

## Evidence-based Reasoning: Computational Theory and Cognitive Assistants

Gheorghe TECUCI, Mihai BOICU, Dorin MARCU,  
David A. SCHUM

Learning Agents Center, George Mason University  
4400 University Drive, Fairfax, VA, USA

E-mail: {tecuci, mboicu, dmarcu}@gmu.edu,  
dschum398@earthlink.net

**Abstract.** Evidence-based reasoning which is at the core of many problem solving and decision making tasks in a wide variety of domains, requires a complex combination of imaginative reasoning, critical reasoning, and expert knowledge. This paper presents research performed in the Learning Agents Center of George Mason University on developing a computational theory of evidence-based reasoning and on its integration within the Disciple learning agents theory, methodology, and tools for the development of cognitive assistants that are able to learn complex problem solving expertise directly from human experts, assist experts and non-experts with problem solving and decision making, and teach the acquired expertise to students. The paper introduces evidence-based reasoning concepts that need to be understood in order to use such a cognitive assistant, illustrating them in the context of a specific system developed for intelligence analysis, named Cogent. It also illustrates a sample session with Cogent and summarizes the assistance it provides to the user. The paper concludes with the envisioned application of these results to other evidence-based reasoning domains.

**Key-words:** evidence, evidence-based reasoning, cognitive assistant, intelligence analysis, argumentation, mixed-initiative reasoning, learning, ontology, patterns of reasoning, hypotheses analysis.

## 1. Introduction

Evidence is any observable sign, indicator, or datum that is relevant in deciding whether a hypothesis of interest is true or false. Evidence-based reasoning (EBR) is at the core of many problem solving and decision making tasks in a wide variety of domains, including physics, chemistry, history, archaeology, medicine, law, forensics, intelligence analysis, cybersecurity, etc. It is important to realize that most of the information we use in these tasks does not consist of facts but of evidence about these facts. Indeed, each information item comes from some source, and we should always question the believability of this source. Additionally, the evidence we have is always incomplete (we can look for more, if we have time), usually inconclusive (it is consistent with the truth of more than one hypothesis), frequently ambiguous (we cannot always determine exactly what the evidence is telling us), and commonly dissonant (some of it favors one hypothesis but other evidence favors other hypotheses) [1, 2]. These, as well as other characteristics, make EBR tasks difficult and time-consuming.

In this paper we overview some of our work on a computational theory of evidence-based reasoning and its use for building software agents that assist end-users with EBR tasks. This theory has been integrated within the Disciple learning agents theory, methodology, and tools for the development of cognitive assistants that are able to:

- learn complex problem solving expertise directly from human experts;
- assist non-expert users solve problems requiring subject matter expertise;
- assist human experts in complex problem solving and decision making;
- teach problem solving and decision making to students.

The Disciple approach relies on developing a powerful learning agent shell that can be taught by a subject matter expert (who does not have computer science or knowledge engineering experience) in ways that are similar to how the expert would teach a student or an apprentice, by explaining problem solving examples to it, and by supervising and correcting its problem solving behavior. In this way, the resulting agent learns to replicate the problem-solving behavior of its human expert [3-9].

A main result of incorporating the computational theory of evidence-based reasoning into the Disciple learning agents theory is that the generic learning agent shell now contains a significant amount of knowledge about the properties, uses, discovery, and marshaling of evidence from the science of evidence [2, 10]. Therefore, when developing a cognitive assistant for a specific EBR task, such as intelligence analysis or cybersecurity, the agent already contains a significant amount of general evidentiary knowledge and only needs to be trained with domain-specific knowledge.

In the next section we overview the current computational theory of evidence-based reasoning embedded into the Disciple approach. Then, in Section 3 we

introduce some of the EBR concepts that need to be understood in order to use a cognitive assistant, illustrating them in the context of a specific cognitive assistant, Cogent [11], that we have developed for intelligence analysis. After that, we describe a sample session with Cogent and summarize the assistance it provides to the user. We conclude with the envisioned application of these results to other evidence-based reasoning domains.

## 2. Computational Theory of Evidence-based Reasoning and Disciple-EBR

Developed in the framework of the scientific method, the computational theory of evidence-based reasoning views EBR as *ceaseless discovery of evidence, hypotheses, and arguments* in a non-stationary world, involving collaborative computational processes of *evidence in search of hypotheses, hypotheses in search of evidence*, and *evidentiary assessment of hypotheses*, performed jointly by a person and his or her knowledge-based cognitive assistant (see Fig. 1):

- ***Evidence in search of hypotheses.*** From the observations they make or the questions they ask about a situation of interest, they generate alternative hypotheses. The question is, *What hypotheses would explain this observation?* or *What are possible answers to this question?* Through *abductive (imaginative)* reasoning, which shows that something is *possibly* true [12-16], they generate a set of *alternative hypotheses* that may explain the observation, or a set of hypothesized answers to the question (see the red, left-hand side, of Fig. 1).
- ***Hypotheses in search of evidence.*** The user and the cognitive assistant put each of the generated hypothesis to work, guiding the collection of relevant evidence. The question is, *What evidence would be observable if this hypothesis were true?* Through *deductive* reasoning, which shows that something is *necessarily* true, they decompose the hypothesis into simpler and simpler hypotheses, and use the simplest hypotheses to generate new lines of inquiry and discover new evidence. The reasoning might go as follows: If H were true then the sub-hypotheses  $H_1$ ,  $H_2$ , and  $H_3$  would also need to be true. But if  $H_1$  were true then one would need to observe evidence  $E_1$ , and so on (see the blue, middle side, of Fig. 1). This process leads both to the discovery of new evidence and to the construction of a tree-like Wigmorean argumentation structure [1, 2, 17, 18] that can now be used to assess the probability of the hypothesis, based on the discovered evidence.
- ***Evidentiary assessment of hypotheses.*** The user and the cognitive assistant use the obtained evidence to assess the generated alternative hypotheses. The question is, *What is the evidence-based probability of each hypothesis?* Through *inductive (probabilistic)* reasoning, which shows that

something is *probably* true, they assess the *relevance*, the *believability* or *credibility*, and the *inferential force* or *weight* of evidence. Then they combine these evidence credentials in complex ways to assess the probability of the hypothesis [1, 2]. This is a process of *multi-INT fusion* since, in general, the assessment of a hypothesis involves fusing different types of evidence from a variety of sources (see the green, right-hand side, of Fig. 1).

Notice that, as shown at the bottom of Fig. 1, this is a recursive process where, for example, the discovery of new evidence may lead to the generation of new hypotheses or the modification of the existing ones which, in turn, may lead to the discovery of new evidence.

