

Recognizing and Countering Biases in Intelligence Analysis with TIACRITIS

Gheorghe Tecuci, David Schum, Dorin Marcu, Mihai Boicu

Learning Agents Center, Volgenau School of Engineering, George Mason University, Fairfax, VA 22030, USA

Abstract— This paper discusses different biases which have been identified in Intelligence Analysis and how TIACRITIS, a knowledge-based cognitive assistant for evidence-based hypotheses analysis, can help recognize and partially counter them. After reviewing the architecture of TIACRITIS, the paper shows how it helps recognize and counter many of the analysts' biases in the evaluation of evidence, in the perception of cause and effect, in the estimation of probabilities, and in the retrospective evaluation of intelligence reports. Then the paper introduces three other types of bias that are rarely discussed, biases of the sources of testimonial evidence, biases in the chain of custody of evidence, and biases of the consumers of intelligence, which can also be recognized and countered with TIACRITIS.

Bias, cognitive assistant, intelligence analysis, evidence-based reasoning, argumentation, symbolic probabilities.

I. INTRODUCTION

Intelligence analysts face the difficult task of drawing defensible and persuasive conclusions from masses of evidence, requiring the development of often stunningly complex arguments that establish and defend the three major credentials of evidence: relevance, believability, and inferential force [1]. This highly complex task is affected by various biases which are inclinations or preferences that interfere with impartial judgment. Some of the biases are due to our simplified information processing strategies that lead to consistent and predictable mental errors. These errors remain compelling even when one is fully aware of their nature, and are therefore exceedingly difficult to overcome [2, p.111-112].

In this paper we propose an approach to the identification and countering of the biases in intelligence analysis. The approach is based on the observation that the best protection against biases comes from the collaborative effort of teams of analysts, who become skilled in the evidential and argumentational elements of their tasks, and who are willing to share their insights with colleagues, who are also willing to listen. As we discuss in this paper, this could be achieved by employing an intelligent analytic tool like TIACRITIS [3] which helps the analyst perform a rigorous evidence-based hypothesis analysis that makes explicit all the reasoning steps, probabilistic assessments, and assumptions, so that they can be critically analyzed and debated. The name TIACRITIS is an abbreviation of Teaching Intelligence Analysts Critical Thinking Skills, which was the initial motivation of developing this system. The system was later extended to also support its use for regular analysis.

In the next section we introduce the architecture of the TIACRITIS cognitive assistant which is based on semantic technologies for knowledge representation, reasoning, and

learning. Then, in Section III, we address the analysts' biases discussed by Heuer [2, pp.111-171]: biases in the evaluation of evidence, in the perception of cause and effect, in the estimation of probabilities, and in the retrospective evaluation of intelligence reports. After that we address three other origins of bias that are rarely discussed, even though they may be at least as important on occasion as any analysts' biases.

II. THE TIACRITIS COGNITIVE ASSISTANT

TIACRITIS is a knowledge-based system that supports an intelligence analyst in performing evidence-based hypothesis analysis in the framework of the scientific method. It guides the analyst to view intelligence analysis as ceaseless discovery of evidence, hypotheses, and arguments in a non-stationary world, involving collaborative processes of *evidence in search of hypotheses*, *hypotheses in search of evidence*, and *evidentiary testing of hypotheses* [1, 3]. Fig.1 is an abstract illustration of this astonishingly complex process. First we search for possible hypotheses that would explain a surprising observation E^* (see the left side of Fig.1): It is possible that F might be true. Therefore G might be true. Therefore H , a hypothesis of high interest, might be true. The problem with drawing this conclusion, however, is that there are other hypotheses that also explain E^* , such as F' , G' , and H' . To conclude H we would need to assess all the competing hypotheses, showing that F , G , and H are more likely than their competitors.




Fig. 1. Scientific method framework of TIACRITIS.

Let us assume that we have shown that F and G are more likely than their corresponding competing hypotheses. Next we have to assess H , H' , To assess H we need additional evidence which is obtained by successively decomposing H into simpler and simpler hypotheses, as shown by the blue tree in the right part of Fig.1. H would be true if G and M would be true. Then M would be true if N , Q , and S would be true. But if N would be true, then we would need to observe evidence E_n^* . So we look for E_n^* and we may or may not find it. This is the

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

process of hypotheses in search of evidence that guides the evidence collection task. Now some of the newly discovered items of evidence (e.g. E_n^*) may trigger new hypotheses, or the refinement of the current hypotheses. Therefore, as indicated at the bottom part of Fig.1, the processes of evidence in search of hypotheses and hypotheses in search of evidence take place at the same time, and in response to one another.

Then we use all the collected evidence to assess the hypothesis H . This assessment is probabilistic in nature because the evidence is always *incomplete*, usually *inconclusive*, frequently *ambiguous*, commonly *dissonant*, and has various degrees of *believability* [1]. In the computational theory of intelligence analysis we have developed [3], hypotheses assessment is based on a combination of ideas from the Baconian probability system [4] and the Fuzzy probability system [5], and uses a symbolic probability scale. In particular, in the latest version of TIACRITIS, the likeliness of a hypothesis may have one of the following ordered values:

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

In this scale, “no support” means that our evidence does not support the conclusion that the hypothesis is true. This may, however, change if new evidence favoring the hypothesis is later discovered. The likeliness of an upper-level hypothesis (e.g., H) is obtained from the likeliness of its sub-hypotheses (i.e., G and M) by using min or max Baconian and Fuzzy combination functions, depending on whether the sub-hypotheses G and M represent necessary and sufficient conditions for the hypothesis H , sufficient conditions, or just indicators. Competing hypotheses (e.g., H') are assessed in a similar way and the most likely hypothesis is selected. But if no hypothesis is more likely than all its competitors, then the processes of hypotheses in search of evidence, and evidence in search of hypotheses have to be resumed.

