Presented at the 58th Annual Convention of the International Studies Association (ISA): Understanding Change in World Politics, February 22nd - 25th, 2017, Baltimore, Maryland

Structured Analysis with Cogent: Integrating Analyst's Imaginative Reasoning with Computer's Critical Reasoning

Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, David Schum
Learning Agents Center, George Mason University, Fairfax, VA 22030
{tecuci, dmarcu, mboicu}@gmu.edu, dschum398@earthlink.net

ABSTRACT

Intelligence analysts face the astonishingly complex task of drawing defensible and persuasive conclusions from masses of evidence of all types, in a world that is changing all the time. This requires a complex combination of imaginative reasoning, critical reasoning, and expert knowledge, as well as a significant amount of time and effort. This paper presents a systematic approach to intelligence analysis grounded in the scientific method, and its implementation into Cogent, an intelligent analytic tool that incorporates a significant amount of knowledge about evidence and its credentials, and about evidence-based reasoning. Starting from a situation of interest, the analyst imagines an intelligence question and hypothesizes possible answers. Then Cogent guides the collection of evidence and the development of arguments linking it to the hypotheses. The analyst assesses the credibility and relevance of evidence, and Cogent determines its inferential force and the probabilities of the alternative hypotheses. This 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 analysts, subjected to critical examination, and correspondingly improved.

ASTONISHING COMPLEXITY OF INTELLIGENCE ANALYSIS

The task of the intelligence analyst has often been described by using the *"Connecting the Dots"* metaphor. This metaphor may have gained its current popularity following the terrorist attacks in New York City and Washington, DC, on September 11, 2001. It was frequently said that the intelligence services did not connect the dots appropriately in order to have possibly prevented the catastrophes that occurred. Since then, we have seen and heard this metaphor applied in the news media to inferences in a very wide array of contexts in addition to intelligence analysis. For example, we have seen it applied to allegedly faulty medical diagnoses; allegedly faulty conclusions in historical studies; allegedly faulty or unpopular governmental decisions; and in discussions involving the conclusions reached by competing politicians. Listening to, or seeing the media accounts of this process may lead one to believe that it resembles the simple tasks we performed as children when, if we connected some collection of *numbered* dots correctly, a figure of Santa Claus, or some other familiar figure, would emerge.

What is also true is that the commentators on television and radio, or the sources of written accounts of inferential failures, never tell us what they mean by the phrase "connecting the dots." A natural explanation is that they have never even considered what this phrase means and what it might involve. Our belief is that people employing this metaphor to criticize intelligence analysts have very little awareness of how astonishingly difficult the process of connecting the dots can be.

So, what is meant by the "connecting the dots" metaphor when it is applied to the evidence-based reasoning tasks performed by the intelligence analysts?

"Connecting the Dots" refers to the task of marshaling thoughts and evidence in the generation or discovery of productive hypotheses and new evidence, and in the construction of defensible and persuasive arguments on hypotheses we believe to be most favored by the evidence we have gathered and evaluated.

We have found this metaphor very useful, and quite intuitive, in illustrating the extraordinary complexity of the evidential and inferential reasoning required to draw defensible and persuasive conclusions from masses of evidence of all kinds from a variety of different sources. An account of seven complexities of "connecting the dots" is provided in (Tecuci et al., 2016a, pp.1-12).

The conclusions drawn in intelligence analysis must rest on arguments that are defensible and persuasive. Since arguments rest on evidence, these conclusions are necessarily probabilistic in nature because our evidence is always *incomplete* (we can look for more, if we have time), usually *inconclusive* (it is consistent with the truth of more than one hypothesis or possible explanation), frequently *ambiguous* (we cannot always determine exactly what the evidence is telling us), commonly *dissonant* (some of it favors one hypothesis or possible explanation but other evidence favors other hypotheses), and with various degrees of *believability* shy of perfection (Schum, 1994/2001a; Tecuci et al., 2016a, pp. 159-172). Thus, one thing necessary in the study of argument construction is substantial information concerning the evidential foundations of arguments. Careful study of these evidential foundations requires consideration of the properties, uses, discovery, and marshaling of evidence. Moreover, each analytic task is unique and always requires mixtures of *imaginative reasoning* and *critical reasoning*. Indeed, the hypotheses must be generated by imaginative thought and then subjected to critical evidence-based analysis. Not only is this process stunningly complex, but it often needs to be performed in a very short amount of time.

This paper presents a systematic approach to dealing with the astonishing complexity of connecting the dots. This approach, which is grounded in the scientific method, is implemented in Cogent, an intelligent analytic tool that incorporates a significant amount of knowledge about evidence and its credentials, and about evidence-based reasoning.

BACONIAN PROBABILITIES WITH FUZZY QUALIFIERS

As indicated in the previous section, hypothesis assessment is necessarily probabilistic in nature. Unfortunately, none of the probability views known to us can cope with all the mentioned characteristics of evidence. For example, only the Baconian view 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 Fuzzy view can cope with the ambiguity of evidence, while a Subjective Bayesian view cannot (Tecuci et al., 2016a, pp. 205-208).

A probabilistic system based on ideas from both the Baconian view and the Fuzzy view may potentially cope with all the five characteristics of evidence: *incompleteness*, *inconclusiveness*, *ambiguity*, *dissonance*, and various degrees of *believability*. Moreover, the use of similar min/max probability combination rules by the Baconian and the Fuzzy views facilitates the development of such an integrated system (Cohen, 1977, pp.167-187; Zadeh, 1965, pp.340-341; Schum, 1979, pp.460-463; Schum, 1994/2001a, p.255, p.266; Tecuci et al., 2016a, pp.201-204). These rules are much more simple than the Bayesian probability combination rules, which is important for the understandability of the analyses performed by the analysts.

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 (Tecuci et al., 2016a, p.204-205). While a numerical probability is much more precise, it is not at all clear how an analyst 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 analysts, 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* (Cohen, 1977; 1989) with *Fuzzy qualifiers* (Zadeh, 1965). This symbolic probability system may be used with different assessment scales, such as the following one:

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

In this scale, the considered hypothesis may be "lacking support (LS)" from the available evidence, it may have some level of support, such as "very likely (VL)."

INTELLIGENCE ANALYSIS IN THE FRAMEWORK OF THE SCIENTIFIC METHOD

The analysis framework implemented in Cogent is illustrated in Figure 1: From questions we ask about a situation of interest we generate alternative hypotheses as possible answers, we use these hypotheses to guide the collection of evidence, and we use the found evidence to assess the hypotheses and determine the most probable answer. This framework will be illustrated in the following sections.

