Ontology of Evidence

Kathryn B. Laskey, David A. Schum, Paulo C. G. Costa
The Volgenau School of Information Technology and Engineering,
George Mason University, Fairfax, VA 22030 USA
[klaskey, dschum, pcosta]@ gmu.edu

Terry Janssen
Lockheed Martin, IS&GS/GSS,
Herndon, VA 20171 USA
terry.janssen@lmco.com

Abstract— **Intelligence analysts rely on reports that are subject to many varieties of uncertainty, such as noise in sensors; deception or error by human sources; or cultural misunderstanding. To be effective, intelligence analysts must understand the relationship between reports, the events or situations reported upon, and the hypotheses of interest to which those events or situations are evidential. Computerized support for intelligence analysts must provide assistance for managing evidential reasoning. For this purpose, computational representations are needed for categories and relationships related to evidential reasoning, such as hypotheses, evidence, arguments, sources, and credibility. This paper describes some of the entities and relationships that belong in an ontology of evidence, and makes the case for the fundamental importance of a carefully engineered ontology of evidence to the enterprise of intelligence analysis.**

Index Terms— **Evidence, probabilistic ontologies, intelligence analysis, inferential reasoning, source credibility**

I. Introduction

Evidential reasoning is fundamental to the practice of intelligence analysis. Much of an intelligence analyst's time is spent constructing complex chains of argument from evidence to conclusion, weighing the force of each argument and the credibility of its component sources, and arriving at overall judgments that, while falling short of certainty, provide useful inputs to decision makers. Reports that give rise to intelligence assessments are characterized by many varieties of uncertainty: noise in sensors; deception or error by human sources; poor understanding of situation or context. To be effective, intelligence analysts must understand the relationship between reports, the events or situations reported upon, and the hypotheses of interest to which those events or situations are evidential.

It follows that effective computerized support for intelligence analysts must support processes of evidential reasoning. For this purpose, computational representations are needed for categories and relationships related to evidential reasoning, such as hypotheses, evidence, sources, credibility, and the like.

Some have argued that computational representations of evidential categories and relationships, while necessary to intelligence analysis, do not belong in an ontology. Ontology, the argument goes, is the systematic study of existence: the categories of things that can exist and the relationships they can bear to one another. In the field of information systems, the term has come to mean the engineering discipline of constructing computational representations of various domains of application. By contrast, epistemology is the study of knowledge: how agents come to know about things that exist. The ontologies we construct, the argument goes, should be about what *is*, not what *might or might not be*, or what agents can reasonably *infer* from available evidence.

Computational support for intelligence analysts requires the ability to represent, store, and manipulate evidence, hypotheses, and arguments relating evidence to hypotheses. Such representations must be stored in a computational structure, which, for want of a better term, we might call an epistemological repository. Let us consider what such an epistemological repository might contain. It would represent concepts such as hypothesis, evidence, source, and report. It would contain relationships such as relevance of evidence to hypothesis, or the source-of relationship connecting a source with a report produced by the source. It would be quite natural to construct the representation using the languages and tools commonly applied in the discipline of ontological engineering. In other words, this epistemological repository would look rather like a domain ontology, where the domain being represented is epistemology – the field devoted to how we use evidence obtained from the world around us to arrive at knowledge about the world. The natural person to build this repository would be someone schooled in constructing such representations – that is, an ontological engineer. To call such a repository an ontology of evidence would hardly seem unreasonable.

In this paper, we argue for the fundamental importance of a carefully engineered ontology of evidence to the enterprise of intelligence analysis for the need for an ontology of evidence, and describe some of the entities and relationships that such an ontology would represent.

II. Evidence and Argument

Schum [1] has written a systematic treatise on evidence and its role in constructing arguments. All evidence, according to Schum, has three major credentials: relevance, credibility, and inferential force or weight. Relevance concerns the degree to which the evidence bears upon the hypothesis under consideration. Credibility means the degree to which the evidence is believable; whether or not the evidence is

Manuscript received November 1, 2008.

trustworthy. Inferential force concerns the strength of the relationship between evidence and hypothesis – the degree to which the evidence sways our belief in the hypothesis.

Evidence can come from diverse types of sources (e.g. physical sensors, human reports, direct tangible evidence such as objects or documents), each with different degrees of relevance, levels of credibility, and force.