TIACRITIS was developed by first customizing the Disciple learning agent shell (a general agent building tool [6, 7]) into a learning agent shell for intelligence analysis, and then by training it with analysis knowledge from several domains [8]. The overall architecture of the Disciple learning agent shell for intelligence analysis is shown in Fig. 2. It contains integrated modules for ontology development, rule learning, problem solving and evidence-based reasoning, mixed-initiative interaction, and tutoring, as well as a hierarchically organized repository of knowledge bases (KB). At the top level of this repository is the general knowledge base for intelligence analysis (IA KB) which

contains knowledge applicable to the evidence-based analysis of any type of intelligence hypothesis, from any domain. Under it, and inheriting from it, are domain-specific knowledge bases. Each such Domain KB contains knowledge specific to a particular type of IA problems, such as 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 particular Scenario KB contains specific knowledge about the United States, as well as items of evidence to make the corresponding analysis. The actual analysis is done by using this knowledge as well as more general knowledge inherited from the corresponding Domain KB and from the IA KB.

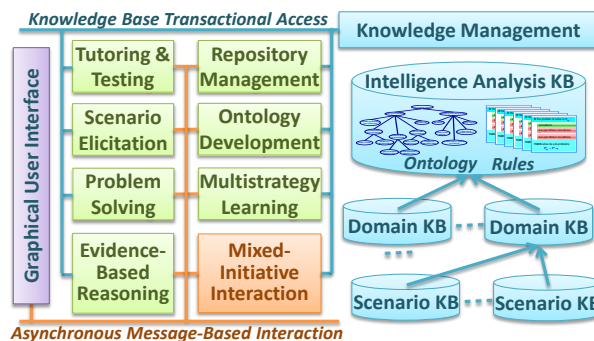


Fig. 2. Learning agent shell for intelligence analysis.

Each of these knowledge bases is structured into an ontology of concepts and a set of general problem solving rules expressed with these concepts. The rules are learned from specific examples of reasoning steps, by using the ontology as a generalization hierarchy [7]. The learning agent shell for intelligence analysis was obtained by training the Disciple learning agent shell with general intelligence analysis know-

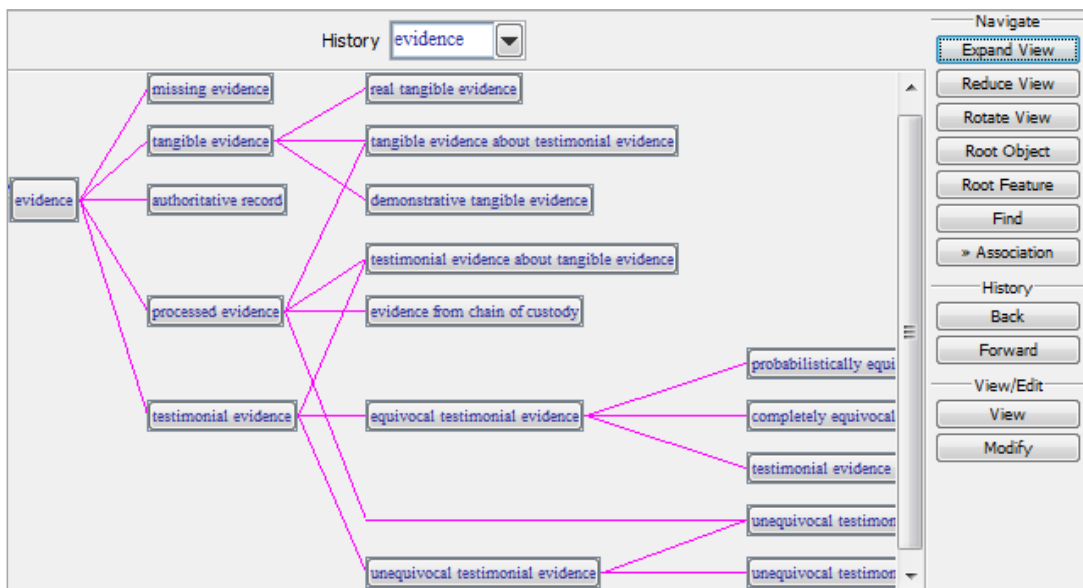


Fig. 3. Ontology fragment showing various types of evidence.

ledge resulting in the development of the IA KB. The IA KB contains both a general ontology and a set of general reasoning rules which are necessary for any Disciple agent for intelligence analysis, as we will briefly present in the following. For example, Fig. 3 shows a general ontology of evidence. It includes both basic types (e.g., testimonial evidence and tangible evidence), as well as evidence mixtures (e.g., testimonial evidence about tangible evidence). The ontology language of Disciple is an extension of RDFS [9] with additional features to facilitate learning [6, 7, 10].

