Credibility Models for Multi-Source Fusion

Edward J. Wright
Information Extraction and Transport, Inc.
1911 N Ft Myer Dr, Suite 600
Arlington, VA, 22209 U.S.A.
ewright@iet.com

Kathryn Blackmond Laskey
Department of Systems Engineering and
Operations Research,
George Mason University
Fairfax, VA 22030 U.S.A.
klaskey@gmu.edu

Abstract *– This paper presents a technical approach for fusing information from diverse sources. Fusion requires appropriate weighting of information based on the quality of the source of the information. A credibility model characterizes the quality of information based on the source and the circumstances under which the information is collected. In many cases credibility is uncertain, so inference is necessary. Explicit probabilistic credibility models provide a computational model of the quality of the information that allows use of prior information, evidence when available, and opportunities for learning from data. This paper provides an overview of the challenges, describes the advanced probabilistic reasoning tools used to implement credibility models, and provides an example of the use of credibility models in a multi-source fusion process.*

Keywords: fusion, Bayesian networks, pedigree, credibility models, Multi-Entity Bayesian Networks.

1 Introduction

In traditional multi-source fusion applications, fusion is performed using data from sensors that are reasonably well understood, and that provide error models where the parameters of the error models are known (e.g., [1]). In today's operational asymmetric warfare environment, there are requirements to integrate / fuse information from a much broader set of sources including HUMINT, open source (web pages, news reports), and communications intercepts. [2]. Multi-source fusion in this expanded problem space requires the ability to integrate information even when the sources are not well characterized.

Ceruti, et al. [3] identify both the potential benefits and the challenges of incorporating more complex pedigree information into the fusion process. They recommend that all source information be included as part of the pedigree, and that a computational model of information quality be provided. The idea is that the quality model can be used to integrate information from diverse sources, characterize the quality of the result, and to automatically update the results when new information becomes available.

The purpose of this paper is to demonstrate that probabilistic models based on Bayesian Networks can be used to implement a computational model that characterizes the quality of information. Such a computational model can be used to: (*i*) incorporate credibility information in the fusion process and characterize the quality of the results; (*ii*) deal with missing or uncertain credibility information; and (*iii*) update results of previous inference to make use of new information on source credibility when it becomes available.

This paper is organized as follows. Section 2 describes fusion challenges in a heterogeneous, net-centric, distributed environment. Section 3 describes the scenario we use to illustrate our approach. Section 4 provides an overview of Bayesian Networks and Multi-Entity Bayesian Networks. Section 5 describes how MEBNs can be used to create credibility models for some non-traditional information sources. Section 6 illustrates the use of probabilistic credibility models to extend the scenario from section 3. Section 7 describes some further applications of the technology. Finally, Section 8 presents our conclusions.

2 Fusion Challenges

Fusion in an asymmetrical warfare environment, especially higher level fusion (levels 2 & 3), requires the ability to integrate information and draw inferences in a complex problem space. Inferences often involve multiple level of abstraction, and draw on diverse information from sources that are not well known. Sources may include HUMINT gathered from agents or collected by interviewing civilians on the street; news or magazine articles; web sites that draw information from a wide variety of sources, etc. Reasoning about the quality of information requires reasoning about the competence and veracity of the sources, as well as reasoning about the capability or opportunity of the source to observe what was reported.

This kind of reasoning about credibility is similar in many respects to reasoning about legal evidence. Levitt and Laskey [4] describe computational techniques for evidential reasoning in legal settings, and illustrate their approach with a model for a French murder case. The methodological approach taken in this paper is similar to that of [4].

To be practical, an automated fusion system must be flexible enough to respond to a wide variety of specific situations, which cannot be defined in advance. This requires that a practical system be capable of using

modular, reusable, components that can be assembled when necessary to build complex models needed to reason about specific situations.

3 Scenario

Laskey and Levitt [5] described a scenario involving a coordinated Biological attack by a terrorist organization on the US. They developed a probabilistic model to perform multi-source fusion of diverse evidence to infer the existence and type of biological attack. We employ this scenario to illustrate our technology for representing and reasoning with source credibility.

The scenario unfolds as follows.

Day 1: Infiltrated stockyard operatives in Chicago infect cattle herds at target stockyards with cutaneous anthrax by sprinkling several grams of it in the cattle feed.