As examples of the factors bearing the credibility of a source, evidence coming from physical sensors needs to be evaluated with respect to environmental conditions, distance from observer, and physical characteristics of the respective sensor. Human sensors, on the other hand, must be scrutinized with respect to opportunity, competence, and veridicality. Opportunity concerns whether the person was in a position to have observed the event or verified the fact. Competence concerns whether the source was capable of making the distinction in question. Veridicality concerns whether the source is telling the truth. Clearly, there may be complex chains of inference involved in ascertaining any of these factors influencing credibility. Approaches for dealing with the weight or strength of evidence include both qualitative and quantitative aspects of the reasoning process adopted to draw inferences from it (e.g. probability theory, logical reasoning, etc).

A vital (and too often overlooked) distinction to be made is the difference between an event and evidence that the event occurred, or between a fact and evidence that the fact obtains. Schum uses the notational device of an asterisk to make the distinction between event or fact E and evidence E^* relating to E. It is important to note that E^* does not entail E; the inference to E depends on the credibility of the source of E^*.

We do not always have the luxury of a direct report E^* on an event or fact E of interest. We may need to reason indirectly from a report R^* to an event or proposition R whose truth bears on the truth of E, and from there to E itself. Collections of interrelated propositions can be chained together into complex arguments. We often think of an argument as a linear chain from evidence through a collection of intermediate conclusions to a final conclusion. However, each link in such a chain must be justified. A judgment must be made that each antecedent in the chain is relevant to its consequent. The evidential force of each link must also be established. These judgments often require evidential reasoning in their own right. Schum uses the term *ancillary evidence* to refer to evidence about the nature and force of an evidential relationship. Intelligence analysts require support for keeping account of chains of argument and the ancillary evidence on which their force depends.

III. Probabilistic Treatments of Evidence

The past century has brought broad appreciation of the statistical regularities underlying the seeming complexity of physical, biological, psychological, and societal phenomena [2]. Computational advances are enabling automated and semi-automated support for many "knowledge tasks" once thought to be the exclusive province of human cognition. Intelligence analysts increasingly rely upon computerized systems that allow them to catalog, organize, and explore the implications of large collections of reports and other evidence. Quantitative measures of the strength of evidence are useful as a way to summarize and communicate the implications of large bodies of evidence. A natural candidate for such summarization, with a long and respected intellectual tradition behind it, is probability. Systematic deviations of intuitive human reasoning from the tenets of probability theory (*e.g.*, [3]) have been cited as justification for heuristic approaches to combining strength of evidence (e.g., [4]). Nevertheless, naturalistic human reasoning can usefully be treated as a computationally bounded approximation to a probabilistic norm (*c.f.*, [5], [6]). There is a robust literature on the use of probability and decision theory to support human inference and decision making, and to protect against errors that can occur in naïve human reasoning (e.g., [7], [8]). Furthermore, heuristic techniques introduced as cognitively natural ways to overcome perceived disadvantages of probability theory have been shown to admit a probabilistic interpretation (e.g., [9]). When the independence conditions justifying the probabilistic interpretation are met, such heuristic weighting factors can work well, but they can produce disastrous results when applied without regard to whether these conditions are met. There is no match for probability theory in its generality, logical coherence, and well-developed methodological base. For this reason, we focus on probability theory as a logically justified approach to combining numerical measures of evidential force.

We provide several examples to illustrate how probability can be used to represent and reason about credibility, to combine reports from different sources, and to handle subtleties such as dependence relationships that can stymie naïve heuristic weighting schemes. Our examples are deliberately kept simple to illustrate the key points. They are not intended to represent the full complexity of the evidential reasoning problems faced in real applications. Nevertheless, they illustrate the building blocks from which a more sophisticated reasoning capability can be constructed.

Figure 1 shows a Bayesian network that illustrates the combination of three independent pieces of evidence regarding the whereabouts of Osama bin Laden. Prior to receiving the reports, the probability is 3% that he is in Kandahar. After receiving the first report, the chance increases to 11%. After a