Learned general rules from the IA KB include those for directly assessing a hypothesis based on evidence. These rules automatically reduce the assessment of a leaf hypothesis, such as **Q** in Fig. 1, to assessments based on favoring and disfavoring evidence and, further down, to the assessment of the *relevance* and the *believability* of each item of evidence with respect to **Q**. Once these assessments are made, they are combined, from bottom-up, to obtain the *inferential force* of all the items of evidence on **Q**, which results in the *likeliness* of **Q**.

An example of a learned rule is shown in Fig. 4. It is an if-then problem reduction rule that expresses how and under what conditions a generic hypothesis can be reduced to simpler generic hypotheses. The conditions are represented as first-order logical expressions [7]. In particular, this rule states that, in order to assess the believability of unequivocal testimonial evidence obtained at second hand, one needs to assess both the believability of our source, and the believability of the source of our source. It is by the application of such rules that an agent can generate the reduction part of the trees in Fig. 1 and Fig. 5.

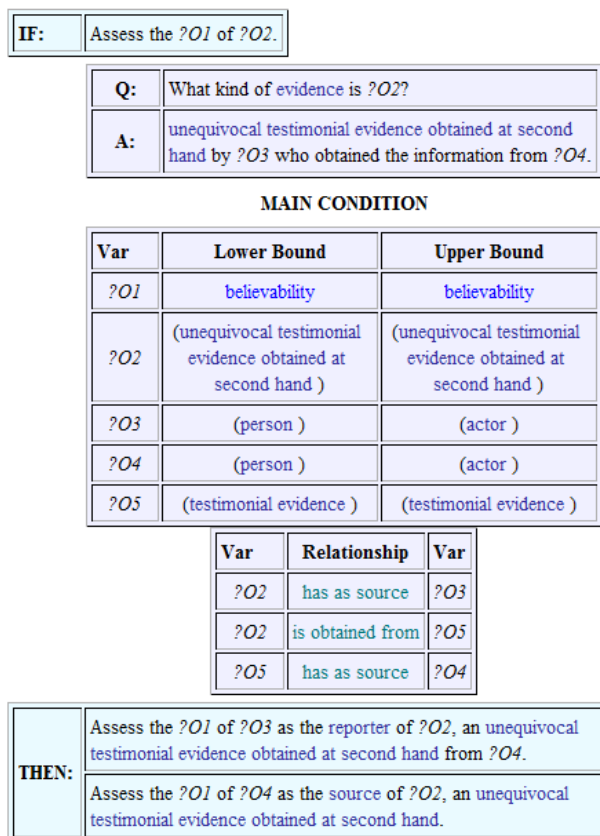


Fig. 4. Learned rule for believability analysis.

The ontology and the rules from the knowledge repository of TIACRITIS allow it to support the analyst in formulating hypotheses, developing arguments that reduce complex hypotheses to simpler and simpler ones (as discussed above), collecting evidence relevant to the simplest hypotheses, and finally assessing the relevance, the believability, and the inferential force of evidence, and the likeliness of the hypotheses. Additionally, TIACRITIS continuously learns from the performed analyses.

As discussed in the rest of this paper, TIACRITIS has one additional important capability. It supports the analysts in recognizing and countering many of their biases. Because Heuer has made a detailed and very well-known analysis of biases in intelligence analysis [2, pp.111-171], we follow his classification and identified characteristic of biases to show how TIACRITIS helps recognizing and countering many of them.

III. BIASES OF THE ANALYST

A. Biases in the Evaluation of Evidence

Heuer first mentions *vividness of evidence* as a necessary criterion for establishing its force. Analysts, like other persons, have preferences for certain kinds of evidence and these preferences can induce biases. In particular, analysts can have a distinct preference for vivid or concrete evidence when less vivid or concrete evidence may be more inferentially valuable. In addition, their personal observations may be over-valued.

First, as discussed in the previous section, the hypothesis in search of evidence phase of the analysis helps identify a wide range of evidentiary needs. For example, the argumentation in Fig. 1 shows that we need evidence relevant to N, evidence relevant to Q, evidence relevant to S, etc. It is unlikely that we would have vivid evidence for each basic hypothesis. So we would be forced to use less vivid evidence as well.

Second, as illustrated by the abstract analysis example in Fig. 5 and discussed in the following, TIACRITIS guides us to assess a simple hypothesis **Q** by performing a uniform, detailed, and systematic evaluation of each item of evidence, *regardless of its "vividness"*, helping us be more objective in the evaluation of the force of evidence.

Let us first consider how to assess the probability of **Q** based only on one item of favoring evidence E_k^* (see the bottom of Fig. 5). First notice that we call this *likeliness* of **Q**, and not *likelihood*, because in classic probability theory likelihood is $P(E_k^*|Q)$, while here we are interested in $P(Q|E_k^*)$, the posterior probability of **Q** given E_k^* . With TIACRITIS, to assess **Q** based only on E_k^* , we have three judgments to make by answering three questions:

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

The *believability* question is: *How likely is it that E_k^* is true?* Here the answer should be one of the values from "no

support” to “certain.” “Certain” means that we are sure that the event E_k reported in E_k^* did indeed happen. “No support” means that E_k^* provides us no reason to believe that the event E_k reported in E_k^* did happen. For example, we believe that the source of E_k^* has lied to us.