Day 3: Same operatives infect herds with foot-and-mouth disease by direct application of pus onto multiple cattle.

Day 5: First reports of anthrax and foot-and-mouth symptoms in herds occur. Confusion of symptoms delays cause identification. At end of shift, operatives spray multiple grams of inhalation anthrax at herd with hand held spray device. Infiltrated stockyard operatives in Kansas City infect cattle herds at target stockyards with cutaneous anthrax by sprinkling several grams of it in the cattle feed.

Day 7: Crop duster sprays Chicago with 50kg anthrax aerosol. Kansas City stockyard infected with foot-and-mouth.

Day 8: Cutaneous anthrax confirmed in Chicago stockyard.

Day 9: Kansas City Stockyard sprayed with inhalation anthrax. Denver stockyard infected with cutaneous anthrax.

Day 11: Crop duster sprays Kansas City with 50kg anthrax aerosol.

Day 12: Inhalation anthrax detected at Chicago stockyard. Foot-and-mouth confirmed at Chicago stockyard. Cutaneous anthrax detected in Kansas City stockyard.

Day 13: Unlikeliness of multiple outbreaks in disparate areas triggers concern about possible multi-city biowarfare attack. Analysis of anthrax in Kansas City and Chicago shows weapons grade inhalation anthrax. Alerts are issued to all cities with major cattle stockyards; local law enforcement engaged for extreme surveillance. Crop dusting alert nationwide.

Day 14: Crop dusting alert finds suspicious operatives planning run in Denver. Dallas/Fort-Worth operation subsequently found and shut down.

For this paper we extend the scenario to incorporate information from an unknown HUMINT source. The problems to be addressed are: (*i*) How to exploit pedigree information and characterize the results of the fusion; (*ii*); How to deal with missing and uncertain pedigree information; and (*iii*) How to update the inference as new information becomes available.

4 Bayesian Networks and MEBNs

The acknowledged standard for logically coherent reasoning under uncertainty is Bayesian probability theory. Bayesian theory provides a principled representation for degrees of plausibility, a logically justified calculus for combining prior knowledge with observations, and a learning theory for refining degrees of plausibility as evidence accrues. Bayesian reasoning has become quite popular since the advent of *Bayesian networks*, a graphical paradigm for representing and computing with large numbers of interrelated uncertain hypotheses [6,7].

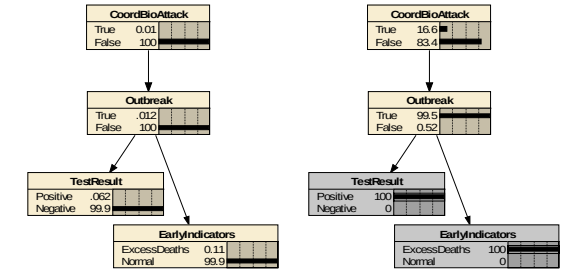

a. Prior Probabilities *b. Posterior Probabilities*

Figure 1. Bayesian Network

A Bayesian network is both a formal language for representing probabilistic knowledge and a computational architecture for drawing inferences about some hypotheses based on information about others. Figure 1 shows a Bayesian network for a highly simplified bioterrorism example. Four hypotheses are shown: whether or not a biological attack is occurring, whether or not there is an outbreak (natural or terrorist-initiated) of anthrax; whether there have been excess deaths in humans and/or livestock, and whether specimens test positive for anthrax. Arcs in the graph show dependence relationships. Dependence relationships may represent cause and effect relationships, statistical association, or other relationships that create an evidential association. In this example, a biological attack causes a disease outbreak, which in turn causes both excess deaths and positive test results for anthrax. The occurrence of excess deaths and the positive test for anthrax increase the probability of a disease outbreak from a hundredth of a percent to over 99%. The probability of a biological attack has increased dramatically from a hundredth of a percent to about 17%, but in the absence of other evidence for a biological attack, a natural outbreak is the most likely explanation for the evidence.

