grow toxigenic C. botulinum and to create fully virulent B. anthracis also relied on...</p> <p>EVD-005-scale-up: A second factor, closely related to but distinguishable from the first, is that movement to scale seems notably easier for chemical than for...</p> <p>EVD-006-significant-funds: This second factor was mitigated for</p>	<p>Selected item of evidence: EVD-002-material-acquisition [RENAME] [DELETE EVIDENCE]</p> <p>Description: A third factor that made chemical weapons easier to develop was that chemicals of the required purity were readily accessible at low visibility for plausibly legitimate business purposes.</p> <p>Richard Danzig, Marc Sageman, Terrance Leighton, Lloyd Hough, Hidemi Yuki, Rui Kotani and Zachary M. Hosford, Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons, Center for a New American Security, July 2011, p.34 [EDIT]</p> <p>Extracted from: not specified (select a collected information to link with)</p> <p>Type: tangible evidence about testimonial evidence [CHANGE]</p> <p>Source of testimony: Danzig et al [RENAME] [CHANGE]</p> <p>Favors:</p> <ul style="list-style-type: none"> it is relatively easy for a legitimate business to acquire production material for sarin based weapons [REMOVE] [REASONING] [COLLECTION] <p>Irrelevant to:</p> <ul style="list-style-type: none"> Aum Shinrikyo preaches that there is an imminent apocalypse [FAVORS] [DISFAVORS] [REASONING] [COLLECTION] Aum Shinrikyo preaches that the apocalypse will start with chaos and destruction [FAVORS] [DISFAVORS] [REASONING] [COLLECTION]
--	---

6 Evidence-based Assessment

Now we present how Disciple-CD assesses the leaf hypotheses of the argumentation structure, based on the identified relevant evidence. Before continuing the concrete example from the previous section, we present an abstract example where the hypothesis to be directly analyzed by evidence is Q (see Figure 6).

Figure 6 The relevance, believability, and inferential force of evidence



Let us begin by discussing how to assess the probability of hypothesis Q based only on one item of favoring evidence E_k^* (see the bottom of Figure 6). First notice that we call this *likeliness* of Q , and not *likelihood*, because in classic probability theory likelihood is $P(E^*|Q)$, while here we are interested in $P(Q|E^*)$, the posterior probability of Q given E^* . With Disciple-CD, to assess Q based only on E_k^* , we have three judgments to make by answering three questions:

The *relevance* question is: *How likely is Q , based only on E_k^* and assuming that E_k^* is true?* If E_k^* tends to favor Q , then our answer should be one of the values from *likely* to *certain*. If E_k^* is not relevant to Q , then our answer should be *no support*, because E_k^* provides no support to the truthfulness of Q . If, however, E_k^* tends to disfavor Q , then it should be moved under the negation, or the complement of Q , that is under Q^c , because it provides support for the truthfulness of Q^c .

The *believability* question is: *How likely is it that E_k^* is true?* Here the answer should be one of the values from **no support** to **certain**. The maximal value, **certain**, means that we are sure that the event E_k reported in E_k^* has indeed happened. The minimal value, **no support**, means that E_k^* provides us no reason to believe that the event E_k reported in E_k^* has happened. For example, we believe that the source of E_k^* has lied to us.

The *inferential force or weight* question is: *How likely is Q based only on E_k^* ?* Disciple-CD automatically determines this answer as the minimum of the relevance and believability answers. Indeed, to believe that Q is true *based only on E_k^** , E_k^* should be both relevant to Q and believable.

When we assess a hypothesis Q , we may have several items of evidence, some favoring Q and some disfavoring Q . Disciple-CD uses the favoring evidence to assess the likeliness of Q , and the disfavoring evidence to assess the likeliness of Q^c . Because the disfavoring evidence for Q is favoring evidence for Q^c , the assessment process for Q^c is similar to the assessment for Q .

When we have several items of favoring evidence, we evaluate Q based on each of them (as was explained above), and then we compose the obtained results. This is illustrated in Figure 6 where the assessment of Q *based only on E_i^** (**almost certain**) is composed with the assessment of Q *based only on E_k^** (**likely**), through the maximum function, to obtain the assessment of Q *based only on favoring evidence* (**almost certain**). In this case the use of the maximum function is justified because it is enough to have one item of evidence that is both very relevant and very believable to persuade us that the hypothesis Q is true.

Let us now assume that Q^c *based only on disfavoring evidence* (for Q) is **likely**. How should we combine this with the assessment of Q *based only on favoring evidence*? As illustrated at the top of Figure 6, Disciple-CD uses an *on balance* judgment: Because Q is **almost certain** and Q^c is **likely**, it concludes that, *based on all available evidence*, Q is **very likely**.

