

---

# High-Level Information Fusion with Bayesian Semantics

---

**Paulo C. G. Costa, Kathryn Laskey, Kuo-Chu Chang, Wei Sun, Cheol Park, Shou Matsumoto**  
Center of Excellence in C4I & Department of Systems Engineering and Operations Research  
George Mason University  
Fairfax, VA 22030  
(pcosta,klaskey,kchang)@gmu.edu; wsun@c4i.gmu.edu; cparkf@gmu.edu; smatsum2@masonlive.gmu.edu

## Abstract

In an increasingly interconnected world information comes from various sources, usually with distinct, sometimes inconsistent semantics. Transforming raw data into high-level information fusion (HLIF) products, such as situation displays, automated decision support, and predictive analysis, relies heavily on human cognition. There is a clear lack of automated solutions for HLIF, making such systems prone to scalability issues. In this paper, we propose to address this issue with the use of highly expressive Bayesian models, which can provide a tighter link between information coming from low-level sources and the high-level information fusion systems, and allow for greater automation of the overall process. We illustrate our ideas with a naval HLIF system, and show the results of a preliminary set of experiments.

## 1 INTRODUCTION

Information fusion is defined as:

“... the synergistic integration of information from different sources about the behavior of a particular system, to support decisions and actions relating to the system.”<sup>1</sup>

A distinction is commonly made between *low-level* and *high-level* fusion. Low-level fusion combines sensor reports to identify, classify, or track individual objects. High-level fusion combines information about multiple objects, as well as contextual information, to characterize a complex situation, draw inferences about the intentions of actors, and support process refinement.

In current information fusion systems, lower-level data fusion is typically accomplished by stove-piped sys-

tems that feed information directly to human users. Subsequent generation of high-level information fusion (HLIF) products, such as situation displays, automated decision support, and predictive analysis, relies heavily on human cognition. The tacit underlying assumption is that humans are still the most efficient resource for translating low-level fusion products into decision-relevant knowledge. While the current approach works well for many purposes, it cannot scale as the data influx grows. Automated assistance for HLIF tasks is urgently needed to mitigate cognitive overload and achieve the necessary throughput.

Stove-piped systems can be extremely efficient at exploiting a specific technology applied to a limited and well defined set of problems. Air Traffic Control Systems, for instance, employ radar technology in a very effective way to provide reliable situation awareness for radar controllers via sophisticated low-level information fusion (LLIF) techniques. The synthetic radar screen shown to traffic controllers in a sector of an Area Control Center (ACC) fuses multiple radar tracks. Data association algorithms infer whether geographically close signals captured by various radars are coming from a single or multiple aircraft. Despite the sophistication of its low-level fusion components, the ATC system relies heavily on humans for HLIF products. For instance, controllers rely on their own understanding of the overall picture to decide how to drive their tracks; area coordinators rely on their knowledge to decide whether the outbound traffic to a given airport should be redirected due to an upcoming storm; and so on.

The ATC system is a good example of a highly sophisticated stove-piped system that relies on human cognition for its major purpose: to ensure that thousands of airplanes in the US can share the airspace in a safe and effective way. As the volume of airplanes increases, more humans are needed to perform HLIF tasks. After a point, the overhead of transferring between ever-smaller control sectors becomes a major

---

<sup>1</sup>The International Society for Information Fusion, <http://isif.org>

scalability issue. That is, cognitive limitations (each human can control only airplanes at once) together with the added complexity of adding extra cognitive units (a.k.a. traffic controllers) become a major obstacle to growth. This scalability problem is common to HLIF systems in other domains as well.

In this paper, we propose to address the issue with the use of highly expressive Bayesian models. Such systems provide a tighter link between low-level and high-level information fusion systems. Because they are sufficiently expressive to reason about high-level information, they provide a coherent framework for expressing and reasoning with uncertain information that spans low and high level systems.

This paper describes our approach by way of a case study in information fusion for Maritime Domain Awareness. Section 2 motivates the use of explicit probabilistic semantics and explains the main concepts behind our approach. Section 3 introduces the Maritime Domain Ontology we used in our experiments. The experiments are described in Section 4. Section 5 concludes with a discussion.

## 2 Semantics in HLIF

Humans are more effective than computers at gathering various pieces of information and correlating them into a coherent picture, but still have a high error rate. For example, an intelligence analyst can correlate images and videos of a road with observers reports that a convoy has passed in the early afternoon, and conclude that this was the same convoy that participated in a terrorist attack 10 miles down that road. These conclusions are based on an implicit understanding of how trucks and cars are represented in each type of media (video, imagery, human reports), as well as the temporal and spatial relationships between cars, roads, convoys, etc. For a computer program to perform the same inferences from the same set of sources, it must possess the same kind of knowledge. Conveying such knowledge to a computer program requires a means to make the humans tacit knowledge explicit and formal, so it can be retrieved and used when needed.

Ontologies are the current paradigm for specifying domain knowledge in an explicit and formal way. One of the most cited definitions of ontologies is the specification of a conceptualization. [1] To perform automated fusion in the above example, concepts such as cars, roads, convoys, people, etc., as well as their relationships, must be formalized in a way that computers can store, retrieve, and use. Not surprisingly, ontologies have been widely considered in the domain of information fusion as a means to enable automated systems to perform HLIF tasks (e.g. [2, 3, 4]).

