

MEBN LOGIC: A KEY ENABLER FOR NETWORK CENTRIC WARFARE

Student Paper

Modeling and Simulation

Paulo Cesar G da Costa*^[1]

Francis Fung^[2]

Kathryn B. Laskey^[1]

Michael Pool^[3]

Masami Takikawa^[2]

Edward J. Wright^[3]

[1] George Mason University
System Engineering & Operations Research Department
4400 University Drive
Fairfax, VA 22030-4444
(703) 993-1644
[pcosta; klaskey]@gmu.edu

[2] Information Extraction and Transport, Inc.
1600 SW Western Blvd., Suite 300
Corvallis, OR 97333
[fung; takikawa]@iet.com

[3] Information Extraction and Transport, Inc.
1911 N. Ft. Myer Dr., Suite 600
Arlington, VA 22209
[mpool; ewright]@iet.com

Point of Contact: Paulo C G da Costa
pcosta@gmu.edu and/or (703) 426-4565

Abstract

Among the lessons learned from recent conflicts stands the dramatic change in the very way wars are fought. There are no more clear-cut enemies or allies; rules of engagement have become increasingly fuzzy; guerrilla and insurgent tactics are now commonplace: in short, the battlespace is a very different place from what it used to be. Furthermore, advances in sensor technology and network computing have brought a new element to the complex equation of warfare: information overload. Nowadays, instead of merely gathering information and displaying assets, command and control systems must be able to fill the gap between the glut of information arriving from a networked grid of sensors and the capacity of human commanders to make sense of it. In short, the quest today is for systems that work under the knowledge paradigm. Systems must automatically provide decision makers with a clear picture of what is happening, how it relates to the current situation, and what are the options and their respective consequences. Facing this challenge with technologies of the past is a recipe for failure. New, more powerful approaches are needed. The objective of this paper is to argue for two claims: (1) Bayesian decision theory is an appropriate technology for modeling human decision-making in complex, ambiguous scenarios; and (2) Bayesian reasoning technology is a promising enabler for Network Centric Warfare. To support both claims, we have applied Multi-Entity Bayesian Networks (MEBN) to model a historical tactical decision from the naval domain. MEBN is a breakthrough Bayesian reasoning system in which complex probabilistic models are constructed from modular components that can be replicated and combined in an infinite variety of ways. MEBN allows models to capture important and subtle aspects of objects and their interrelationships that would be impossible to model using existing technologies. We provide a brief overview of modeling in MEBN and then present our model and the outcome of applying it to a historical scenario. Our results clearly support the validity of our approach.

1. Introduction

In order to meet the new challenges faced by American forces in a changing battlespace environment, the DOD has been aggressively pursuing a doctrinal and operational transformation focused on taking full advantage of Information Age technologies. According to the Secretary of Defense Donald Rumsfeld, "U. S. Forces must leverage information technology and innovative network-centric concepts of operations to develop increasingly capable joint forces. New information and communications technologies hold promise for networking highly distributed joint and multinational forces."¹ Behind these transformations lie concepts such as Network Centric Warfare (NCW), which translates information superiority into combat power by effectively linking knowledge entities in the battlespace [1]. NCW and other transforming concepts are driving warfare toward greater

¹ Source: Department of Defense Office of Force Transformation's Network Centric Warfare Primer, available at http://www.oft.osd.mil/library/library_files/document_318_NCW_GateFold-Pages.pdf

levels of situational awareness, greater autonomy and increased freedom of action at ever lower levels of the command chain.

To realize these and similar concepts stated in the Joint Vision 2010, massive investments have been made to achieve sensor interoperability and information sharing between combatant platforms in a tactical environment. These investments are consistent with the general perception within the DOD that the key to successful tactical decision support systems resides in their ability to cope with increasing amounts of asynchronous data arriving from many diverse sensors, sometimes presenting ambiguous, contradictory, or uncertain evidence to support decision-makers. More important, coping in this context requires not only the capacity to collect, store, and access data, but also the ability to make sense of it, and to sift through it to identify the data most relevant to a decision maker's problem. This is no minor challenge, and as we will see in the next section, combining the ability to express intricate situations while also being able to reason with uncertainty has been a quandary that most technologies in use today are not well suited to solve.

2. Background

Most currently fielded command and control systems use a rule-based methodology for storing expert knowledge and using it to guide tactical decisions according to pre-established policies. This technology has been widely used because of its flexibility and relative computational efficiency, but its simplicity renders the approach incapable of coping with the increasing complexity of modern warfare.