In general, as indicated in the right and upper side of Table 4, if the assessment of Q^c (based on disfavoring evidence for Q) is higher than or equal to the assessment of Q (based on favoring evidence), then we conclude that, based on all the available evidence, there is no support for Q . If, on the other hand, the assessment of Q is strictly greater than the assessment of Q^c , then the assessment of Q is decreased, depending on the actual assessment of Q^c (see the left and lower side of Table 4).

One important aspect to notice is that the direct assessment of hypotheses based on favoring and disfavoring evidence is done automatically by Disciple-CD, once the relevance and the believability of evidence are assessed by the analyst.

Another important aspect to notice is that the evaluation of upper level hypotheses (such as those from Figure 2), requires the analyst to indicate what function to use when composing the assessments of their direct sub-hypotheses. This was discussed in Section 3.

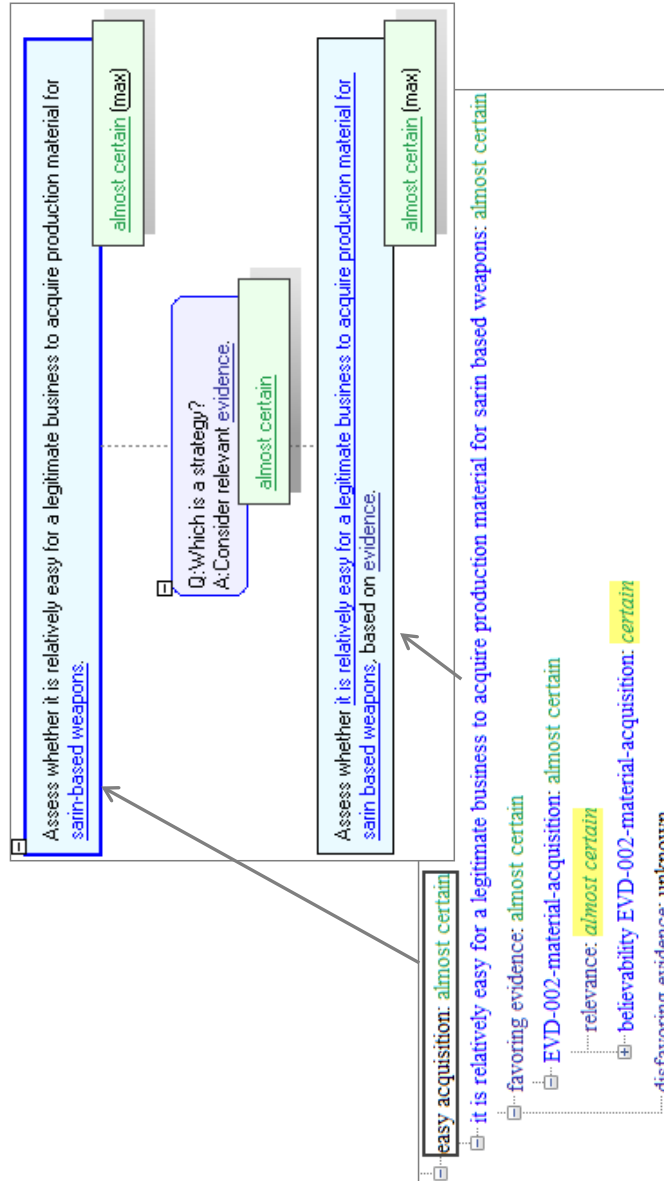
Table 4 The on balance function in Disciple-CD

		Q ^c based only on disfavoring evidence				
		no support	likely	very likely	almost certain	certain
Q based only on favoring evidence	Q based on all evidence	no support	no support	no support	no support	no support
	no support	no support	no support	no support	no support	no support
	likely	likely	no support	no support	no support	no support
	very likely	very likely	likely	no support	no support	no support
	almost certain	almost certain	very likely	likely	no support	no support
	certain	certain	almost certain	very likely	likely	no support

Let us now briefly illustrate the actual use of Disciple-CD to directly assess hypotheses based on relevant evidence. When the analyst associates an item of evidence with a hypothesis, as illustrated in Figure 5, Disciple-CD automatically generates a decomposition tree like the one in Figure 6. The bottom part of Figure 7 shows the abstraction of the tree that is automatically generated by Disciple-CD when the analyst indicates that the item of evidence [EVD-002-material-acquisition](#) favors the leaf hypothesis “it is relatively easy for a legitimate business to acquire production material for [sarin-based weapons](#).”

Disciple-CD automatically generates the reduction from the top of Figure 7, where the leaf hypothesis is reduced to the elementary hypothesis with the name “[it is relatively easy for a legitimate business to acquire production material for sarin based weapons](#),” to be directly assessed based on evidence. This elementary hypothesis corresponds to the hypothesis **Q** in Figure 6. Then Disciple-CD decomposes it as shown in the bottom part of Figure 7. This corresponds to the tree in Figure 6 except that there is only one item of favoring evidence, namely [EVD-002-material-acquisition](#). After that the analyst assesses the relevance of this item of evidence to the considered hypothesis (e.g., [almost certain](#)), as well as its believability (e.g., [certain](#)), and Disciple-CD automatically composes them, from bottom-up, to obtain the assessment of the leaf hypothesis. When the analyst adds additional items of evidence as either favoring or disfavoring evidence, Disciple-CD extends the reasoning tree from Figure 7 as indicated in Figure 6.

