Combining Facts and Expert Opinion in Analytical Models via Logical and Probabilistic Reasoning

Marco Valtorta, John Byrnes, Michael Huhns, and Richard Rohwer

University of South Carolina CSE TR 2006-012

Summary

This report describes a proposal to develop computationally feasible technology to support optimal management of uncertainty, inconsistency and disagreement in collaborative intelligence analysis by importing semantically guided proof search techniques in development at HNC into Bayesian network techniques in development at USC.

Uncertainty is pervasive in intelligence analysis. Support systems for intelligence must be able to quantify and track uncertainties in evidence findings, in data used by inferential processes, in the imperfect theories that emerge from the individual and collective experience of analysts, and from other sources. Bayesian probability theory defines the ***unique*** paradox-free method for reasoning with uncertainty, a proven result [Van Horn 2003] that is less widely known than it deserves to be. Although they enjoy certain advantages in versatility and computational complexity, logical knowledge bases are ill-suited to represent uncertainty and then reason about it correctly, because knowledge representation languages based on classical logic do not provide facilities for representing and reasoning about uncertainty expressed in a probabilistic form.

Recent work shows that, in principle, such facilities can be provided by extending the logical framework to support such representations as multiple-entity Bayesian networks and probabilistic relational models [Bangsø and Wuillemin 2000; Getoor et al. 2002; Laskey and da Costa 2005], but the scalability of such approaches is questionable [Jaeger 2000]. We propose to overcome this problem by developing approximate methods that automatically convert logical proofs into Bayesian networks. A proof is derived from the application of a logical knowledge base to a particular situation, leveraging knowledge summarization techniques developed by HNC that guide and accelerate proof search. The Bayesian network can then be used to reason about the uncertainty of data sources, the uncertainty associated with expert judgment, conflicting data, and conflicting judgments. Conflicting data will be a major issue as larger knowledge bases are used, and particularly as more of their content is extracted automatically from text, because logic engines fail catastrophically upon encountering a contradiction. Our approach will provide the main advantages of a full integration of logical knowledge bases with Bayesian networks without facing the computational complexity of such a project.

The results of our effort will be a logical and probabilistic reasoning system (BALER, for *Bayesian And Logical Engine for Reasoning*) that (1) can be incorporated into other CASE projects, (2) will be used, demonstrated, and validated within USC's Magellan system for generating and evaluating hypotheses, and (3) will enable analyst teams to collaborate on large-scale tasks. We will evaluate the reasoning system on realistic intelligence problems and, uniquely, by using large groups of students at USC acting as teams of novice analysts or, similarly, analysts working in a domain that is new to them. We expect to be able to demonstrate our claim that the amount of information that an analyst must process is greatly reduced because the combined reasoner will have removed the irrelevant portions, while bringing to the attention of an analyst the most relevant information.

Having constructed this advanced reasoner and integrated it into the CASE environment, we will study its scalability properties and extend its design for operability with large teams of analysts working on large

data sets during the two option years (FY09 and FY10). We will provide documentation and training in support of field testing of our tools and will evaluate the results of field tests in order to continue to make the developed tools more usable and more powerful. We will develop scenarios, tutorials, and demonstrations that provide for transfer of technology into operational environments within the intelligence community.

The intelligence community has long had a need for intelligent tools to support analysts and decision makers. Reasoning in formal logical systems, at various levels, has been able to provide some of these tools in the past. Statistical and probabilistic systems have also contributed valuable tools that have enhanced an analyst's command over data. These distinct approaches have yet to be fruitfully merged, and doing so requires a significant improvement in the state of the art. It also requires collaboration between communities that struggle to find a common vocabulary and are not often motivated to seek each other out. The practical tools that we will generate through participation in the CASE program and the improved understanding of the theory relating Bayesian and logical techniques will be a significant step toward the synergistic merger of these historically disparate approaches to reasoning over massive data sets and other information sources. The fruit of this integrated approach will be tools for analysts providing far higher intelligence (and hence utility) than any that can be constructed under the current state of the art.

1.0 Technical Approach

Innovative Claims

1. We will extend classical logic formalisms to support reasoning over uncertain information. Classical knowledge representation formalisms provide for reasoning from certain assumptions to certain conclusions. Uncertainty in any assumption entails uncertainty in the conclusion, but there is no mechanism for quantifying the degree of uncertainty. Applying Bayesian networks to formal logical arguments provides such a mechanism, allowing for robust reasoning from uncertain assumptions.

2. We will provide for meaningful reasoning over inconsistent information. As larger knowledge bases become available, as their content is derived automatically from text, and as collaboration results in the integrated use of theories from multiple analysts, it will become increasingly difficult to maintain the logical consistency that today's knowledge engines require. Inconsistency-tolerant reasoning is naturally supported by probabilistic modeling. Since our assumptions may be uncertain already, the appearance of a logical contradiction only reduces the degree of certainty held in some of those assumptions entailing the contradiction. The Bayesian framework provides the unique theoretic foundation for appropriate propagation of the resulting changes in certainty.

3. We will develop a tool that identifies significant agreements and disagreements with potential collaborators and identifies the most significant contributors to disagreements. We will create tools that support different analysts working on the same problem, with similar models, and with different beliefs. These tools will find and analyze disagreements between analysts and determine whether differences in the models are consequential or inconsequential for the hypotheses considered by analysts. These tools will also be able to lead the analysts to understand the source of disagreements, such as crucial pieces of evidence that only one of the potential collaborators has seen.

4. We will develop a tool that considers an analyst's existing prior/tacit knowledge before reporting newly discovered data, prioritizing that data which will be most surprising to the analyst. The Bayesian framework naturally updates degrees of certainty upon the receipt of new data. Data which does not significantly confirm any working hypotheses or alternative competing hypotheses need not be brought to the analyst's attention.

5. We will provide a well-founded theoretical framework in which a new family of techniques can be developed which allow robust reasoning over uncertainty and inconsistency and which directly address a number of challenges facing the intelligence community. The current proposed research is a necessary precursor to the types of applications described below. Any non-Bayesian approach to modeling uncertainty can be shown to lead the reasoner into a paradoxical state of beliefs in which the reasoner becomes vulnerable to arbitrage [Van Horn 2003]. Non-Bayesian software could tell the analyst that there is a 70% certainty that a given conclusion is correct, while at the same time behaving as though there is a 95% certainty that the same conclusion is incorrect; the latter evaluation may or may not be available to the analyst.

1.1 Technical Discussion

Provenance and collaboration. When information sources are made explicit in the reasoning fragments, one can model the degree to which particular sources of information contribute to strongly trusted conclusions, and which sources tend to yield conclusions which are discarded. A similar analysis of information from collaborators will provide further insight into collaboration work habits and help analysts find appropriate collaborators.

Detecting analyst bias. By having a model of data generation from an analyst's assumptions, we can estimate elements of the source data that are most consistent with and most inconsistent an analyst's beliefs. If a conclusion is trusted more strongly than is warranted by the argument for it, we have an indication that the analyst may be relying on a prior prejudice toward the conclusion or may be incorporating unconscious knowledge into the reasoning. This bias can be tracked, isolated, and presented to the analyst for evaluation.

Study of the relationship between the prior/tacit knowledge of individuals and that of groups. We will develop techniques for assembling larger Bayesian networks from smaller fragments. When the smaller fragments become sufficient to represent the beliefs of individual analysts, the assembled networks represent the beliefs of the groups. This provides for both a theoretical and empirical investigation into the relationship between the prior and tacit knowledge of individuals and that of groups.