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




Fig. 5. The relevance, believability, and inferential force of evidence.

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

When we have several items of favoring evidence, we evaluate Q based on each of them (as was explained above), and then we compose the obtained results. This is illustrated in Fig.5 where the assessment of Q based only on E_k^* (almost certain) is composed with the assessment of Q based only on E_k^* (likely), through the maximum function, to obtain the assessment of Q based only on favoring evidence (almost certain). In this case the use of the maximum function is justified because it is enough to have one item of evidence that is both very relevant and very believable to make us believe that the hypothesis is true.

Let us now assume that Q^c based only on disfavoring evidence is “likely.” How should we combine this with the assessment of Q based only on favoring evidence? As shown at the top of Fig.5, TIACRITIS uses an *on balance* judgment: Because Q is “almost certain” and Q^c is “likely,” it concludes that, *based on all available evidence*, Q is “very likely.”

Heuer also mentions the *absence of evidence* as another origin of bias. The bias here concerns a failure to consider the degree of completeness of available evidence. Consider again the argumentation from Fig. 1 which decomposes complex hypotheses into simpler sub-hypotheses that are assessed based

on evidence. This argumentation structure makes very clear that S is not supported by any evidence. Thus the analyst should lower her confidence in the final conclusion, countering the *absence of evidence* bias.

The next source of bias mentioned by Heuer is a related one: *oversensitivity to evidence consistency, and not enough concern about the amount of evidence we have*. This kind of bias can easily manifest when using an analytic tool like Heuer’s ACH [11] where the analyst judges alternative hypotheses based on evidence, without building any argumentation. With TIACRITIS, the argumentation will reveal if most of the evidence is only relevant to a small fraction of sub-hypotheses, while many other sub-hypotheses have no evidentiary support. For example, the argumentation from Fig. 1 shows that most of the evidence is related to hypothesis Q.

According to Heuer [2, pp. 121-122]: “When working with a small but consistent body of evidence, analysts need to consider how representative that evidence is of the total body of potentially available information.” The argumentation from Fig. 1 makes very clear that the available evidence is not representative of all the potentially available information. We have no evidence relevant to S. If we would later find such evidence which would indicate “no support” for S, then the considered argumentation would provide “no support” for the top-level hypothesis H. When faced with sub-hypotheses for which there is no evidence (e.g., S in Fig. 1), TIACRITIS allows the analyst to consider various what-if scenarios, making alternative assumptions with respect to the likeliness of S, and determining their influence on the likeliness of H. This should inform the analyst on how to adjust her confidence in the analytic conclusion, to counter the oversensitivity to evidence consistency bias.

Finally, Heuer lists the *persistence of impressions based on discredited evidence* as an origin of bias. If Heuer had written his book in 2003, he might have used the case of Curveball as a very good example [12]. In this case, Curveball’s evidence was discredited on a number of grounds but was still believed and taken seriously by some analysts as well as many others.

TIACRITIS helps countering this bias by incorporating in the argumentation an explicit analysis of the believability of evidence, especially for key evidence that has a direct influence on the analytic conclusion. When such an evidence item is discredited, specific elements of its analysis are updated, and this leads to the automatic updating of the likeliness of each hypothesis to which it is relevant. For example, as shown in the left hand side of Fig. 6, the believability of the observations performed by a source (such as Curveball) depends on source’s *competence* and *credibility*. Moreover, competence depends on *access* and *understandability*. Credibility depends on *veracity*, *objectivity*, and *observational sensitivity under the conditions of observation*. Thus, the bias that would result from the *persistence of impressions based on discredited evidence* is countered in TIACRITIS with a rigorous, detailed and explicit believability analysis.

But there are additional biases in the evaluation of evidence that Heuer does not mention, particularly with respect to establishing the credentials of evidence: relevance, believability, and inferential force or weight. An analyst may

confuse the competence of a HUMINT source with his/her credibility. Or, the analyst may focus on the veracity of the source and ignore source's objectivity and observational sensitivity. Analysts may fail to recognize possible synergisms in convergent evidence, as happened in the 9/11/2001 disaster. Analysts may even overlook evidence having significant inferential force.




Fig. 6. Believability of testimonial and tangible evidence.

B. Biases in the Perception of Cause and Effect

As noted by Heuer, analysts seek explanations for the occurrence of events and phenomena. These explanations involve assessments of causes and effects. But biases arise when analysts assign causal relations to those that are actually accidental or random in nature. One related consequence is that analysts often overestimate their ability to predict future events from past events, because there is no causal association between them. One major reason for these biases is that analysts may not have the requisite level of understanding of the kinds and amount of information necessary to infer a dependable causal relationship.

According to Heuer, when feasible, the “increased use of scientific procedures in political, economic, and strategic research is much to be encouraged”, to counter these biases [2, p.128]. Because TIACRITIS makes all the judgments explicit, they can be examined by other analysts to determine whether they contain any mistakes or are incomplete. Because different people have different biases, comparing and debating analyses of the same hypothesis made by different analysts can also help identify individual biases. Finally, as a learning system, TIACRITIS can acquire correct reasoning patterns from expert analysts which can then be used to analyze similar hypotheses.