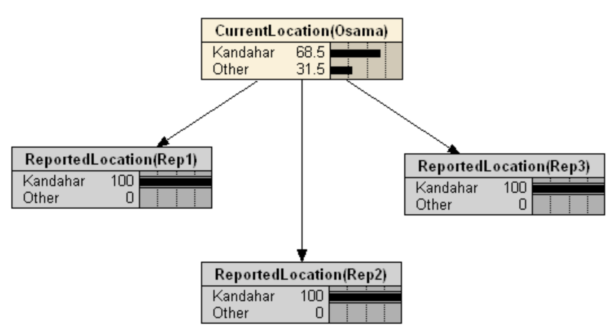

Figure 1: Three Independent Reports Increase Probability of Hypothesis from 3% to 69%

second report, the probability is 35%; the third report brings the probability to 69%. The figure shows the situation after the third report has been received. The top rectangle represents hypotheses about bin Laden's location and their probabilities (Kandahar at 69%; Other at 31%). The three reports are shown below the location hypotheses. The gray color indicates that they have been specified as evidence, with 100% probability assigned to the actual reported location. Figure 2 extends this example to explicitly represent report credibility. The figure now shows credibility hypotheses (low, moderate and high) for the three reports. If we had specified no evidence about the credibility values, the results would have been the same as Figure 1. But if we specify that the credibility of the third report is low, then the probability decreases to 55% that bin Laden is in Kandahar. That is, lowering the credibility of a report decreases its evidential force, resulting in less change in belief when the report is received.

Our final example illustrates an issue not easily accounted for by heuristic methods for assigning and combining evidential weights. Suppose we discover that two of the reports, which we had originally treated as independent, may have actually come from the same informant. We can treat this case by explicitly representing a hypothesis for whether the reports came from the same source. In Figure 3*a*, we indicate that the sources of the two reports are different. In this case, they can be treated as independent evidence items, and the resulting belief in bin Laden's location is the same as in Figure 1. However, if we specify that the sources are the same (Figure 3*b*), the probability that bin Laden is in Kandahar is reduced to 35%, the same as if we had received only two independent reports. The structural assumptions (the independence relationships represented in the graphs) together with the numerical probability values ensure that subtleties such as source credibility and common sources are properly accounted for in evidential reasoning.

Additional treatments of probabilistic representations of relevance and credibility in evidential reasoning can be found in [10] and [11].

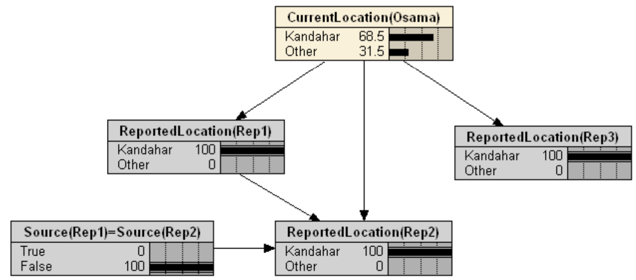

a. Sources for Rep1 and Rep2 are Different

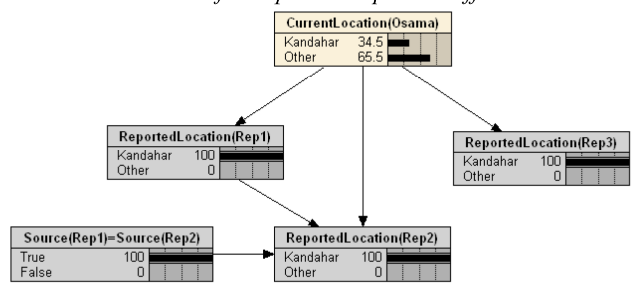

b. Sources for Rep1 and Rep2 are the Same

Figure 3: Common Source Reduces Force of Report

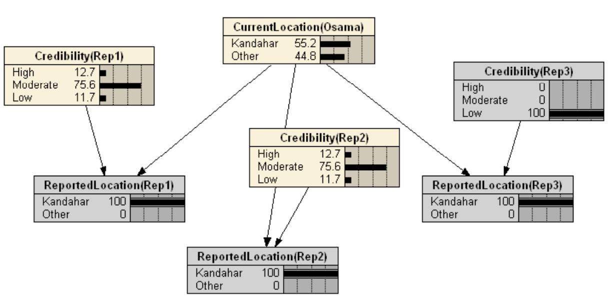

Figure 2: Low Credibility Reduces Force of Report

IV. A Probabilistic Ontology of Evidence and Inferential Reasoning

The above concepts pertain to the use of evidence as an informational asset and to the inferential process that transforms it into knowledge. This is clearly a multi-disciplinary subject. Practitioners from many disciplines can profit from a formalization of the discipline of evidential reasoning. Due to its heavy dependence on evidence in almost every aspect of its operations, the domain of intelligence analysis would be a prime beneficiary of an ontology of evidence. Benefits of an ontology of evidence include a common, shared vocabulary for important features and relationships that occur across different applications of evidential reasoning, as well as the ability to share information among diverse systems.

