

Anticipatory Intelligence Analysis with Cogent: Current Status and Future Directions

Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, Chirag Uttamsingh

Learning Agents Center, Volgenau School of Engineering, George Mason University, Fairfax, VA 22030, USA
{ tecuci; dmarcu; mboicu; cuttamsi }@gmu.edu

Abstract

Anticipatory intelligence analysis is the complex task of drawing defensible and persuasive conclusions about future events or states based on current information. This paper presents a systematic approach to anticipatory intelligence analysis with Cogent, a software cognitive assistant that enables a synergistic integration of the analyst's imagination and expertise with the computer's knowledge and critical reasoning. It shows how current, as well as envisioned capabilities of Cogent, help alleviate many of the complexities of this task.

Introduction

Anticipatory intelligence analysis is the complex task of drawing defensible and persuasive conclusions about future events or states based on current information of all kinds that come from a variety of different sources. It addresses new and emerging trends, changing conditions, and underappreciated developments (ODNI, 2019).

The prevailing approach to anticipatory intelligence analysis and intelligence analysis in general is the holistic approach where the analysts, after reviewing large amounts of information and performing the reasoning in their heads, reach a conclusion (Marrin, 2011). A complementary approach uses structured analytic techniques, such as those described by Heuer and Pherson (2011) that guide the hypothesis generation and testing process performed by the analysts. Some of these methods and more advanced ones based on Bayesian probabilistic inference networks are implemented in analytical tools, such as Netica (2019).

We have developed a sequence of cognitive assistants based on Wigmorean networks (Wigmore, 1937). The first of these systems, Disciple-LTA, integrates capabilities for analytic assistance, learning, and tutoring (Tecuci et al., 2008). TIACRITIS and its subsequent version, Disciple-CD (Tecuci et al, 2016a), were developed primarily for teaching intelligence analysis and were experimentally used in many

IC and DOD organizations. While praising their theoretical framework and evidentiary knowledge, the analysts desired a simplified interface and interaction, which has led to the development of Cogent (Tecuci et al., 2015; 2018).

In this paper we discuss how Cogent supports an analyst in performing anticipatory analysis. We start with a brief account of the complexity of this task. Then we introduce a systematic approach to anticipatory analysis which is grounded in the science of evidence (Schum, 2009). Following that we present two examples of anticipatory analysis with Cogent, and discuss how it assists the analysts in coping with their complexity. Finally we discuss future developments of Cogent.

Complexity of Anticipatory Analysis

Evidence-based Reasoning

The evidence upon which the anticipation of possible future states or events eventually rests has five major characteristics that make these anticipations necessarily probabilistic in nature (Tecuci et al., 2016a, pp.159-167). The evidence is always *incomplete* no matter how much we have. It is commonly *inconclusive* in the sense that it is consistent with more than one future state or event. Further, the evidence is frequently *ambiguous* and, in most situations, *dissonant*, some of it favoring one future state or event while other evidence favoring others. Finally, the evidence comes from sources having different levels of *credibility*. Arguments to test the hypothesized future states or events are necessary in order to establish and defend the three major credentials of evidence: its *relevance*, its *credibility*, and its *inferential force* or *weight*. These arguments rest upon both imaginative and critical reasoning.

Assessing the Credibility of Evidence

Evidence credibility assessments form the very foundation for all arguments we make from evidence to possible future states or events. The different types of tangible, testimonial,

and mixed evidence have many credibility indicators and all would need to be assessed based on ancillary evidence in order to have high confidence in the accuracy of our anticipations (Tecuci et al., 2016a, pp. 118-133).

Limits of Individual Probability Views

While anticipatory intelligence analysis is probabilistic in nature, none of the non-enumerative probability views known to us (Subjective Bayesian, Belief Functions, Baconian, and Fuzzy) can optimally cope with all the five characteristics of evidence mentioned above (Schum, 2001a; Tecuci et al., 2016a, pp.173-208). For example, the conventional Subjective Bayesian view cannot cope well with ambiguities or imprecision in evidence. On the other hand, the Fuzzy view can naturally cope with such imprecisions. But neither the Bayesian view nor the Fuzzy view can account for the incompleteness of the coverage of evidence. The only view that can account for this is the Baconian view where the probability of a future state or event depends on how complete the evidence is, or how many questions recognized as being relevant remain unanswered by the evidence we have. This is in contrast with the Bayesian, Belief Functions, and Fuzzy views that all rest on how strong is the evidence we have about the considered future state or event. While on the Bayesian probability scale “0” means disproof, on the Baconian scale, “0” simply means lack of proof. A future state/event currently having “0” Baconian probability can be revised upward in value as soon as we have some evidence for it.

Time Constraints

A major objective of anticipatory intelligence analysis is to help insure that the policies and decisions reached by the governmental and military leaders, at all levels, are well informed. In many cases 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.