Hypothesis generation and tracking. Testing and comparing alternative hypotheses is addressed by comparing the degree to which each hypothesis explains observations. The degree by which each hypothesis can be derived from observations and background assumptions can also be measured numerically. When multiple hypotheses are considered within a group of collaborators, they can be assessed not only in terms of the degree to which they are consistent with the fragments contributed by each collaborator but also by the variation in acceptance. Key reasons for the stronger acceptance of a hypothesis by one collaborator than by another could be automatically worked out in terms of different prior knowledge held by the collaborators. As with all assertions, the uncertainty in a given hypothesis would be quantified precisely given any body of background knowledge.

Data integration tools. Finally, we observe that the ability to integrate BN fragments is a step toward the ability to integrate data, information, knowledge, and people. Bayesian reasoning provides the only well-founded theory for the science of uncertain reasoning, and this science would certainly be furthered by our ability to apply the proposed techniques to the challenges facing the intelligence community.

1.1.1 Core Scientific Principles

The core scientific principles underlying our proposal are that (1) it is possible to extend first-order logic theories to Bayesian networks for the purpose of addressing the uncertainty inherent in many phases of analytical activities, (2) such extension is feasible in a practical way for particularly salient situations, and (3) it allows the resolution of key problems that occur in collaboration.

1.1.2 Scenarios

Our scientific program extends classical knowledge representation and reasoning systems. We begin by considering the use of such a classical system in order to illustrate the extensions that will be made available by our proposed research. We borrow from the Sign of the Crescent [Hughes 2003] for familiarity. We present two scenarios. The first one emphasizes use of an argument graph. The second scenario emphasizes the use of causal information in the construction of probabilistic models.

1.1.2.1 Scenario 1: Extending an Argument Graph

Anna was using her new CASE environment to put together an argument that terrorists had access to the New York Stock Exchange. Initially, she applied techniques she had learned in Professor Hughes' class. Anna structured the argument and checked it for completeness: Sahim Albakri was a terrorist and was the roommate of Hamid Alwan, therefore Albakri and Alwan collaborated on terrorist activities. Alwan had access to the NYSE, therefore a terrorist has access to the NYSE. She thought that this should be enough, but the argument checker told her that it wasn't formally complete. It asked whether people who collaborated on terrorist activities were necessarily terrorists. She was amazed that this wasn't already in the knowledge base, but she replied yes. The environment asked her how certain she was about this, and she replied that this was a hard fact that should get absolute certainty. Now came the part that hadn't been in class. She asked the system for a probability estimate in the truth of the conclusion that a terrorist had access to the NYSE. The system replied that the likelihood was 64%. Anna had expected a stronger result, so she asked why it was so low. The system indicated that the primary uncertainty drew from the fact that Anna had indicated only a 75% probability that Albakri and Alwan were actually roommates. Figure 1 diagrams Anna's argument, where the assertions are [Hughes 2003, p 41]:

144. Hamid Alwan has access to the NYSE.

151. Sahim Albakri belongs to a terrorist organization in the USA.

152. Sahim Albakri shares an apartment with Hamid Alwan.

154. Sahim Albakri collaborates with Hamid Alwan in terrorist activities.

128. A terrorist operating in the New York City area has access to the NYSE.

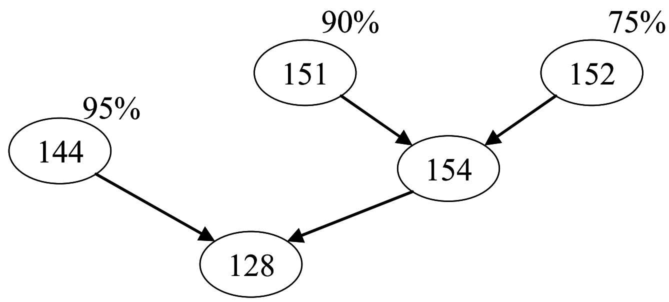

Figure 1. Part of the argument diagram from Sign of the Crescent, Chart H. Here we have annotated the leaves with probability assessments.

Anna was about to write up her report when an update came in. She had turned off updates that weren't related to her current work because she didn't want to get distracted, so she decided to look at this update right away. It informed her that the system had gotten new information that indicated with 70% certainty that Alwan lived on 10th street. Since Anna was considering Albakri's address on 50th street to be certain, the system warned her about the inconsistency: the system indicated 75% certainty that the two were roommates and 70% certainty that they were not. She was glad to get the warning---she had heard about systems in the past that would have held both probabilities simultaneously without warning the analyst. She had even heard that someone had used that kind of data to make a case for a conclusion with 110% certainty. The system asked whether she wanted to manually track down the inconsistency or just let the system handle it. She recalled that her previous system used to handle this type of thing by holding on to one or the other claim and throwing away everything that depended on the claim not chosen. But that wasn't a probabilistic system, so she was curious what would happen. She let the system resolve the inconsistency itself, and it suggested setting the probability that the two were roommates to 56%, resulting in a 48% certainty in the conclusion. Anna decided that this was too low, and asked the system to find someone for her to collaborate with.

Anna expected the system to send her to Bill, who she knew had been doing some work on Alwan. (Bill, in fact, had pointed her to the information that Albakri and Alwan were roommates.) Instead, the system recommended talking to Chris and Dan. The system said that Chris had some information on Alwan having received weapons training in Afghanistan, and that Dan had the same information as Anna about the roommates, but had assigned different probabilities. Anna asked the system to explain the relevance of Chris's information, and it responded that since Chris had independent intelligence indicating with 75% certainty that Hamid Alwan was a terrorist, and this additional evidence of Alwan's being a terrorist boosted confidence in her conclusion to 87.6%. Anna decided that she was likely to agree with Chris's assessment of the reliability of this information, so she decided to trust it for now and go see why Dan had so much more confidence in the information about roommates than she did. Dan explained that the informant who had supplied the 10th Street address for Alwan had turned out to have supplied a large amount of misinformation. Anna returned to her desk and told her system to throw out the 10th Street address. The system correctly returned to the previous estimate of 75% for the likelihood that Alwan and

Albakri were roommates. The conclusion was now held with probability 92%. Anna was pleased that it was higher than before and realized that this was because of Chris's additional information. She decided that she should have asked for potential collaborators initially, even before getting the bogus address information, and decided she would have to remember to do that next time before writing her analysis. She then went to talk to Chris to verify the information and its probability assessment.

1.1.2.2 Scenario 2: Exploiting Causal and Temporal Information

The Sign of the Crescent example contains several messages that, when taken together, lead one to consider the likelihood that Faysal Goba could be attending a meeting in Springfield and, at approximately the same time, be on a train in Atlanta. Formalization of this project was taken up under NIMD. One approach is to include among the axioms assumptions about typicality. In a pure logical system, these axioms are sufficient to entail with certainty the conclusion that Faysal Goba cannot be present at the Springfield meeting. The proof is illustrated in 2. Note that a proof is a directed acyclic graphs (DAG).

This example is typical of envisioned applications of knowledge representation and reasoning systems. Our research will address two well-known weaknesses of this approach to reasoning. Firstly, the degree of uncertainty in the assumptions of the system (the "typicality axioms") is never quantified, and the resulting degree of uncertainty in the conclusion is also not quantified. Note that depending on the structure of the argument, a conclusion may have either higher or lower certainty than the assumptions from which it is drawn. Secondly, it is quite reasonable that inconsistent assumptions are made (such as "trains follow published schedules" and "trains often run late"). Classical logical reasoning in inconsistent systems is meaningless, as all conclusions can be reached with equal validity.