Most languages for expressing ontologies, such as the most popular variant of the W3C Recommendation OWL [5], are based on Description Logic [6]. Other ontology languages such as KIF<sup>2</sup> and the ISO Standard Common Logic<sup>3</sup>, are based on first-order logic. Classical logic has no standard means to represent uncertainty. This is a major drawback for HLIF systems, which must operate in environments in which uncertainty is pervasive. Inputs from LLIF systems come with uncertainty, as do the high-level domain relationships that analysts use to draw inferences about a complex situation.

Although LLIF algorithms often have a basis in probability theory, it is common to report results according to a threshold rule without any confidence qualifier. This is often justified by cognitive limitations of human decision makers. As an example, suppose a video analysis report assigns 86% probability of person Joe being inside a car driving towards place A. If the threshold for the input source was 85%, a LLIF system might simply report the statement without qualification, and a HLIF system might treat this as a true statement. Such threshold rules lose uncertainty information. Other information sources, each with its own internal processing and threshold rules, might provide additional reports about Joe, A, and other aspects of the situation relevant to inferences about Joes destination. Without uncertainty qualifiers, it is difficult for the HLIF system to draw sound inferences about Joes destination. Other limitations of HLIF with respect to the handling of uncertainty are discussed within the context of the International Society of Information Fusions working group on Evaluation of Technologies for Uncertainty Reasoning (ETURWG)<sup>4</sup> [7, 8].

Representing uncertainty with ontologies is an active area of research, especially in the area of the Semantic Web (e.g., [9, 10]). HLIF requires reasoning with uncertain information about complex situations with many interacting objects, actors, events and processes. Automating HLIF therefore requires expressive representation formalisms that can handle uncertainty. Probabilistic ontologies [11, 12], extend traditional ontologies to capture both domain semantics and associated uncertainty about the domain. The probabilistic ontology language PR-OWL [12] is based on multi-entity Bayesian Networks [13].

### 2.1 Multi-Entity Bayesian Networks

MEBNs represent the world as a collection of inter-related entities and their respective attributes. Knowl-

<sup>2</sup><http://www-ksl.stanford.edu/knowledge-sharing/kif/>

<sup>3</sup><http://www.iso-commonlogic.org/>

<sup>4</sup><http://eturwg.c4i.gmu.edu/>

edge about attributes of entities and their relationships is represented as a collection of repeatable patterns, known as *MEBN Fragments* (MFragments). A set of MFragments that collectively satisfies constraints ensuring a unique joint probability distribution is a *MEBN Theory* (MTheory).

An MFragment is a parametrized fragment of a directed graphical probability model. It represents probabilistic relationships among uncertain attributes of and relationships among domain entities. MFragments are templates that can be instantiated to form a joint probability distribution involving many random variables. Such a ground network is called a *situation-specific Bayesian network* (SSBN).

MEBN provides a compact way to represent repeated structures in a Bayesian Network. There is no fixed limit on the number of random variable instances, which can be dynamically generated as needed. The ability to form a consistent composition of parametrized model fragments makes MEBN well suited for knowledge fusion applications [14]. MEBN inference can be performed by instantiating relevant MFragments and assembling them into SSBNs to reason about a given situation. As evidence arrives, it is fused into the SSBN to provide updated hypotheses with associated levels of confidence. These are very convenient features for representing diverse information coming from various sensors, which make MEBN attractive as a logical basis for probabilistic ontologies.

## 2.2 PR-OWL Probabilistic Ontologies

There are basically three aspects that must be addressed for a representational and reasoning framework in support of effective higher-level knowledge fusion:

1. A rigorous mathematical foundation,
2. The ability to represent intricate patterns of uncertainty, and
3. Efficient and scalable support for automated reasoning.

Current ontology formalisms deliver a partial answer to items 1 and 3, but lack a principled, standardized means to represent uncertainty. This has spurred the development of palliative solutions in which probabilities are simply inserted in an ontology as annotations (e.g. marked-up text describing some details related to a specific object or property). These solutions address only part of the information that needs to be represented, and too much information is lost to the lack of a good representational scheme that captures structural constraints and dependencies among probabilities. A true probabilistic ontology must be capable of properly representing those nuances. More formally:

Definition 1 (from [11]): A probabilistic ontology (PO) is an explicit, formal knowledge representation that expresses knowledge about a domain of application. This includes:

- Types of entities that exist in the domain;
- Properties of those entities;
- Relationships among entities;
- Processes and events that happen with those entities;
- Statistical regularities that characterize the domain;
- Inconclusive, ambiguous, incomplete, unreliable, and dissonant knowledge related to entities of the domain; and
- Uncertainty about all the above forms of knowledge;

where the term entity refers to any concept (real or fictitious, concrete or abstract) that can be described and reasoned about within the domain of application. ■

POs provide a principled, structured, sharable formalism for describing knowledge about a domain and the associated uncertainty and could serve as a formal basis for representing and propagating fusion results in a distributed system. They expand the possibilities of standard ontologies by introducing the requirement of a proper representation of the statistical regularities and the uncertain evidence about entities in a domain of application. POs can be implemented using PR-OWL<sup>5</sup>, a Probabilistic Web Ontology Language that extends OWL with constructs for expressing first-order Bayesian theories. PR-OWL structures map to MEBN structures, so PR-OWL provides a means to express MEBN theories in OWL.