Non-Stationary World

As outlined above, anticipatory intelligence analysis has many inherent difficulties, but none seem more difficult than the fact that analysts must assess future states or events in a non-stationary world that keeps changing as analysts are trying to understand it. As a result, we have continuing streams of new information, some items of which are relevant evidence regarding our anticipations. 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 future event regarded as highly likely today may be overtaken by events we will learn about tomorrow. In fact, the very questions we asked yesterday may need to be revised or may even seem

unimportant in light of what we learn today. The consequence is that the complex process of discovery or investigation in anticipatory analysis is a ceaseless activity.

Anticipatory Intelligence Analysis in the Framework of the Scientific Method

Following the framework of the scientific method, we model anticipatory analysis as *ceaseless discovery of evidence, anticipations, and arguments*, in a non-stationary world, involving collaborative computational processes of *evidence in search of anticipations, anticipations in search of evidence, and evidentiary testing of anticipations*, as represented in Figure 1 (Tecuci et al., 2016a).

First, through *abductive (imaginative) reasoning* that shows that something is *possibly* true (Schum, 2001b), we generate alternative future events or states that may explain an intelligence alert. If, instead of an intelligence alert, the starting point is an intelligence question, the alternative anticipations are the possible answers to this question. Next, through *deductive reasoning* that shows that something is *necessarily* true, we use the hypothesized future states/events to generate new lines of inquiry and discover new evidence. After that, through *inductive reasoning* that shows that something is *probably* true (Schum, 2001a), we test each anticipation by developing an argumentation that shows how the discovered evidence favors or disfavors it.

As shown at the bottom of Figure 1, these are collaborative processes that support each other in recursive calls. For example, the discovery of new evidence may lead to the modification of the existing hypotheses of future events/states or the generation of new ones that, in turn, lead to the search and discovery of new evidence. Also, inconclusive testing of the considered anticipations requires the discovery of additional evidence. The next sections illustrates this process using Cogent in two analyses.

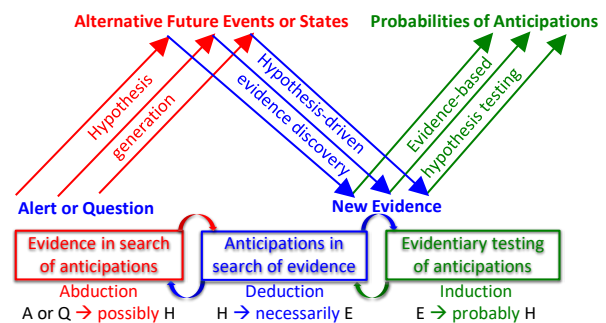


Figure 1: A framework for anticipatory analysis.

Alert-Driven Anticipatory Analysis

The following example of anticipatory analysis with Cogent shows how evidence about a missing cesium-137 canister

leads to anticipating that a dirty bomb will be set off in the Washington, D.C., area (Tecuci et al., 1016a). Note that this scenario and all the entities involved are fictitious.

Mavis, a counterterrorism analyst, reads in today's Washington Post that a canister containing cesium-137 is missing from the warehouse of the XYZ Company in Maryland (see E* at the bottom of Figure 2). The question is: *What hypothesis would explain this observation?*

Through abductive (imaginative) reasoning, Mavis infers that a dirty bomb will be set off in the Washington, D.C., area (see H₅ at the top of Figure 2). However, no matter how imaginative or important this future event is, no one will take it seriously unless Mavis and her cognitive assistant, Cogent, are able to justify it. So they develop the chain of abductive inferences shown in the left side of Figure 2:

We have evidence that the cesium-137 canister is missing. Therefore it is possible that it is indeed missing. It is possible that it was stolen. It is possible that it was stolen by a terrorist organization. It is possible that the terrorist organization will use the cesium-137 canister to build a dirty bomb. It is possible that the dirty bomb will be set off in the Washington, D.C., area.

But these are not the only hypotheses that may explain the evidence. Just because there is evidence that the cesium-137 canister is missing does not mean that it is indeed missing. At issue here is the credibility of the source of this information. Thus an alternative hypothesis is that the cesium-137 canister is not missing. But let us assume that it is missing. Then it is possible that it was stolen, but it is also possible that it was misplaced, or maybe it was used in a project at the XYZ Company. Now let us assume that the

cesium-137 canister was stolen. Then it is possible that it was stolen by a terrorist organization, or by a competitor of XYZ, or by an employee. Upper level hypotheses concern possible future events, such as, the terrorist organization will use the cesium-137 canister to build a dirty bomb, the dirty bomb will be set off in the Washington, D.C., area, or it will be set off in the New York area.