Bayesian networks use a simple attribute-value representation. That is, each problem instance involves reasoning about a fixed set of hypotheses, each of which can take on a fixed set of possible values. Only the observations change from problem instance to problem instance. This representation is not expressive enough for

our application. In the above scenario, the attack targets multiple cities and involves several biological agents. Reasoning about multi-city, multi-agent problems requires replicating portions of the Bayesian network of Figure 1. More expressive knowledge representation formalisms are needed to handle this kind of repeated structure.

Figure 2 illustrates how *multi-entity Bayesian networks* can be applied to extend our example to multiple cities and agents [8]. A multi-entity Bayesian network encodes domain knowledge as parameterized argument structures called *MEBN Fragments (MFrags)*. An MFrag is a modular component representing a fairly small, separable, and conceptually meaningful part of the total argument structure supporting or denying a given hypothesis. MFrags can represent alternative hypothetical world states, evidence that bears upon which hypotheses are true, and chains of argument relating evidence to hypotheses. MFrags can be combined to build models relating complex configurations of many features, can be repeatedly instantiated to represent multiple related entities of a given type (such as multiple biological agents, attack locations, or information sources), and can be re-used across multiple scenarios. Multi-entity Bayesian networks have sufficient expressive power to represent a logically coherent probability assignment to any collection of hypotheses that can be expressed in the language of first-order logic. Thus, MEBN provides a synthesis of classical logic and Bayesian probability theory.

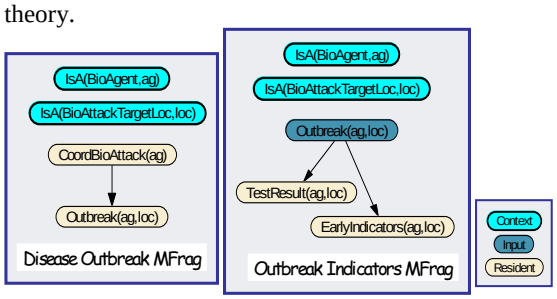

Figure 2. MEBN Fragments (MFrags)

To reason about a particular scenario, *instances* are created of the MFrags of Figure 2 by substituting relevant entities for the variables. For example, our scenario involves the hypothesis of a coordinated anthrax attack launched against Chicago and Kansas City. To reason about this scenario, the agent *Anthrax* is substituted for the placeholder variable *ag*, and the cities *Chicago* and *KC* are substituted for the placeholder variable *loc* in the MFrags of Figure 2. This results in a single instance of *CoordBioAttack*(*ag*) and two instances of each of the random variables *Outbr*(*ag*, *loc*), *EarlyIndicators*(*ag*, *loc*), and *TestResult*(*ag*, *loc*). These instances are combined into the *situation-specific Bayesian network* shown in Figure 3.

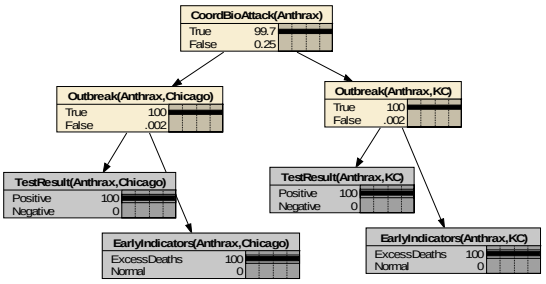

Figure 3. Situation-Specific Bayesian Network

An implementation of MEBN logic is available in IET's Quiddity*Suite, a knowledge-based probabilistic reasoning, tool-building toolkit [9].

5 Credibility models

Characterizing the quality of information based on its source and the conditions under which it was collected requires building a credibility model. This is particularly challenging when the fusion is directed at inference of higher level cognitive states such as intent and goals of an entity. Difficult challenges in Multi-INT fusion arise from the need to incorporate unstructured text information from HUMINT, COMINT, or from open sources. Information from these sources is characterized by uncertainty in the credibility of the source. This will be addressed by automatically generating an explicit credibility model for each source.

The primary evidence for fusion of cognitive states will come from a variety of sources that generate information in plain language. Examples are HUMINT, COMINT intercepts, and open source. Fusing information from plain language sources is complex because the sources may use language in a way that is different from the receiver's usage, rendering the meaning unclear. Even if Blue forces are using a common syntax to facilitate automated processing of textual information, automated fusion of textual inputs will not be possible unless the system can accommodate the different possible meanings of the potential sources.