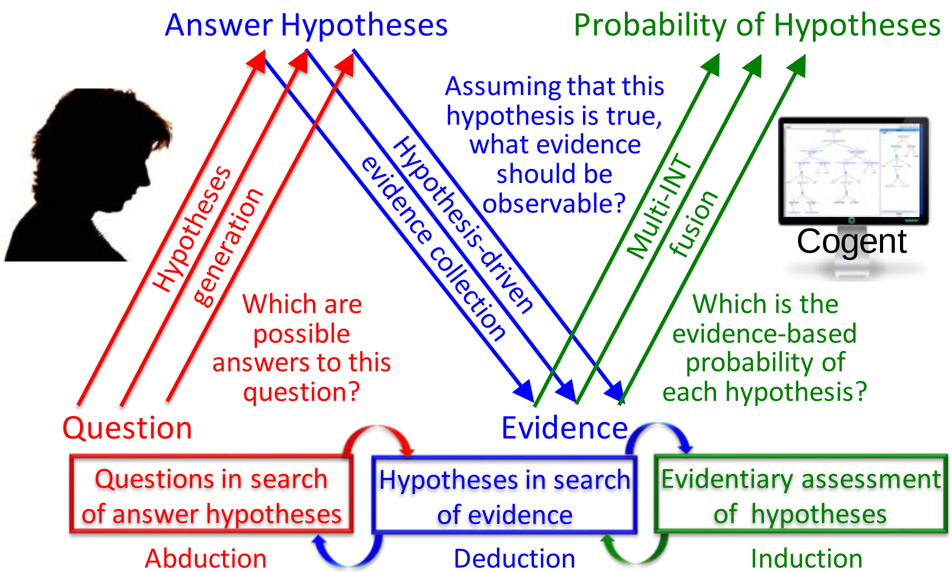

Figure 1: Discovery-based analysis framework.

1) Questions in Search of Answers

As shown in the top part of Figure 2, our analysis with Cogent starts with specifying a situation of interest: Aum Skinrikyo, a Japanese apocalyptic sect suspected of terrorist activities (Danzig at al., 2011).

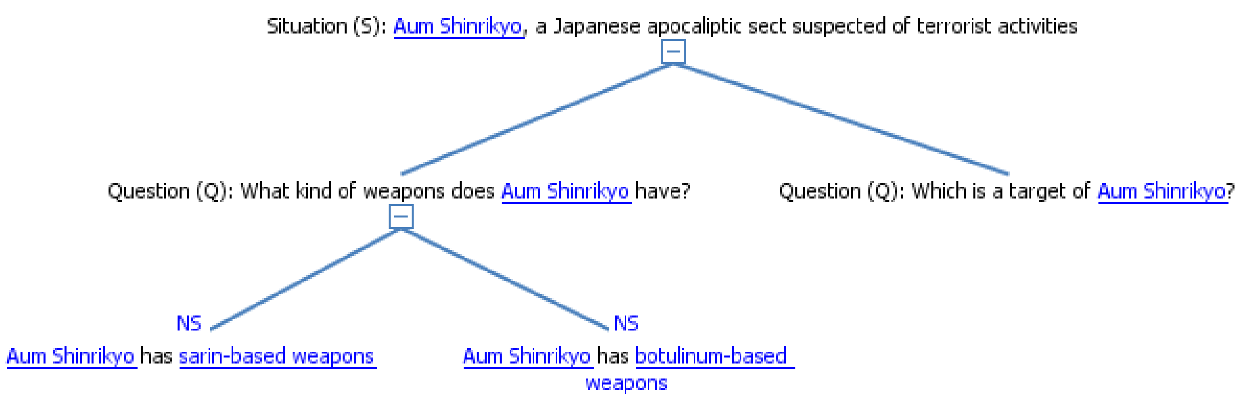

Figure 2: Hypotheses formulation.

We may now formulate any number of intelligence questions we would like to answer in our analysis, such as "What kind of weapons does Aum Shinrikyo have?"

For each intelligence question we may formulate possible answers as hypotheses to assess, such as "Aum Shinrikyo has sarin-based weapons" and "Aum Shinrikyo has botulinum-based weapons."

The considered hypotheses are not required to be disjoint. For example, both hypotheses from the bottom of Figure 2 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 later, each of the considered hypotheses is analyzed independent of one another.

For the same reason, *the set of the considered hypotheses is not required to be complete* either. At each moment we may formulate additional hypotheses, for example "Aum Shinrikyo has nuclear weapons."

These features significantly facilitate hypothesis analysis in the context of a dynamic world that is changing all the time, where new evidence may suggest new questions and the formulation of additional hypotheses or the modification of the existing ones, as indicated by the red and blue arrows from the bottom of Figure 1.

2) Hypotheses in Search of Evidence

Now that multiple hypotheses have been formulated, we need to assess the probability of each of them based on evidence. *But how can you find evidence to assess such complex hypotheses?* Cogent guides 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 that show very clearly what evidence to look for.

Figure 3 provides 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 (Wigmore 1913, 1937, 1940; Schum, 1987, 1994/2001a; Tecuci et al., 2016a, pp.82-90):

- It is almost certain (AC) that Aum Shinrikyo has sarin-based weapons if it develops sarin-based weapons
- It is very likely (VL) that Aum Shinrikyo develops sarin-based weapons if it has expertise, production material and funds
- It is certain (C) that Aum Shinrikyo has production material if it has a legitimate business and it is safe for such a business to acquire it.