Alternative schemes to represent complex, intricate situations are usually based on classical logic systems. More specifically, many systems are based on some variation of First-Order Logic (FOL), which according to Sowa "has enough expressive power to define all of mathematics, every digital computer that has ever been built, and the semantics of every version of logic, including itself" ([2], page 41). For this reason, FOL has become the *de facto* standard for logical systems from both a theoretical and practical standpoint. However, systems based on classical first-order logic lack a theoretically principled, widely accepted, logically coherent methodology for reasoning under uncertainty, thus limiting their suitability for open, uncertain environments where tactical decision support systems operate.

A promising approach for handling uncertainty in NCW decision support systems is Bayesian Inference. Its advantages over rule-based systems have been widely acknowledged (e.g.[3; 4]). So far, the dominant technology for Bayesian inference is Bayesian Networks (BN) [5; 6], a very capable probabilistic representational scheme that has been successfully used in many different areas such as language understanding [7; 8], visual recognition [9], medical diagnosis [10], and search [11]. Covering BNs is not in the scope of this work, but the interested reader seeking for a review of recent applications of Bayesian Networks will find it in Heckerman [12].

Unfortunately, BNs are not expressive enough for many real-world applications. More specifically, Bayesian Networks assume a simple attribute-value representation – that is, each problem instance involves reasoning about the same fixed number of attributes, with only the

evidence values changing from problem instance to problem instance. Present day tactical situations involve intricate relationships among many variables, rendering techniques with limited representational power such as BNs unsuitable for building useful, detailed models.

As a result, a number of languages have appeared that extend the expressiveness of standard BNs in various ways (e.g. [13], [14], [15]). Although these methods provide significant advances over traditional BNs, they do not achieve the full representational power of First-Order Logic (FOL).

In order to develop a tactical decision system that combines the advantages of FOL with the power of Bayesian Inference, we propose to employ Multi-Entity Bayesian Networks (MEBN, [16]). An implementation of MEBN logic is IET's Quiddity*Suite, a knowledge-based probabilistic reasoning toolkit. As a means of demonstrating the potential of MEBN logic to represent complex, real-world tactical decision environments we built a prototype model of a tense tactical situation involving a possible attack on a U.S. Aegis cruiser by a Libyan gunboat.

3. The Gunboat Scenario

This scenario was presented by Cohen, et al. ([17]) to illustrate a model of human cognition that accounts for the decision making process of the Commanding Officer (CO) and the Tactical Action Officer (TAO) of an U.S. Aegis vessel. In that article, the authors describe a historical situation assessment scenario, which they use as a means to illustrate their model of human cognition. The authors argue that their model accounts for the decision making process followed by the U.S. Navy officers whose reasoning they are describing. Their model of human cognition assumes that human decision makers apply plan recognition templates to scenarios they encounter, matching features of the template to the scenario, and then using the template as a guide to future problem solving.



Figure 1 – The Gunboat Scenario

The already tense situation between the U.S. and Libya worsened after the hijacking of a TWA airliner in July 1985, which was followed by bombing attacks at American Airlines' counters at Vienna and Rome in December of that same year, both linked to terrorist Abu Nidal under patronage of Libya. Following the bombing attacks, the U.S. started a series of operations code-named "Attain Document", in which U.S. ships conducted freedom of navigation demonstrations in Libyan-claimed waters, amidst threats from the Libyan leader. The operations lasted until March 1986. During this time, many incidents occurred in which Libyan assets were damaged or destroyed.

Figure 1 depicts the particular situation of interest, which occurred on March 26th when a gunboat emerged from a Libyan port in the Gulf of Sidra, turned toward an Aegis cruiser, and increased speed. The TAO on the cruiser, from now on referred to as *Ownship*, was convinced the gunboat intended to attack. Several factors supported his assessment. The track was inbound and appeared to be a combat vessel from a hostile nation. It was moving fast at 45 knots on a non-cardinal bearing towards *Ownship*, which at this point was a logical target for the Libyans since it was 20 miles within the "Line of Death" (i.e. within Libyan claimed waters).

Furthermore, *Ownship* (but not other members of the battle group) had detected apparent missile launches toward other American ships earlier in the day, indicating that Libya was actively engaging surface vessels. However, other factors complicated this assessment. The Libyans had far better air assault assets than this small gunboat, and had in fact used them earlier in the day; the gunboat stood a slim chance against the U.S. fleet. However, both the *Ownship's* CO and its TAO considered that perhaps Libya was willing to use every available asset to strike the U.S.