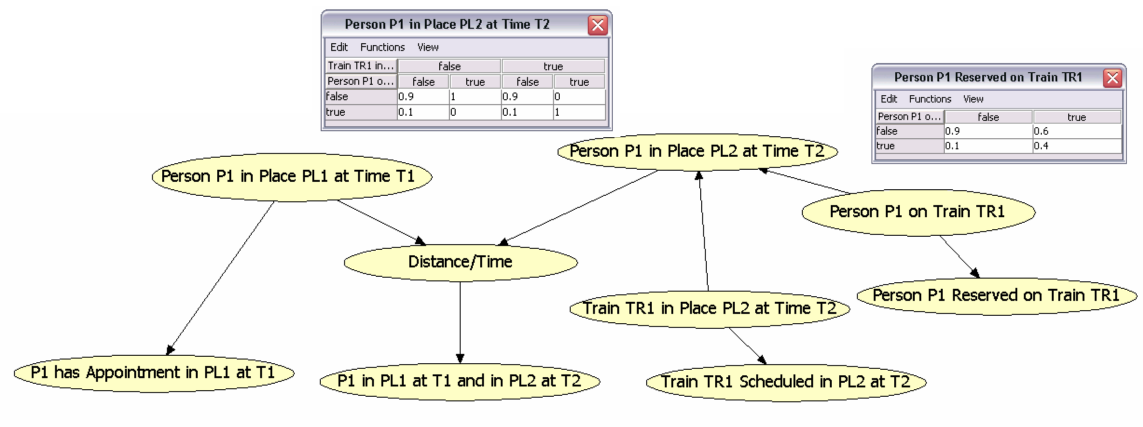

Figure 2. From the logical axioms, we obtain the DAG shown here. The theory is written causally, so edges point from an event to a manifestation that is evidence for the event. The conditional probability assessments shown quantify the "typicality" of axioms informally stated as "if someone is on a train, then that person as a reservation on that train," and "if someone is on a train and that train is in some place, then that person is in that place."

In the ultimate Bayesian framework that we envision enabling, the system knows how strongly the analyst believes in the axioms of the system already. The system estimates the degree of certainty that Goba will be on the train in Atlanta and then reports the degree of certainty that Goba cannot be in a meeting in Springfield. When additional data comes in that confirms Goba's likely attendance at the Springfield meeting (We are now departing from the Sign of the Crescent exercise), confidence that Goba will be on the train drops. The analyst looks for the most weakly supported background assumptions that most strongly cause the two conclusions (attending the Springfield meeting and being on the train in Atlanta) to both have low degrees of confidence. Although there are a large number of factors influencing both conclusions, the system finds that one of the most easily overturned assumptions is the date of the Springfield meeting, for which very little evidence has been gathered. The analyst asks the system to consider the alternative hypothesis that the Springfield meeting will occur one day earlier than the Atlanta train ride, and the system reports that this is sufficient to grant high likelihood that Goba will both attend the meeting and ride the train. The analyst decides to focus on information about the date of the meeting.

1.1.3 Bayesian Extension of Classical Logic

1.1.3.1 Classical logic and uncertainty

Uncertainty is pervasive in the real world. Several types of uncertainty are present in intelligence analysis applications. Some has a linguistic character, some is inherent in the evidence itself, and some reflects the lack of complete knowledge of the true state of the world on the part of even the most informed and experienced decision maker. While there are approaches to the representation of these types of uncertainty in classical logic, they are, at least, cumbersome, *because they require a level of detail equivalent to capturing probability theory in the logical formalism used*, and often wrong, because of the natural tendency to avoid such level of detail.

A brief review of the literature on model-based diagnosis provides an illuminating example of attempts to solve a problem inherently ripe with uncertainty using logic. A seminal paper by Reiter [1987] describes how to apply default logic [Reiter, 1980] to the diagnosis of a system made up of discrete components that can fail and be replaced as a unit. In this approach, a system is at fault if its description is inconsistent with observations about its behavior. A key to this approach is the definition of a component-based logical description of the system to be diagnosed. Some of the logical formulas take the form of rules whose premise includes a normality assumption. A diagnosis is defined as a minimal set of components that, when assumed to be working abnormally, reconcile the observations with the system description. It was soon realized that this conceptually powerful framework leads to a combinatorial explosion in the number of diagnoses. The need to prioritize diagnoses as they are generated, as well as to select observations that are valuable in guiding the diagnostic process, became apparent [de Kleer and Williams 1987]. Some attempts were made to add "ad-hoc" numerical priorities to the components, with varying degrees of success, and a move towards explicit use of probabilities in component-based diagnostic solvers is evident [Jensen et al. 2001b].

The problems that follow from application of classical logic to reasoning under uncertainty are well documented in the literature. Pearl [1988] contains many examples of the paradoxical results that arise from using extensional (also known as truth-functional or compositional) approaches, which "treat uncertainty as a generalized truth value; that is, the certainty of a formula is defined to be a unique function of the certainties of its subformulas." He concludes that probability is a "faithful guardian of common sense." On the theoretical side, the definitive study of compositional systems is in [Hajek et al. 1992]. Their overall conclusion, that "compositional systems seem unlikely to become again a matter of central theoretical interest," corresponds, on the applied side, to a blossoming of applications of graphical probabilistic models.

A major issue to be addressed in making probabilistic reasoning a reality is computational complexity. Worst-case complexity of the three major problems in probabilistic inference, namely probability update (Pr, the computation of posterior probability in the presence of evidence), most probable explanation (MPE, the computation of the most likely state of all variables in a Bayesian network in the presence of evidence), and most likely a posteriori hypothesis (MAP, the computation of the most likely state of selected variables in the presence of evidence), is well known and not encouraging [Park 2002]. One should be discouraged by these apparently negative results. The complexity of the most commonly used solution algorithms depend on graphical parameters, such as treewidth [Dechter 1996; Bodlaender 1998]. "Many probabilistic networks appear to have bounded treewidth in practice" [Bodlaender 2005], and the treewidth is almost invariably very small when the models are obtained from people or human-generated documents (as opposed to models that may arise in, say, bioinformatics). It is therefore possible to conjecture that for all-source intelligence applications, the complexity of probabilistic inference is not a major issue, but we shall be prepared to deal with other sources of complexity in the process of Bayesian network fragment matching and composition. In particular, as the number of fragments in a repository grows to the point that the existing algorithm in Magellan cannot cope, we shall formulate the matching and composition process as a heuristically guided search process.

1.1.3.2 Quantifying uncertainty

Bayesian networks (BNs) provide a means to quantify uncertainty about information and to propagate that uncertainty in a consistent manner. A BN can be constructed over any DAG; the primary component in addition to the DAG is a conditional probability distribution over each node given its parents. Considering the DAG in Figure 1, for example, what we really want to know is the probability that the hypothesis of interest to an analyst is true, given estimates of the probabilities that the axioms are true. The axioms are the leaves of the tree, and the leaves of the tree encode evidence received by the analyst, which may be soft, because of uncertainty associated with the finding in which case the marginal probabilities associated with the leaves are a special case of the conditional distributions defined by the BN extension of the DAG [Valtorta et al. 2002; Vomlel 2004; Kim et al. 2004; Chan and Darwiche 2005]. Any such BN provides an assessment of the probability that the conclusion is true. Not only is this of value when we initially construct our analysis, but it can be used to automatically update the assessment of the conclusion in light of any updated assessment of the assumptions, such as finding out that the information about the train schedule came from an unreliable source. Sensitivity analysis can also be applied to determine those assumptions which have the most impact on the current assessment of the conclusion. These may be given a high prioritization for further analysis.