## 2.3 The UnBBayes MEBN/PR-OWL Plugin

In order to develop and use POs, we have developed a MEBN/PR-OWL plugin to the graphical probabilistic package UnBBayes<sup>6</sup>, an open source, Java<sup>TM</sup>-based application developed at the University of Brasilia. The plugin provides both a GUI for building probabilistic ontologies and a reasoner based on the MEBN/PR-OWL framework [15, 16]. Reasoning in the UnBBayes MEBN/PR-OWL plugin involves SSBN construction, which can be seen type of proposition-alization, and the subsequent inferential process over the resulting SSBN. Figure 1 shows a screenshot of the UnBBayes MEBN/PR-OWL plugin.

<sup>5</sup><http://www.pr-owl.org>

<sup>6</sup><http://unbbayes.sourceforge.net>

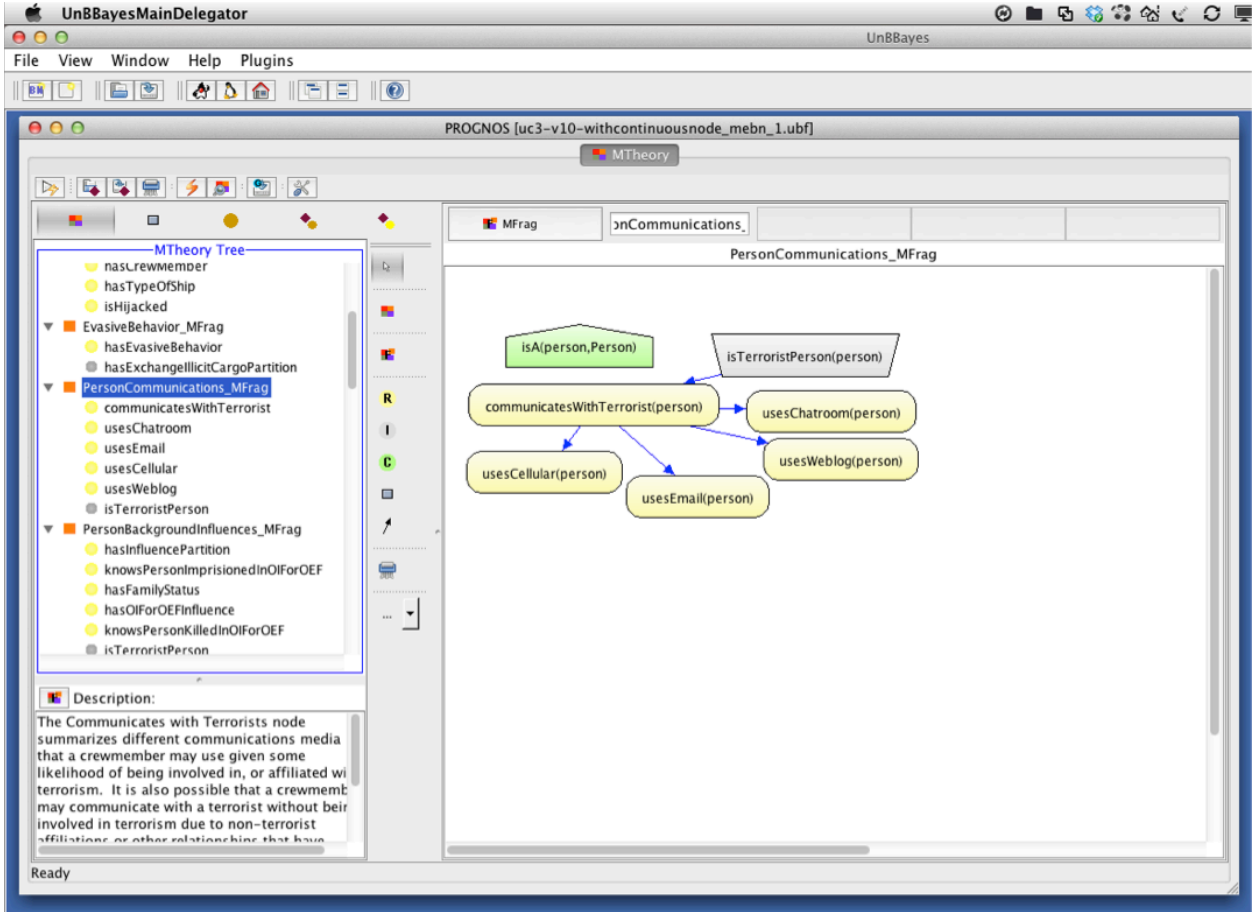


Figure 1: The UnBBayes MEBN/PR-OWL plugin

Many HLIF problems involve spatio-temporal entities, and require reasoning with discrete and continuous, possibly non-Gaussian, variables. To support this requirement, a capability for hybrid MTheories was added to the UnBBayes MEBN/PR-OWL plugin. The plugin can handle MTheories in which continuous variables can have discrete or continuous parents, but no discrete variable is allowed to have a continuous parent. To specify hybrid models, constructs were added to the local distribution scripting language for continuous distributions. The SSBN construction algorithm for building the ground model is basically unchanged except that local distributions can be continuous.