The analyst and Cogent need to assess each of these competing hypotheses, and determine which of them are likely. Starting from bottom-up, each hypothesis is put to work to guide the collection of additional evidence:

Assuming that the cesium-137 canister is indeed missing, what other things should be observable? What are the necessary conditions for an object to be missing from a warehouse? It was in the warehouse, it is no longer there, and no one has checked it out.

As a result, the analyst contacts Ralph, the supervisor of the warehouse, who reports that the cesium-137 canister is registered as being in the warehouse, that no one at the XYZ Company had checked it out, but it is not located anywhere in the hazardous materials locker. He also indicates that the lock on the hazardous materials locker appears to have been forced. Ralph's testimony provides several items of relevant evidence, and the question is: *What is the probability that the cesium-137 canister is missing, based on this evidence?*

To answer this question, Mavis and Cogent build the Wigmorean probabilistic inference network from Figure 3 (Wigmore, 1937; Tecuci et al., 2018). It integrates logic and Baconian probabilities with Fuzzy qualifiers, and uses the min/max probability combination rules common to the Baconian and Fuzzy views of probability (Cohen, 1977; 1989; Zadeh, 1983). That is, the probability of a conjunction of hypotheses is the minimum of their probabilities, and the probability of a disjunction of hypotheses is the maximum of their probabilities.

First Mavis and Cogent have to assess the probabilities of the bottom hypotheses in Figure 3, based on the corresponding relevant evidence. Then these probabilities are composed to produce the probability of the top hypothesis. The probability of a hypothesis, like the one from the bottom left of Figure 3, shown also at the top of Figure 4, is assessed based on the three credentials of evidence: *credibility, relevance, and inferential force* (Tecuci et al., 2016a, pp. 62-73). They are assessed by using the ordered symbolic probability scale from the upper right of Figure 3. As in the Baconian system, "lacking support" for a hypothesis means that we currently have no basis to consider that the hypothesis might be true. However, we may later find evidence to infer that the hypothesis is, for instance, "likely." Figure 3 shows a favoring argument for the top hypothesis and therefore it appears under the left (green) square. Disfavoring arguments (if any) appear under the right (pink) square.

The credibility of the evidence item E2 in Figure 4 is

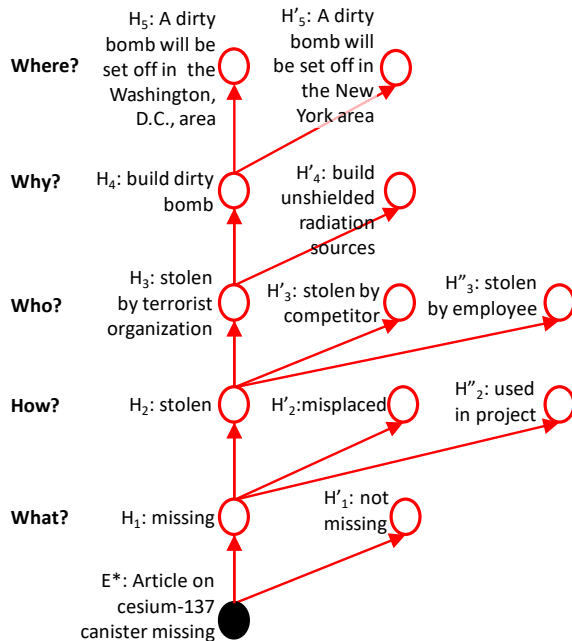


Figure 2: Multi-step abduction and competing explanations.