Now, here is something that can occur in any analysis concerning chains of reasoning. It is always possible that an analyst's judgment will be termed biased or fallacious, on structural grounds if it is observed that this analyst frequently leaves out important links in his/her chains of reasoning. This is actually a common occurrence since, in fact, there is no such thing as a uniquely correct or perfect argument. Someone can always find alternative arguments to the same hypothesis; what this says is that there may be entirely different inferential routes to the same hypothesis. Another possibility is that someone may find arguments based on the same evidence that lead to different hypotheses. This is precisely why there are trials at law; the prosecution and defense will find different arguments, and tell different stories, from the same body of evidence.

C. Biases in Estimating Probabilities

There are different views among probabilists on how to assess the force of evidence [1]. The view of probability that

Heuer assumes is the conventional view of probability which might be best called the Kolmogorov view of probability since the Russian mathematician was the first one to put this view of probability on an axiomatic basis [13, 14]. This is also the only view of probability considered by Heuer's sources of inspiration on biases: Daniel Kahneman, Amos Tversky, and their many colleagues in psychology [15, 16]. In his writings, Kolmogorov makes it abundantly clear that his axioms apply only to instances in which we can determine probabilities by counting. But Heuer also notes that intelligence analysis usually deals with one-of-a-kind situations for which there are never any statistics. In such cases, analysts resort to subjective or personal numerical probability expressions. He discusses several reasons why verbal assessments of probability are frequently criticized for their ambiguity and misunderstanding. In his discussion he recalls Sherman Kent's advice that verbal assessments should always be accompanied by numerical probabilities [17].

Since Heuer only considers numerical probabilities conforming to the Kolmogorov axioms, any biases associated with them (e.g., using the availability rule, the anchoring strategy, expressions of uncertainty, assessing the probability of a scenario) are either irrelevant or not directly applicable to a type of analysis that is based on different probability systems, such as the one performed with TIACRITIS, which is based on the Baconian and Fuzzy probability systems. Indeed, analysts using TIACRITIS never assess any numerical probabilities.

Heuer [2, p.122] mentions *coping with evidence of uncertain accuracy* as an origin of bias: “The human mind has difficulty coping with complicated probabilistic relationships, so people tend to employ simple rules of thumb that reduce the burden of processing such information. In processing information of uncertain accuracy or reliability, analysts tend to make a simple yes or no decision. If they reject the evidence, they tend to reject it fully, so it plays no further role in their mental calculations. If they accept the evidence, they tend to accept it wholly, ignoring the probabilistic nature of the accuracy or reliability judgment.” He then further notes [2, p.123]: “Analysts must consider many items of evidence with different degrees of accuracy and reliability that are related in complex ways with varying degrees of probability to several potential outcomes. Clearly, one cannot make neat mathematical calculations that take all of these probabilistic relationships into account. In making intuitive judgments, we unconsciously seek shortcuts for sorting through this maze, and these shortcuts involve some degree of ignoring the uncertainty inherent in less-than-perfectly-reliable information. There seems to be little an analyst can do about this, short of breaking the analytical problem down in a way that permits assigning probabilities to individual items of information, and then using a mathematical formula to integrate these separate probability judgments.”

First, as discussed in the previous section, concerning the believability of evidence, there is more than just its accuracy to consider. Second, as discussed above, Heuer only considers the *conventional view of probability* which, indeed, involves complex probability computations. With TIACRITIS, the analyst does precisely what Heuer imagined that could be done for countering this bias. It breaks a hypothesis into simpler hypotheses (see Fig.1), and assesses the simpler hypotheses

based on evidence (see Fig.5). Also, TIACRITIS allows the analyst to express probabilities in words rather than numbers, and to employ simple min/max strategies for assessing the probability of interim and final hypotheses that do not involve any full-scale and precise Bayesian or other methods that would require very large numbers of probability assessments.

There are many places to begin a defense of verbal or fuzzy probability statements. The most obvious one is law. All of the forensic standards of proof are given verbally: “beyond reasonable doubt”, “clear and convincing evidence”, “balance of probabilities”, “sufficient evidence”, and “probable cause”. Over the centuries attempts have been made to supply numerical probability values and ranges for each of these standards, but none of them have been successful. The reason, of course, is that every case is unique and rests upon many subjective and imprecise judgments. Wigmore [18] understood completely that the catenated inferences in his Wigmorean networks were probabilistic in nature. Each of the arrows in the chain of reasoning describe the force of one hypothesis on the next one, e.g., $E \rightarrow F$. Wigmore graded the force of such linkages verbally using such terms as “strong force”, “weak force”, “provisional force”, etc. Toulmin [19] also used fuzzy qualifiers in the probability statements of his system which grounds Rationale [20]. There are many other examples of situations in which it is difficult or impossible for people to find numerical equivalents for verbal probabilities they assess. Intelligence analysis so often supplies very good examples in spite of what Sherman Kent said some years ago.

We conclude this discussion by recalling what the well-known probabilist Professor Glenn Shafer said years ago [21]: *Probability is more about structuring arguments than it is about numbers. All probabilities rest upon arguments. If the arguments are faulty, the probabilities however determined, will make no sense.* In TIACRITIS, the structure of the bottom-up argument is given by the logical top-down decomposition, and the conclusions are hedged by employing rigorous Baconian operations with fuzzy qualifiers, leading to a defensible and persuasive argument.