Figure 7 Evidence-based assessment of an elementary hypothesis



7 Believability Assessments based on a Substance-Blind Classification of Evidence

Figures 6 and 7 have illustrated how an analyst may directly assess the believability of an item of evidence. However, Disciple-CD has a significant amount of knowledge about the various types of evidence and its believability credentials, enabling an analyst to perform a much deeper believability analysis, as will be discussed in this section. The analyst may wish to perform such a detailed believability analysis for those items of evidence that are critical to the final result of the analysis. We will start with presenting a classification of evidence which is at the basis of the believability analysis.

Attempts to categorize evidence in terms of its substance or content would be a fruitless task, the essential reason being that the substance or content of evidence is virtually unlimited. What we have termed a *substance-blind* classification of evidence refers to a classification of recurrent forms and combinations of evidence, based not on substance or content, but on the inferential properties of evidence (Schum, 2001a, pp. 114 – 130). In what follows, we identify specific attributes of the believability of various recurrent types of evidence without regard to their substance or content.

Here is an important question we are asked to answer regarding the individual kinds of evidence we have: *How do you, the analyst, stand in relation to this item of evidence?* Can you examine it for yourself to see what events it might reveal? If you can, we say that the evidence is *tangible* in nature. But suppose instead you must rely upon other persons, assets, or informants, to tell you about events of interest. Their reports to you about these events are examples of *testimonial evidence*. Figure 8 shows a substance-blind classification of evidence based on its believability credentials. This classification is discussed in the following sections.

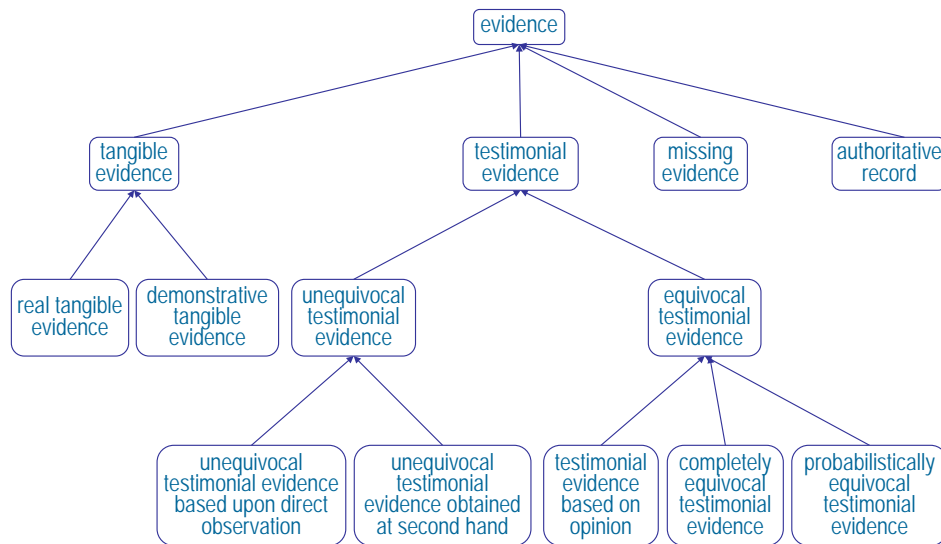
7.1 Tangible Evidence

There is an assortment of tangible items we might encounter and that could be examined by an intelligence analyst. Both IMINT and SIGINT provide various kinds of sensor records and images that can be examined. MASINT and TECHINT provide various objects such as soil samples and weapons that can be examined. COMINT can provide audio recordings of communications that can be overheard and translated if the communication has occurred in a foreign language. Documents, tabled measurements, charts, maps and diagrams or plans of various kinds are also tangible evidence.

There are two different kinds of tangible evidence: *real tangible evidence* and *demonstrative tangible evidence* (Lempert et al., 2000, pp. 1146–1148). Real tangible evidence is an actual thing and has only one major believability attribute: *authenticity*. *Is this object what it is represented as being or is claimed to be?* There are as many ways of generating deceptive and inauthentic evidence as there are persons wishing to generate it. Documents or written communications may be faked, captured weapons may have been

tampered with, and photographs may have been altered in various ways. One problem is that it usually requires considerable expertise to detect inauthentic evidence.