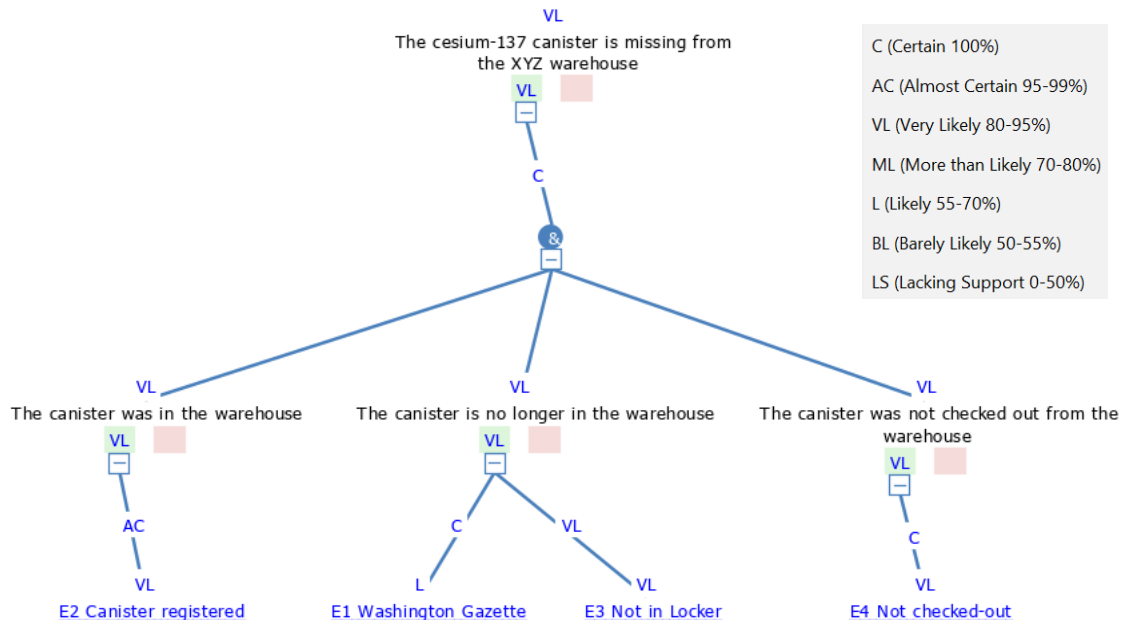


Figure 3: Wigmorean probabilistic inference network.

assessed as very likely because its source, Ralph, has access to the reported information and has a reputation for honesty. The relevance of E2 is assessed as almost certain because the records of the XYZ Company are almost certainly correct. Consistent with both the Baconian and the Fuzzy min/max probability combination rules, the inferential force of E2 on the hypothesis H is determined as the minimum between the credibility of E2 (very likely) and the relevance of E2 (almost certain). Thus, the inferential force of E2 on H is very likely. 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 credible will have no inferential force either. Only an item of evidence that is both relevant and credible supports the truthfulness of a hypothesis.

Because in the argumentation from Figure 4 there is only one item of favoring evidence, E2, its inferential force on the hypothesis is also the probability of the hypothesis. In general, however, the probability of the hypothesis would be the result of the balance of probabilities between the combined inferential force of the favoring evidence items and the combined inferential force of the disfavoring items.

As shown at the top of Figure 3, it is very likely that the cesium-137 canister is missing, this being the minimum between the probabilities of the three sub-hypotheses and the relevance of their conjunctive argument.

Some of the newly discovered evidence may trigger new hypotheses or the refinement of the current hypotheses. For example, during her initial investigation, Mavis discovered a video segment from a security camera at the warehouse showing a person loading a container into a U-Haul truck, leading her to refine the “stolen” hypothesis to indicate that

the cesium-137 canister was stolen with a U-Haul truck.

Having concluded that the cesium-137 canister is missing, Mavis and Cogent have now to establish whether it was stolen with a truck, it was misplaced, or it was used in a project at the XYZ Company. Each of these hypotheses is put to work to guide the collection of evidence for assessing it: *If the cesium-137 canister was stolen with a truck, what other things should be observable?*

Based on the current evidence, Mavis imagines the following scenario on how the cesium-137 canister might have been stolen: *The truck entered the company, the canister was stolen from the locker, the canister was loaded into the truck, and the truck left with the canister.*

Such scenarios have enormous heuristic value in advancing the investigation because they consist of mixtures of what is taken to be factual and what is conjectural. Conjecture is necessary in order to fill in natural gaps left by the absence of existing evidence. Each such conjecture, however, opens up new avenues of investigation, and the

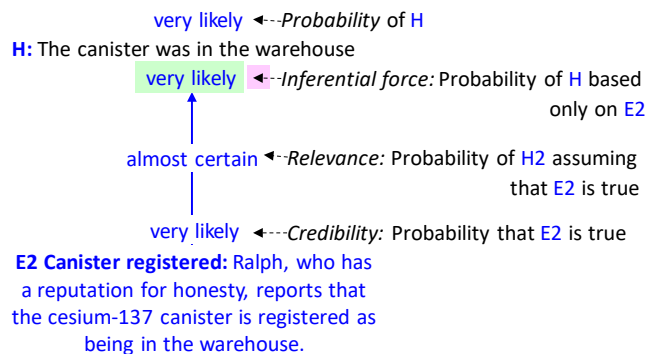


Figure 4: Evidence credentials.

discovery of additional evidence, if the scenario turns out to be true. This scenario, for instance, leads Mavis to check the records of the security guard and they show that a panel truck bearing Maryland license plate number MDC-578 was in the XYZ parking area on the day before the discovery that the cesium-137 canister was missing.

Fusing all the discovered evidence, Mavis and Cogent conclude that it is very likely that the cesium-137 canister was stolen with the MDC-578 truck. After further investigation, they also conclude that the two competing hypotheses, “misplaced” and “used in a project,” lack evidentiary support.