There is also a challenge with credibility of plain text sources. Human sensors cannot be characterized in the same way as physical sensors, so it is difficult to evaluate the quality of the information provided. The credibility issue extends well beyond the traditional concerns of sensor accuracy, and must include multiple attributes of the source's credibility and competence. The attributes of credibility are veracity – whether or not the source is telling the truth, objectivity – whether or not the report is based on received sensory evidence or on prior expectations of beliefs, and observational sensitivity or accuracy – an assessment of the quality of the sensory evidence. Credibility refers to the extent that a source is believable about the event(s) reported, while competence refers to the person's capability to understand what they observed. For example a source who cannot distinguish between the types of enemy vehicles is not competent to report the presence of a particular type of vehicle.

The objective is to make an assessment of the credibility of the source of the HUMINT, COMINT or

open source report, resulting in a probability that the report is true. These all come from a human source.

6 Fusion

This section provides an example of the use of credibility models in a multi-source fusion scenario. The example extends the coordinated biological attack scenario from section 3 to add a HUMINT source.

The extended scenario begins at Day 8 of the original scenario, after cutaneous anthrax has been confirmed at the Chicago stockyard, but prior to any confirmation of inhalation anthrax. For simplicity, we use the simple model of Figure 2. Our SSBN will be like Figure 3 but with only Chicago. The probability of a coordinated biological attack, as shown in Figure 1, is about 17%. Events in the enhanced scenario unfold from this point as follows:

E1. Agent X, an unknown source, reports that the terrorist organization is carrying out a coordinated biological attack.
E2. Agent X reports that his knowledge of the attack comes from a meeting he attended, where details of the attack were discussed.
E3. Agent Y, a trusted known source, reports that Agent X has successfully infiltrated the terrorist organization.
E4. A SIGINT report positively identifies agent X at a time and place where he could not have attended the attack planning meeting.

Figure 4. Evidence for a Coordinated Biological attack, from Figure 2, abstracted to a single evidence node.

The following example demonstrates how the evidence from the above enhanced scenario can be included in the multi-source fusion process.

Figure 4 shows the result from the original scenario. For simplicity, the previous evidence has been abstracted into one evidence node, which when true, results in a posterior probability of 17% for the hypothesis *Coord Bio Attack.*

Figure 5 shows the result of scenario step E1, which includes HUMINT from Agent X. The source generates a report which is evidence for the top level hypothesis. The strength of the report is moderated by a credibility model for Agent X. For traditional sensors, a credibility model would be an error model that defines the detection and false alarm probabilities for a sensor. In a Bayesian network, this error model might be encoded directly into the local probability distribution of the evidence node or may be represented explicitly by an error node as an additional parent to the evidence node. For a nontraditional source such as HUMINT, the credibility model is more complex and may require multiple nodes in the Bayesian network. In a MEBN representation, we would include one or more credibility MFrags, which would depend on characteristics of the source and the event being reported on.

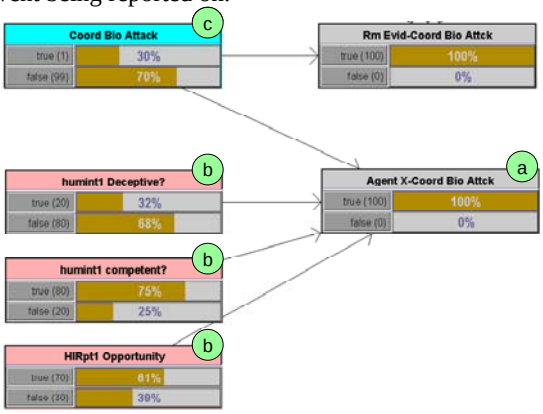

a) HUMINT report; b) Credibility nodes for this report; c) P(coord Bio Attack) has increased to 30%.

Figure 5. HUMINT report from Agent X

In this simplified credibility model, one node represents source competence and one node represents deceptiveness of the source. An additional node represents the opportunity the source had to obtain the reported information.

Figure 5 shows the default prior information for the credibility model for Agent X.