Figure 8 Substance-blind classification of evidence



Demonstrative tangible evidence does not concern things themselves but only representations or illustrations of these things. Examples include diagrams, maps, scale models, statistical or other tabled measurements, and sensor images or records of various sorts such as IMINT, SIGINT, and COMINT. Demonstrative tangible evidence has three believability attributes. The first concerns its *authenticity*. For example, suppose we obtain a hand drawn map from a captured insurgent showing the locations of various groups in his insurgency organization. Has this map been deliberately contrived to mislead our military forces or is it a genuine representation of the locations of these insurgency groups?


The second believability attribute is *accuracy* of the representation provided by the demonstrative tangible item. The *accuracy question* concerns the extent to which the device that produced the representation of the real tangible item had a degree of sensitivity (resolving power or accuracy) that allows us to tell what events were observed. We would be as concerned about the accuracy of the hand-drawn map allegedly showing insurgent group locations as we would about the accuracy of a sensor in detecting traces of some physical occurrence. Different sensors have different resolving power that also depends on various settings of their physical parameters (e.g., the settings of a camera).

The third major attribute, *reliability*, is especially relevant to various forms of sensors

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

The left side of Figure 9 shows how Disciple-CD assesses the believability of an item of demonstrative tangible evidence E_i^* as the minimum of its authenticity, accuracy, and reliability.

Figure 9 Assessing the believability of evidence with Disciple-CD



7.2 Testimonial Evidence

For *testimonial evidence* we have two basic sources of uncertainty: *competence* and *credibility*. The reason why it is more appropriate to talk about the *believability* of testimonial evidence (rather than *credibility*) is that *believability* is a broader concept that includes both *competence* and *credibility* considerations. The first question to ask related to *competence* is whether this source actually made the observation the source claims to have made or had access to the information the source reports. The second *competence* question concerns whether this source understood what was being observed well enough to provide us with an intelligible account of what was observed. Thus *competence* involves *access* and *understandability*.

Assessments of human source *credibility* require consideration of entirely different attributes: *veracity* (or *truthfulness*), *objectivity*, and *observational sensitivity under the conditions of observation* (Schum, 1989). Here is an account of why these are the major attributes of testimonial *credibility*. First, is this source telling us about an event this source believes to have occurred? This source would be untruthful if he/she did not believe the reported event actually occurred. So, this question involves the source's *veracity*. The second question involves the source's *objectivity*. The question is: Did this source base a belief on sensory evidence received during an observation, or did this source believe the reported event occurred either because this source expected or wished it to occur? An objective observer is one who bases a belief on the sensory evidence

instead of desires or expectations. Finally, if the source did base a belief on sensory evidence, how good was this evidence? This involves information about the source's relevant *sensory capabilities and the conditions under which a relevant observation was made*.

As indicated in Figure 8, there are several types of testimonial evidence. If the source does not hedge or equivocate about what the source observed (i.e., the source reports that he/she is certain that the event did occur), then we have *unequivocal testimonial evidence*. If, however, the source hedges or equivocates in any way (e.g., "I'm fairly sure that E occurred") then we have *equivocal testimonial evidence*. The first question we would ask this source of *unequivocal testimonial evidence* is: How did you obtain information about what you have just reported? It seems that this source has three possible answers to this question. The first answer is: "I made a *direct observation* myself." In this case we have *unequivocal testimonial evidence based upon direct observation*. The second possible answer is: "I did not observe this event myself but heard about its occurrence (or nonoccurrence) from another person." Here we have a case of secondhand or hearsay evidence, called *unequivocal testimonial evidence obtained at second hand*. A third answer is possible: "I did not observe event E myself nor did I hear about it from another source. But I did observe events C and D and inferred from them that event E definitely occurred." This is called *testimonial evidence based on opinion* and it raises some very difficult questions. The first concerns the source's credibility as far as his/her observation of events C and D; the second involves our examination of whether we ourselves would infer E based on events C and D. This matter involves our assessment of the source's *reasoning ability*. It might well be the case that we do not question this source's credibility in observing events C and D, but we question the conclusion that event E occurred the source has drawn from his/her observations. We would also question the certainty with which the source has reported the opinion that E occurred. Despite the source's conclusion that "event E definitely occurred", and because of many sources of uncertainty, we should consider that *testimonial evidence based on opinion* is a type of *equivocal testimonial evidence*.

There are two other types of equivocal testimonial evidence. The first we call *completely equivocal testimonial evidence*. Asked whether event E occurred or did not, our source says: "I don't know", or "I can't remember".

But there is another way a source of HUMINT can equivocate; the source can provide *probabilistically equivocal testimonial evidence* in various ways: "I'm 60 percent sure that event E happened"; or "I'm fairly sure that E occurred"; or "It is very likely that E occurred". We could look upon this particular probabilistic equivocation as an assessment by the source of his/her own observational sensitivity.

The right side of Figure 9 shows how Disciple-CD assesses the believability of an item of testimonial evidence based upon direct observation by a source, as the minimum of source's competence and credibility. Source's competence is assessed as the minimum of his/her access and understandability, while source's credibility is assessed as the

minimum of his/her veracity, objectivity, and observational sensitivity.