Continuing the investigation with the rental company owning the truck, it is discovered that Omar al-Massari rented the MDC-578 truck giving his alias, Omer Riley, and a false address, and that the truck is now contaminated because cesium-137 is radioactive. These lead to the conclusion that Omar used the truck to steal the cesium-137 canister. It is further discovered that Omar al-Massari has ties with terrorist organizations, and that he has given the cesium-137 canister to Saeed al-Nami, alias Kenny Derwish, who is a member of the terrorist organization Jihad Bis Sayf. These discoveries lead to the specializations of the hypotheses from Figure 2 as shown in Figure 5.

Figure 6 shows the analysis of the top anticipatory hypothesis “Jihad Bis Sayf will set off a dirty bomb in the Washington, DC, area.” It shows that Jihad Bis Sayf has reasons, desire, and capability to set off the dirty bomb. It has reasons because it is a terrorist organization opposed to the United States, it has a presence in the Washington, DC, area, and a dirty bomb in this area would have a very high impact. Furthermore, Jihad Bis Sayf has both the ability to build the bomb and to set it off. In particular, it has the radioactive material from the stolen cesium-137 canister, and further investigation has determined that it has both the necessary explosive material (Saeed al-Nami has stolen 2 pounds of RDX explosive) and expertise to build the bomb (Saeed al-Nami has expertise in explosives and has received training in the building of dirty bombs).

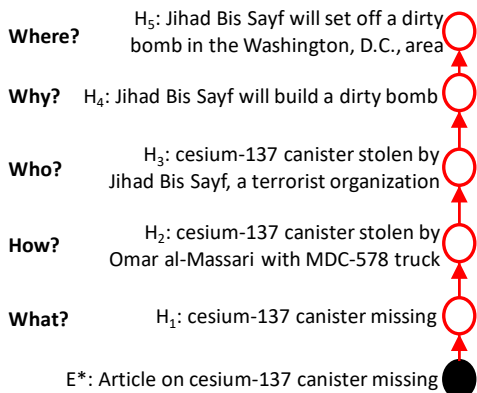


Figure 5: Evidence-based hypothesis specialization.

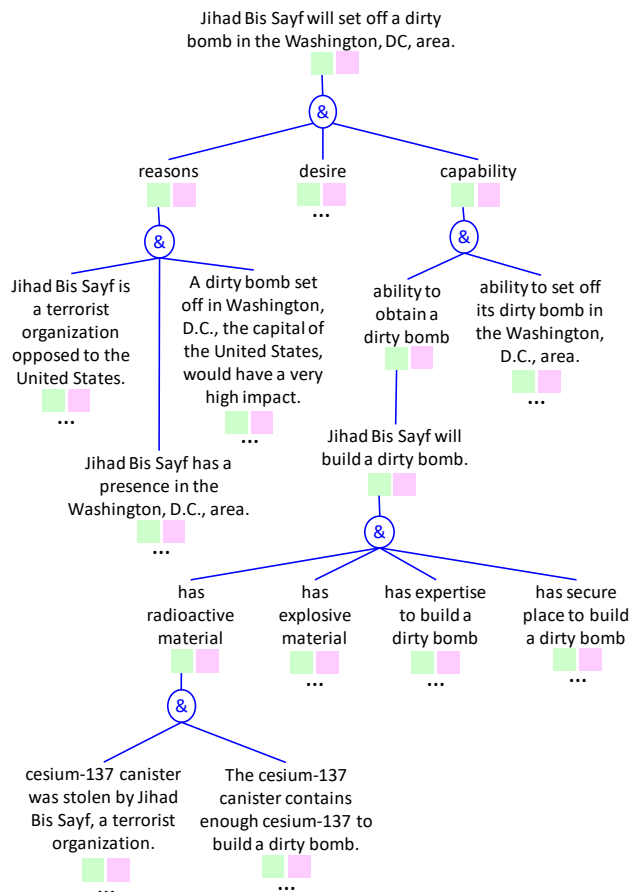


Figure 6: Anticipatory analysis.

Through such spiral hybrid reasoning, where abductions, deductions, and inductions feed on each other in recursive calls, Mavis and Cogent continuously generate and update intermediate alternative hypotheses, use these hypotheses to guide the collection of relevant evidence, and use the evidence to test these hypotheses, until the probability of the top-level hypotheses are assessed, ultimately anticipating that *Jihad Bis Sayf will likely set off a dirty bomb in the Washington, D.C., area.*

Note, however, that performing this analysis is not as simple as one may infer from this presentation. It is the methodology from Figure 1 and Cogent that guide the analyst and simplify it. Many things can and will indeed go wrong. But Cogent provides the means to deal with them. Based on evidence, you come up with some hypotheses, but then you cannot find more evidence to support any of them. So you need to come up with other hypotheses, and you should always consider alternative hypotheses. The deduction-based decomposition approach guides you on how to look for evidence, but your knowledge and imagination also play a crucial role. As illustrated here, Mavis imagined a scenario where the cesium-137 canister was stolen with a truck. But let us now assume that she did not find supporting evidence for this scenario. Therefore,

Mavis has to imagine other scenarios. Maybe the cesium canister was stolen by someone working at the XYZ Company, or maybe it was stolen by Ralph, the administrator of the warehouse. The important thing is that each such scenario opens up a new line of investigation.