Despite considerable diversity and individual variation in the conduct of investigation and analysis, there are fundamental common structures and processes. Examples include assessing the credibility and relevance of individual items or of masses of evidence, or constructing reasoning chains to connect evidence to hypothesis. A formal representation of evidence and evidential relationships provides the obvious benefit of allowing analysts to query a knowledge base not just for conclusions (e.g., "Where is Osama bin Laden?"), but also for the evidence on which the conclusions are based (e.g., "What is the evidence that bin Laden is in Kandahar?") Analysts can reason about the relevance of evidence to hypotheses, the credibility of sources, errors that may be common to several evidential reasoning chains, and other subtleties of evidential reasoning.

There has been an increasing emphasis in recent years in sharing knowledge among intelligence applications. An ontology of evidence and inferential reasoning is a first step in that direction. Ontologies provide shared representations of the entities and relationships characterizing a domain, into which vocabularies of different systems can be mapped so to provide interoperability among them. Techniques for making semantic information explicit and computationally accessible are key to effective exploitation of evidence from diverse sources, with distinct grades of credibility and relevance. Shared formal semantics enables systems with different internal representations to exchange information, and provides a means to enforce business rules such as access controls for security.

However, traditional ontologies do not provide a principled means to ensure semantic consistency with respect to issues of uncertainty related to credibility of sources, relevance of evidence, and other aspects of the evidential reasoning process. Because uncertainty is a fundamental aspect of evidential reasoning, this is a serious deficiency.

When faced with the challenge of representing uncertainty in an ontology, the natural tendency is to introduce a means to annotate property values with information regarding their level of confidence. This approach addresses only part of the information that needs to be represented in a full ontology of evidence. To understand why more is needed, consider the example from Section II above, in which evidence from several sources is combined to draw an inference about the current location of Osama bin Laden. We saw that the inferential force of each report depended not only on that report's credibility, but also on whether the information from which it was derived overlapped with the information on which another report was derived. In other words, we need to represent not just a single credibility number, but information about how that credibility was derived. An assessment from source x, in order to be used in conjunction with evidence coming from other sources would not only state that (say) "with 75% probability, Osama bin Laden is in Kandahar." To be part of a comprehensive probabilistic model capable of performing sophisticated evidential reasoning, such a statement would have to include the supporting evidence and how its credibility affects the overall assessment. A simple example would be "with 75% probability, *given* reports that his physician was spotted in a local market (evidence E1) and that a radio communication regarding his whereabouts was intercepted (evidence E2)," accompanied by information clarifying how this number changes as the credibility of E1 and E2 varies. Further, as new evidence accrues, a sophisticated evidential reasoning system must be capable of capturing the impact of additional evidence on the body of evidence being analyzed. As an example, if a source were found to be a double agent, the credibilities of all reports from that agent would need to be called into question. A system that relies on or can represent only numerical weights of individual arguments cannot cope with the complexity and dynamic aspect of real world multi-source evidential reasoning.

In short, annotating a standard ontology with numerical probabilities is not sufficient, as too much information is lost due to the lack of a good representational scheme that captures structural constraints and dependencies among probabilities. Over the past several decades, semantically rich and computationally efficient formalisms have emerged for representing and reasoning with probabilistic knowledge (e.g., [12]). A true probabilistic ontology must be capable of properly representing the nuances these more expressive languages were designed to handle. We have argued elsewhere (e.g. [5]) that for domains characterized by uncertainty, probabilistic ontologies ([13], [14]) provide a principled means to represent the structural and numerical aspects of evidential reasoning. Indeed, many researchers have pointed out the importance of structural information in probabilistic models (e.g. [15], [16]), and it is well known that some questions about evidence can be answered entirely in structural terms ([1], page 271). Shafer ([17], pages 5-9) argues that probability is more about structure than it is about numbers. Numerical probabilities enable quantitative assessment of the force of evidence, which depends on the strength of relevance and credibility arguments. Exploring the details of probabilistic ontologies is not in the scope of this work, but the interested reader is referred to http://www.pr-owl.org.

Finally, apart from the advantages of knowledge sharing tools to the Intelligence Analysis domain, it is important to foresee the institutional and cultural implications of systematizing and standardizing vocabulary and semantics of evidential reasoning. The very difficulties an effective information-sharing scheme is meant to overcome can become obstacles to its widespread adoption. Given the nature of the field, with highly personalized approaches to analysis, a knowledge tool may encounter resistance if it is perceived as threatening deeply ingrained processes. Yet, the increasing demands within the Intelligence community for effective exchange create an opportunity for developing standardized representations and approaches. This is an important and difficult issue. A probabilistic ontology of evidence is a promising first step to provide a structure for knowledge sharing that is sufficiently flexible to address the demands of the multiple approaches currently used to handle evidential reasoning.