7.3 Missing Evidence

To say that evidence is missing entails that we must have had some basis for expecting we could obtain it. There are some important sources of uncertainty as far as missing evidence is concerned. In certain situations missing evidence can itself be evidence. Consider some form of tangible evidence, such as a document, that we have been unable to obtain. There are several reasons for our inability to find it, some of which are more important than others. First, it is possible that this tangible item never existed in the first place; our expectation that it existed was wrong. Second, the tangible item exists but we have simply been looking in the wrong places for it. Third, the tangible item existed at one time but has been destroyed or misplaced. Fourth, the tangible item exists but someone is keeping it from us. This fourth consideration has some very important inferential implications including denial and possibly deception. An adverse inference can be drawn from someone's failure to produce evidence.

7.4 Accepted Facts

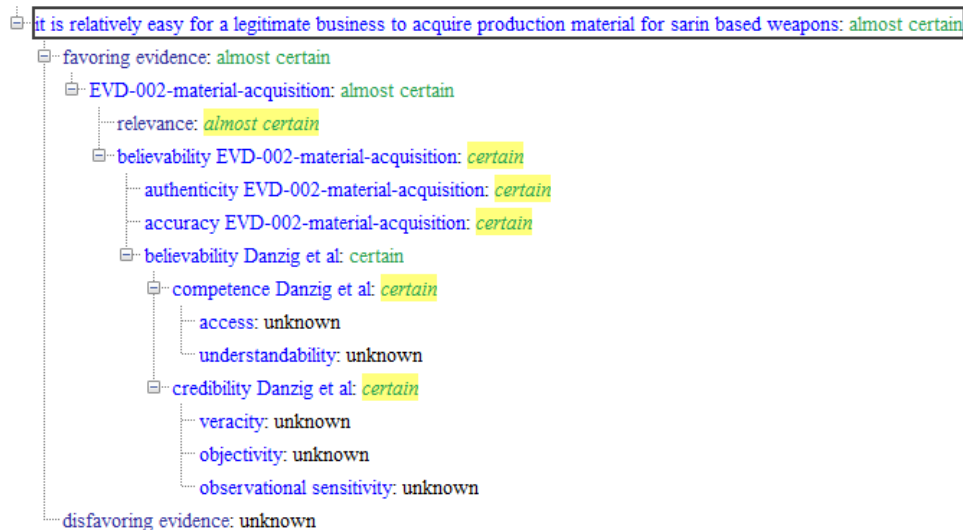
There is one final category of evidence about which we would never be obliged to assess its believability. Tabled information of various sorts such as tide tables, celestial tables, tables of physical or mathematical results such as probabilities associated with statistical calculations, and many other tables of information we would accept as being believable provided that we used these tables correctly. For example, an analyst would not be obliged to prove that temperatures in Iraq can be around 120 degrees Fahrenheit in summer months, or that the population of Baghdad is greater than that of Basra.

7.5 Mixed Evidence and Chains of Custody

We have just described a categorization of individual items of evidence. But there are situations in which individual items can reveal various mixtures of the types of evidence shown in Figure 8. One example is testimonial evidence about tangible evidence where a source describes a weapon observed at a scene of a crime. Another example is a tangible document containing a testimonial assertion based on other alleged tangible evidence. As a concrete example, [EVD-002-material-acquisition](#) in Figure 5 is a statement in the copy of a research report on Aum Shinrikyo by the authors of that report (Danzig et al., 2011). This is an item of tangible evidence about testimonial evidence and Disciple-CD analyzes its believability as shown in Figure 10. It considers the authenticity and the accuracy of the copy of the report, and the believability of the authors of the report. Further on, Disciple-CD reduces the authors' believability to their competence and credibility. Finally it reduces the authors' competence to their access and understandability, and their

credibility to their veracity, objectivity, and observational sensitivity. The analyst may now assess these lowest level credentials, and Disciple-CD automatically determines the believability of EVD-002-material-acquisition. Alternatively, the analyst may assess some upper-level credentials, such as competence and credibility, as illustrated in Figure 10.

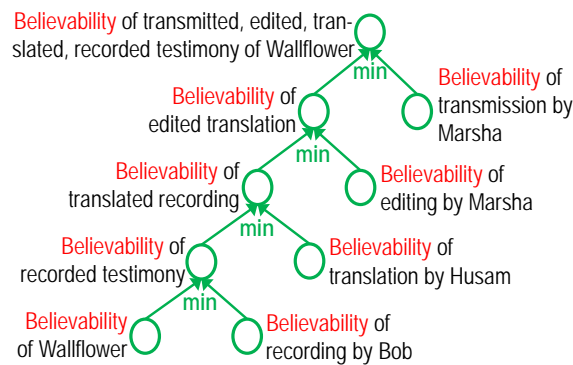
Figure 10 Believability analysis of tangible evidence about testimonial evidence