1.1.3.3 Reasoning in the presence of logical inconsistency

Reasoning from inconsistent assumptions means that certain nodes of the DAG will contain contradictory assertions. Classical logic approaches to this problem include default logic [Reiter 1980] and belief revision [Arlo-Costa et al. 2004; Levi 1991], both of which result in systems that need to remove large amounts of argumentation, because logic is “all or nothing” and cannot simultaneous consider two inconsistent propositions in a meaningful way. In the BN framework, on the other hand, recognizing inconsistency only imposes a new set of constraints. In their simplest form, these modifications amount to requiring that the probability of *not*(*A*) is 1 minus the probability of *A*. This allows us to have a consistent probability distribution in which to assess our confidence in a given conclusion. Minimum entropy techniques provide one reasonable approach to finding distributions satisfying these kinds of constraints [Vomlel 1999]. A special case of this approach is to define the given state which is logically inconsistent to be statistically consistent by considering contradictory assertions to be independent variables. This is what was done by Anna’s system, in the scenario, when it had a 75% certainty that Albakri and Alwan were roommates and a 70% certainty that they were not.

Let X and Y represent “roommates” and “not roommates”, respectively. X=TRUE if the two are roommates and FALSE otherwise; Y=TRUE if the two are not roommates and FALSE otherwise. The original distribution, P_0, holds that P_0(X=TRUE)=0.75 and P_0(Y=FALSE)=0.70. This is a consistent probability distribution on independent variables. The logical inconsistency arises from the recognition that these variables should not be modeled as independent, and in fact is a special case of learning new dependencies in general. What we desire is a probability distribution P for which P(X=TRUE)=P(Y=FALSE) (since these events are now recognized as being identical). Of all possible such P, we will choose the one that lies the closest to P_0 under the Kullback-Leibler divergence. (A number of alternative techniques may be considered, but this one offers a convenient starting point.) We choose p=P(X=TRUE) to minimize D[P, P_0]:

$$D[P,P_0] = \sum_{x,y \in \{True,False\}} P(X=x,Y=y) \ln \frac{P(X=x,Y=y)}{P_0(X=x,Y=y)} = p \ln \frac{p}{0.75*0.3} + (1-p) \ln \frac{1-p}{0.25*0.7}$$

Numerical techniques can be applied to find that $D[P, P_0]$ is minimized when p=0.5625 (reported as 56% in the scenario).

This approached generalizes to the more typical case where a large number of formulas are involved in generating an inconsistency. This is especially problematic for classical logic systems, which need to choose which of the many formulas to eliminate.

1.1.3.4 Incorporation of new information

We now provide another short example from the Sign of the Crescent case study that illustrates *explaining away*, a powerful pattern of uncertain reasoning that is well supported by Bayesian networks, but hard to represent using a logical rule representation [Wellman and Henrion 1993]. This is the case of Shiela [sic] Watson, a person who (in the fictitious example) gives an incorrect address on a job application. We would believe that she is a liar and, therefore, a suspect in the context of the case study. However, we are later told that "she simply made a mistake in listing her home address on her NYSE vendors' ID application. She had recently moved and gave her earlier address by mistake." The conclusion we would reach then is that the mistake is an alternative explanation for the incorrect address, and we would therefore reduce our belief in Shiela's guilt. The situation is described in Figure 3.

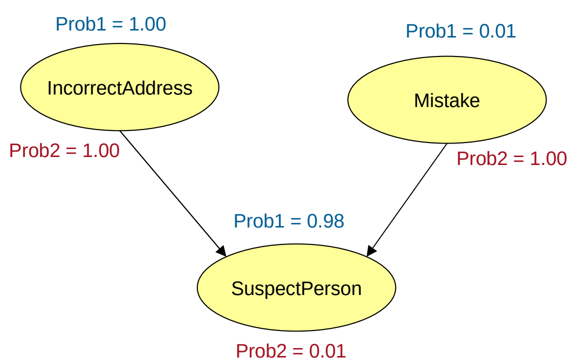

Figure 3. The initial probability assessments are shown above the nodes of the DAG. The probability of a mistake is low and the resulting probability of being a suspect is high. The probability assessments below each node are those made after the mistake is understood. The probability of an incorrect address does not change, but is explained by the new high probability of having made a mistake. The resulting probability of being a suspect is reduced from a very high value to a very low one.

1.1.4 Technical Challenges and Detailed Approach

Our approach to the difficult problem of integrating logical knowledge bases and Bayesian networks builds on the strengths of the two partners and previous work in the NIMD program.

1.1.4.1 Construction of Bayesian Networks over proofs

We intend to start with proofs generated through the use of proof planning based on *association-grounded semantics* (AGS). AGS is a class of techniques for representing the meaning of a data object by a probability distribution over the contexts in which the data object may occur. AGS can hierarchically cluster symbols in a formal knowledge representation system, yielding a hierarchy of abstraction layers in which formal reasoning may occur. A proof in one of these layers becomes a plan for finding proofs in lower layers. Given a low-layer proof, a higher layer abstract proof may easily be derived, allowing for summarization. Choosing a greater degree of summarization should yield more feasible BN computation, but choosing a greater degree of specificity yields probability assessments of more precise conclusions.

The proof instances arrived at will be mapped to Bayesian networks. The precise details of this mapping shall be carefully researched, beyond the illustrations given in this proposal. (Cf. [Getoor 2002; Jaeger 2002; Laskey and da Costa 2005; Poole 2003].) While the structure is given at this point, assessing the needed parameters, which are probabilities for particular variables and conditional probabilities for families of variables, is a more challenging task. We will experiment with several techniques, including using a Kroneker delta with smudge, noisy OR and related functional models [Vomlel 2006], counting explicit mentions of contradictions, and counting references both in the raw input data and in finalized intelligence products. We also anticipate being able to make use of data made available by those contractors that model the analyst or capture analyst activity. Given this starting point, we will extend Laskey and Mahoney's framework of Bayesian network fragment composition, which we already use in the Magellan system within NIMD [Cheng et al. 2005; Laskey 1997].

Sets of proofs or partial proofs for related problems may be considered to be fragments of an overall investigation or of an analyst's overall hypothetical beliefs. For the purpose of fragment composition, the fragments (DAGs) are represented as terms, the matching of fragments with evidence consists of a restricted form of unification, and composition of fragments into larger DAGs (which we call situation-specific scenarios) also consists of a restricted form of unification. See Figure 4 for an illustration and [Cheng et al. 2005] for algorithms. We shall extend the existing implementation to take advantage of some of the features of newer versions of Laskey's system [Laskey and da Costa 2005], and to review related approaches [Bangsø and Wuillemin 2000; Getoor et al. 2002]. In the implementation of [Cheng et al. 2005], the probabilities of each fragment do not depend on the attributes of evidence items, but only on their values. While this is a convenient simplification, we have experienced in the Magellan system that it is a limitation on what can be expressed conveniently. We shall therefore allow prior marginal and conditional probability tables to be affected by the particular details of the evidence items that trigger the use of fragments.