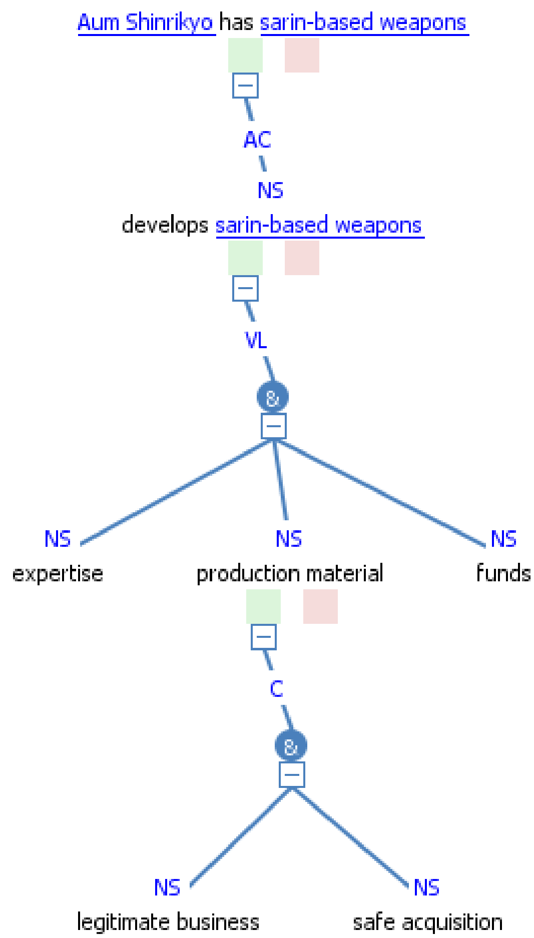

Figure 3: Hypotheses-driven evidence collection.

This guides us to look for evidence on whether Aum Shinrikyo has created any legitimate business that uses the production material needed for sarin-based weapons. This leads to the identification of the following item of evidence: "Aum created two dummy companies – both run by Niimi – under Hasegawa Chemical, an already existing Aum shell company, which justifies the purchase of the required technical equipment and substantial amounts of chemicals" (Danzig et al., 2011).

The other bottom leaf hypothesis in Figure 3 guides us to look for evidence on whether it is safe for a legitimate business to acquire sarin, by simply checking the Japanese regulations on buying it. This leads to the identification of the following item of evidence: "Chemicals of the purity required for sarin-based weapons are readily accessible at low visibility for plausibly legitimate business purposes" (Japanese regulations).

3) Evidentiary Assessment of Hypotheses

Once we have obtained evidence, we have to assess the hypotheses considered. We will show first how to directly assess a simple hypothesis based on evidence. Then we will show how to assess a complex hypothesis that was decomposed into simpler hypotheses.

Hypothesis assessment based on the credentials of evidence

One can directly assess a hypothesis based on an item of evidence by assessing the credentials evidence, as illustrated in Figure 4 and explained in the following.

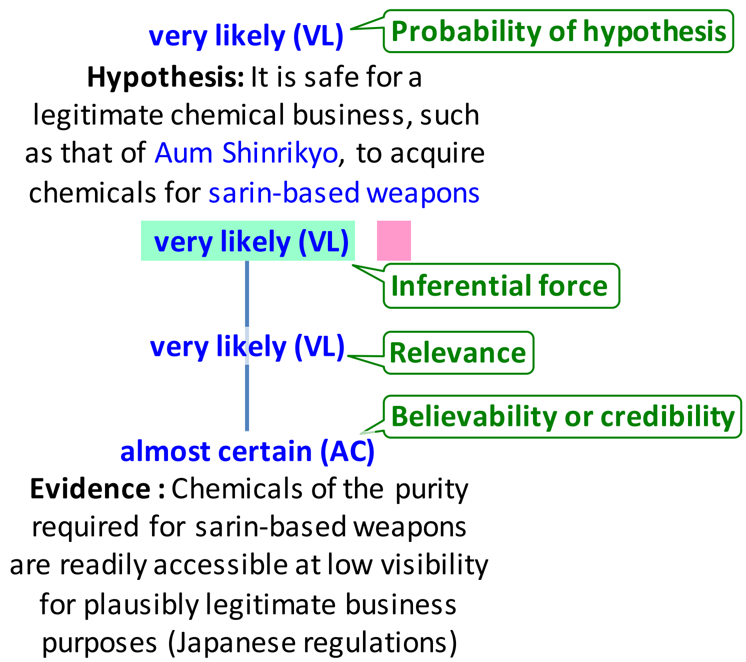

Figure 4: Evidence-based hypothesis assessment.

First we assess the *believability or credibility* 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's assess this probability to be almost certain (AC), as shown at the bottom-left of Figure 4.

Next we assess 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's assess this probability to be very likely (VL).

Finally, Cogent determines the *inferential force or weight* of the item of evidence on the hypothesis. The inferential force answers the question, *"What is the probability of the hypothesis, based only on this item of evidence?"* Obviously, an irrelevant item of evidence will have no inferential force, and will not convince us that the hypothesis is true. An item of evidence that is not believable will have no inferential force either. Only an item of evidence that is both very relevant and very believable may convince us that the hypothesis is true. Consistent with both the Baconian min/max probability combination rules (Cohen, 1977, pp.167-187; Schum, 1979, pp.460-463; Schum, 1994/2001a, p.255; Tecuci et al., 2016a, pp.201-202) and the with the Fuzzy min/max probability combination rules (Zadeh, 1965, pp.340-341; Schum, 1994/2001a, p.266; Tecuci et al., 2016b, p.204), Cogent determines the inferential force of an item of evidence on a hypothesis as the minimum between its believability and its relevance:

Inferential force(Evidence) = min(Believability(Evidence), Relevance(Evidence))

Therefore, the inferential force of "Evidence" from the bottom of Figure 4 is:

Inferential force(Evidence) = min(almost certain, very likely) = very likely.

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. Figure 5 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 these four items of evidence, for example,

Believability(Evidence1) = almost certain, and Relevance(Evidence1) = very likely.

Then Cogent automatically determines their individual and collective inferential forces on Hypothesis1.

The individual inferential force of each item of favoring evidence on Hypothesis1 is determined as discussed above:

Inferential force(Evidence1) = min(Believability(Evidence1), Relevance(Evidence1)) =
= min(almost certain, very likely) = very likely

Inferential force(Evidence2) = min(Believability(Evidence2), Relevance(Evidence2)) =
= min(almost certain, almost certain) = almost certain

Then the inferential force of all the favoring evidence (i.e., Evidence1 and Evidence2) on Hypothesis1 is computed as the maximum of the inferential forces of the individual items of favoring evidence, and displayed in the left (green) box under Hypothesis1 (see the upper-left side of Figure 5):

Inferential force(Evidence1, Evidence2) = max(Inferential force(Evidence1), Inferential force(Evidence2)) = max(very likely, almost certain) = almost certain

The inferential force of the disfavoring evidence is computed in a similar way and displayed in the right (pink) box under Hypothesis1 (see the upper-right side of Figure 5):

Inferential force(Evidence3) = min(Believability(Evidence3), Relevance(Evidence3)) =
= min(likely, very likely) = likely

Inferential force(Evidence4) = min(Believability(Evidence4), Relevance(Evidence4)) =
= min(very likely, likely) = likely