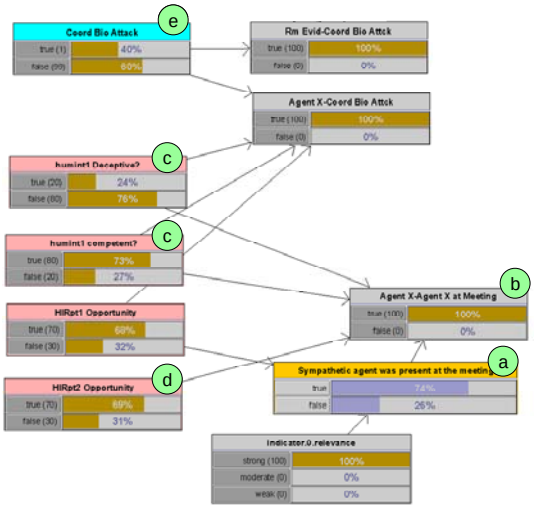

a) New hypotheses that Agent X was present at the meeting; b) Agent X's report that he has at the meeting; c) This report shares the competence and deception nodes with the original report; d) Report has a new opportunity node; e) P(Coord Bio Attck) has increased to 40%.

Figure 6. Additional nodes to reason about Agent X's opportunity to make the observation.

Additional information, if it is available can be added to reason about any of the nodes in the credibility model. For example, Agent X may report that his information about the coordinated biological attack resulted from his attendance at a meeting where the details of the attack were discussed. This information is represented in the model as an additional hypothesis about Agent X's attendance at the meeting, and his report becomes

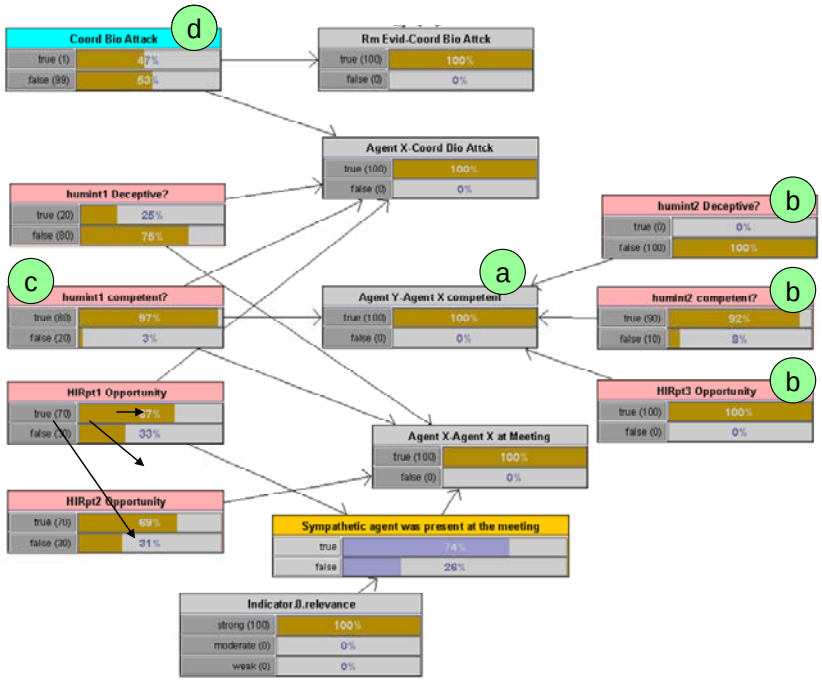

a) Evidence provided by Agent Y; b) Agent Y, a trusted friendly agent, is extremely credible; c) belief in the competence of agent X has increased from 73% to 97%; d) P(Coord Bio Attack) has increased to 47%.

Figure 7. The network after Agent Y verifies Agent X's competence.

evidence for this hypothesis. In Figure 6, this information is added to the network as an additional report from Agent X. Because this report is from the same source, it shares the credibility nodes for competence and deception. Because it is a separate report, it has a new source opportunity node.

An advantage of the MEBN representation is that it automatically keeps track of which nodes must be replicated for a new report and which are shared with other reports The node *Competence*(*src*) has only one argument, the source, whereas the node *Opportunity*(*src*, *evnt*) has two arguments, the source and the event reported upon. Because competence is a property of a source, the situation-specific Bayesian network contains only one instance of the competence node per source. Because opportunity depends on both the source and the event, the situation-specific Bayesian network contains a an instance of the opportunity node for each event / source combination for which we have a report.