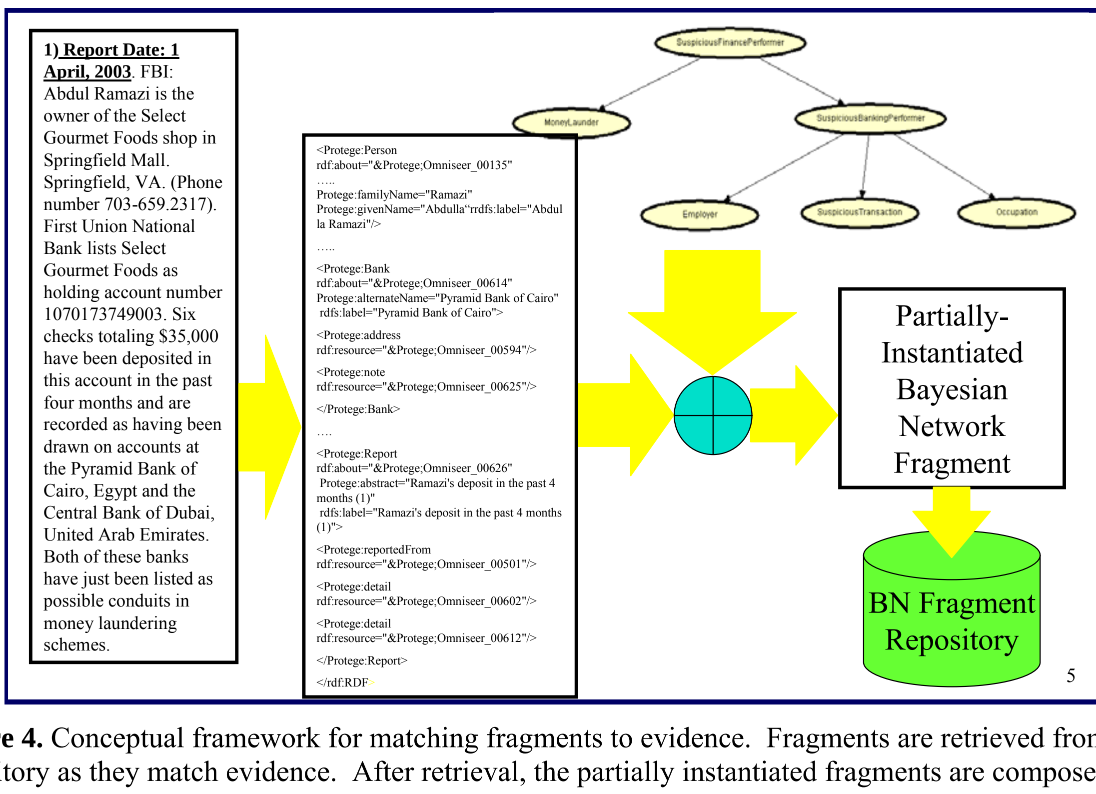

Figure 4. Conceptual framework for matching fragments to evidence. Fragments are retrieved from a repository as they match evidence. After retrieval, the partially instantiated fragments are composed into situation-specific scenarios.

Prior and Tacit Knowledge

Part of the process of inducing the Bayesian network over a given argument tree involves determining the prior beliefs of the analyst with regard to those assertions occurring at the leaves of the argument. The estimation of priors involves a number of technical difficulties in the realm of continuous, parametric approaches to statistical inference and modeling that are the subject of much work in the statistical literature. The use of graphical probabilistic models eliminates many of these difficulties, because the only priors to be estimated are the probabilities of nodes without parents in the Bayesian network and the conditional probabilities of all the other nodes. An illustration is given below.

A Bayesian network (BN) is a graphical representation of the joint probability distribution for a set of discrete variables. The representation consists of a directed acyclic graph (DAG), prior probability tables for the nodes in the DAG that have no parents and conditional probabilities tables (CPTs) for the nodes in the DAG given their parents. More formally, a Bayesian network is a pair composed of: (1) a multivariate probability distribution over n random variables in the set $V = \{V_1, \ldots, V_n\}$, and (2) a directed acyclic graph (DAG) whose nodes are in one-to-one correspondence with $V_1;\ldots, V_n$. (Therefore, for the sake of convenience, we do not distinguish the nodes of a graph from variables of the distribution.)

Bayesian networks allow specification of the joint probability of a set of variables of interest in a way that emphasizes the qualitative aspects of the domain. The defining property of a Bayesian network is that the conditional probability of any node, given any subset of non-descendants, is equal to the conditional probability of that same node given the parents alone. In most applications, this property is insured by qualitative causal or independence relations that form the basic structure of a domain of interest.

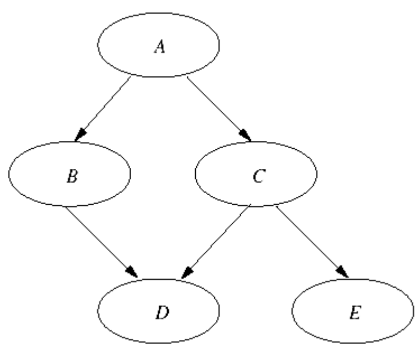

	A_1	A_2
B_1	0.2	0.1
B_2	0.6	0.6
B_3	0.2	0.3

Figure 5. A sample Bayesian network structure with the conditional probability table P(B|A).

Three features of Bayesian networks are worth mentioning. First, the directed graph constrains the possible joint probability distributions represented by a Bayesian network. For example, in any distribution consistent with the graph of Figure 5, D is conditionally independent of A given B and C. Also, E is conditionally independent of any subset of the other variables given C. Second, the explicit representation of constraints about conditional independence allows a substantial reduction in the number of parameters to be estimated. In the example, assume that the domain of the values of the variables have size 2, 3, 2, 4, and 4, in order. Then, the joint probability table P(A,B,C,D,E) has 2*3*2*4*4 = 192 entries. It would be very difficult to assess 191 independent parameters. However, the independence constraints encoded in the graph permit the factorization P(A,B,C,D,E) = P(A)*P(B|A)*P(C|A)*P(D|B,C)*P(E|C), which reduces the number of parameters to be estimated to 1+ 4+ 2+18 +6 = 31. The second term in the sum is the table for the conditional probability of B given A. This probability is shown in Figure 5; note that there are only four independent parameters to be estimated since the sum of values by column is one. Thirdly, the Bayesian network representation allows a substantial (usually, dramatic) reduction in the time needed to compute marginals for each variable in the domain. The explicit representation of constraints on independence relations is exploited to avoid the computation of the full joint probability table in the computation of marginals both prior and conditioned on observations. See [Dechter 1996; Bloemeke and Valtorta 1998; Jensen 2001] for a discussion of some algorithms for this computation.

Some prior information may be elicited from the analyst directly, for example, if, as in the Magellan system, an ACH front-end is used to obtain estimates of the credibility of evidence items and of conditional probabilities. Research will need to be carried out on whether such information can be used reliably in the assessment process, since ACH is not, in its basic form, a probabilistic system. In a way somewhat related to [Nielsen and Jensen 2004], simply by observing which conclusions the analyst trusts the most, we can start to get estimates of the analyst's confidence in a variety of assumptions on which the arguments are based. Any other tools that monitor the analyst's knowledge can also provide input for our estimation of these prior beliefs. By having a model of data generation from these assumptions, we can estimate elements of the source data that are most consistent with or most inconsistent with an analyst's beliefs. If a conclusion is trusted more strongly than is warranted by the argument for it, we have an indication that the analyst may be relying on a prior prejudice toward the conclusion or may be incorporating unconscious knowledge into the reasoning.