For inference in a hybrid SSBN, the plugin implements the direct message passing (DMP) algorithm [17] to compute, propagate, and integrate messages. DMP combines the unscented transformation [18] and the traditional message-passing algorithm to deal with arbitrary, not necessary linear, functional relationships between continuous variables in the network. DMP gives exact results for polytree conditional linear Gaussian (CLG) networks and approximate results for networks with loops (via loopy propagation), networks with non-linear relationships (via the unscented transformation) and networks with non-Gaussian variables

(via mixtures of Gaussians). Mixtures of Gaussian distributions are used to represent continuous messages. The number of mixture components can be as large as the size of the joint state space of all discrete parents. To achieve scalability, the algorithm can restrict the number of mixture components in the messages to satisfy a predefined error bound [19].

Specifically, without loss of generality, suppose a typical hybrid model involving a continuous node  $X$  with a discrete parent node  $D$  and a continuous parent node  $U$ . As shown in Figure 2, messages sent between these nodes are: (1)  $\pi$  message from  $D$  to  $X$ , denoted as  $\pi_X(D)$ ; (2)  $\pi$  message from  $U$  to  $X$ , denoted as  $\pi_X(U)$ ; (3)  $\lambda$  message from  $X$  to  $D$ , denoted as  $\lambda_X(D)$ ; and (4)  $\lambda$  message from  $X$  to  $U$ , denoted as  $\lambda_X(U)$ .




Figure 2: Example hybrid model

In general, for a polytree network, any node  $X$   $d$ -separates evidence into  $\{\mathbf{e}^+, \mathbf{e}^-\}$ , where  $\mathbf{e}^+$  and  $\mathbf{e}^-$  are evidence from the sub-network “above”  $X$  and “below”  $X$  respectively. The  $\lambda$  and  $\pi$  message maintained in each node are defined as,

$$\lambda(X) = P(\mathbf{e}_X^- | X) \quad (1)$$

and

$$\pi(X) = P(X | \mathbf{e}_X^+) \quad (2)$$

With the two messages, it is straightforward to see that the belief of a node  $X$  given all evidence is just the normalized product of  $\lambda$  and  $\pi$  values, namely,

$$\begin{aligned} BEL(X) &= P(X | \mathbf{e}) = P(X | \mathbf{e}_X^+, \mathbf{e}_X^-) \\ &= \alpha \lambda(X) \pi(X) \end{aligned} \quad (3)$$

where  $\alpha$  is a normalizing constant. It can be shown that for a hybrid network, the  $\pi$  message can be recursively computed as,

$$\begin{aligned} \pi(X) &= \sum_D \int_U P(X | D, U) \pi_X(D) \pi_X(U) dU \\ &= \sum_D \left[ \pi_X(D) \int_U P(X | D, U) \pi_X(U) dU \right] \end{aligned} \quad (4)$$

where the integral of  $P(X | D = d, U) \pi_X(U)$  over  $U$  is equivalent to a functional transformation of  $\pi_X(U)$ , which is a continuous message in the form of a Gaussian mixture.

Similarly, the  $\lambda$  message for the discrete parents can be obtained as

$$\lambda_X(D = d) = \int_X \lambda(X) \int_U P(X | D = d, U) \pi_X(U) dU dX \quad (5)$$

where  $\int_U P(X | D = d, U) \pi_X(U) dU$  is a functional transformation of a distribution over  $U$  to  $X$ .

On other hand, the  $\lambda$  message for continuous parent  $U$  can be computed as

$$\begin{aligned} \lambda_X(U) &= \int_X \lambda(X) \sum_D P(X | D, U) \pi_X(D) \pi_X(D) dX \\ &= \sum_D \left[ \pi_X(D) \int_X \lambda(X) P(X | D, U) dX \right] \end{aligned} \quad (6)$$

Equations (3) to (6) form a baseline for computing direct messages between mixed variables.

As mentioned earlier, with the unscented transformation, this method can be modified for arbitrary non-linear non-Gaussian hybrid models. In addition, the algorithm is scalable by combining the mixture components in the messages with any given error bound

### 3 Maritime Domain Awareness PO

In 2008, the Department of Defense issued a directive to establish policy and define responsibilities for Maritime Domain Awareness (MDA).<sup>7</sup> The directive defines MDA as the “effective understanding of the global maritime domain and its impact on the security, safety, economy, or environment of the United States.” The ability to automatically integrate information and recommendations from multiple intelligence sources in a complex and ever-changing environment to produce a dynamic, comprehensive, and accurate battlespace picture is a critical capability for MDA. This section reports on a prototype probabilistic ontology for maritime domain awareness (MDA-PO). This model, which is depicted in Figure 3, was developed as part of the PROGNOS project [20, 21], with the assistance of two retired Navy officers who served as subject-matter experts.

Figure 4 depicts one of the MFrag of the MDA-PO, the *AggressiveBehavior* MFrag. As the name implies, this is a chunk of knowledge that captures some of the concepts and relationships that are useful to infer whether a ship is displaying aggressive behavior. The three different types of MFrag nodes can be seen: *Context*, *Input*, and *Resident* nodes.