In scenario step 3, Agent Y, a trusted friendly agent, provides information that Agent X has in fact infiltrated the terrorist organization. This information can be applied as evidence for the competence node of the Agent X's credibility model. The result is shown in Figure 7.

At this point the inference for P(Coord Bio Attack)= 47%, which is significantly higher than the probability without the HUMINT report.

In the final step of the scenario, a SIGINT report (with known high credibility) establishes that Agent X was at a different location at the time of the meeting, and could not have been present. Because the SIGINT is a source with high credibility, this evidence overwhelms the evidence from Agent X. This causes a change in the credibility model for Agent X. Belief that X is deceptive has increased. This weakens the force of his evidence about the Coordinated Biological Attack. The final inference is shown in Figure 8.

At each step of this fusion process, if it was necessary to make the inference results available to other analysts or decision makers, the results should be packaged with metadata that characterizes the quality of the inference. One opportunity is to include the current inference Bayesian network (or a link to it), as part of the metadata. Then at any time in the future if additional evidence becomes available – either direct evidence for the coordinated attack hypotheses or evidence about the credibility of any of the sources, then the Bayesian network can be used as a computational model to update the inference to include the new information.

This example has shown:

- The ability to build a credibility model for a HUMINT source which allows information about the source credibility to be factored into the fusion results. This simplified credibility model can easily be extended to more realistic complexity.
- The ability to reason about information from sources with incomplete credibility model.
- The ability to include credibility information in the fusion results. Use of credibility information causes belief changes in intuitive ways.
- The potential to use the BN (or a link to it) as a part of the metadata that provides a computational model that characterizes the quality of the fused results.

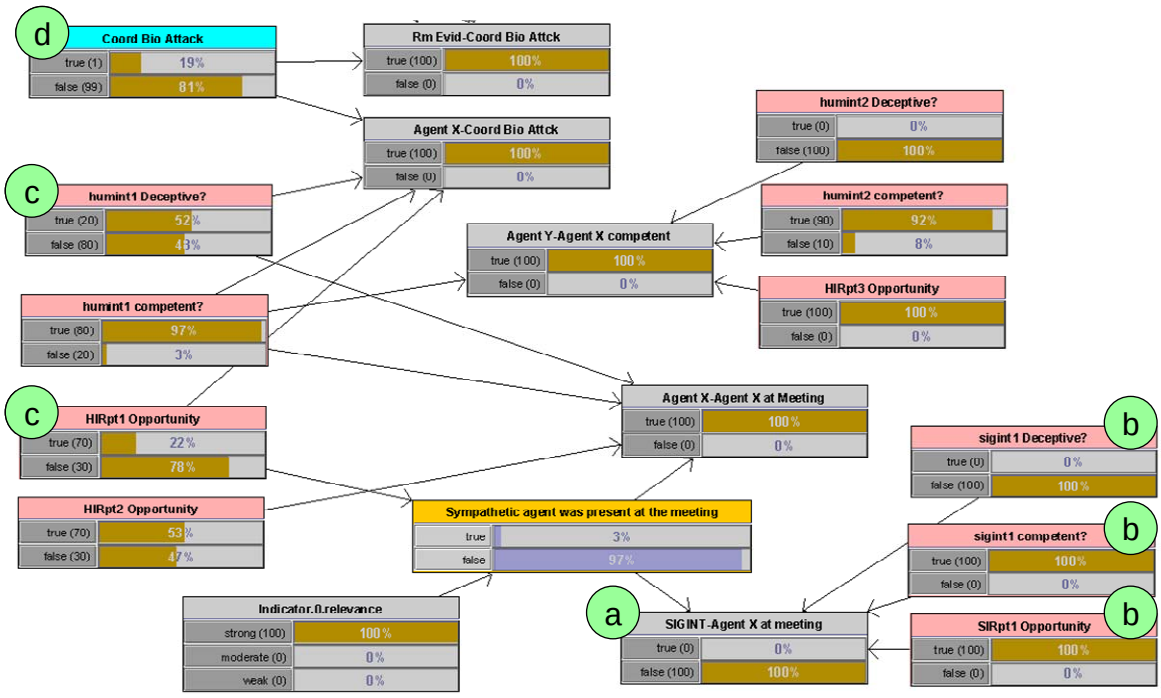

a) SIGINT report has provided evidence the Agent A was somewhere else – and so could not have been at the meeting; b) SIGINT report is extremely credible; c) Credibility model for Agent X is updated -- belief that he had the opportunity to make the original report decreases; belief that he is deceptive increases;, d) Agent X's original report is discounted, decreasing P(Coord Bio Attck) to 19%.