Question-Driven Anticipatory Analysis

The previous section illustrated a situation where the anticipatory analysis was driven by an intelligence alert. This section illustrates a situation where the analysis is driven by the intelligence question: *Who will be the world leader in wind power within the next decade?*

The top level of the corresponding anticipatory analysis is shown in Figure 7. It is likely that the United States will be the world leader in wind power within the next decade because almost certainly they have reasons, likely they have the desire, and almost certainly they have the necessary capability. A reason is that significant production of wind power will reduce the current need of the United States to consume huge quantities of oil that represent a danger to the environment. The desire of the United States, which is a representative democracy, is determined by the desire of the people (almost certain), the desire of the major political parties (very likely), and the desire of the energy industries

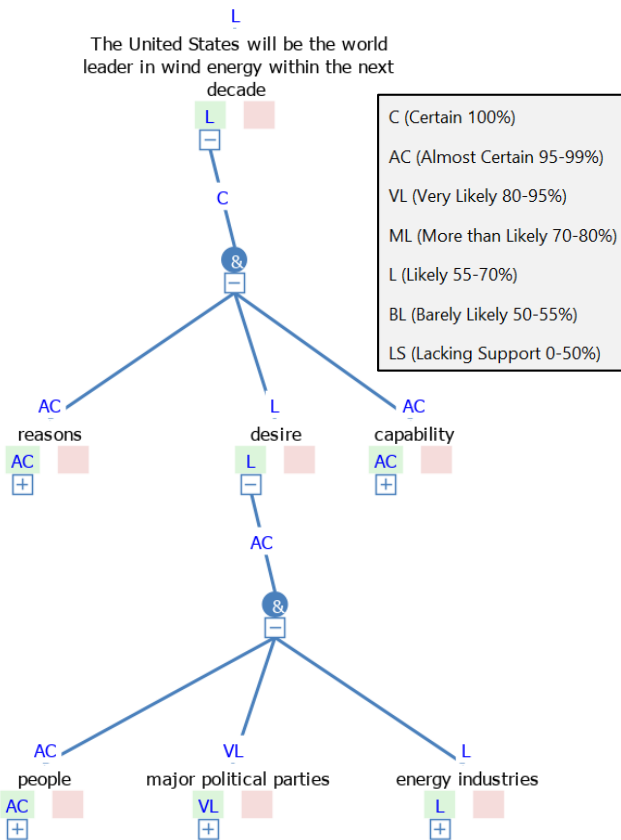


Figure 7: Another example of anticipatory analysis.

(likely). They combine into an overall desire of likely. Finally, almost certainly the United States have the required capabilities that consist of required homegrown scientific knowledge, technical knowledge, economic resources, and natural resources.

Current Cognitive Assistance

Consider again the process described in Figure 1. With the current version of Cogent, the analyst has to imagine the possible future states or events. However, Cogent helps with developing argumentations that lay out the underlying analytical framework for every anticipation, including the connection between the evidence and various intermediate hypotheses in the analysis, the evaluation of the credibility of evidence and its strength in supporting a hypothesis, and the role of assumptions in addressing missing information.

Anticipatory analysis may be affected by the analyst's biases. Cogent can detect several of them, such as the *confirmation bias* (building an argumentation and/or only searching for evidence that confirms the analyst's beliefs while dismissing or ignoring evidence to the contrary), the *satisficing bias* (choosing the first hypothesis that appears good enough, rather than carefully identifying all possible hypotheses and determining which one is the most consistent with the evidence), and potential *absence of evidence bias* (failure to consider the degree of completeness of the available evidence). Many other biases are avoided because explicit argumentations are developed that employ an intuitive system of symbolic probabilities.

Additionally, Cogent facilitates the analysis of what-if scenarios, where the analyst may make various assumptions and Cogent automatically determines their influence on the analytic conclusion. It also automatically updates the computed probabilities based on new or revised evidence.

Once the analysis is finalized, Cogent generates a structured report that the analyst then transforms into a more understandable and persuasive report that includes argumentation fragments and evidence, can be shared with other analysts, subjected to critical analysis, and correspondingly improved.