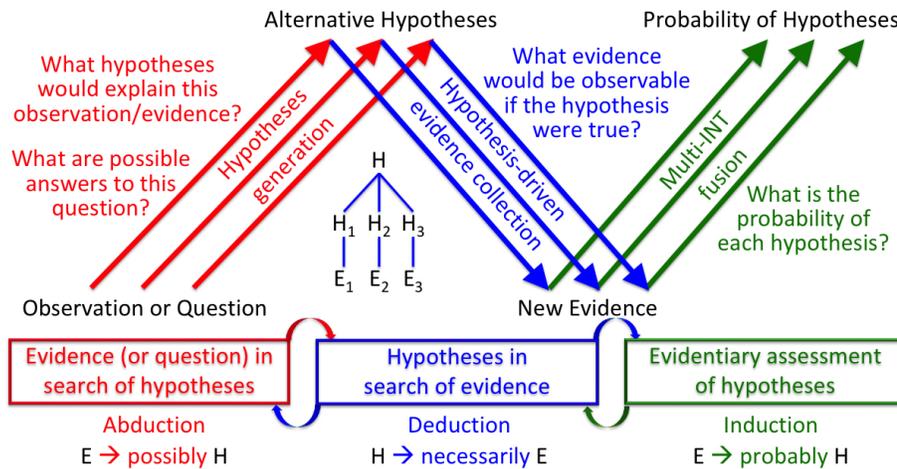


Fig. 1. Evidence-based reasoning framework.

In addition to providing a general framework for evidence-based reasoning, the computational theory also provides integrated computational models for essential reasoning tasks: hypotheses generation (guided by intelligence questions or evidence), evidence collection (driven by the generated hypotheses), and hypotheses assessment (based on the obtained evidence, through a combination of logic and Baconian probabilities [19, 20] with Fuzzy qualifiers [21, 22]).

The integration of some of these models with the knowledge representation, reasoning, and learning methods and modules of the Disciple learning agent shell have led to the development of Disciple-EBR, the Disciple learning agent shell for evidence-based reasoning [23]. The overall architecture of this shell is shown in Fig. 2. The shell includes multiple modules for problem solving, learning, evidence-based reasoning, mixed-initiative interaction, evidence management, and repository management. It also includes a hierarchically organized knowledge base (KB) additionally structured into an ontology of concepts and a set of general problem solving rules expressed with these concepts. At the top

level of the hierarchy is the general knowledge base for evidence-based reasoning (EBR KB), containing knowledge applicable to evidence-based reasoning in any domain, including a general ontology of evidence, and general believability analysis patterns. Under the EBR KB and inheriting from it, are domain-specific knowledge bases. Each such Domain KB contains knowledge on a specific EBR domain (*e.g.*, intelligence analysis) and the type of problems in that domain (*e.g.*, predictive analysis related to energy sources, or assessments related to the current production of weapons of mass destruction by various actors). Under each Domain KB there are several Scenario KBs, each corresponding to an instance of a problem pattern from that domain, such as: “Assess whether the United States will be a world leader in wind power within the next decade.” This Scenario KB will contain specific knowledge about the United States, as well as items of evidence to make the corresponding analysis. The actual analysis will be done by using this knowledge as well as more general knowledge inherited from the corresponding Domain KB and from the EBR KB.

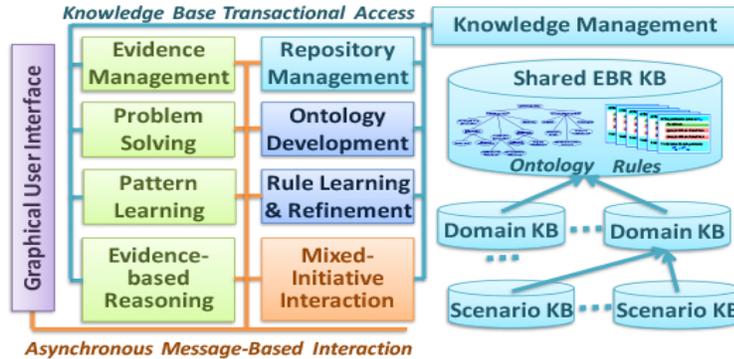


Fig. 2. The overall architecture of Disciple-EBR.

The Disciple-EBR shell can be customized to a specific EBR domain, by possibly extending it with special interfaces and domain-specific modules, and by training it on how to solve representative problems in that domain. This will result in the population of the Domain KB with a domain-specific ontology and with reasoning rules.

In the next section we introduce some of the EBR concepts that need to be understood in order to use a cognitive assistant, illustrating them in the context of a specific cognitive assistant, Cogent, that we have developed for intelligence analysis.

### 3. Evidence-based Reasoning Concepts

We will illustrate how Cogent might have been used by an intelligence analyst who was monitoring *Aum Shinrikyo*, a Japanese apocalyptic sect [24, 25]. As

shown in Figure 3, the analysis starts with a situation of interest, such as “*Aum Shinrikyo*”. For this situation, several intelligence questions are asked, such as “Is *Aum Shinrikyo* a threat?” Then, for each intelligence question, several answers are hypothesized, such as “*Aum Shinrikyo* has *sarin-based weapons*.” Finally, each of these hypotheses is assessed, based on evidence. Notice also that some of the words, such as *Aum Shinrikyo*, appear in blue and are underlined. This is because they are part of Cogent’s knowledge base, and are recognized by it.

The considered *hypotheses are not required to be disjoint*. For example, both hypotheses from the bottom of Fig. 3 may be true at the same time since *Aum Shinrikyo* may have both *sarin-based weapons* and *botulinum-based weapons*. This is because, as will be discussed in the following, each of the considered hypotheses is analyzed independently of one another. The probability of each hypothesis depends only on the developed argumentation structure for that hypothesis and the evaluation of the available evidence, as illustrated in Fig. 5. Moreover, if there are disfavoring arguments and evidence, they must be made explicit and indicated for each hypothesis in part. For the same reason, *the set of the considered hypotheses is not required to be complete* either. At each moment the analyst may formulate additional hypotheses, for example “*Aum Shinrikyo* has *nuclear weapons*.” This will not change the probability of the other hypotheses, nor will the probability of this hypothesis be changed because of changes in the probabilities of the other hypotheses considered. These features significantly facilitate hypotheses analysis in the context of a dynamic world that is changing all the time, where new evidence may suggest the formulation of additional hypotheses or the modification of the existing ones, as indicated by the red and blue arrows from the bottom of Fig. 1.