D. Hindsight Biases in Evaluating Intelligence Reporting

As Heuer notes, analysts often overestimate the accuracy of their past judgments; customers often underestimate how much they have learned from an intelligence report; and persons who conduct post-mortem analysis of an intelligence failure will judge that events were more readily foreseeable than was in fact the case. “The analyst, consumer, and overseer evaluating analytical performance all have one thing in common. They are exercising hindsight. They take their current state of knowledge and compare it with what they or others did or could or should have known before the current knowledge was received. This is in sharp contrast with intelligence estimation, which is an exercise in foresight, and it is the difference between these two modes of thought—hindsight and foresight—that seems to be a source of bias. ... After a view has been restructured to assimilate the new information, there is virtually no way to accurately reconstruct the pre-existing mental set.” [2, p.162]

Apparently Heuer did not envision the use of a system like TIACRITIS that keeps track of the performed analysis, what evidence we had, what assumptions we made and what were

their justifications, and what was the actual logic of our analytic conclusion. We can now add additional evidence and use our hindsight knowledge to restructure the argumentation and re-evaluate our hypotheses, and we can compare the hindsight analysis with the foresight one. But we will not confuse them. As indicated by Heuer [2, pp.166-167]: “A fundamental question posed in any postmortem investigation of intelligence failure is this: Given the information that was available at the time, should analysts have been able to foresee what was going to happen? Unbiased evaluation of intelligence performance depends upon the ability to provide an unbiased answer to this question.” We suggest that this may be accomplished with a system like TIACRITIS.

IV. SOME FREQUENTLY OVERLOOKED ORIGINS OF BIAS

So much of the discussion of bias in intelligence analysis is directed at intelligence analysts themselves. But we have identified three other origins of bias that are rarely discussed, even though they may be at least as important on occasion as any analysts’ alleged biases. The three other origins of bias we will consider are: (1) persons who provide testimonial evidence about events of interest (i.e. HUMINT sources); (2) other intelligence professionals having varying capabilities who serve as links in what we term “chains of custody” linking the evidence itself, as well as its sources, with the users of evidence (i.e. the analysts); and (3) the “consumers” of intelligence analyses (government and military officials who make policy and decisions regarding national security).

A. HUMINT Sources

Our concern here is with persons who supply us with testimonial evidence consisting of reports of events about matters of interest to us. Heuer [2, p.122] does mention the “bias on the part of the ultimate source,” but he does not analyze it. In our work on evidence in a variety of contexts, we have always been concerned about establishing the believability of its sources, particularly when they are human witnesses, sources, or informants [1]. In doing so, we have made use of the 600 year-old legacy of experience and scholarship in the Anglo-American adversarial trial system concerning witness believability assessments. We have identified the three major attributes of the credibility of ordinary witnesses: *veracity*, *objectivity*, and *observational sensitivity* (see Fig. 6). We will show how there are distinct and important possible biases associated with each such believability attribute. These biases are recognized in the MACE system (Method for Assessing the Credibility of Evidence), developed for the IC [22]. This system incorporates both Baconian and Bayesian methods for combining evidence about our source.

As discussed above, assessing the credibility of a human source *S* involves assessing *S*’s veracity, objectivity, and observational sensitivity. We have to consider that source *S* can be biased concerning any of these attributes. On *veracity*, *S* might prefer to tell us that event *E* occurred, whether *S* believed *E* occurred or not. As an example, an analyst evaluating *S*’s evidence *E** might have evidence about *S* suggesting that *S* would tell us that *E* occurred because *S* wishes to be the bearer of what *S* believes we will regard as good news that event *E* occurred. On *objectivity*, *S* might choose to believe that *E*

occurred because it would somehow be in S's best interests if E did occur. On *observational sensitivity*, there are various ways that S's senses could be biased in favor of recording event E; clever forms of deception supply examples.

These three species of bias possible for HUMINT sources must be considered by analysts attempting to assess the credibility of source S and how much weight or force S's evidence E* should have in the analyst's inference about whether or not event E did happen. The existence of any of these three biases would have an effect on an analyst's assessment of the weight or force of S's report E*. As we know, all assessments of the credibility of evidence rest upon available evidence about its sources. In the case of HUMINT we need ancillary evidence about the veracity, objectivity, and observational sensitivity of its sources. In the process, we have to see whether any such evidence reveals any of the three biases just considered. TIACRITIS supports the analyst in this determination by guiding her to answer specific questions based on ancillary evidence. For instance, the veracity questions considered are shown in Table 1.

Table 1. Questions concerning the veracity of human sources.

1. <i>Goals of this source?</i> Does what this source tells us support any of his or her goals?
2. <i>Present influences on this source?</i> Could this source have been influenced in any way to provide us with this report?
3. <i>Exploitation potential?</i> Is this source subject to any significant exploitation by other persons or organizations to provide us this information?
4. <i>Any contradictory or divergent evidence?</i> Is there any evidence that contradicts or conflicts with what the source has reported to us?
5. <i>Any corroborative or confirming evidence?</i> Is there any other evidence that corroborates or confirms this source's report?
6. <i>Veracity concerning collateral details?</i> Are there any contradictions or conflicts in the collateral details provided by this source that reflect the possibility of this source's dishonesty?
7. <i>Source's character?</i> What evidence do we have about this source's character and honesty that bears upon this source's veracity?
8. <i>Reporting record?</i> What does the record show about the truthfulness of this source's previous reports to us?
9. <i>Source expectations about us?</i> Is there any evidence that this source may be reporting events he/she believes we will wish to hear or see?
10. <i>Interview behavior?</i> If this source reported these events to us, what was this source's demeanor and bearing while giving us this report?