*Resident* nodes are the random variables that form the core subject of an MFrag. The MFrag defines a local distribution for each resident node as a function of the parents of the resident node in the fragment graph. They can be discrete or continuous. There are three discrete nodes in this MFrag, which are depicted as yellow rounded rectangles in the picture, and five continuous nodes, depicted as rounded rectangles with double lines.

As an example of how the representation works, reports on the propeller turn count of a ship will be an indicator of whether the ship speed is changing or not. Also, there will be different probability distributions for *speedChange(ship)* if the ship is behaving aggressively or not (i.e. if the state of node *hasAggressiveBehavior(ship)* is true or false).

*Input* nodes, depicted as gray trapezoids in the figure, serve as “pointers” referring to resident nodes in other MFrag. Input nodes influence the local distributions of resident, but their own distributions are defined in the MFrag in which they are resident.

In a complete MTheory, every input node must point to a resident node in some MFrag. For instance, the *hasBombPortPlan(ship)* input node influences the distribution of all the *hasAggressiveBehavior(ship)* nodes

<sup>7</sup> www.dtic.mil/whs/directives/corres/pdf/200502p.pdf

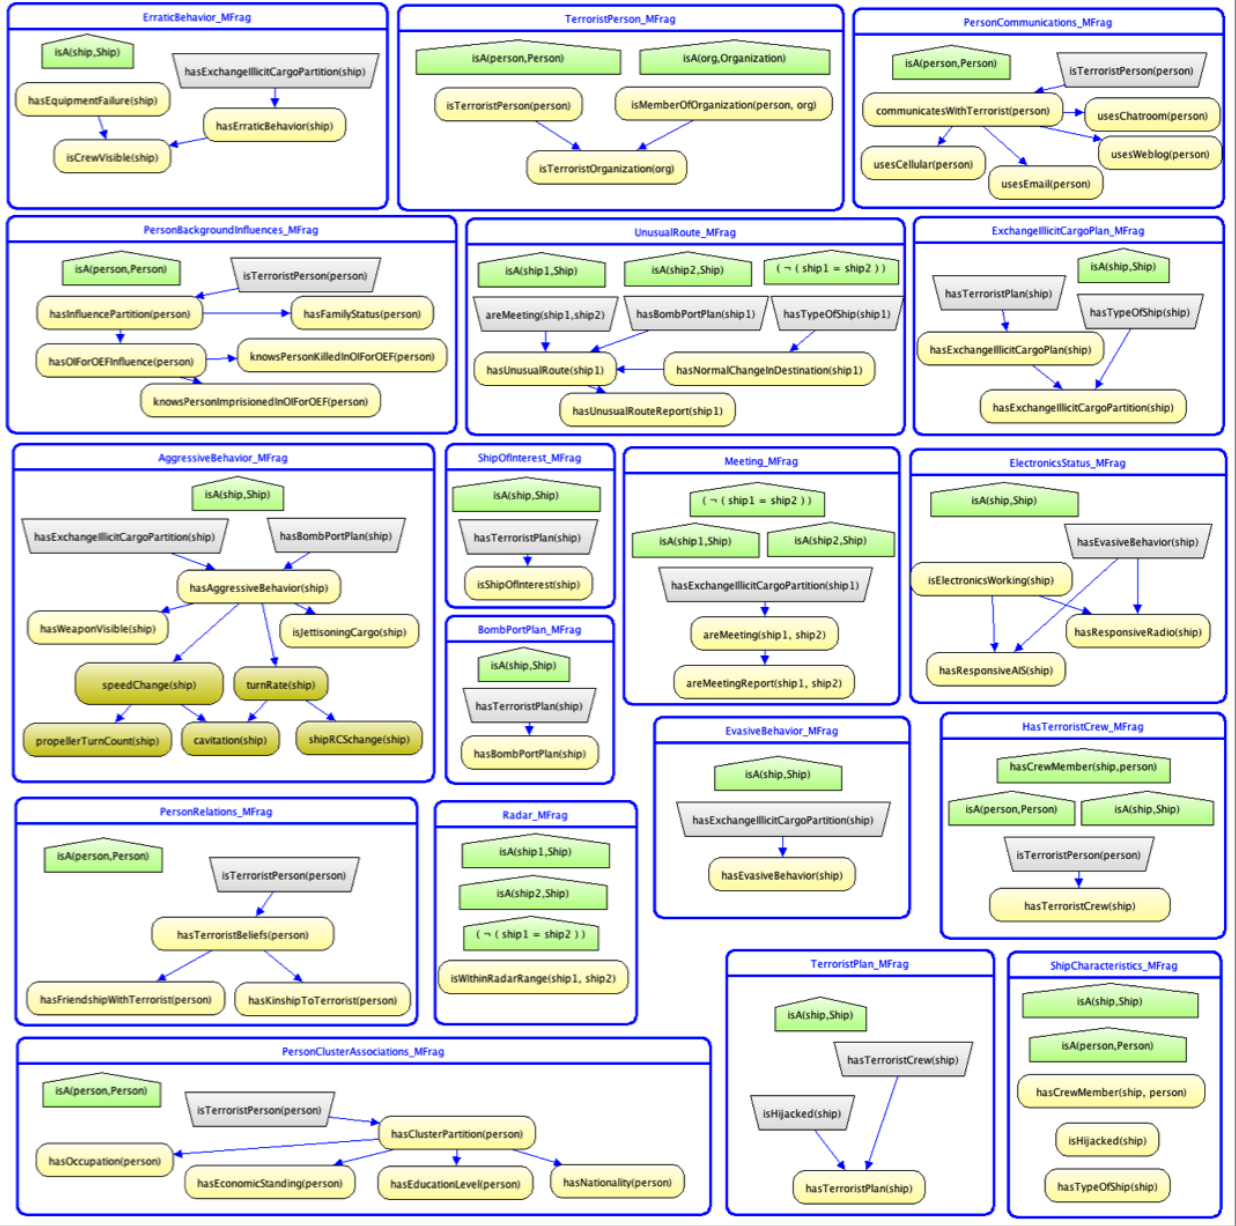


Figure 3: The MDA-PO

that would be instantiated in an SSBN construction process.

*Context* nodes are Boolean (i.e., true/false) random variables representing conditions that must be satisfied for the probability distribution of an MFrag to apply. Like input nodes, context nodes also have distributions defined in other MFraGs.

By allowing uncertainty on context nodes, MEBN can represent several types of sophisticated uncertainty patterns, such as relational uncertainty or existence uncertainty. There is only one context node in the *AggressiveBehavior* MFrag, seen in the figure as a green pentagon.

The MDA-PO is described in detail in [22]. In PROG-NOS, the MDA-PO was also used to build the model used to run the test and evaluation process, which we explain in the next Section.

## 4 Experimental Results

The main objective of the set of experiments presented in this paper was to assess the accuracy, scalability, and overall performance of the SSBN construction and DMP algorithms combined. As a benchmark, we used the UnBBayes implementation of the Junction Tree (JT) algorithm, which is a well-known belief propagation method for Bayesian networks [23] and the fo-




Figure 4: The Aggressive Behavior MFRag

cus of various efforts on algorithm optimization (e.g. [24, 25]).

#### 4.1 Setup and Metrics

Obtaining a real data set for maritime HLIF was not an option for our team. Therefore, we generated various synthetic datasets through an agent-based simulation module, depicted in Figure 5. The module generates simulated scenarios, including entities (e.g., ships, people) and their features. The simulated scenarios serve as the ground truth for evaluating performance. The simulation module also generates reports of the kind the eventual operational system is expected to receive, thus exercising the interfaces and the reasoning module in a realistic manner [21].