For the first scenario, we used an extremely simple first algorithm for constructing the BN. We use the DAG of the argument itself. Each node is a formula that is either true or false. The probability of truth of a node given the truth of all of its parents is 1. Given any other combination of truth values for the parents, we take the probability of the conclusion being true to be the same as the probability of its truth without any argument. Since the probability of an intelligence claim such as "Hamid Alwan is a terrorist" with no supporting evidence of any kind is extremely small, we set the probability to 0 for simplicity. Looking at nodes 151, 152, and 154 in Figure 1, we have assessed P(151)=0.9 (i.e., we believe that there is a 90% chance that claim #151 is true) and P(152)=0.75. We construct the following conditional probability table for P(154|151,152):

	151 True, 152 True	151 True, 152 False	151 False, 152 True	151 False, 152 True
154 True	1	0	0	0
154 False	0	1	1	1

An incorrect application of this table allows us to conclude with certainty that an assertion is false whenever the premises are false. This is not what is desired, which is why use of the table is restricted to the context of the argument: we assume that an argument based on false assumptions offers no support for its conclusion. This is not to be used to reject the conclusion independently of the argument; rather it serves to quantify the value of the argument itself. This restricted applicability is what gives us the simplicity of the above table.

In this context, "incorporation of new information" becomes incorporation of additional arguments for a conclusion (where a conclusion might be any node in the original argument tree). In the scenario, Chris had an argument for Alwan being a terrorist that was independent of Anna's original argument. Internally, Anna's original argument was represented by several more steps than were included in Figure 1; this was the job of the logic assistant as developed under NIMD and IKRIS. One of these additional nodes asserted that Alwan was a terrorist---this node is marked "A is T" in Figure 6. Chris's argument (or an argument constructed by the logic assistant from Chris's information) was a formal proof of this particular claim; it is sufficient to represent the conclusion of the argument as a single node, labeled "Chris" in Figure 6. "A is T" is represented differently than the other nodes because it will be justified by either of its parents, rather than requiring both parents as the other nodes do. This is a familiar graphical structure from logical reasoning and from AI in general: the search space is an "AND/OR tree". A single proof is a fragment of this tree that does not include OR nodes, but multiple partial proofs are naturally joined by the OR nodes, so that the result of the proof search may be this more informative structure.

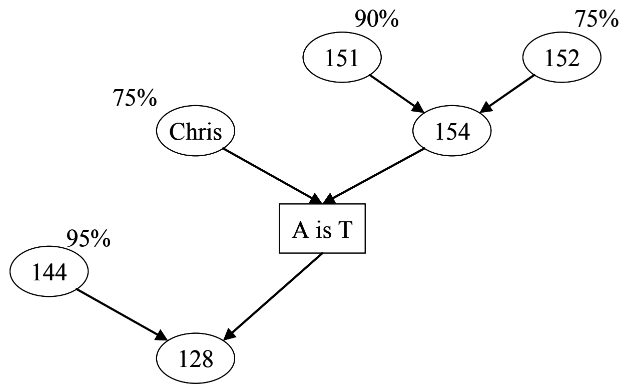

Figure 6. An augmented argument tree containing an "OR" node. The claim "Alwan is a terrorist" is true if *either* Chris's conclusion *or* claim 154 are true. This is in contrast to claim 154, for example, which is justified in the argument only if *both* claims 151 *and* 152 are true.

Accommodating this join of two arguments is done simply by creating the appropriate conditional proof table for the "A is T" node. Prob(A is T | Chris, 154) is defined by the following table:

	Chris True, 154 True	Chris True, 154 False	Chris False, 154 True	Chris False, 154 False
A is T True	1	1	1	0
A is T False	0	0	0	1

This is the table used in the scenario to compute both the 87.6% and 92% confidence in the conclusion of Anna's argument, depending on the different probabilities assigned to the two suspects being roommates.

We described above how BNs can produce meaningful inference in the presence of contradictory assertions, and this includes contradictory prior and tacit knowledge. Prior knowledge that is found to be inconsistent with the knowledge of other analysts or with data can be highlighted and possibly removed from consideration. One of the central properties of BNs is their ability to quantify uncertainty in light of additional information. The discovery of additional information can easily lead to increased doubt, as in the "explaining away" example in section 1.1.3.4 above. In this example, if we had some doubt about Shiela Watson's guilt originally, we had less doubt when we thought she had falsified a document. The additional information that an honest mistake was made returns us to our original position of doubt.

A particularly exciting extension of the compositional fragment model arises when we consider each analyst to be represented by a set of fragments. We can then attempt to fit large sets of fragments together to begin both a theoretical and empirical investigation into the relationship between the prior and tacit knowledge of individuals and that of groups.

1.1.4.2 Collaboration

The use of Bayesian network fragments as a representation of prior knowledge supports inference and analytical functions, as we have shown in the Magellan system [Cheng et al 2005]. We will use this same representation to support collaboration among analysts. In particular, the use of probabilistic models greatly simplifies the construction of more detailed and complete situation models, even when they are constructed from fragments obtained from several analysts, and it supports the integration of new knowledge (possibly obtained from a collaborator) into an existing model. Suppose that an analyst using a probabilistic situation model obtains from another analyst a more precise model relating some of the events of the original model. By using probability distribution integration techniques, this new model can be integrated into the larger model [Kim et al. 2004; Valtorta et al. 2002; Vomlel 1999].

The viewpoints, assumptions, and biases shared (or not) by a pair of potential collaborators can be quantified and isolated when they are represented as beliefs in BN fragments. The degree of agreement or disagreement between any two analysts on any set of modeled beliefs can be quantified immediately. Given a strong disagreement, it is sensible to ask *why* the two analysts have come to such different conclusions about a given assertions. Although computationally difficult, it should be feasible in many cases to be able to find a key set of background beliefs or key body of evidence which explains the difference in beliefs about the conclusion of concern. This is the type of information that can take a long time to elicit through human discussion and may well block much potential collaboration.

Our research will be guided by concepts provided by the Bayesian and information-theoretic frameworks, such as: nearest-consistent theory, model averaging, probabilistic conflict detection, adaptation, soft-evidential update, and the use of the iterative proportional fitting procedure. Differences of opinion may be due to differences in evidence or in parameter values or in structure. The alignment matters only on variables that are of interest. The interest comes from some source, such as existing profiles, recently read documents, or explicit statements of what an analyst specifies. Trust and provenance will be evaluated by modeling multiple unreliable sources of data or expert judgment as evidence nodes in a data fusion problem [Jensen 2001].

Multiple Experts Represented Explicitly in Bayesian Networks

There are several approaches to accommodating the differing opinions of multiple sources of expertise or multiple analysts within a Bayesian network framework. Our proposed approach allows analysts to share the common scenario fragments on which they agree and to maintain separate fragments on which they disagree. Special model-selection nodes will be introduced into the composed situations that can attach the fragments on which the analysts disagree. When evidence is placed on these nodes, the fragment from a particular analyst is emphasized and the model reflects the analyst's judgment, while enabling the differing opinions to be shown and considered. When evidence is not placed on such nodes, the system can reveal which analyst's judgment is most consistent with the available evidence. The outcomes provided by the system are

- Consensus views can be computed
- The range of views is maintained
- Situations are identified in which certain analyst's judgments are supported
- Based on the range of judgments, the system can produce a range of conclusions for the hypotheses of interest.