B. Persons in Chains of Custody of Evidence

Unfortunately, there are other persons, apart from HUMINT sources, whose possible biases need to be carefully considered. We know that analysts make use of an enormous variety of evidence that is not testimonial or HUMINT, but is *tangible* in nature. Examples include objects, images, sensor records of various sorts, documents, maps, diagrams, charts, and tabled information of various kinds.

But the intelligence analysts only rarely have immediate and first access to HUMINT assets or informants. They may only rarely be the first ones to encounter an item of tangible evidence. What happens is that there are several persons who have access to evidence between the times the evidence is first acquired and when the analysts first receive it. These persons may do a variety of different things to the initial evidence during the time they have access to it. In law, these persons constitute what is termed a "*chain of custody*" for evidence.

Heuer [2, p.122] mentions the "distortion in the reporting chain from subsource through source, case officer, reports officer, to analyst" but he does not analyze it. In criminal cases in law, there are persons identified as "evidence custodians", who keep careful track of who discovered an item of evidence, who then had access to it and for how long, and what if anything they did to the evidence when they had access to it.

These chains of custody add three major additional sources of uncertainty for intelligence analysts to consider, that are associated with the persons in chains of custody whose competence and credibility need to be considered. The first and most important question involves *authenticity*: *Is the evidence received by an analyst exactly what the initial evidence said and is it complete?* The other questions involve assessing the *reliability* and *accuracy* of the processes used to produce the evidence if it is tangible in nature (see the right side of Fig. 6), or also used to take various actions on the evidence in a chain of custody, whether the evidence is tangible or testimonial. As an illustration, consider an item of testimonial HUMINT coming from a foreign national whose code name is "Wallflower", who does not speak English [23]. Wallflower gives his report to *case officer* Bob. This report is *recorded* by Bob and then *translated* by Husam. Then, Wallflower's translated report is *transmitted* to a *report's officer* Marsha who *edits* it and *transmits* it to the analyst Clyde who evaluates it and assesses its weight or force.

Now, here is where forms of bias can enter that can be associated with the persons involved in these chains of custody. The case officer Bob might have intentionally overlooked details in his recording of Wallflower's report. The translator Husam may have intentionally altered or deleted parts of this report. The report's officer Marsha might have altered or deleted parts of the translated report of Wallflower's testimony in her editing of it. The result of these actions is that the analyst Clyde receiving this evidence almost certainly did not receive an authentic and complete account of it, nor did he receive a good account of its reliability and accuracy. What he received was the transmitted, edited, translated, recorded testimony of Wallflower. Fig. 7 shows how TIACRITIS may determine the believability of the evidence received by the analyst. Although the information to make such an analysis may not be available, the analyst should adjust the confidence in his conclusion, in recognition of these biases.

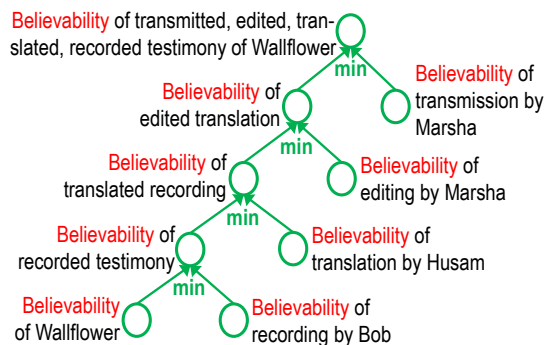


Fig. 7. Chain of custody of Wallflower's testimony.

C. Consumers of Intelligence Analyses

The policy-making consumers or customers of intelligence

analysts are also subject to a variety of inferential and decisional biases that may influence the reported analytic conclusions. As is well known, the relationships between intelligence analysts and governmental policy makers are much discussed and involve considerable controversy [24, 25]. On the one hand we hear intelligence professionals say that they do not make policies but only try to help policy makers be as informed as they can be when they do form policies and make decisions in the nation's best interests. But we also learn facts about the intelligence process that complicate matters. An intelligence analysis is usually a hierarchical process involving many intelligence officers, at various grade levels, who become involved in producing an intelligence "product". At the most basic level of this hierarchy are the so-called "desk analysts" who are known and respected experts in the specific subject matter of the analysis at hand. An analysis produced by one or more desk analysts is then passed "upward" through many administrative levels, at each of which persons at these higher levels can comment on the desk analysts' report. It is often recognized that the higher an editor is in this hierarchy, the more political his/her views and actions become that may affect the content and conclusions of the analysis at hand. As this "upward" process continues, the analysis that results may be quite different from the one produced by the desk analysts, reflecting the biases of those who have successively edited it. In some cases, these editing biases are the direct result of the consumer's biases who may wish to receive a certain analytic conclusion. Using a system like TIACRITIS that shows very clearly how the analytic conclusion is rooted in evidence would significantly help in reducing the above biases.