The hierarchical structure of the Wigmorean argumentation enables the analyst and Cogent to perform the analysis at different levels of abstraction. Moreover, the analyst may drill down on selected sub-hypotheses as much as allowed by the available time and evidence. Consider, for example, the dirty bomb anticipated event from Figure 6. In time-limiting situations the analyst may assume that Jihad Bis Sayf has the reasons and desire to set off the dirty bomb, and focus the investigation on its capability.

Note that an anticipated future state made at a given moment in time may change afterwards because many of its indicators are dynamic, such as those that determine the

desire of the United States in the wind power scenario (see Figure 7). Therefore evidence of these indicators needs to be continuously monitored and updated. In the current version of Cogent, this monitoring has to be done by the analyst who also needs to insert the new evidence in the analysis. After that Cogent automatically updates the analysis. However, as discussed in the next section, continuous monitoring and updating of evidence can also be automated.

Cogent has a knowledge base that includes an ontology of evidence and rules for assessing its credibility. Figure 8 shows a fragment of this ontology (Schum et al., 2009). For each type of evidence from this ontology, Cogent has a procedure for assessing its credibility. For example, as illustrated in the left hand side of Figure 8, the credibility of an item of demonstrative tangible evidence (e.g., a map, a sound recording, or a satellite image) depends on its authenticity, its accuracy, and the reliability of the instrument the produced it. The credibility of a human source depends on the source's competence, veracity, and accuracy. These indicators depend on lower level indicators. For example, the competence depends on the source's access and expertise, while the accuracy depends on the source's objectivity and observational sensitivity.

As discussed in the next section, future research will address the problem of learning general analysis rules from an expert analyst, which will speed-up and improve the development of new analyses. Note, for example, that the analyses from the previous sections, although very different, both make use of the following reasoning pattern: *An actor will perform a certain action or achieve a certain state if it has reasons, desire, and capability.* The rules learned from the analysis in Figure 7 will enable Cogent to automatically generate analyses of future states such as: *China will be the world leader in solar energy within the next decade.*

Mixed-Initiative and Automatic Anticipatory Analysis

We plan to significantly extend Cogent with knowledge-based reasoning and learning capabilities, effectively evolving it into the multi-agent architecture from Figure 9, similar to that described in (Tecuci et al., 2019). It will have two complementary functions:

(1) *Analysis and Learning*, shown in the upper part of Figure 9, where the Analyst and Cogent (through its Mixed-Initiative Learning and Reasoning component) will rapidly develop complex, logical, and compelling argumentations in a transparent manner. At the same

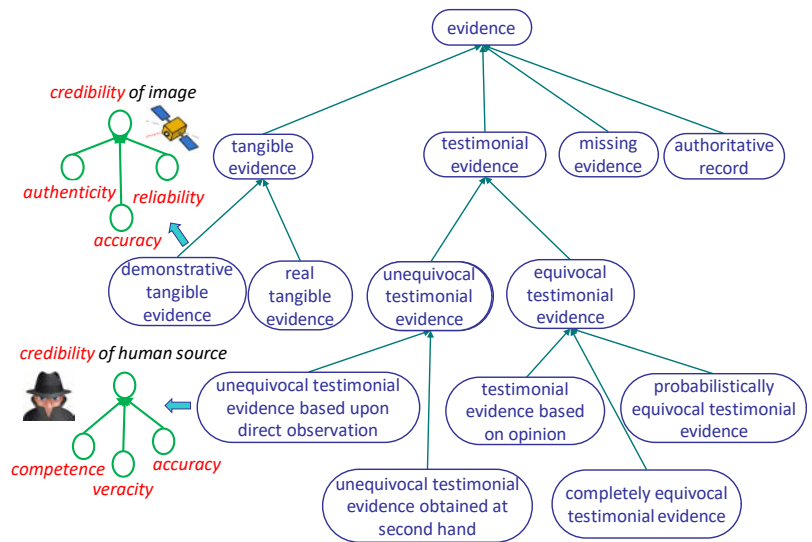


Figure 8: Evidence ontology and credibility patterns.

time, Cogent will learn domain analysis rules from the contributions of the Analyst, through the employment of the Disciple-EBR multistrategy learning approach, which integrates learning from examples, learning from explanations, and learning by analogy and experimentation, in a mixed-initiative interaction with the expert. Successive versions of this learning approach are presented in (Tecuci 1998; Tecuci et al. 2002; 2005; 2008; 2016b). Cogent will also facilitate the development of analyses in collaboration with other analysts, each contributing sub-arguments based on their expertise, reviewing and commenting on each-other's contributions, and sharing previously learned domain analysis rules, sources, and evidence.