Because multiple experts, with their unique opinions and conclusions, are represented, the presence of group think can be made evident and then mitigated, where appropriate.

1.1.5 Performance Evaluation

Initial evaluation of our system will involve work on synthetic problems of the kind described in, e.g., [Pearl 1988; Shenoy and Biswas 1990; Bezdek 1992; Halpern 2003]. We want to insure that the system does not fail in addressing paradoxical situations that fool some of the naïve systems for reasoning under uncertainty, such as compositional systems, and that it has sufficient expressive power.

A second kind of evaluation will involve examples such as those in [Heuer 2001] and case studies, such as Frank Hughes's well-known "Sign of the Crescent." The advantage of the synthetic problems above is that they have correct answers. Case studies such as the "Sign of the Crescent" provide good illustrative examples, but do not provide "correct answers" for probability assessments of arguments. We will need to make subjective assessments, but would like to avoid assessing our tools directly. This will be achieved by constructing a variety of arguments based on the suggested examples and attempting to rank the relative strengths of the arguments based on natural intuition. This will be carried out by the principle researchers, by students, and possibly by experienced analysts or other subject matter experts. We will compare the rankings with the rankings produced by our tools as an evaluation of how well the tool captures human intuition about what constitutes strong evidence for a conclusion.

In order to test the ability of our machinery to reason in the presence of inconsistent assertions, we will reason over consistent case study examples, and then add inconsistent assumptions and evaluate the degree of stability of the system. Theoretical analyses should also be possible that are able to quantify the sensitivity of the reasoning system to "noise" such as false information. Our research will also be presented at peer review conferences and submitted to peer-reviewed journals.

The initial task we propose is dedicated to surveying work in the field, and we will give special attention to evaluation techniques that have been proposed and applied in the past. We anticipate that part of the overall contribution we make to the state of the art will be both in the development of evaluation techniques and in the development of some detailed examples for which we obtain ground-truth.

We propose to exploit the CASE analytic environment to evaluate utility towards intelligence analysis. We will collaborate with other contractors to integrate our technologies at appropriate levels into analyst-facing tools. Taking advantage of the NIST evaluations and the CASE Challenge Problems, our technologies will be evaluated by the changes in analyst productivity metrics similar to those used in NIMD, with and without our various technical components enabled. Sample tasks for testing the capabilities of the proposed tools include those in which an analyst is required to evaluate quickly which evidence to use when arguing for a conclusion, or when an analyst is faced with explicitly contradictory intelligence reports. We would expect our tools to have access to the same data as the analysts can reach, as well as any products that the analysts return to the environment.

We propose to use logical theories from case studies developed by us, other program participants, activities such as IKRIS, and other parties. We will also use open source data, especially for parameter estimation. We will actively seek out additional data sources, especially those most useful for further evaluation and for technology transfer.

1.1.6 Deliverables

The project will deliver:

- Design documents and code for successive versions of the Magellan system for logical and Bayesian reasoning about hypotheses and evidence. We plan to deliver a new version of the software every six months. Regular software releases will be packaged with detailed technical documentation; user documentation, libraries, sources and executable code, and delivered to the CASE-enabled analytic environment.
- Management reports as required by the program office. We expect to provide monthly and quarterly reports in softcopy. Monthly expenditure and technical status reports will be submitted via the JIFFY system or by any means reasonably requested by the Government.
- Presentations and demonstrations at bi-annual PI meetings and site visits.
- Successively more robust and sophisticated versions of the BALER reasoning system. These will be made available to other project teams in the CASE program, to evaluators at NIST and PNNL, and to analysts.

1.2 Technical Program Summary

We propose to develop computationally feasible technology to support optimal management of uncertainty, inconsistency and disagreement in collaborative intelligence analysis by importing semantically guided proof search techniques in development at HNC into Bayesian network techniques in development at USC. Although they enjoy certain advantages in versatility and computational complexity, logical knowledge bases are ill-suited to represent uncertainty and then reason about it correctly, because knowledge representation languages based on classical logic do not provide facilities for representing and reasoning about uncertainty expressed in a probabilistic form.

Although such facilities can be provided by extending the logical framework to support such representations as multiple-entity Bayesian networks and probabilistic relational models, the scalability of such approaches is problematic [Jaeger 2000]. We propose to overcome this problem by developing approximate methods that automatically convert logical proofs into Bayesian networks. HNC techniques will guide and accelerate proof search, while the Bayesian network will then be used to reason about the uncertainty of data sources and judgments and conflicting data and judgments. We will avoid the problems of catastrophic failure in classical logic engines due to conflicting data, while also avoiding the computational complexity of default and nonmonotonic logic approaches.

The results of our effort will be a logical and probabilistic reasoning system that (1) can be incorporated into other CASE projects, (2) will be used, demonstrated, and validated within USC's Magellan system for generating and evaluating hypotheses, and (3) will enable analyst teams to collaborate on large-scale tasks. We will evaluate the reasoning system on realistic intelligence problems and, uniquely, by using large groups of students at USC acting as teams of novice analysts or, equivalently, analysts working in a domain that is new to them. We expect to be able to demonstrate our claim that the amount of information that an analyst must process is greatly reduced because the combined reasoner will have removed the irrelevant portions, while bringing to the attention of an analyst the most relevant information.

1.2.1 Option 1 (FY'09): Reasoning Extensions and Field Testing

In the first option year, the goal is to provide support for the field testing of our reasoner. To achieve this goal, we will significantly harden the implementation of our prototype in terms of system performance and robustness. We will also extend the evaluation tools and methodology, and develop training and documentation material for outside administrators and users. One important research and development activity is for our system to accommodate the sponsor's real-world data sets and the target testing environment.

Major research challenges include the definition and implementation of the evaluation tools and methodologies as well as the use of real-world data sets. The modular nature of our reasoner supports efficient integration with other tools, selective hardening of the software, and easy extension or transition to other data sources.

1.2.2 Option 2 (FY'10): CASE Integration and Technology Insertion

We devote the final year of the CASE program to technology insertion and transition activities. To achieve this goal, our work will focus on system tuning and configuring to maximize performance and usability. The development of new software or data is restricted to the implementation of new tools for the evaluation of the reasoning prototype and tuning of this software. We will promote it to RDEC and any other potential transition partner. We will aggressively advertise its newly developed capabilities in the intelligence, defense, and homeland-security communities. Our strategy for marketing the capabilities of the logic plus Bayesian reasoner is to raise its attractiveness by lowering its deployment threshold and increasing its usability.

Research challenges include (1) collecting empirical and formal rules and metrics to guide the tuning and deployment of the reasoning prototype, and (2) exploring new application areas for our technology.

1.3 Risk Analysis and Alternatives

The argumentation DAGs corresponding to the proofs generated by the HNC logic summarization engine might not exhibit the independence properties needed for belief update, efficient assessment of probabilities, sensitivity analysis, and value-of-information computations, as needed by the Bayesian reasoner. A secondary risk is that the concepts identified within the HNC proofs might not match events in the probability space of the analytical task under consideration. These risks are reduced by the use of Bayesian network fragments with a rich vocabulary of attributes and the redundancy afforded by the availability of multiple overlapping fragments.