The simulation is based on maritime activities (regular and suspicious) with the objective of prevention and disruption of terrorist attacks, sabotage, espionage, or subversive acts. Therefore, the agents on the simulation tool simulate commercial, fishing, recreational, and other types of ships in their normal and suspicious behaviors. Suspicious behaviors are characterized by ships that do not follow their regular or most probable routes according to their origin and destination, by ships that meet in the middle of the ocean for no apparent reason, etc.

For the experiments, the simulation engine generated 9 types of scenarios with different combinations of the number of ships and the associated entities (e.g. organizations, people, etc.). The maximum size of the dataset was limited to 10K entities. After generating each dataset, we added noise as to assess robustness to model misspecification. The scenario for the experiments emulates a U.S. Navy destroyer conducting Maritime Security Operations in a relatively busy area. This is a “needle in a haystack” type of problem, in which the destroyer has information about dozens of ships within a certain radius, but can only verify a few

of them. In this case, the HLIF system must integrate the data coming from the ship sensors with information coming from other sources, such as intelligence reports, signals intelligence, HUMINT, and others. We emulate this in the experiments with “area queries” in which all ships within a 60NM radius are queried by the system. More precisely, information on all  $n$  known ships within that radius trigger the instantiation of MFragS storing pertinent knowledge, including the one containing the  $shipOfInterest(ship)$  node. The system would then query the  $n$   $shipOfInterest(ship)$  nodes that were instantiated.

All experiments were performed in a dedicated computer with an Intel quad-core i7<sup>TM</sup> processor with 8 GB of RAM, and running MS Windows 7<sup>TM</sup> 64bit.

Accuracy is assessed by the quadratic scoring rule [26]:

$$B(r, i) = \sum_{j=1}^C (y_j - r_j)^2$$

where  $y_j = 1$  when the  $j^{th}$  event is correct and 0 otherwise.  $C$  is the number of classes. This is a proper scoring rule, i.e., the score is minimized when the assessed probability is equal to the actual frequency.

To assess scalability, we measured the computation time as a function of the generated SSBN size and the number of ships involved in a query.

#### 4.2 Preliminary Results

The results for accuracy are depicted in Table 1. From the obtained scores, it is clear that the Hybrid system performed better in capturing both the cases in which the  $shipOfInterest(ship)$  node state was true ( $\approx 10.36\%$  better) in the ground truth, as well as those in which the node state was false ( $\approx 14.02\%$  better).