Inferential force(Evidence3, Evidence4) = max(Inferential force(Evidence3), Inferential force(Evidence4)) = max(likely, likely) = likely

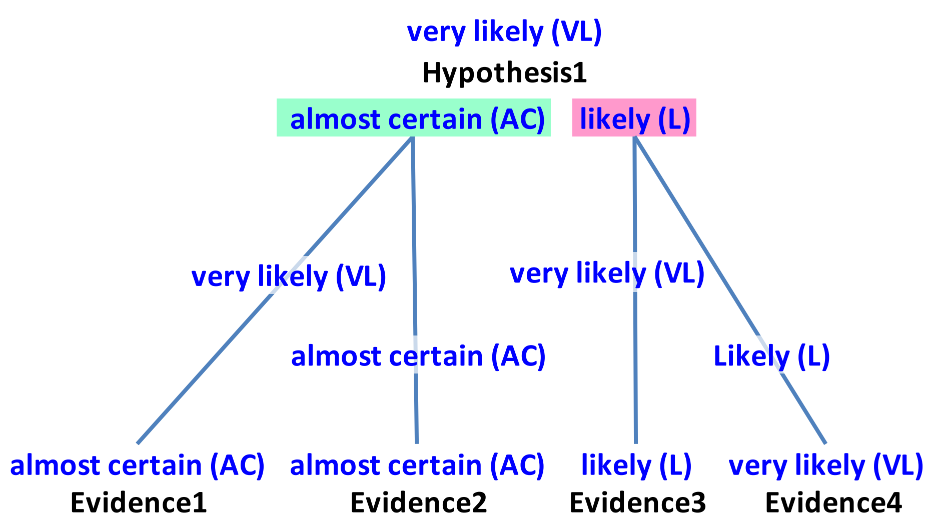

Figure 5: Hypothesis assessment based on multiple items of evidence.

The Baconian probability system (Cohen, 1977; 1989) requires considering either **H** or **not H** as probably true, but not both at the same time. To assess a hypothesis that has both favoring and disfavoring evidence, such as Hypothesis1 at the top of Figure 5, we have introduced an "on-balance" function that balances the inferential force of the favoring evidence with that of the disfavoring evidence. Figure 6 shows the actual "on-balance" function that we have defined for the current version of Cogent. As indicated in the right and upper side of Figure 5, if the inferential force of the disfavoring evidence is higher than or equal to that of the favoring evidence, then Cogent concludes that, based on all the available evidence, **H** is lacking support. If, however, the inferential force of the favoring evidence is strictly greater than that of the disfavoring evidence (and there is some force of the disfavoring evidence), then the probability of **H** is lowered, based on the inferential force of the disfavoring evidence (see the left and lower side of Figure 6).

Inferential force of the disfavoring evidence and arguments

Inferential force of the favoring evidence and arguments

H	lacking support	likely	very likely	almost certain	certain
lacking support	lacking support	lacking support	lacking support	lacking support	lacking support
likely	likely	lacking support	lacking support	lacking support	lacking support
very likely	very likely	likely	lacking support	lacking support	lacking support
almost certain	almost certain	very likely	likely	lacking support	lacking support
certain	certain	almost certain	very likely	likely	lacking support

Figure 6: On balance function in Cogent.

In the example from Figure 5, the inferential force of the favoring evidence (i.e., Evidence1 and Evidence2) is almost certain, as displayed inside the left (green) box under Hypothesis1. The inferential force of the disfavoring evidence (Evidence3 and Evidence4) is likely, as displayed inside the right (pink) box under Hypothesis1. Therefore, using the "on-balance" function from Figure 6, Cogent automatically determines the inferential force of all these items of evidence on Hypothesis1 as very likely, by balancing the inferential force of the favoring items (almost certain) with that of the disfavoring items (likely), as shown by the cells with white background from Figure 6.

In the current version of Cogent, the on-balance function from Figure 6 was defined as a uniform distribution of the values from the probability scale, and is used for all the hypotheses in an analysis. However, comparing the force of the favoring evidence (or arguments) with that of the disfavoring one is a subjective and context-dependent judgement, and an analyst may wish to select a different on-balance value than that proposed by the system. Therefore we envision that, in future versions of Cogent, each type of hypothesis will have a specific on-balance function that will be learned from the intelligence analysts.

Assessing Complex Hypotheses

Once the evidence is found, the probabilities of the simplest hypotheses can be assessed, as discussed in the previous sections. Then the probabilities of the upper-level hypotheses are automatically computed by Cogent using the logic embedded in the argumentation, in accordance with Baconian and Fuzzy min/max probability combination rules: AND (&) structures in the argumentation require the use of the minimum (min) function, while OR (alternative) structures require the use of the maximum (max) function. Figure 7 illustrates this process.

Based on the items of evidence E1 Chemical business and E2 Chemicals acquisition, the probabilities of the "legitimate business" and "safe acquisition" are determined as almost certain (AC) and very likely (VL), respectively (see the bottom of Figure 7). Then the probability of "production material" is automatically determined as min(min(AC, VL)), C) = VL because "production material" requires both "legitimate business" and "safe acquisition," and their combined relevance is certain (C).

However, we may not always find evidence for each hypothesis in the argumentation, or we may not have time to look for it. In such cases we may make *assumptions* with respect to the probabilities of these hypotheses. For example, we may assume that it is almost certain (AC) that Aum Shinrikyo has expertise to develop sarin-based weapons because all that is required is a degree in chemistry.

If we also make the assumption that it is very likely that Aum Shinrikyo has funds, then Cogent assesses that it is very likely that it develops sarin-based weapons.