(2) *Automatic Analysis Updating*, shown in the bottom part of Figure 9, where Cogent (through its collaborative autonomous agents) continuously monitors the Multi-INT Environment and updates the performed analysis based on the newly discovered evidence. These agents are copies of the corresponding modules of the Learning and Reasoning

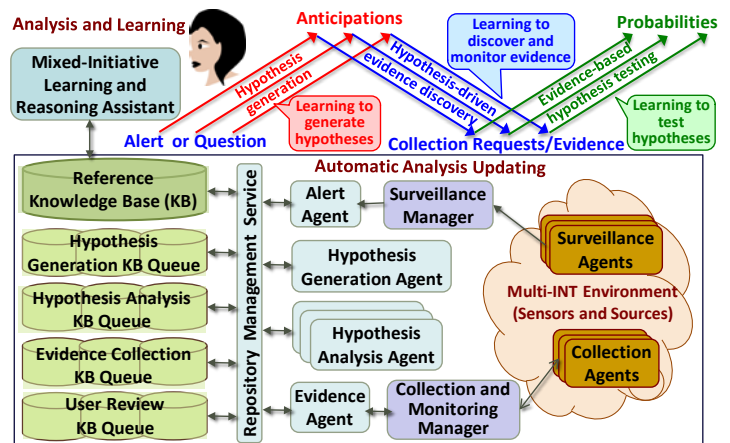


Figure 9: Envisioned extended architecture of Cogent.

Assistant except that they are configured to run autonomously and communicate by developing and exchanging Knowledge Bases. The component agents of Cogent will be connected to the application environment through a Surveillance Manager and a Collection and Monitoring Manager, the latter continuously monitoring the results returned by the Collection Agents that operate on the Multi-INT Environment. Once a new or updated evidence item is detected, it is introduced in the analysis by the Evidence Agent, and the analysis is updated by the Hypothesis Analysis Agent.

The automatic hypothesis generation is performed by the Alert Agent in collaboration with the Hypothesis Generation Agent. Next, the hypothesis-driven evidence discovery is performed by Hypothesis Analysis Agents in collaboration with the Evidence Agent. Then evidence requests are issued to collection agents through the Collection and Monitoring Manager. After that, evidence-based hypothesis testing is performed by the Evidence Agent in collaboration with the Hypothesis Analysis Agents.

Conclusions

After reviewing several of the complexities of anticipatory intelligence analysis, this paper illustrated how they can be alleviated through the use of the Cogent cognitive assistant within a systematic approach based on the science of evidence and the scientific method. Key to overcoming these complexities and performing more accurate anticipatory analyses based on imperfect information in a dynamic world, is the synergistic integration of the analyst's imagination and expertise with the computer's knowledge and critical reasoning.

Acknowledgements

David Schum has significantly influenced this research that was supported in part by NSF under grant number 1611742, by AFRL under contract number FA8750-17-C-0002, by IARPA under contract number 2017-16112300009, and by George Mason University. 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 any organization of 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.
- Heuer, R. J., and Pherson, R. H. 2011. *Structured Analytic Techniques for Intelligence Analysis*, CQ Press, Washington, DC.
- Marrin, S. 2011. *Improving Intelligence Analysis: Bridging the gap between scholarship and practice*. Routledge, London & New York.
- Netica, 2019. <https://www.norsys.com/>
- ODNI, 2019. *The National Intelligence Strategy of the United States of America*, Office of the Director of National Intelligence.
- Schum, D. A. 2009. Science of Evidence: Contributions from Law and Probability. *Law Probab Risk* **8** 197–231.
- 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):1645–1681.
- Schum, D., Tecuci, G., Boicu, M., Marcu, D., Substance-Blind Classification of Evidence for Intelligence Analysis, in *Proceedings of the Conference "Ontology for the Intelligence Community"*, Fairfax, Virginia, 20-22 October 2009.
- Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, San Diego: Academic Press.
- 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. *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, 853-860. Edmonton, Alberta: AAAI Press.
- Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M. 2005. The Disciple-RKF Learning and Reasoning Agent. *Computational Intelligence*, 21 462-479.
- Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M. 2008. Disciple-LTA: Learning, Tutoring and Analytic Assistance. *Journal of Intelligence Community Research and Development*, July.
- 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"*, 58-65, Arlington, VA, Technical Report FS-15-02, AAAI Press, Palo Alto, CA.
- 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.
- Tecuci, G., Kaiser, L., Marcu, D., Uttamsingh, C., Boicu, M. 2018. Evidence-based Reasoning in Intelligence Analysis: Structured Methodology and System, *Computing in Science and Engineering*, 20(6) 9-21, November/December.
- Tecuci, G., Meckl, S., Marcu, D., Boicu, M. 2019. Instructable Cognitive Agents for Autonomous Evidence-Based Reasoning. *Seventh Annual Conference on Advances in Cognitive Systems*, August 2-5, MIT, Cambridge, MA.
- 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.
- Zadeh, L. 1983. The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems, *Fuzzy Sets and Systems*, 11 199-227.