Table 1: Results for Accuracy

Queries on the “Ship of Interest” node	Hybrid System	Discrete System
Ship of Interest = true (ground truth)	0.88203	0.79947
Ship of Interest = false (ground truth)	0.89439	0.78439

These results are consistent with expectations, given the inherent inaccuracies in discretizing the continuous random variables in the MDA-PO. The 10 to 14% improvement from the hybrid with respect to the discrete model was consistent all over the 9 datasets. However, since the datasets were all generated from the same model, it is difficult to assess robustness with this run of experiments. In any case, more complex relationships between nodes are likely to increase the difference




Figure 5: PROGNOS simulation module

in accuracy between the JT and the DMP systems.

Figure 6 below shows the results of the area query experiments. The x-axis conveys the number of nodes generated by each query, which tends to be correlated with the number of ships. However, it was not uncommon to see a few ships generating a large network or vice-versa. The y-axis depicts query time in milliseconds.

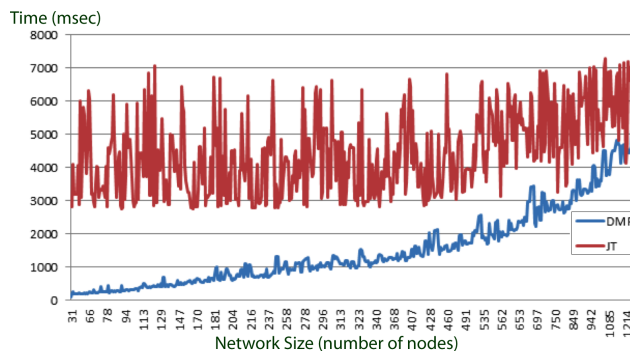


Figure 6: Area Query Time vs. Network Size

Regarding performance, most of the generated networks were between 100 and 500 nodes, and generally yielded a query time below 2 seconds for DMP and 5 seconds for JT. The maximum query time for JT was 7.3 seconds, while the DMP system worst case was 6.4 seconds for a query. The results also show lower variance for DMP query times for a given network size.

Figure 7 shows results for query time vs. number of ships within the 60NM area.

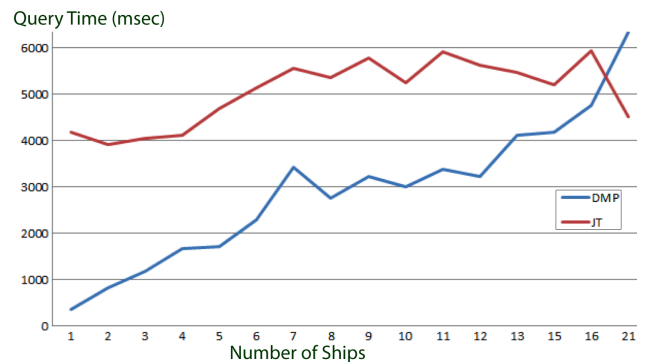


Figure 7: Query Time vs. Number of Ships

This was a different run of experiments, in which the focus was on keeping a controlled number of ships within the query area. This allowed an assessment of how each system reacted to the controlled increase in that number. The performance of JT stayed relatively steady, while the DMP system performed much better for simpler problems but approached the performance of JT when the number of ships was above twenty.

These results were also expected, since our implementation of DMP was not as optimized as the JT implementation in UnBBayes. More specifically, the DMP algorithm was initially implemented in MATLAB™, and the translation to Java™ was not tuned for performance. Yet, the graph suggests a linear increase in the range considered.



## 5 Discussion

The experiments were meant to simulate an HLIF system within a relatively simple scenario. In spite of the overall size of the experiments and the fact that it was conducted within a controlled environment, the performance figures are promising. There remain many ways to improve the efficacy of the algorithm. As previously mentioned, the main objective of the testing and evaluation was to assess the gains in accuracy, which clearly lived up to our expectations.

The results also show promise for the feasibility of using probabilistic ontologies as a driver for HLIF systems. Our future steps towards this goal are to continue the optimization of the algorithms, and to seek out new forms of knowledge acquisition techniques. The latter involves automated learning, which has been the subject of our latest research efforts. We also plan to address the research on rare events, and to work with other datasets.

### Acknowledgements

The PROGNOS project was partially funded by the Office of Naval Research. The authors acknowledge Richard Haberlin and Michael Lehocky who served as subject-matter experts for developing the MDA-PO, and Rommel Carvalho, for his various contributions to the PROGNOS project.