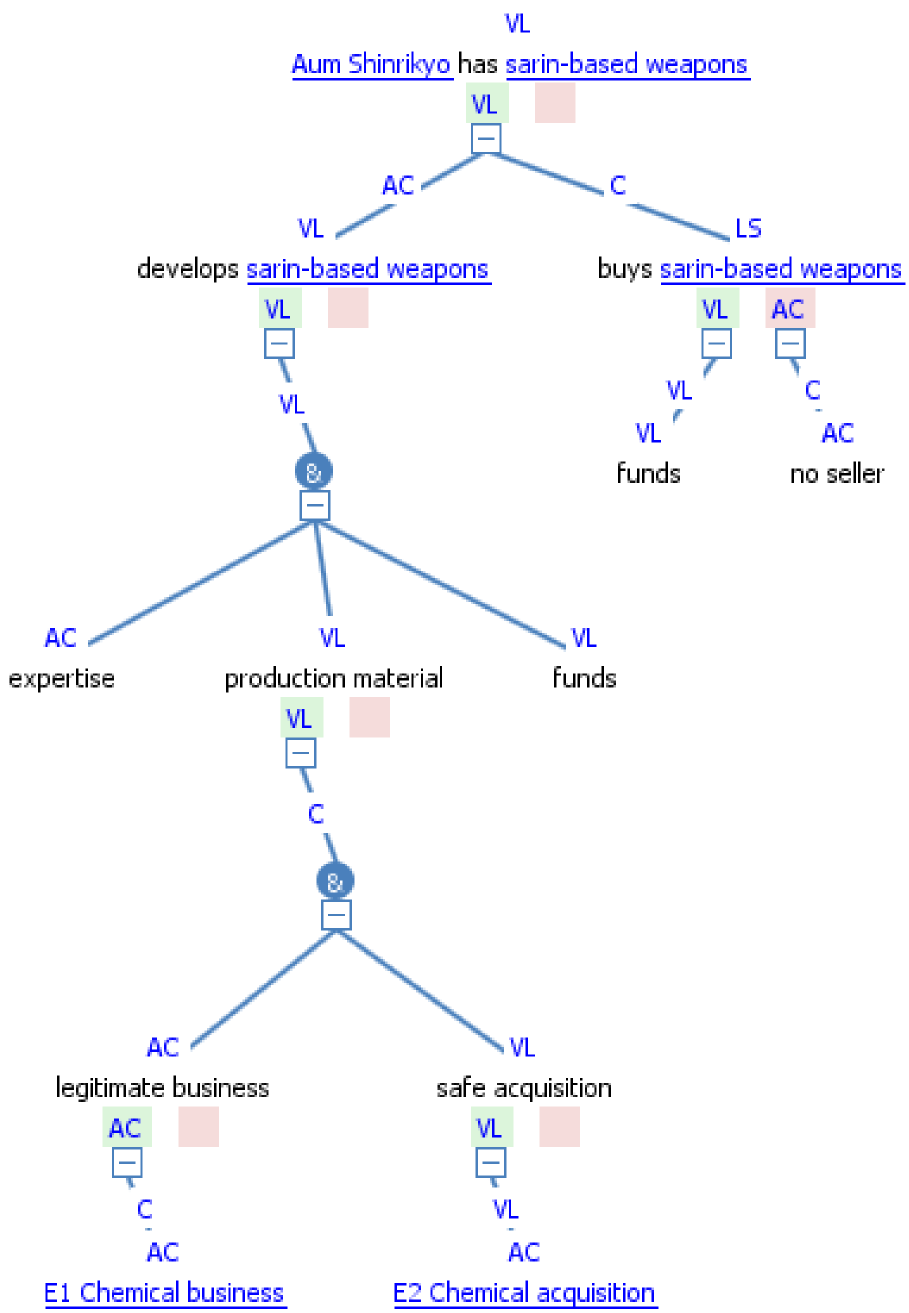

Figure 7: Wigmorean network with probabilities assessed.

Notice that the top-level hypothesis in Figure 7 has two favoring arguments:

- It is almost certain (AC) that Aum Shinrikyo has sarin-based weapons if it develops sarin-based weapons
- It is certain (C) that Aum Shinrikyo has sarin-based weapons if it buys sarin-based weapons

Because it is very likely that Aum Shinrikyo develops sarin-based weapons, while Aum Shinrikyo buys sarin-based weapons is lacking support, Cogent concludes that:

- It is very likely (VL) that Aum Shinrikyo has sarin-based weapons

Indeed, the inferential force of the "develops sarin-based weapons" argument is min(very likely, almost certain) = very likely, the inferential force of the "buys sarin-based weapons" argument is min(lacking support, certain) = lacking support, and their maximum is very likely.

Why is the hypothesis "Aum Shinrikyo buys sarin-based weapons" lacking support? Because the inferential force of its disfavoring argument ("no seller") is min(almost certain, certain) = certain, and this is higher than the inferential force of its favoring argument ("funds"), which is min(very likely, very likely) = very likely.

SEARCHING, REPRESENTATION, AND USE OF EVIDENCE

Once a top-level hypothesis is defined, the analyst interacts with Cogent to develop an argumentation like that from Figure 7, through easy operations, for example by dragging and dropping building blocks from the Argument assistant under, above, or next to existing hypotheses, and by editing them. Then the analyst looks for evidence and attaches it to the corresponding hypotheses. This process is illustrated in Figure 8.

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) box under the hypothesis. Similarly, disfavoring evidence is dropped on the right (pink) box. As a result, Cogent automatically defines the item of evidence in the Evidence assistant (see the upper right side of Figure 8), and attaches it to the hypothesis (see the bottom left side of Figure 8). 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 by double-clicking on the corresponding labels and selecting the desired value from the displayed list, 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 can be greatly facilitated by the reuse of learned patterns, as will be discussed in the next section.