**Fig. 3.** Hypotheses as possible answers to intelligence questions.

Hypothesis assessment is necessarily probabilistic in nature because, as already noticed, our evidence is always incomplete, usually inconclusive, frequently ambiguous, commonly dissonant, and with various degrees of believability [1, 2]. Unfortunately, none of the probability systems known to us can optimally cope with all these features. Table 1, for instance, shows the most well-known non-enumerative probability systems and which of the above features are best dealt with. For example, only the Baconian system [19, 20] can

account for the incompleteness of the coverage of evidence, but it cannot cope with the ambiguities or imprecision in evidence. On the other hand, the Belief Functions [26] and the Fuzzy system [21, 22] can optimally cope with the ambiguity of evidence, while a Subjective Bayesian [1] has difficulties with it. What Table 1 suggests is that an approach that combines or allows different probability views may deal with all the five characteristics of evidence. This is the approach taken by the computational theory of evidence-based reasoning and Cogent, which uses a probability system that integrates elements of the Baconian and Fuzzy systems, within a Wigmorean probabilistic inference network [1]. This was possible because of the use of similar min/max probability combination rules by the Baconian and the Fuzzy systems ([1, p.255, p.266], [19, pp.167-187], [21, pp.340-341]). These rules are also much simpler than the Bayesian probability combination rules, which is important for the understandability of the agent’s reasoning.

**Table 1.** Non-enumerative probability systems and what they best capture

Evidence characteristic	Subjective Bayes	Belief Functions	Baconian	Fuzzy
Incompleteness			☑	
Inconclusiveness	☑	☑	☑	☑
Ambiguities		☑		☑
Dissonance	☑	☑	☑	☑
Source believability	☑		☑	

There is also the issue of using a numerical probability scale, which is required by a Bayesian view, as opposed to a symbolic scale required by a Fuzzy view. While a numerical probability is much more precise, it is not at all clear how a user would be able to defend his or her subjective assessment that, for instance, the probability of the hypothesis  $H_k$  is exactly 77%. Such precise assessments would necessarily lead to variations in the assessment of probabilities by different users, which would impede their collaboration. Because words are less precise than numbers, there will often be less disagreement about a verbal or fuzzy probability.

Starting from such considerations, we have defined an intuitive and easy to use system of *Baconian probabilities* [19, 20] with *Fuzzy qualifiers* [21, 22]. This symbolic probability system may be used with different assessment scales, such as the following one:

lack of support(LS) < likely(L) < very likely(VL) < almost certain(AC) < certain(C)

In this scale, there may be a “lack of support (LS)” from the available evidence to the considered hypothesis, or the evidence may indicate some level of support, such as “very likely (VL)”. Each symbolic probability value (*e.g.*,

“very likely”) is abbreviated (“VL”) in the Cogent analysis whiteboard in order to reduce space usage and facilitate the visualization of larger argumentations.

### 3.1. Evidence-based Hypothesis Assessment

One can directly assess a hypothesis based on an item of evidence by assessing the credentials of this item of evidence, as illustrated in the left-hand side of Fig. 4, and explained in the following.

First the analyst assesses the *believability* of the item of evidence by answering the question, “What is the probability that what this item of evidence is telling us is true?” Let us assess this as “almost certain (AC)”, as shown at the bottom-left of Fig. 4. Next the analyst assesses the *relevance* of the item of evidence to the hypothesis by answering the question, “What would be the probability of the hypothesis if this item of evidence were true?” That is, assuming that “Chemicals of the purity required for sarin-based weapons are readily accessible at low visibility for plausibly legitimate business purposes,” what is the probability that “It is safe [*i.e.*, not suspicious] for a legitimate chemical business, such as that of Aum Shinrikyo, to acquire chemicals for sarin-based weapons?” Let us assess this as *very likely* (VL). The minimum of these answers (*i.e.*, *very likely*) is the *inferential force* of this item of evidence on the hypothesis.

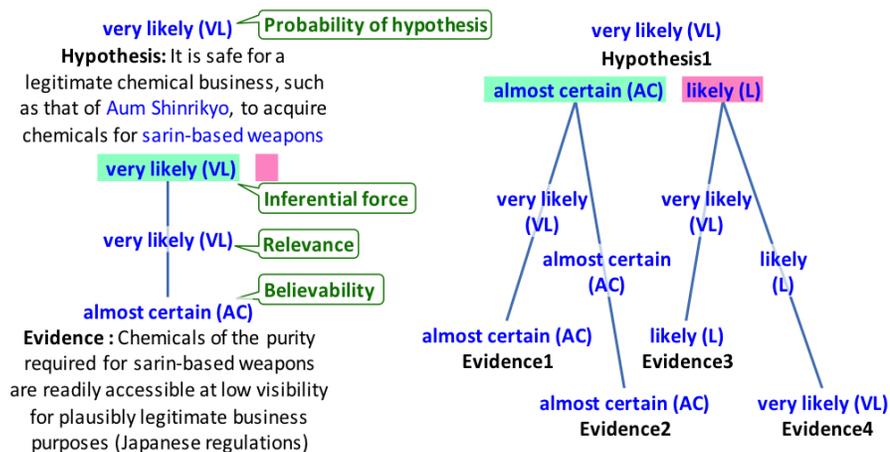


Fig. 4. Hypothesis assessment based on the credentials of evidence.

In general, there may be several items of evidence that are relevant to a given hypothesis, some favoring it and some disfavoring it, and each with a specific believability, relevance, and inferential force. The right-hand side of Fig. 4 illustrates a situation where there are two items of evidence favoring Hypothesis1 (Evidence1 and Evidence2), and two items of evidence disfavoring it (Evidence3 and Evidence4). The analyst needs to assess the believability and

relevance of each of the two favoring items of evidence. Then Cogent computes their individual inferential force, and takes their maximum as the combined inferential force of the favoring evidence (*i.e.*, **almost certain**). The combined inferential force of the disfavoring evidence is determined in a similar way as **likely**. Finally, Cogent determines the inferential force of all the evidence by employing an *on balance function* [11], concluding either some level of support for the hypothesis (if the force of the favoring evidence is greater than that of the disfavoring evidence), or **lack of support** otherwise.

### 3.2. Assessing Complex Hypotheses

The previous section presented a simple way of directly assessing a hypothesis based on evidence. This works well when it is easy to assess the relevance of the evidence to the hypothesis. But it does not work well for a complex hypothesis, such as “*Aum Shinrikyo* has *sarin-based weapons*” from the bottom-left of Fig. 3.