But the believability analysis of an item of evidence can be even more complicated. Very rarely, if ever, has the analyst access to the original evidence. Most often, what is being analyzed is an item of evidence that has undergone a series of transformations through a *chain of custody* (Schum, et al., 2009). Here we have borrowed an important concept from the field of law where a chain of custody refers to the persons or devices having access to the original source evidence, the time at which they had such access, and what they did to the original evidence when they had access to it. These chains of custody add three major sources of uncertainty for intelligence analysts to consider, that are associated with the persons in chains of custody whose competence and credibility need to be considered. The first and most important question involves *authenticity*: Is the evidence received by the analyst exactly what the initial evidence was and is it complete? The other questions involve assessing the *reliability* and *accuracy* of the processes used to take various actions on the evidence in the chain of custody. As an illustration, consider an item of testimonial HUMINT coming from a foreign national whose code name is “Wallflower”, who does not speak English. Wallflower gives his report to the *case officer* Bob. This report is *recorded* by Bob and then *translated* by Husam. Then Wallflower’s

translated report is *transmitted* to the *report's officer* Marsha who *edits* it and *transmits* it to the analyst Clyde who evaluates it. Figure 11 shows how Disciple-CD may determine the believability of the evidence received by the analyst. A more detailed discussion is provided in (Schum et al., 2009).

Figure 11 Chain of custody of Wallflower's testimony



The case officer might have intentionally overlooked details in his recording of Wallflower's report. Thus, as shown at the bottom of Figure 11, the believability of the recorded testimony of Wallflower is the minimum between the believability of Wallflower and the believability of the recording. Then Husam, the translator, may have intentionally altered or deleted parts of this report. Thus, the believability of the translated recording is the minimum between the believability of the recorded testimony and the believability of the translation by Husam. Then Marsha, the report's officer, might have altered or deleted parts of the translated report of Wallflower's testimony in her editing of it, and so on. The result of these actions is that Clyde, the analyst receiving this evidence, almost certainly did not receive an authentic and complete account of it, nor did he receive a good account of its reliability and accuracy. What Clyde received was the transmitted, edited, translated, recorded testimony of Wallflower. Although the information to make such an analysis will not, in most cases, be available, the analyst should adjust the confidence in his conclusion, in recognition of these factors.

8 Drill-down Analysis, Assumption-based Reasoning, and What-If Scenarios

A major objective of intelligence analysis is to help insure that the policies and decisions reached by the governmental and military leaders, at all levels, are well informed. The policy-relevance of analytic "products" is a goal routinely kept in mind. Analysts face

different requirements in their efforts to serve these policy and decision-making "customers". In some cases current analyses are required to answer questions that are of immediate interest and that do not allow analysts time for extensive research and deliberation on available evidence regarding the questions being asked. In other cases, teams of analysts participate in more lengthy analyses that combine evidence from every available source to make long-term assessments on matters of current and abiding interest.

An important feature of Disciple-CD is that it allows an analyst to perform analyses at different levels of detail. What this means is that a hypothesis may be reduced to many levels of sub-hypotheses or just a few levels that are then assessed based on relevant evidence. The same applies to assessing the believability of evidence. The analyst may directly assess it, as was illustrated in Figure 7 where the believability of [EVD-002-material-acquisition](#) was assessed as *certain*. But if [EVD-002-material-acquisition](#) has an important influence on the result of the analysis, then the analyst may perform a deeper believability analysis, as was illustrated in Figure 10, where the analyst assessed lower level believability credentials: the authenticity and the accuracy of [EVD-002-material-acquisition](#), and the competence and the credibility of its source. The analyst could have drilled even deeper to assess the source's veracity, objectivity, and observational sensitivity, instead of source's credibility.

It may also happen that the analyst does not have the time or the evidence to assess a sub-hypothesis, in which case he/she may make various assumptions with respect to its likeliness. Consider, for example, the analysis from Figure 2 and the four sub-hypotheses of the top-level hypothesis. Two of these sub-hypotheses have been analyzed as discussed in the previous sections. However, for the other two sub-hypotheses the analyst has made the following assumptions:

It is *certain* that [Aum Shinrikyo](#) desires to acquire and use [sarin-based weapons](#).

It is *very likely* that [Aum Shinrikyo](#) is developing capabilities to use [sarin-based weapons](#).

Assumptions are distinguished from system-computed assessments by the fact that the assumed probabilities have a yellow background.

The analyst may provide justifications for the assumptions made. He/she may also experiment with various what-if scenarios, where he/she makes different assumptions to determine their influence on the final result of the analysis.

Thus Disciple-CD gives the analyst the flexibility to perform the analysis that makes the best use of his/her time constraints and available evidence.