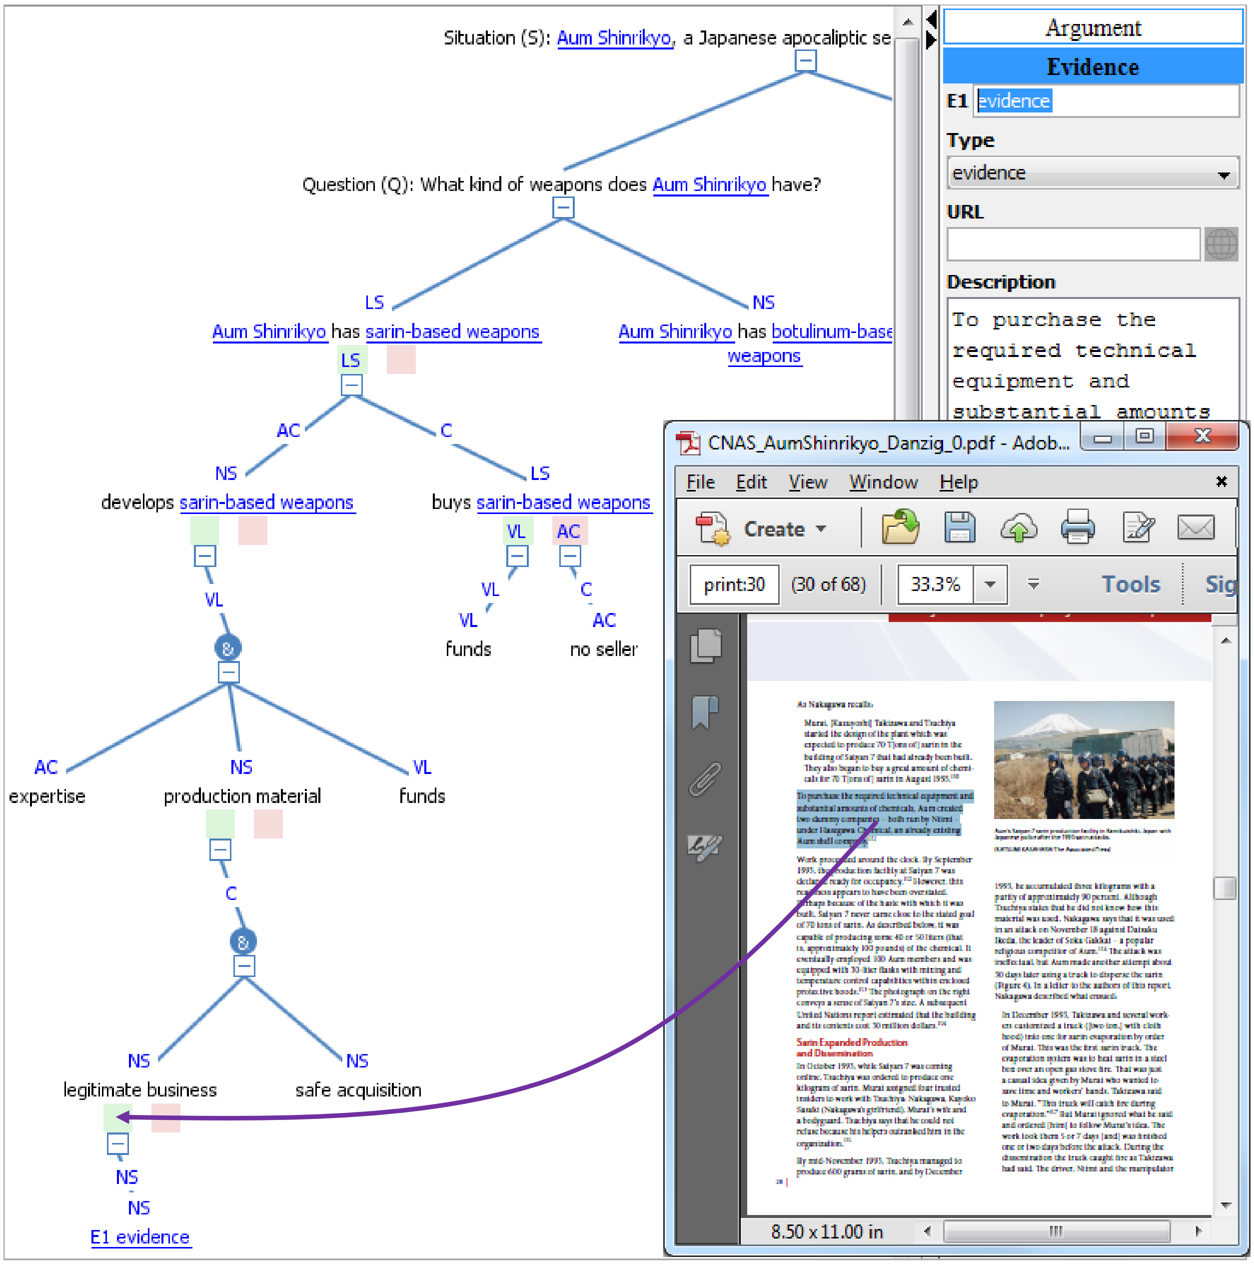

Figure 8: Creating and attaching an item of evidence to a hypothesis.

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 Figure 7, and selecting "Learn." As a result, Cogent learns the argument patterns shown at the right hand side of Figure 9. The patterns are obtained by using Cogent's ontology as a generalization hierarchy, where individual entities from the analysis in Figure 7, such as sarin-based weapons, are replaced with more general concepts from the ontology, such as weapon (Tecuci et al., 2005; 2008; 2016b).