One question is, *How confident would you be in assessing the relevance of an item of evidence to a complex hypothesis?* For example, how confident would you be in assessing the relevance of the item of evidence from the bottom-left of Fig. 4 (*i.e.*, “Chemicals of the purity required for sarin-based weapons are readily accessible at low visibility for plausibly legitimate business purposes”) to the hypothesis “*Aum Shinrikyo* has *sarin-based weapons*”? If you would know that this item of evidence is true, would you be able to assess the probability of the hypothesis “*Aum Shinrikyo* has *sarin-based weapons*”? Probably not, because there are other factors to consider, such as the availability of funds and expertise, or even the possibility of buying sarin-based weapons on the black market. Depending of these ignored factors, the probability of the hypothesis may take different values. And there is an additional question, *How can you find evidence to assess such complex hypotheses?*

Solutions to both these difficulties are provided by the computational theory and by Cogent, which guide the analyst in developing an argumentation structure that successively reduces a complex hypothesis to simpler and simpler hypotheses, down to the level of very simple hypotheses, such as that from the top-left of Fig. 4. Then the analyst has to search for evidence that is relevant to these very simple hypotheses, and to assess their believability and relevance, as was illustrated in the previous section. Once this is done, Cogent automatically composes these assessments, from bottom-up, based on the logic embedded in the argumentation, finally assessing the probability of the top-level hypothesis.

The next section presents the structure of the argumentation in the computational theory and in Cogent.

### 3.3. Argumentation Structure

Figure 5 shows an example of an argumentation for the hypothesis “*Aum Shinrikyo* has *sarin-based weapons*.” This is a type of Wigmorean logic and

probabilistic inference network, where the top-level hypothesis is successively decomposed into simpler and simpler hypotheses [1, 2, 17, 18].

Notice that only the top-level hypothesis is completely specified, while the sub-hypotheses are abstracted, and are understood in the context of their upper level hypotheses. For example, “develops *sarin-based weapons*” is understood as “*Aum Shinrikyo* develops *sarin-based weapons*.” Similarly, “production material” is understood as “*Aum Shinrikyo* has production material to develop *sarin-based weapons*.” Also, “legitimate business” is understood as “*Aum Shinrikyo* has a legitimate business for acquiring production material to develop *sarin-based weapons*.” The use of context-dependent hypotheses enables a very succinct representation of an argumentation, helping the analyst to visualize a larger portion of it in the Cogent whiteboard.

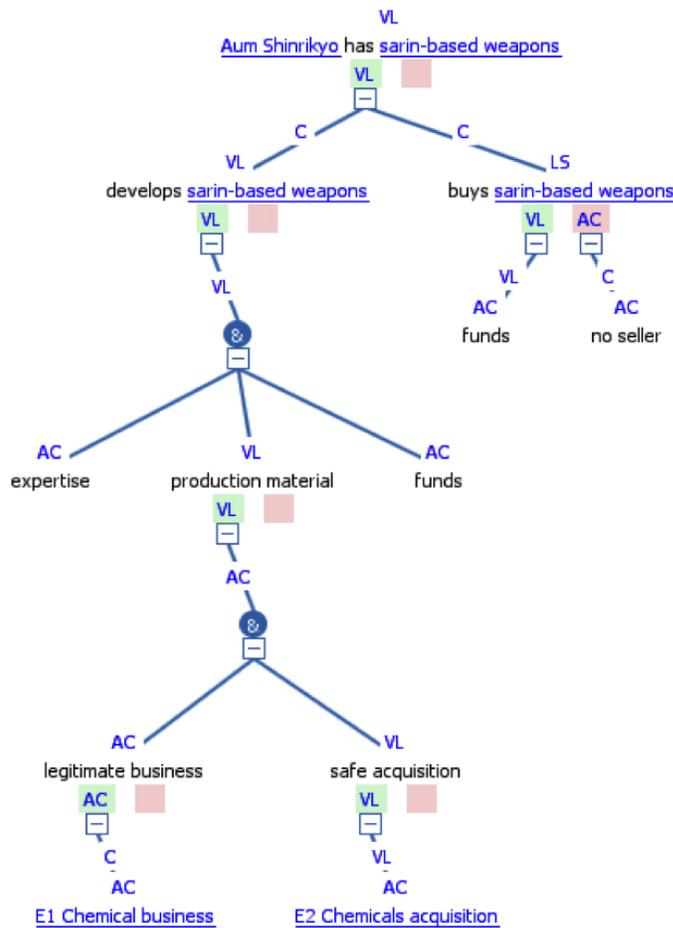


Fig. 5. Argumentation structure.

The top-level hypothesis is assessed considering both favoring arguments (under the left, green square) and disfavoring arguments (under the right, pink square). There may be one argument of each type, several, or none. Each argument reduces the top-hypothesis to simpler hypotheses for which favoring and disfavoring arguments are again considered.

For example, there are two favoring arguments and no disfavoring argument for the top hypothesis “*Aum Shinrikyo* has *sarin-based weapons*”, the first one being: “It is **certain (C)** that *Aum Shinrikyo* has *sarin-based weapons* if it develops *sarin-based weapons*”, where **certain (C)** is the relevance of this argument.

Further down in the argumentation: “It is **very likely (VL)** that *Aum Shinrikyo* develops *sarin-based weapons* if it has expertise, production material, and funds.” And further down: “It is **almost certain (AC)** that *Aum Shinrikyo* has production material for *sarin-based weapons* if it has a legitimate business for acquiring the production material, and it is safe for this business to acquire it.”

Now it should be quite easy to search for evidence that is relevant to these simpler hypotheses (*i.e.*, “legitimate business” and “safe acquisition”). For example, it would be enough to check whether Aum has created any legitimate business that uses the production material needed for sarin-based weapons. One may easily discover that it has, in fact, created two dummy chemical companies, Beck and Belle Epoc, – both run by Niimi – under Hasegawa Chemical, an already existing Aum shell company [24]. Similarly, one may easily find evidence for the hypothesis that it is safe for a legitimate chemical business, such as those of Aum Shinrikyo, to acquire chemicals for sarin-based weapons, by simply checking the Japanese regulations on the buying of chemicals.

Thus, the reduction of complex hypotheses (such as “*Aum Shinrikyo* has *sarin-based weapons*”) to simpler ones (*e.g.*, “legitimate business” and “safe acquisition”), guides the analyst to collect relevant evidence. Once the evidence is found, the probabilities of the simplest hypotheses can be determined, as discussed in the previous sections. Then the probabilities of the upper-level hypotheses are automatically assessed, based on their favoring and disfavoring arguments, since each subhypothesis in an argument will have a probability, and the entire argument will have a relevance.

When there is no evidence relevant to a given subhypothesis, the analyst may make an assumption with respect to its probability. When new items of evidence are added to the argumentation, and their believability and relevance are assessed, the probabilities of all the hypotheses are automatically updated. The same kind of automatic updating takes place when the assessment of any item of evidence is changed.