V. CONCLUSIONS

A wide variety of biases affect the correctness of intelligence analyses. In this paper we have shown how the use of TIACRITIS, a knowledge-based cognitive assistant, helps analysts recognize and counter many of them. TIACRITIS integrates several semantic technologies (knowledge representation through ontologies and rules, evidence-based reasoning, machine learning and knowledge acquisition). It can run in a browser as a web-based system, or it can be installed locally, and has been used in many civilian, military, and intelligence organizations.

There are two complementary ways by which TIACRITIS helps mitigate biases. First, as a cognitive assistant, it helps automate many parts of the analysis process, making this task much easier for the analyst. Thus it alleviates one of the main causes of biases, which is the employment of simplified information processing strategies on the part of the analyst. Second, TIACRITIS performs a rigorous evidence-based hypothesis analysis that makes explicit all the reasoning steps, evidence, probabilistic assessments, and assumptions, so that they can be critically analyzed and debated. Indeed, the best protection against biases comes from the collaborative effort of teams of analysts, who become skilled in solving their analytic tasks through the development of sound evidence-based arguments, and who are willing to share their insights with colleagues, who are also willing to listen. TIACRITIS makes all this possible.

Finally, this paper adds a strong argument in favor of using

structured analytic methods, in the debate on how to significantly improve intelligence analysis [26].

REFERENCES

- [1] Schum D.A. (2001). *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press.
- [2] Heuer R.J. (1999). *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington, DC.
- [3] Tecuci, G., Marcu, D., Boicu, M., Schum, D.A., Russell K. (2011). Computational Theory and Cognitive Assistant for Intelligence Analysis, in *Proc. 6th Int. Conf. on Semantic Technologies for Intelligence, Defense, and Security*, pp. 68-75, Fairfax, VA, 16-18 November.
- [4] Cohen L.J. (1977). *The Probable and the Provable*, Clarendon Press, Oxford.
- [5] Zadeh L. (1983). The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems, *Fuzzy Sets and Systems*, vol.11, pp.199-227.
- [6] Tecuci, G. (1998). *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, San Diego: Academic Press.
- [7] Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M. (2005). The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence*, Vol.21, No.4, pp. 462-479.
- [8] Tecuci, G., Boicu, M., Marcu, D., Schum, D. (2013). How Learning Enables Intelligence Analysts to Rapidly Develop Practical Cognitive Assistants, in *Proc. 12th International Conference on Machine Learning and Applications (ICMLA'13)*, Miami, Florida, December 4-7.
- [9] W3C (2004). <http://www.w3.org/TR/rdf-schema/>, accessed 10/11/13.
- [10] Boicu M., Tecuci G., Schum D. (2008). Intelligence Analysis Ontology for Cognitive Assistants, in *Proc. of Conf. "Ontology for the Intelligence Community"*, Fairfax, VA, 3-4 December.
- [11] Heuer R.J. (2008). Computer-Aided Analysis of Competing Hypotheses, in George R.Z., Bruce J.B., eds., *Analyzing Intelligence: Origins, Obstacles, and Innovations*, Georgetown Univ. Press, Washington, DC.
- [12] Drogin B. (2007). *CURVEBALL: Spies, Lies, and the Con Man Who Caused a War*. Random House, New York, NY.
- [13] Kolmogorov A.N. (1933). *Foundations of a Theory of Probability* (1933), 2nd English edition, Chelsea Publishing, New York, NY., 1956
- [14] Kolmogorov A.N. (1969). The Theory of Probability. In: Aleksandrov, A. D., Kolmogorov, A. N., Lavrentiev, M. A. (eds) *Mathematics: Its Content, Methods, and Meaning*. MIT Press, Cambridge, MA.
- [15] Kahneman, D., Tversky, A. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, Vol, 185, 1124-1131.
- [16] Kahneman, D., Slovic, P., Tversky, A. (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- [17] Kent S. (1994). Words of Estimated Probability, in Steury D.P., ed., *Sherman Kent and the Board of National Estimates: Collected Essays*, Center for the Study of Intelligence, CIA, Washington, DC.
- [18] Wigmore J.H. (1937). *The Science of Judicial Proof*. Boston, MA: Little, Brown & Co.
- [19] Toulmin S.E. (1963). *The Uses of Argument*. Cambridge Univ. Press.
- [20] van Gelder T.J. (2007). The Rationale for Rationale, *Law, Probability and Risk*, 6, pp. 23-42.
- [21] Shafer G. (1988). Combining AI and OR. *University of Kansas School of Business Working Paper No. 195*.
- [22] Schum D.A., and Morris J. (2007). Assessing the Competence and Credibility of Human Sources of Evidence: Contributions from Law and Probability, *Law, Probability and Risk*, Vol. 6, pp. 247-274.
- [23] Schum D.A., Tecuci G, and Boicu M. (2009). Analyzing Evidence and its Chain of Custody: A Mixed-Initiative Computational Approach. *Int. Journal of Intelligence and Counterintelligence*, Vol. 22, pp. 298-319.
- [24] George, R., Bruce, J. (eds) (2008). *Analyzing Intelligence: Origins, Obstacles, and Innovations*. Georgetown Univ. Press, Washington, DC.
- [25] Johnston, R. (2005). *Analytic Culture in the U.S. Intelligence Community*. Central Intelligence Agency, Washington, DC.
- [26] Marrin, S. (2011). *Improving Intelligence Analysis: Bridging the gap between scholarship and practice*. Routledge, London and New York.