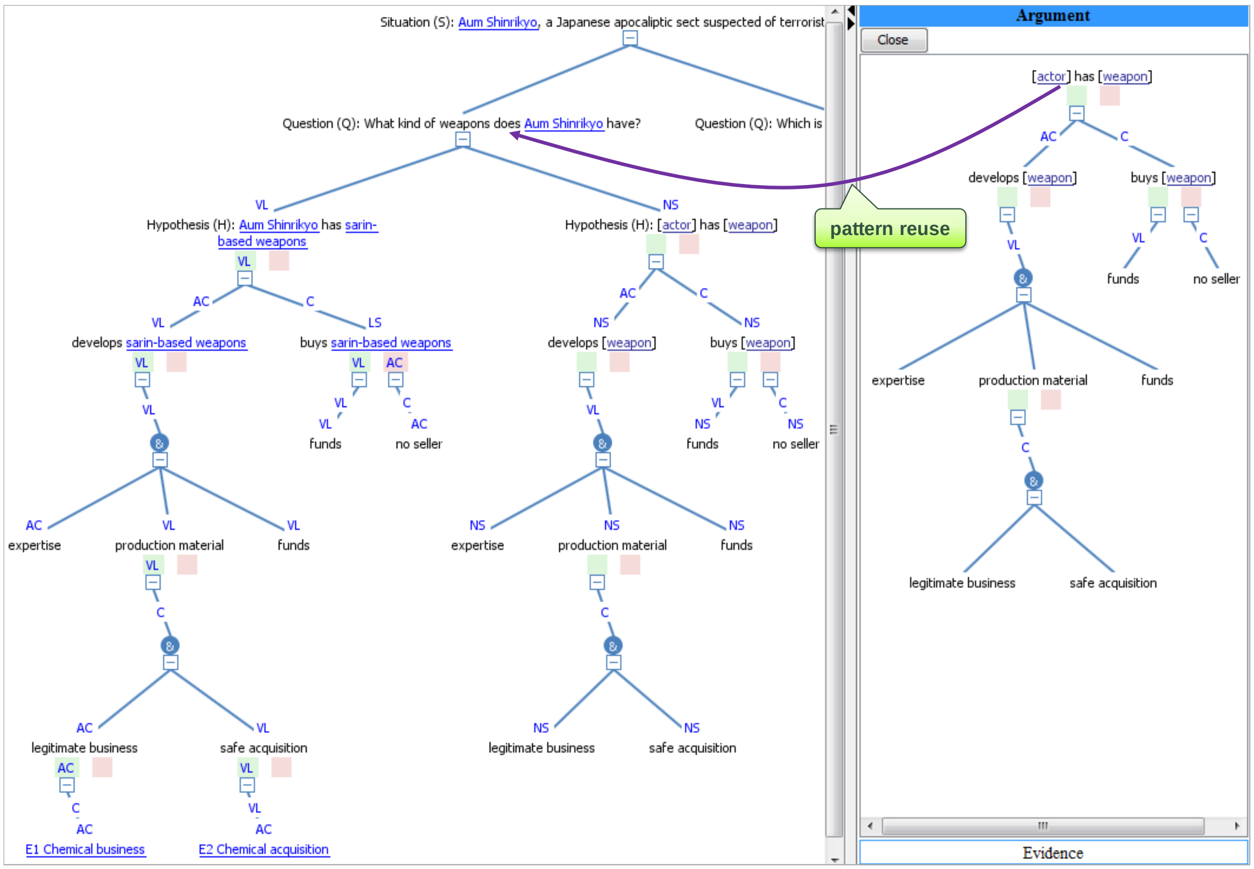

Figure 9: Argument development through pattern reuse.

Figure 9 also illustrates the reuse of the learned hypothesis patterns. The analyst simply drags "[actor] has [weapon]" from the Argument assistant, and drops it on "Question(Q): "What kind of weapons does Aum Shinrikyo have?" As a result, Cogent generates the analysis from the middle part of Figure 9. Then the analyst clicks on each concept in the generated analysis (e.g., [weapon]) and selects an entity to instantiate it from a list provided by Cogent, or defines a new entity (e.g., "botulinum-based weapons"). Cogent automatically replaces each occurrence of that concept in the analysis generated from the pattern with the selected entity.

To complete the analysis of the new hypothesis, the analyst needs to search for evidence relevant to the leaf hypotheses, and to assess it.

Notice in this illustration how much faster the analysis can be completed through the reuse of the learned patterns.

DEEPER BELIEVABILITY/CREDIBILITY 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 indicated above, the knowledge base of Cogent include a general ontology of evidence, a fragment of which is shown in Figure 10.

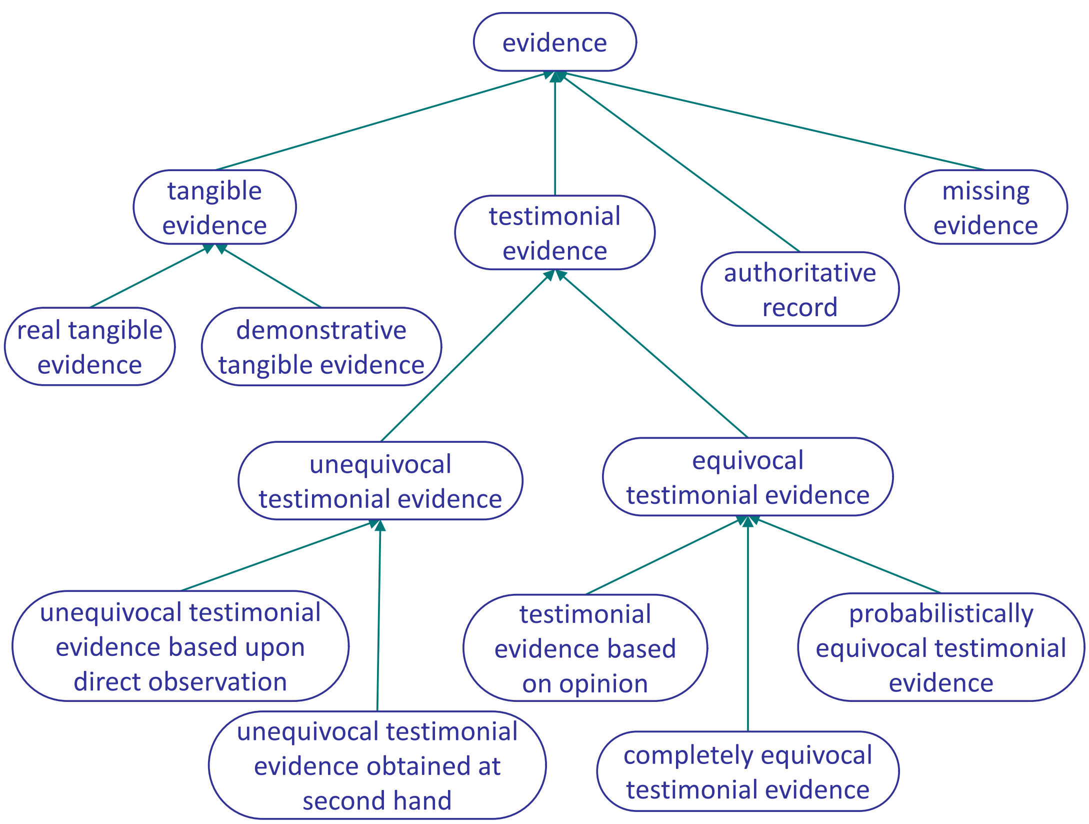

Figure 10: Fragment of the ontology of evidence.

For each evidence type in this ontology there is a general pattern for assessing its believability based on lower-level credentials [Tecuci et al., 2016a, pp.118-138]. Thus, when a deeper, more detailed believability analysis is needed, Cogent can generate the corresponding Wigmorean network.

Figure 11, for example, shows the argument pattern for assessing the believability of *demonstrative tangible evidence*.

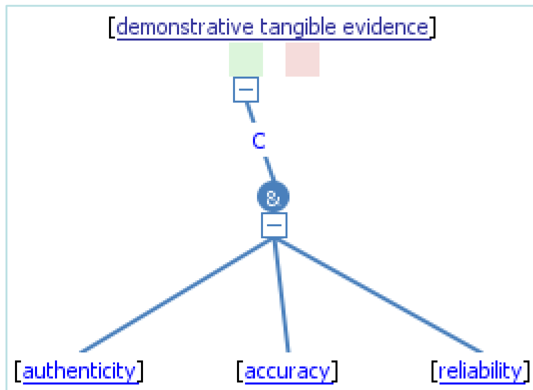

Figure 11: Pattern for assessing the believability of demonstrative tangible evidence.