### 3.4. Deeper Believability Analysis

The previous sections have discussed the process of evidence-based hypothesis assessment where the analyst directly assesses the believability of each item of evidence. However, if the probability of the top-level hypothesis changes with

the believability of an item of evidence, then that item is *key evidence*, and its believability has to be more carefully assessed.

As discussed in Section 2, the knowledge base of Cogent include a general ontology of evidence, a fragment of which is shown in Fig. 6. For each evidence type in this ontology there is a general pattern for assessing its believability based on lower-level credentials, as established in the computational theory of intelligence analysis [2]. Thus, when a deeper, more detailed believability analysis is needed, Cogent can generate the corresponding Wigmorean network.

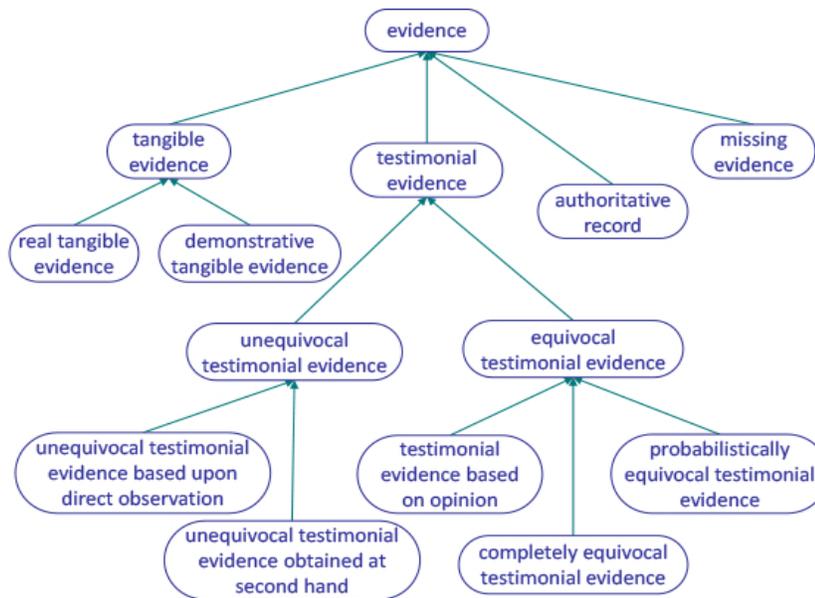
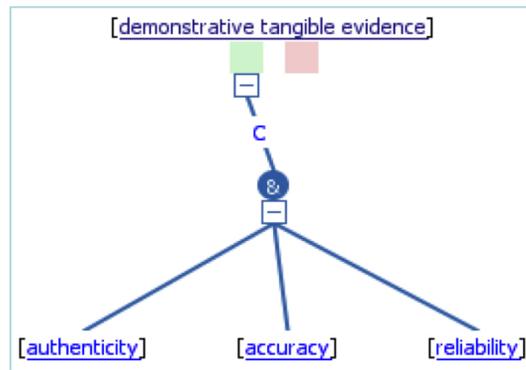


Fig. 6. Fragment of the ontology of evidence.

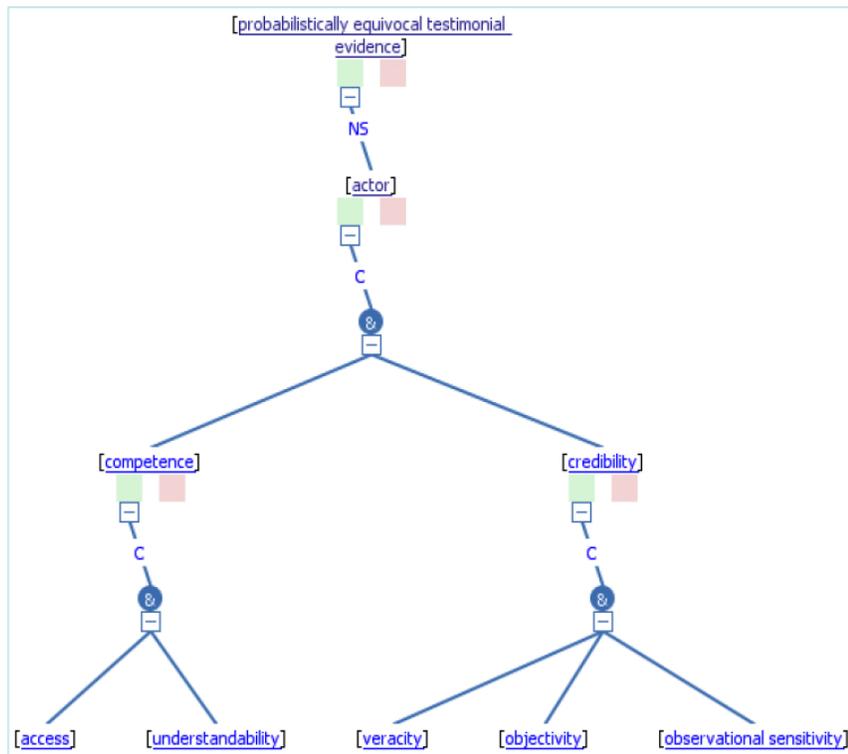
Figure 7, for example, shows the argument pattern for assessing the believability of *demonstrative tangible evidence*.

*Demonstrative tangible evidence* is a representation or illustration of a thing, such as a diagram, map, scale model, or sensor image. It has three believability attributes: (1) *authenticity* (Is this representation or illustration of a thing what it is claimed to be?), (2) *accuracy* (If the tangible item was produced by a sensing device or a person, how accurate was this representation or illustration?), and (3) *reliability* (If the tangible item was produced by a sensing device of some sort, how reliable was this device and the processes used to display the results?).

Figure 8 shows the argument pattern for assessing the believability of *probabilistically equivocal testimonial evidence* provided by a certain actor (or source). Asked whether event E occurred, the actor might say: “I’m almost certain that E occurred.” In this case, “almost certain” is the actor’s own assessment of his or her believability. But how to assess his or her believability?



**Fig. 7.** Pattern for assessing the believability of demonstrative tangible evidence.



**Fig. 8.** Pattern for assessing the believability of probabilistically equivocal testimonial evidence.

As shown in the pattern from Fig. 8, the actor's believability depends on the actor's *competence* and *credibility*. The first question to ask related to *competence* is whether this actor actually made the observation he or she claims to have made or had access to the reported information. The second competence question concerns whether this actor 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*. Here is an account of why these are the major attributes of testimonial credibility. First, is this actor telling us about an event he or she believes to have occurred? This actor would be untruthful if he or she did not believe the reported event actually occurred. So, this question involves the actor's *veracity*. The second question involves the actor's *objectivity*. The question is, did this actor base a belief on sensory evidence received during an observation, or did this actor believe the reported event occurred either because this actor 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 actor did base a belief on sensory evidence, how good was this evidence? This involves information about the actor's relevant *sensory capabilities and the conditions under which a relevant observation was made*.