### References

- [1] T. R. Gruber, "A translation approach to portable ontology specifications," *Knowledge acquisition*, vol. 5, no. 2, pp. 199–200, 1993.
- [2] D. McGuinness, "Ontologies for information fusion," in *Information Fusion, 2003. Proceedings of the Sixth International Conference of*, vol. 1, pp. 650 – 657, 2003.
- [3] E. G. Little and G. L. Rogova, "Designing ontologies for higher level fusion," *Information Fusion*, vol. 10, pp. 70–82, Jan. 2009.
- [4] E. Blasch, E. Dorion, P. Valin, E. Bosse, and J. Roy, "Ontology alignment in geographical hard-soft information fusion systems," in *Information Fusion (FUSION), 2010 13th Conference on*, pp. 1 –8, July 2010.
- [5] P. F. Patel-Schneider, P. Hayes, and I. Horrocks, "OWL web ontology language semantics and abstract syntax." <http://www.w3.org/TR/owl-semantics/>, Feb. 2004. W3C Recommendation.
- [6] F. Baader, I. Horrocks, and U. Sattler, "Description logics as ontology languages for the semantic web," in *Mechanizing Mathematical Reasoning*, pp. 228–248, 2005.
- [7] E. Blasch, J. Llinas, D. Lambert, P. Valin, S. Das, C. Chong, M. Kokar, and E. Shahbazian, "High level information fusion developments, issues, and grand challenges: Fusion 2010 panel discussion," in *Information Fusion (FUSION), 2010 13th Conference on*, pp. 1 –8, July 2010.
- [8] P. C. G. Costa, R. N. Carvalho, K. B. Laskey, and C. Y. Park, "Evaluating uncertainty representation and reasoning in HLF systems," in *Proceedings of the Fourteenth International Conference on Information Fusion*, (Chicago, Illinois, USA), July 2011.
- [9] K. Laskey and K. Laskey, "Uncertainty reasoning for the world wide web: Report on the URW3-XG incubator group," URW3-XG, W3C, 2008.
- [10] L. Predoiu and H. Stuckenschmidt, "Probabilistic extensions of semantic web languages - a survey," in *The Semantic Web for Knowledge and Data Management: Technologies and Practices*, Idea Group Inc, 2008.
- [11] P. C. G. Costa, *Bayesian semantics for the Semantic Web*. PhD dissertation, George Mason University, Fairfax, VA, USA, July 2005. Brazilian Air Force.
- [12] P. C. G. Costa and K. B. Laskey, "PR-OWL: a framework for probabilistic ontologies," in *Proceedings of the International Conference on Formal Ontology in Information Systems (FOIS 2006)* (B. Bennet and F. Christiane, eds.), vol. 150 of *Frontiers in Artificial Intelligence and Applications*, (Baltimore, MD, USA), pp. 237–249, IOS Press, Nov. 2006.
- [13] K. B. Laskey, "MEBN: a language for first-order Bayesian knowledge bases," *Artificial Intelligence*, vol. 172, pp. 140–178, Feb. 2008.
- [14] P. C. G. Costa, K. Chang, K. B. Laskey, and R. N. Carvalho, "High level fusion and predictive situational awareness with probabilistic ontologies," (George Mason University, Fairfax, VA, USA), May 2010.
- [15] R. N. Carvalho, L. L. Santos, M. Ladeira, and P. C. G. Costa, "A GUI tool for plausible reasoning in the semantic web using MEBN," (Los Alamitos, CA, USA), pp. 381–386, IEEE Computer Society, Oct. 2007.

- [16] P. Costa, M. Ladeira, R. N. Carvalho, K. Laskey, L. Santos, and S. Matsumoto, "A first-order Bayesian tool for probabilistic ontologies," in *Proceedings of the Twenty-First International Florida Artificial Intelligence Research Society Conference (FLAIRS 2008)*, (Coconut Grove, FL, USA), pp. 631–636, AAAI Press, May 2008.
- [17] W. Sun and K. Chang, "Direct message passing for hybrid Bayesian network and its performance analysis," in *SPIE Defense and Security Symposium*, (Orlando, FL, USA), Apr. 2010.
- [18] S. J. Julier, "The scaled unscented transformation," in *Proceedings of the American Control Conference*, vol. 6, p. 45554559, 2002.
- [19] H. Chen, K. Chang, and C. J. Smith, "Constraint optimized weight adaptation for gaussian mixture reduction," in *SPIE Defense and Security Symposium*, (Orlando, FL, USA), Apr. 2010.
- [20] P. C. G. Costa, K. B. Laskey, and K. Chang, "PROGNOS: applying probabilistic ontologies to distributed predictive situation assessment in naval operations," in *Proceedings of the Fourteenth International Command and Control Research and Technology Conference*, (Washington, DC, USA), CCRP Publications, June 2009. Best paper award of the Collaborative Technologies for Network-Centric Operations Track.
- [21] R. N. Carvalho, P. C. G. Costa, K. B. Laskey, and K. Chang, "PROGNOS: predictive situational awareness with probabilistic ontologies," (Edinburgh, UK), July 2010.
- [22] R. Carvalho, R. Haberlin, P. Costa, K. Laskey, and K. Chang, "Modeling a probabilistic ontology for maritime domain awareness," in *Information Fusion (FUSION), 2011 Proceedings of the 14th International Conference on*, pp. 1–8, July 2011.
- [23] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1 ed., Sept. 1988.
- [24] L. Zheng, O. J. Mengshoel, and J. Chong, "Belief propagation by message passing in junction trees: Computing each message faster using GPU parallelization," *Proc. of the 27th Conference on Uncertainty in Artificial Intelligence (UAI-11)*, 2011.
- [25] A. L. M. F. V. Jensen, A. L. Madsen, A. L. Madsen, F. V. Jensen, and F. V. Jensen, "Lazy propagation in junction trees," in *In Proc. 14th Conf. on Uncertainty in Artificial Intelligence*, p. 362369, Morgan Kaufmann Publishers, 1998.
- [26] G. W. Brier, "Verification of forecasts expressed in terms of probability," *Monthly Weather Review*, vol. 78, pp. 1–3, Jan. 1950.