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 12 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?

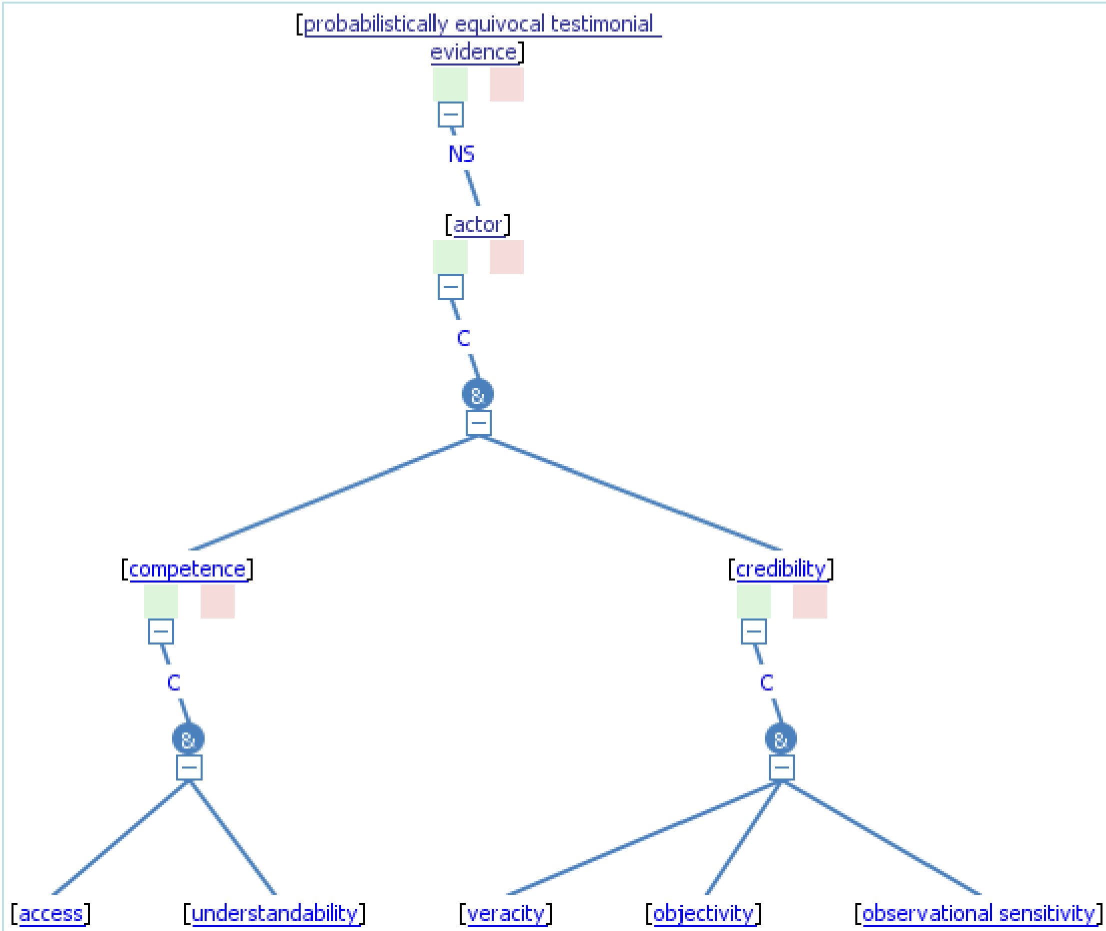

Figure 12: Pattern for assessing the believability
of probabilistically equivocal testimonial evidence.

As shown in the pattern from Figure 12, 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*.

COGNITIVE ASSISTANCE

As discussed in the previous sections, 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 process leads to the development of defensible and persuasive conclusions.

Cogent also enables 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 (or revised) evidence and assumptions.

CONCLUSIONS

This paper has presented a brief overview of Cogent, emphasizing the integration of an analyst's imaginative reasoning with the computer's critical reasoning.

Future work will continue the development of Cogent to make it easier to use. It will integrate several advanced analytic capabilities to automatically analyze the developed argumentation, identify critical elements, detect potential problems, suggest improvements, and generate specialized views of the analysis, as well as analysis reports.

ACKNOWLEDGEMENTS

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA). Work on a previous version of Cogent (Tecuci et al., 2015) was supported in part by the Intelligence Community (IC), the Department of Defense (DOD), and George Mason University (GMU). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, IC, DOD, GMU, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

Cohen L. J. 1977. *The Probable and the Provable,* Clarendon Press, Oxford.

Cohen L. J. 1989. *An Introduction to the Philosophy of Induction and Probability,* Clarendon Press, Oxford.

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.

Schum D. A. 1987. *Evidence and Inference for the Intelligence Analyst* (2 volumes). University Press of America.

Schum D. A. 1979. A Review of a Case Against Blaise Pascal and his Heirs, *Michigan Law Review*, 77(Jan.-Mar.): 446-483.

Schum D. A. 1994/2001a. *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.

Tecuci G., Boicu M., Boicu C., Marcu D., Stanescu B., and Barbulescu M. 2005. The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence, Vol.21, No.4*, pp. 462-

479. http://lac.gmu.edu/publications/2005/TecuciG_Disciple_RKF_CI.pdf

Tecuci G., Boicu M., Marcu D., Boicu C., and Barbulescu M. 2008. Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development*, July. http://lac.gmu.edu/publications/2008/Disciple-LTA08.pdf

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 November 2015, Technical Report FS-15-02, AAAI Press, Palo Alto, CA. http://lac.gmu.edu/publications/2015/Cogent-overview.pdf

Tecuci G., Schum D. A., Marcu D., Boicu M. 2016a. *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*, Cambridge University Press.

Tecuci G., Marcu D., Boicu M., Schum D. A. 2016b. *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, Cambridge University Press.

Wigmore J. H. 1913. The Problem of Proof, *Illinois Law Review*, Vol. 8, No. 2, pp. 77 – 103.

Wigmore J. H. 1937. *The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience and Illustrated in Judicial Trials*, 3rd edition, Little, Brown & Co, Boston, MA.

Wigmore J. H. 1940. *Wigmore on Evidence,* vol. 7.

Zadeh L.A. 1965. Fuzzy Sets, *Information and Control*, Vol.8, Issue 3, pp. 338 – 353, Elsevier. http://www.sciencedirect.com/science/article/pii/S001999586590241X (accessed January 19, 2016).