## 4. A Sample Session with Cogent

Having presented the evidence-based reasoning concepts, this section illustrates a sample session with the current version of Cogent, to provide an intuitive understanding of its use by a typical analyst.

### 4.1. Starting the Analysis

The analyst starts by defining, in natural language, the situation of interest, the intelligence question(s), and the alternative hypotheses, as was illustrated in Fig. 3. While editing a statement, the analyst may select a phrase (*e.g.*, *sarin-based weapons*) and ask Cogent to introduce it into its knowledge base as a new entity. In general, entities in a hypothesis are those phrases that can be replaced with other phrases to obtain similar hypotheses. Consider, for example, the hypothesis "*Aum Shinrikyo* has *sarin-based weapons*." One may replace *Aum Shinrikyo* and *sarin-based weapons* with *Al Qaeda* and *botulinum-based weapons*, respectively, to obtain a similar hypothesis: "*Al Qaeda* has *botulinum-based weapons*." When the analyst starts typing a word, Cogent proposes its completion with known entities.

Recognizing entities enables Cogent to learn and reuse both hypotheses patterns from specific hypotheses, such as "[actor] has [weapon]," and argument patterns from specific arguments, as will be discussed in Section 4.3.

## 4.2. Developing the Argumentation

Once a top-level hypothesis is defined, the analyst interacts with Cogent to develop an argumentation like that from Fig. 5, through easy operations, for example by dragging and dropping building blocks from the Argument assistant under, above, or next to existing hypotheses, and by updating them. Then the analyst looks for evidence and attaches it to the corresponding hypotheses. This process is illustrated in Fig. 9.

The screenshot displays the Disciple software interface. The main window shows an argumentation diagram with a central question: "Question (Q): Is Aum Shinrikyo a threat?". Below this, a hypothesis is stated: "Aum Shinrikyo has sarin-based weapons". This hypothesis is supported by two causal links (C) leading to "develops sarin-based weapons" and "buys sarin-based weapons". The "develops" hypothesis is supported by three causal links (C) leading to "expertise", "production material", and "funds". The "production material" hypothesis is supported by two causal links (C) leading to "legitimate business" and "safe acquisition". The "funds" hypothesis is supported by two causal links (C) leading to "legitimate business" and "safe acquisition". The diagram uses various colored boxes and lines to represent different types of links and evidence.

On the right side, there is an "Argument" panel with an "Evidence" section. It contains a text box with the following description: "To purchase the required technical equipment and substantial amounts of chemicals, Aum created two dummy companies - both run by Niimi - under Hasegawa Chemical, an". Below this, there are fields for "URL" and "Description".

In the foreground, a PDF document titled "CNAS\_AumShinrikyo\_Danzig\_0.pdf" is open. The document contains text about the Aum Shinrikyo group and their activities. A purple arrow points from the "E1 evidence" box in the Disciple interface to a specific paragraph in the PDF document.

Fig. 9. Creating and attaching an item of evidence to a hypothesis.

The analyst selects a paragraph from a document which represents favoring evidence for the “legitimate business” hypothesis. Then it drags and drops it on the left (green) square under the hypothesis. Similarly, disfavoring evidence is dropped on the right (pink) square. As a result, Cogent automatically defines the item of evidence in the Evidence assistant (see the upper right side of Fig. 9), and attaches it to the hypothesis (see the bottom left side of Fig. 9). The automatically generated evidence name in the Evidence assistant is selected in case the analyst desires to replace it with a more suggestive one. Notice that the believability and relevance of the newly created item of evidence are NS (Not Set). Once the analyst assesses them, Cogent automatically assesses the probability of the elementary hypothesis and of the upper-level ones, based on the defined structure of the argumentation, as was previously discussed.

While argument development may seem a laborious process, it is greatly facilitated by the reuse of learned patterns, as will be discussed in the next section.

### 4.3. Pattern Learning and Reuse

Once the analysis is completed, the analyst may wish to request Cogent to learn hypotheses and argument patterns to be reused in future analyses. This is done by simply right-clicking on a hypothesis, such as “Aum Shinrikyo has sarin-based weapons” from the top of Fig. 5, and selecting “Learn.” As a result, Cogent learns both a hypothesis pattern (“[actor] has [weapon]”), and two argument patterns (one for each of the two favoring arguments), as illustrated in Fig. 10.

The patterns are obtained by using Cogent’s ontology as a generalization hierarchy, where individual entities from the analysis in Fig. 5 (*e.g.*, [sarin-based weapons](#)) are replaced with more general concepts from the ontology [23]. The analyst may change the pattern in the Learning assistant by clicking on a concept (*e.g.*, [\[weapon\]](#)) and selecting another concept (*e.g.*, [\[WMD\]](#)) from a list presented by Cogent. Cogent maintains the example from which each pattern has been learned, and will use it to refine the pattern when new examples are encountered (*e.g.*, when the pattern is reused). Patterns are a special type of rules that are not automatically applied by Cogent. To reuse a learned pattern, the analyst simply drags it from the Argument assistant, and drops it on a question or a hypothesis.

## 5. Cognitive Assistance

The computational theory of evidence-based reasoning embedded into Cogent guides the user through a systematic analysis process that synergistically integrates the user’s imaginative reasoning and expertise with the agent’s critical reasoning. For example, the user imagines the questions to ask and hypothesizes possible answers. Cogent helps with developing the arguments by reusing

previously learned patterns, and guides the evidence collection by the user. The user assesses the believability and relevance of the evidence, and Cogent determines its inferential force and the probabilities of the hypotheses. The jointly-developed analysis makes very clear the argumentation logic, what evidence was used and how, what is not known, and what assumptions have been made. It can be shared with other users, subjected to critical analysis, and correspondingly improved. As a result, this systematic and theoretically justified process leads to the development of defensible and persuasive conclusions.