V. Summary and Conclusions

After identifying some concepts regarding the process of transforming masses of evidence into knowledge, we explored the need for formal representations of evidential processes as a means to provide cross-fertilization among domains that depend on processes of evidential reasoning. Among these, intelligence analysis is paradigmatic. We proposed a probabilistic ontology of evidence as a key enabler of this vision. Implementation of this concept must be cognizant of institutional and cultural barriers. In conclusion, we argue that the benefits of effective evidential reasoning and knowledge sharing tools far outpace the difficulties in implementing them.

Acknowledgements

The authors gratefully acknowledge helpful comments from anonymous reviewers of an earlier draft of this paper.

References

[1] D. A. Schum, *Evidential Foundations of Probabilistic Reasoning*. New York: John Wiley & Sons, Inc., 1994.

[2] G. Gigerenzer, Z. Swijtink, T. Porter, L. Daston, J. Beatty, and L. Kruger. *The Empire of Chance: How Probability Changed Science and Everyday Life. Cambridge, Cambridge University Press, 1990*.

[3] D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge Univ. Press, 1982.

[4] S. Bringsjord, J. Taylor, A. Shilliday, M. Clark and K. Arkoudas. "Slate: An Argument-Centered Intelligent Assistant to Human Reasoners," in F. Grasso, N. Green, R. Kibble and C. Reed (*eds.*) *Proceedings of the 8th International Workshop on Computational Models of Natural Argument* (*CMNA 8*), Patras, Greece, 2008. http://kryten.mm.rpi.edu/Bringsjord_etal_Slate_cmna_crc_061708.pdf.

[5] L. Martignon and K. B. Laskey. "Taming Wilder Demons: Bayesian Benchmarks for Fast and Frugal Heuristics." In *Simple Heuristics that Make us Smart.* The ABC Group, Oxford University Press, 1999.

[6] J. R. Anderson and M. Matessa. "Explorations of an Incremental, Bayesian Algorithm for Categorization." *Machine Learning* 9(4), 275-308, 1992.

[7] D. Von Winterfeld and W. Edwards. *Decision Analysis and Behavioral Research*. Cambridge, U.K.: University Press, 1986.

[8] W. Edwards, R. F. Miles, Jr., and D. von Winterfeldt, *Advances in Decision Analysis: From Foundations to Applications*. Cambridge, U.K.: University Press, 2007.

[9] D.E. Heckerman. "Probabilistic interpretations for MYCIN's certainty factors". In J. Lemmer and L. Kanal, editors, *Uncertainty in Artificial Intelligence Vol 1*, pages 167-196, Amsterdam: Elsevier, 1986.

[10] S. Mahoney, D. Buede, and J. Tatman. "Patterns of Report Relevance." *Proceedings of the Third Annual Bayesian Modeling Applications Workshop*, 2005. http://www.intel.com/research/events/bayesian2005/docs/Mahoney-ReportRelevance.pdf.

[11] E. Wright and K. Laskey. "Credibility Models for Multi-Source Fusion." *Proc. 9th International Conf. on Information Fusion*, 2006. http://ite.gmu.edu/~klaskey/papers/Wright_Laskey_Credibility.pdf

[12] K. B. Laskey, "MEBN: A Language for First-Order Bayesian Knowledge Bases" *Artificial Intelligence*, 172(2-3), 200

[13] K. B. Laskey, P. C. G. Costa, and T. Jensen, "Probabilistic Ontologies for Knowledge Fusion," in *Proc. 11th International Conf. on Information Fusion*, Cologne, Germany, 2008.

[14] P. C. G. Costa, "Bayesian Semantics for the Semantic Web," PhD Dissertation, Dept. of Sys. Eng. and Op. Res., George Mason Univ. 315p, Fairfax, VA, USA, 2005.

[15] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA, USA: Morgan Kaufmann Publishers, 1988.

[16] J. B. Kadane, and D. A. Schum, *A Probabilistic Analysis of the Sacco and Vanzetti Evidence*. New York: John Wiley & Sons, 1996.

[17] G. Shafer, "Combining AI and OR," University of Kansas School of Business, Working Paper No. 195, April 1988.