1.4 References

Arlo-Costa, Horacio and Isaac Levi. "A complete characterization of a notion of contraction based on information-value", *Proceedings of the Tenth International Workshop on Non-Monotonic Reasoning (NMR2004)*, Whistler, BC, Canada, 25-40, 2004.

Bangsø, Olav and Pierre-Henri Wuillemin. "Object-Oriented Bayesian Networks: A Framework for Top-Down Specification of Large Bayesian Networks and Repetitive Structures." Technical Report CIT-87.2-00-obphw1, Department of Computer Science, University of Aalborg, Denmark, 2000.

Bezdek, James. "Special Issue on Belief Function Revisited: Questions and Answers." *International Journal of Approximate Reasoning*, 3, 3 May, 1992.

Bodlaender, Hans. "Treewidth: Algorithmic Techniques and Results." *Proceedings of the 22nd International Symposium on Mathematical Foundations of Computer Science* (MFCS'97), Lecture Notes in Computer Science, volume 1295, 1998.

Bodlaender, Hans. "Discovering Treewidth," Technical Report UU-CS_2005-018, Institute of Information and Computing Sciences, Utrecht University, 2005.

Bloemeke, Mark and Marco Valtorta. "A Hybrid Algorithm to Computer Marginal and Joint Beliefs in Bayesian Networks and Its Complexity." *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence* (UAI-98), Madison, Wisconsin, pp. 208-214, 1998.

Chan, Hei and Adnan Darwiche. "On the Revision of Probabilistic Beliefs Using Uncertain Evidence." *Artificial Intelligence*, 163, 67-90, 2005.

Cheng, John, R. Emami, L. Kerschberg, E. Santos, Jr., Q. Zhao, H. Nguyen, H. Wang, M. Huhns, M. Valtorta, J. Dang, H. Goradia, J. Huang, and S. Xi. "OmniSeer: A Cognitive Framework for User Modeling, Reuse of Prior and Tacit Knowledge, and Collaborative Knowledge Services." Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS38), Big Island, Hawaii, January 3-6, 2005.

Dechter, Rina. "Bucket Elimination: A Unifying Framework for Probabilistic Inference." *Proc. 12th International Conf. on Uncertainty in Artificial Intelligence (UA196),* Portland, OR, 1996, pp. 211-219, 1996.

deKleer and Williams. "Diagnosis with Behavioral Modes." *Artificial Intelligence*, 32, 97-130, 1987.

Getoor, Lise, Nir Friedman, Daphne Koller, and Benjamin Taskar. "Learning probabilistic models of Relational Structure." *Journal of Machine Learning Research*, 3, 679-707, 2002.

Hajek, Petr, Tomas Havranek, and Radim Jirousek. Uncertain Information Processing in Expert Systems. Boca Raton, Florida: CRC Press, 1992.

Halpern, Joseph. Reasoning about Uncertainty. Cambridge, MA: Massachusetts University Press, 2003.

Heuer, Richards. *Psychology of Intelligence Analysis*. ISBN 1 929667-00-0 (also available at http://www.cia.gov/csi/books/19104/index.html) Center for the Study of Intelligence, 1999.

Hughes, Francis J. and David A. Schurn. "Case Study #4: The Sign of the Crescent." Joint Military Intelligence College, Defense Intelligence Agency, 2003.

Jaeger, Manfred. "On the complexity of inference about probabilistic relational models**."** *Artificial Intelligence*, 117, 297-308, 2000.

Jaeger, Manfred. "Relational Bayesian Networks: a Survey." *Electronic Transactions in Artificial Intelligence*, 6 , 2002

Jensen, Finn V. Bayesian Networks and Decision Graphs. Springer, 2001.

Jensen, Finn V., Uffe Kjaerulff, B. Christiansen, Helge Langseth, Chris Skaanning, Marta Vomlelova, and Jiri Vomlel. "The SACSO Methodology for Troubleshooting Complex Systems." Special Issue on AI in Equipment Service, *Artificial Intelligence for Engineering Design, Analysis and Manufacturing (AIEDAM),* 15, 321-333, 2001.

Kim, Young-Gyun, Marco Valtorta, and Jiří Vomlel. "A Prototypical System for Soft Evidential Update." *Applied Intelligence*, 21, 81-97, 2004.

Laskey, Katherine and Suzanne Mahoney. "Network Fragments: Representing Knowledge for Constructing Probabilistic Models." *Proceeding of the Thirteenth International Conference on Uncertainty in Artificial Intelligence (UAI-97)*, Providence, Rhode Island, 334-341, 1997.

Laskey, Katherine and Paulo C.G. da Costa. “Of Starships and Klingons: Bayesian Logic for the 23rd Century.” *Proceedings of the Twenty-First International Conference on Uncertainty in Artificial Intelligence (UAI-05)*, Edinburgh, Scotland, 346-353, 2005.

Lauritzen, Steffen. “Causal Inference from Graphical Models.” In O.E. Barndorff-Nielsen, C. Klueppelberg (eds.), *Complex Stochastic Systems*. London, UK: Chapman and Hall, 63-107, 2001.

Levi, Isaac. *The Fixation of Belief and Its Undoing: Changing Beliefs through Inquiry.* Cambridge University Press, Cambridge, England , 1991.

Nielsen, Thomas and Finn V. Jensen. “Learning a Decision Maker’s Utility Function from (possibly) inconsistent behavior.” *Artificial Intelligence*, 160, 53-78, 2004.

Poole, David. “First-order Probabilistic Inference,” *Proceedings of IJCAI-03*. Acapulco, August 2003, pp. 985-991, 2003.

Pearl, Judea. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, California: Morgan Kaufmann, 1988.

Pearl, Judea. *Causality, Modeling, Reasoning, and Inference*. Cambridge, UK: Cambridge Univ. Press, 2000.

Reiter, Raymond. “Default Logic.” *Artificial Intelligence*, 13, 81-132, 1980.

Reiter, Raymond. “A Theory of Diagnosis from First Principles.” *Artificial Intelligence*, 32, 57-96, 1987.

Shenoy, Prakash and Gautam Biswas, eds. “Special Issue on Belief Functions and Belief Maintenance in Artificial Intelligence.” International Journal of Approximate Reasoning, 4, 5/6 (September/November), 1990.

Valtorta, Marco, Young-Gyun Kim, and Jiří Vomlel. “Soft Evidential Update for Probabilistic Multiagent Systems.” *International Journal of Approximate Reasoning*, 29, 71-106, 2002.

Van Horn, Kevin S. 2003. “Constructing a Logic of Plausible Inference: A Guide to Cox’s Theorem.” *International Journal of Approximate Reasoning*, 34, 1 (September 2003), pp.3-24.

Vomlel, Jiří. “Methods of Probabilistic Knowledge Integration.” Ph.D. Thesis. Faculty of Electrical Engineering, Czech Technical University, 1999.

Vomlel, Jiří. “Probabilistic Reasoning with Uncertain Evidence.” Neural Network World, International Journal on Neural and Mass-Parallel Computing and Information Systems, 5, 453-465, 2004.

Vomlel, Jiří. “Noisy-OR Classifier.” *International Journal of Intelligent Systems*, 21, 3, 381-398, 2006.

Wellman, Michael and Max Henrion. “Explaining ‘Explaining Away’.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15, 3, 287-292, 1993.