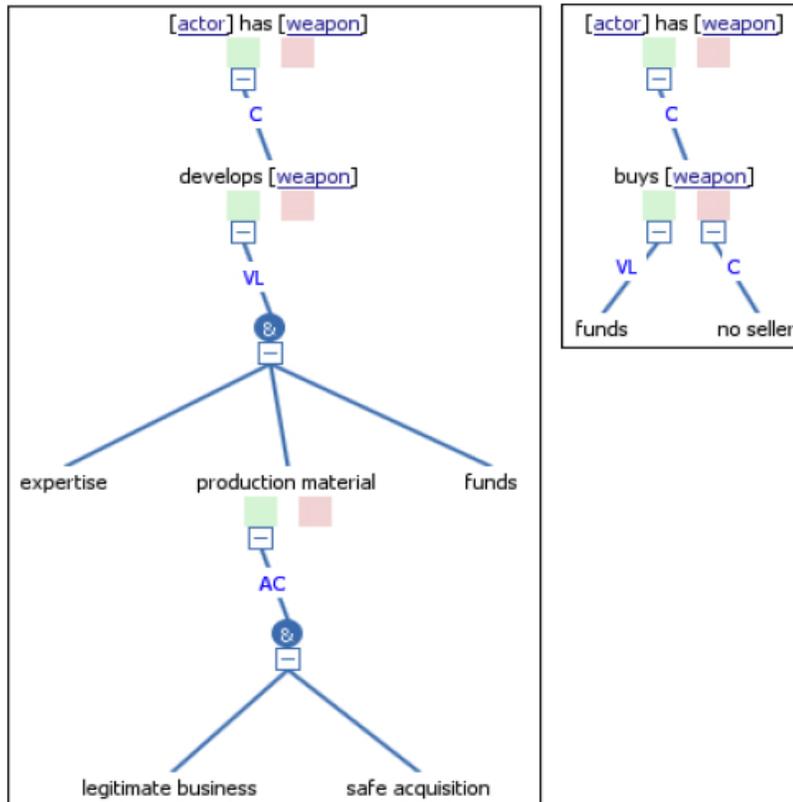


Fig. 10. Learned patterns.

Cogent also enable rapid analysis, not only through the reuse of patterns, but also through a drill-down process where a hypothesis may be decomposed to different levels of detail, depending on the available time. It facilitates the analysis of what-if scenarios, where the user may make various assumptions and the assistant automatically determines their influence on the analytic conclusion. The assistant also makes possible the rapid updating of the analysis based on new evidence and assumptions.

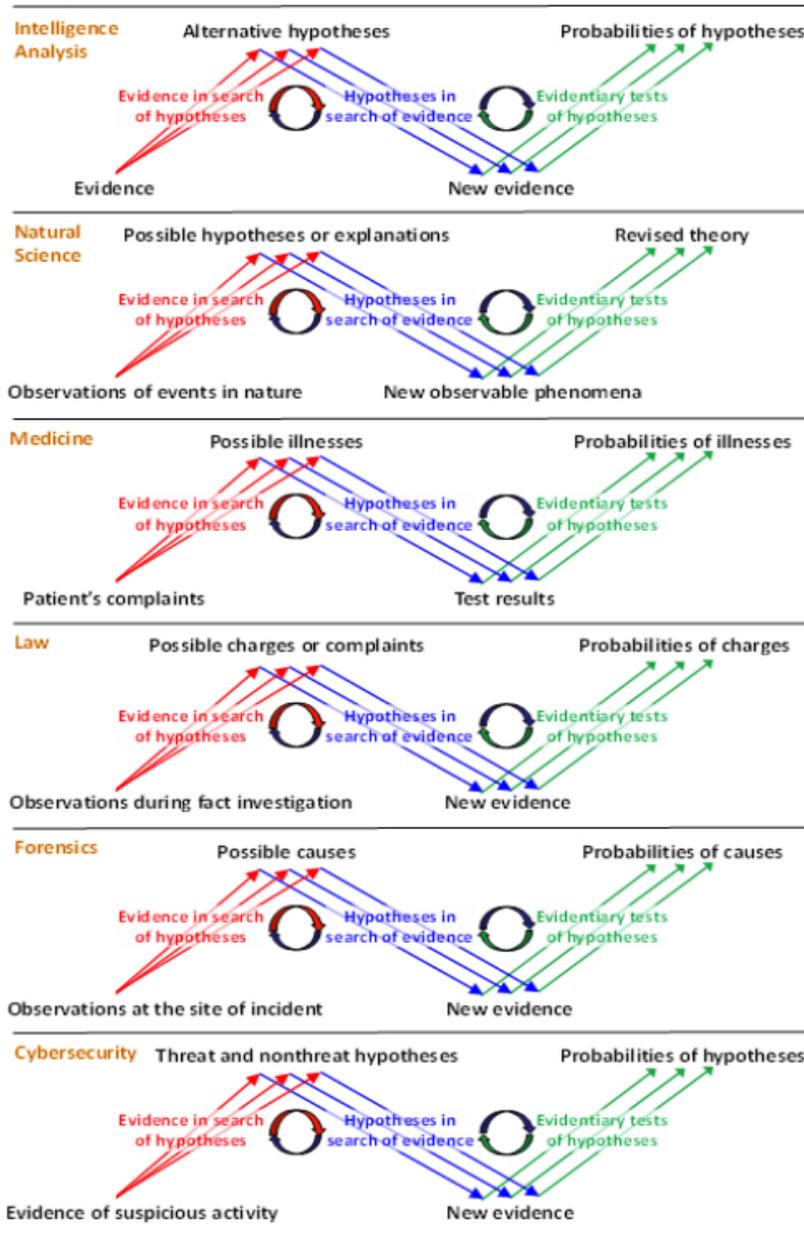


Fig. 11. Evidence-based reasoning everywhere.

## 6. Conclusion

This paper has overviewed research on developing a computational theory of evidence-based reasoning and on its integration within the Disciple learning agents theory, methodology, and tools for the development of cognitive assistants. Although the presentation was done in the context of a cognitive assistant for intelligence analysis, the obtained results are applicable in many other domains, including natural science, medicine, law, forensics, and cybersecurity, as shown in Fig. 11.

For example, in natural science education, a cognitive assistant may facilitate inquiry-based teaching and learning, engaging students in understanding, extending, creating, critiquing, and debating evidence-based scientific argumentations in real-life scientific investigations, giving the students opportunities to exercise imagination and creativity, and develop critical scientific practices, particularly: asking questions; constructing explanations; engaging in argument from evidence; and obtaining, evaluating, and communicating explanations [27, p.3]. Also, in cybersecurity, a Cogent-like agent [28] may be integrated into a Cybersecurity Operations Center (CSOC) to automate the investigation of intrusion alerts from a variety of intrusion detection devices, integrating multiple detection techniques with automated network forensics, to significantly increase the probability of accurately detecting intrusion activity while drastically reducing the workload of the CSOC operators.

**Acknowledgements.** This research was partially supported by George Mason University, the Intelligence Community, and the Department of Defense. 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 any agency of the U.S. Government.