9 Learning and Reuse of Analytic Expertise

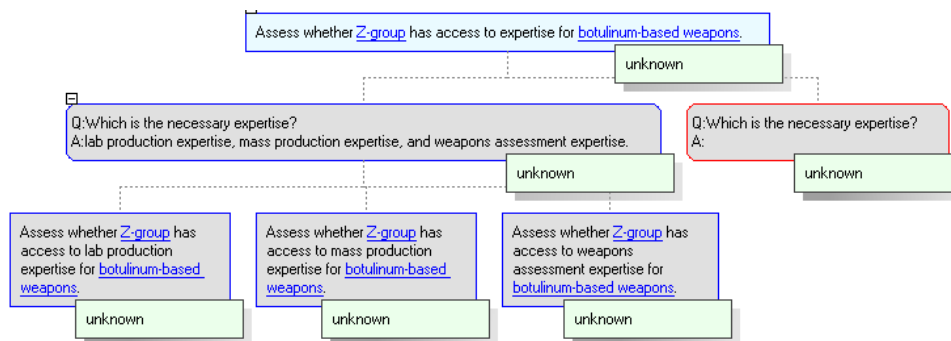
Disciple-CD is a cognitive agent which includes a significant amount of knowledge about evidence and its credentials, enabling it to perform the reasoning discussed in the

previous sections. It has been developed by using the Disciple learning agent shell (Tecuci et al., 2013). As such, it has various learning capabilities, some to be used by expert analysts with limited knowledge engineering support, and some by end-user analysts by themselves, with no support. For example, Disciple-CD may be taught by an expert analyst (with limited support from a knowledge engineer) in a way that is similar to how the expert would teach a student, through examples and explanations, and by supervising and correcting agent's reasoning. In essence, the expert will show the agent a sample reasoning tree like the one in Figure 2, and the agent will learn reasoning rules for developing argumentation structures and assessing similar hypotheses (Tecuci 1998; Tecuci et al., 2005; 2008; Boicu et al., 2011; 2012).

Here we only briefly illustrate how Disciple-CD can easily learn analytic expertise directly from an end-user analyst who does not have any knowledge engineering experience, and receives no knowledge engineering support.

Consider again the analysis from Figure 2. From each reduction step, Disciple-CD learns a reduction pattern. Then, when the analyst analyses a similar hypothesis, such as "Assess whether *Z-group* has access to expertise for *botulinum-based weapons*", Disciple-CD suggests a similar reduction, as shown in Figure 12. The analyst may accept the suggestion or may define a different reduction.

Figure 12 Reduction suggested by Disciple-CD based on a learned analysis pattern



10 Conclusions

This paper presented a computational approach to intelligence analysis and its implementation in the Disciple-CD cognitive assistant. This provides a rigorous systematic framework for intelligence analysis as ceaseless discovery of evidence, hypotheses, and arguments in a complex world that is changing all the time. Disciple-CD incorporates a significant amount of knowledge from the science of evidence and uses it in the analysis. It knows about the substance-blind classification of evidence and about

the ingredients of believability assessments for tangible as well as testimonial evidence, knowledge which allows it to develop theoretically-justified argumentation structures for believability assessments. It supports the development of relevance arguments linking evidence to hypotheses. It provides simple Baconian-Fuzzy methods for dealing with uncertainty as opposed to elaborate equations or extensive probability assessments necessary in approaches based on Bayes' rule or Belief Functions using Dempster's rule (Schum, 2001a). It also knows how to analyze various types of hypotheses and enforces necessary conditions for sound analysis, such as considering both favoring and disfavoring evidence, or qualifying each analytic conclusion with the assumptions made.

Disciple-CD is also concerned about the many demands placed on analysts and does allow for particular simplification methods. However, these simplifications are not mandated but chosen by the analyst. In particular, Disciple-CD allows the analyst to drill-down to various levels in the analysis at hand, to make assumptions concerning various verbal assessments of uncertainty, and to revise these assumptions in light of new evidence.

The use of Disciple-CD for intelligence analysis education in various civilians, military, and intelligence organizations is supported by a textbook which contains numerous case studies of analysis, enabling a quite unique learning by doing experience (Tecuci et al., 2011c). The types of hypotheses that Disciple-CD has been taught to analyze include: A non-state actor has nuclear weapons; A state actor is pursuing a nuclear program for military purposes; A state actor is pursuing a nuclear program for economic purposes; There is an ambush threat at a certain location; A terrorist organization will set-off a dirty bomb in a certain location; A government will crack down on its opposition; The government of a country supports the insurgency in another country, A professor would be a good PhD advisor for a student; A website is believable, etc.

Acknowledgements

This material is based on research performed in the Learning Agents Center of George Mason University and was partially sponsored by several U.S. Government organizations, including the Department of Defense and the Intelligence Community. 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 paper are those of the authors and do not necessarily reflect the official policy or position of any agency of the U.S. government.