Figure 8. The network after adding evidence for deception

Although MFrags were not shown for the credibility model, the hypotheses, sources, and reports are easily represented by MEBN fragments – providing a modular, reusable representation that is easy to maintain, and can be used by automated systems to build situation specific Bayesian Networks.

This example does have some limitations. It is a simplified model, with realistic complexities omitted. An operational credibility model may need to include additional variables to reason about the sources competence for different tasks. For example, a source who is not competent to observe and report on types of military vehicles, may be competent to report on biological agents. In addition it may be necessary to include the source's motivation, as this could influence the conditions under which the source is deceptive. It is possible to include these additional influences in a credibility model, but the result is beyond the scope of this paper.

The parameters used in the local probability distributions for the human credibility models were generated as qualitative assessments based on common sense. The results of inference using these qualitative parameters are consistent with intuition, demonstrating the potential for generating useful results even when knowledge elicitation from an expert or learning from data are not possible. An operational application should provide a way for analysts to review, update and document parameters used by the inference.

7 Applications

The ideas illustrated by the example can be applied in a number of ways to enhance multi-sensor fusion. Bayesian Hierarchical Inference, implemented with Bayesian Networks and MEBNs, provides a common representation of uncertainty that is flexible enough to reason about the complex interactions between factors that influence credibility models for human sources. Use of appropriate credibility models provides a scientific methodology for including HUMINT and other open source information in the fusion process.

The Bayesian Networks also has potential to provide a compact representation of sophisticated metadata to characterize the quality of a fused result. The Bayesian Network provides documentation of the information and assessments used to generate the results, but also provides the computational model of the quality of the result which can be used to automatically propagate changes and updates to derived products when source information changes.

8 Conclusions

Incorporating information from diverse sources is a complex fusion challenge. Using information from human sources is complicated by the need to reason about the credibility of the source, taking into account the

sources competence, veracity, and opportunity to observe the activity reported. Credibility models that reason about these factors can be implemented as probabilistic models using Bayesian Networks and MEBNs. These probabilistic credibility models provide a methodology for integrating information from human sources with information from traditional sensors in Multi-source fusion.

References

1. Hall, D.L., *Handbook of Multisensor Data Fusion*, Artech House, 2001.
2. Cardillo, R., NGA Challenges, briefing presented to the DARPA-NGA Partnership Industry Workshop, 7 Sep 2005, http://dtsn.darpa.mil/ixo/DARPA_NGA/index.html
3. Ceruti, M.G., Ashenfelter, A., Brooks, R. Chen, G., Das, S. Raven, M.S., Wright, E., Pedigree Information for Enhanced Situation and Threat Assessment, 9th International Conference on Information Fusion, Florence, Italy, 10-13 July, 2006
4. Levitt, T. and Laskey, K.B. Computational Inference for Evidential Reasoning in Support of Judicial Proof, *Cardozo Law Review*, 2000.
5. Laskey, K.B. and Levitt, T. S., Multisource fusion for opportunistic detection and probabilistic assessment of homeland terrorist threats. *Proc. SPIE Vol. 4708, p. 80-89.*
6. Pearl, J. *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, 1988.
7. Neapolitan, R. *Learning Bayesian Networks*, Prentice Hall, 2003.
8. Laskey, K.B. *First-Order Bayesian Logic*, Fairfax, VA: Department of Systems Engineering and Operations Research, George Mason University, 2005. http://ite.gmu.edu/~klaskey/papers/Laskey_MEBN_Logic.pdf.
9. Fung, F., Laskey, K. B., Pool, M., Takikawa, M., &Wright, E. J. (2004). PLASMA: combining predicate logic and probability for information fusion and decision support. Paper presented at the AAAI Spring Symposium, Stanford, CA.