## References

- [1] SCHUM, D. A. (1994/2001). *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.
- [2] TECUCI, G., SCHUM, D. A., MARCU, D., BOICU, M. (2016). *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*, Cambridge University Press.
- [3] TECUCI, G. (1988). Disciple: A Theory, Methodology and System for Learning Expert Knowledge, *Thèse de Docteur en Science*, University of Paris-South. <http://lac.gmu.edu/publications/1988/TecuciG.PhD.Thesis.pdf>
- [4] TECUCI, G. (1998). *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. London, England: Academic Press. <http://lac.gmu.edu/publications/1998/TecuciG.Building.Intelligent.Agents/default.htm>

- [5] BOICU, M., TECUCI, G., MARCU, D., BOWMAN, M., SHYR, P., CIUCU, F., LEVCOVICI, C. (2000) Disciple-COA: From Agent Programming to Agent Teaching, in *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, Stanford, California 2000, Morgan Kaufman. [http://lac.gmu.edu/publications/data/2000/2000\\_il-final.pdf](http://lac.gmu.edu/publications/data/2000/2000_il-final.pdf)
- [6] TECUCI, G., BOICU, M., MARCU, D., STANESCU, B., BOICU, C., COMELLO, J., LOPEZ, A., DONLON, J., CLECKNER, W. (2002) Development and Deployment of a Disciple Agent for Center of Gravity Analysis, in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 853-860, Edmonton, Alberta, Canada, AAAI Press/The MIT Press. <http://lac.gmu.edu/publications/data/2002/dddacga.pdf>
- [7] TECUCI, G., BOICU, M., BOICU, C., MARCU, D., STANESCU, B., BARBULESCU, M. (2005). The Disciple-RKF learning and reasoning agent. *Computational Intelligence* 21(4) pp. 462-479. <http://lac.gmu.edu/publications/2005/TecuciG.Disciple.RKF.CI.pdf>
- [8] TECUCI, G., BOICU, M., AYERS, C., CAMMONS, D. (2005). Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving, *Proceedings of the First International Conference on Intelligence Analysis*, McLean, VA, 2-6 May, 2005. <http://lac.gmu.edu/publications/data/2005/Tecuci-Disciple-LTA.pdf>
- [9] TECUCI, G., BOICU, M., MARCU, D., BOICU, C., BARBULESCU, M., AYERS, C., CAMMONS, D. (2007). Cognitive Assistants for Analysts, *Journal of Intelligence Community Research and Development (JICRD)*, 2007. Also published in John Auger, William Wimbish (eds.), *Proteus Futures Digest: A Compilation of Selected Works Derived from the 2006 Proteus Workshop*, Joint publication of the National Intelligence University, Office of the Director of National Intelligence, and U.S. Army War College Center for Strategic Leadership, 2007, pp. 303-329. [http://lac.gmu.edu/publications/2007/TecuciG.Cognitive\\_Assistants.pdf](http://lac.gmu.edu/publications/2007/TecuciG.Cognitive_Assistants.pdf)
- [10] 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.
- [11] TECUCI G., MARCU D., BOICU M., SCHUM D. A. (2015). COGENT: Cognitive Agent for Cogent Analysis. In *Proceedings of the 2015 AAAI Fall Symposium "Cognitive Assistance in Government and Public Sector Applications,"* pp.58-65, Arlington, VA, 12-14 Nov. 2015, Technical Report FS-15-02, AAAI Press, Palo Alto, CA. <http://lac.gmu.edu/publications/2015/Cogent-overview.pdf>
- [12] PEIRCE, C. S. (1898 [1992]). *Reasoning and the Logic of Things*, edited by Ketner, K., Harvard University Press, Cambridge, MA.
- [13] PEIRCE, C. S. (1901 [1955]). Abduction and Induction, in *Philosophical Writings of Peirce*, edited by Buchler, J., Dover, New York, NY, pp. 150-156.
- [14] ECO, U., SEBEOK, T. (eds.) (1983). *The Sign of Three: Dupin, Holmes, Peirce*, Indiana University Press: Bloomington.
- [15] THAGARD P. R. (1993). *Computational Philosophy of Science*, MIT Press: Cambridge, MA.

- [16] SCHUM, D. A. (2001). Species of Abductive Reasoning in Fact Investigation in Law, *Cardozo Law Review*, vol. 22, nos. 5–6, pp. 1645–1681.
- [17] WIGMORE, J. H. (1913). The Problem of Proof, *Illinois Law Review*, vol. 8, no. 2, pp. 77–103.
- [18] WIGMORE, J. H. (1937). *The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience and Illustrated in Judicial Trials*, 3rd ed., Little, Brown, Boston, MA.
- [19] COHEN, L. J. (1977). *The Probable and the Provable*, Clarendon Press, Oxford, UK.
- [20] COHEN, L. J. (1989). *An Introduction to the Philosophy of Induction and Probability*, Clarendon Press, Oxford, UK.
- [21] ZADEH L.A. (1965). Fuzzy Sets, *Information and Control*, Vol. 8, Issue 3, pp.338–353, Elsevier.
- [22] NEGOITA C. V., RALESCU D. A. (1975). *Applications of Fuzzy Sets to Systems Analysis*, Wiley, New York.
- [23] TECUCI G., MARCU D., BOICU M., SCHUM D. A. (2016). *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, Cambridge University Press.
- [24] DANZIG R., SAGEMAN M., LEIGHTON T., HOUGH L., YUKI H., KOTANI R., HOSFORD Z. M. (2011) *Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons*, Center for a New American Security, Washington, DC.
- [25] BOICU M., TECUCI G., MARCU D. (2012). Rapid Argumentation Capture from Analysis Reports: The Case Study of Aum Shinrikyo, in *Proceedings of the Seventh International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS 2012*, Fairfax, VA, 23-26 October. <http://ceur-ws.org/Vol-966/STIDS2012.T05.BoicuEtAl.RapidArgumentation.pdf>
- [26] SHAFER, G. (1976). *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ.
- [27] NATIONAL RESEARCH COUNCIL. (2012). *A Framework for K-12 Science Education: Practices, Crosscutting Concepts, and Core Ideas*. Washington, DC: The National Academies Press.
- [28] MECKL S., TECUCI G., BOICU M., MARCU D. (2015). Towards an Operational Semantic Theory of Cyber Defense Against Advanced Persistent Threats. *Proceedings of the Tenth Conference on Semantic Technologies for Intelligence, Defense, and Security (STIDS 2015)*, Fairfax, VA, 18-20 November, CEUR Workshop Proceedings, Vol. 1523, pp. 58–65. <http://lac.gmu.edu/publications/2015/APT-LAC.pdf>