References

- Boicu, M., Tecuci, G., Marcu, D. (2012) Rapid Argumentation Capture from Analysis Reports: The Case Study of Aum Shinrikyo, in *Proceedings of the 7th International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS 2012*, Fairfax, VA, 23-26 October.
- Cohen, J. (1977) *The Probable and the Provable*, Oxford, UK, Clarendon Press.
- Cooper, J.R. (2008) *Curing Analytic Pathologies*, Center for the Study of Intelligence, Central Intelligence Agency, 2008, <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/curing-analytic-pathologies-pathways-to-improved-intelligence-analysis-1/index.html>
- Danzig, R., Sageman, M., Leighton, T., Hough, L., Yuki, H., Kotani, R. and Hosford, Z.M. (2011) *Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons*, Center for a New American Security, Washington, DC, July.
- Heuer, R.J. (1999) *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington, D.C.
- Heuer, R.J. (2008) Computer-Aided Analysis of Competing Hypotheses, in Roger Z. George and James B. Bruce, eds., *Analyzing Intelligence: Origins, Obstacles, and Innovations*.
- Lempert, R.O., Gross, S.R. and Liebman, J.S. (2000) *A Modern Approach to Evidence*, 3rd ed., St. Paul, MN: West Publishing.
- Nilsson, N.J. (1971) *Problem Solving Methods in Artificial Intelligence*, New York: McGraw-Hill.
- Peirce, C.S. (1898) *Reasoning and the Logic of Things* (Kenneth L. Ketner, ed., Harvard University Press, Cambridge, MA, 1992).
- Peirce, C.S. (1901) Abduction and Induction, in Justus Buchler, ed., *Philosophical Writings of Peirce* (Dover, New York, NY, 1955), pp. 150-156.
- Powell, G.M. and Schmidt, C.F. (1988) A First-order Computational Model of Human Operational Planning, *CECOMTR-01-8*, US Army CECOM, Fort Monmouth, NJ.
- Schum, D.A. (1987) *Evidence and Inference for the Intelligence Analyst* (2 Vols), University Press of America, MD: Lanham.
- Schum, D.A. (1989) Knowledge, Probability, and Credibility, *Journal of Behavioral Decision Making*, Vol. 2, pp. 39-62.
- Schum, D.A. (2001a) *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.
- Schum, D.A. (2001b) Species of Abductive Reasoning in Fact Investigation in Law, *Cardozo Law Review*, 22 (5-6), pp. 1645–1681.
- Schum, D.A. (2009) Classifying Forms and Combinations of Evidence: Necessary in a Science of Evidence, in *Evidence, Inference and Inquiry*, The British Academy, Oxford University Press.
- Schum, D.A., Tecuci, G., and Boicu, M. (2009) Analyzing Evidence and its Chain of Custody: A Mixed-Initiative Computational Approach, *International Journal of Intelligence and Counterintelligence*, Vol. 22, No. 2, pp. 298-319.
- Tecuci, G. (1988) DISCIPLINE: A Theory, Methodology and System for Learning Expert Knowledge, *Doctoral dissertation*, University of Paris-South, France.

- Tecuci, G. (1998) *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, San Diego: Academic Press.
- Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M. (2005) The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence*, Vol.21, No.4, pp. 462-479.
- Tecuci, G., Boicu, M., Marcu, D., Le, V. and Boicu, C. (2008) Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development*, <http://iac.gmu.edu/publications/2008/Disciple-LTA08.pdf>
- Tecuci, G., Schum, D., Boicu, M., Marcu, D., Russell, K. (2011a) Toward a Computational Theory of Evidence-based Reasoning, in *Proceedings of the 18th International Conference on Control Systems and Computer Science*, University Politehnica of Bucharest, Romania, 24-27 May.
- Tecuci, G., Marcu, D., Boicu, M., Schum, D.A. and Russell, K. (2011b) Computational Theory and Cognitive Assistant for Intelligence Analysis, in *Proceedings of the 6th International Conference on Semantic Technologies for Intelligence, Defense, and Security*, pp. 68-75, Fairfax, VA.
- Tecuci, G., Schum, D.A., Boicu, M., Marcu, D. (2011c) *Introduction to Intelligence Analysis: A Hands-on Approach with TIACRITIS*, 220 pages, Learning Agents Center, George Mason Univ.
- Tecuci, G., Boicu, M., Marcu, D., Schum, D.A. (2013) How Learning Enables Intelligence Analysts to Rapidly Develop Practical Cognitive Assistants, in *Proceedings of the 12th International Conference on Machine Learning and Applications (ICMLA'13)*, Miami, Florida, December 4-7.
- Toulmin, S. E. (1963) *The Uses of Argument*, Cambridge University Press.
- Wigmore, J.H. (1937) *The Science of Judicial Proof*, Boston: Little, Brown.
- Zadeh, L. (1983) The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems, *Fuzzy Sets and Systems*, Vol.11, pp. 199-227.