The CO and the TAO believed the gunboat probably did not have the capability to detect *Ownship* at the range at which it had turned inbound, and the dark of night only complicated localization. They reasoned that third party targeting or unusual technology might explain the choice of vectors. Furthermore, *Ownship* was not the only Aegis cruiser below the "Line of Death." Virtually any maneuver by the track would have put it on a vector to a friendly ship. Was the vector to *Ownship* merely a coincidence? There was also another U.S. cruiser farther into Libyan waters and closer to the gunboat. If the gunboat planned to attack, why wasn't the other cruiser the target? If the Libyans wanted to attack, they had an air capability that would have been a more effective means of attack than the gunboat. Thus, there were strong arguments against hostile intent.

This scenario has all the ingredients of a common tactical decision in a modern conflict: a large set of hypotheses that are valid given the ambiguous and conflicting evidence at hand. As a means of both presenting MEBN and showing its ability to deal with such dynamic, complex situation, in the next section we give some background on modeling situations using MEBN logic, while at the same time introducing our model and the process we used to build it.

4. Modeling with Multi-Entity Bayesian Networks

MEBN logic represents the world as comprised of entities that have attributes and are related to other entities. Random variables represent features of entities and relationships among entities. MEBN logic expresses knowledge about attributes and relationships as a collection of *MEBN fragments* (MFrag) organized into *MEBN Theories* (MTheories). An MFrag represents a conditional probability distribution for instances of its resident random variables given their parents in the fragment graph and the context nodes. An MTheory is a set of MFrag that collectively satisfies consistency constraints ensuring the existence of a unique joint probability distribution over instances of the random variables represented in each of the MFrag within the set.

Each MFrag is composed of five components: a fragment graph, a finite set of context nodes, a finite set of input nodes, a finite set of resident nodes, and a set of local and default distributions. The fragment graph is an acyclic directed graph in which the nodes represent random variables and the edges represent conditional probability dependence. Figure 2 shows an example of an MFrag. In this case, the fragment graph has four nodes. The upper yellow node is a context node; the gray node just below it is an input node; and the two white nodes at the bottom are resident nodes.

A node in an MFrag may have a parenthesized list of arguments. These arguments are placeholders for entities in the domain. For example, the argument c to $Aggressiveness(c)$ is a placeholder for an entity that is a combatant, and whose aggressiveness we want to represent. Each actual entity in the domain is assumed to have a *unique identifier*. By convention, unique identifiers begin with an exclamation point, and no two distinct entities can have the same unique identifier. By substituting unique identifiers for a random variable's arguments, we can make *instances* of the random variable. For example, $Aggressiveness(!G1)$ is an instance of the random variable that represents the attribute aggressiveness of entity $!G1$, where $!G1$ is a symbol that refers to a unique entity in the model.

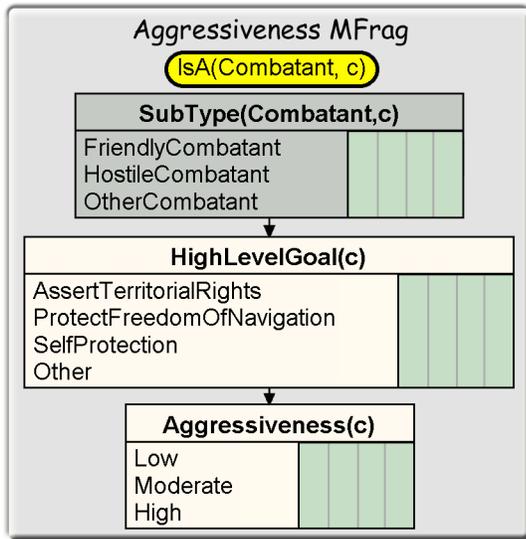


Figure 2 – The Aggressiveness MFrag

The resident nodes of an MFrag have local distributions that define how the probabilities for each of its states depend on the values of their parents in the fragment graph. In the example of Figure 2, the distribution of states *Low*, *Moderate*, and *High* of node $Aggressiveness(c)$ depends on the state of its parent node $HighLevelGoal(c)$. Because $HighLevelGoal(c)$ is a resident node, its distribution is also defined in this MFrag. Its probability is a function of the value of its parent, the input node $SubType(Combatant, c)$. That is, a combatant's high-level goal depends probabilistically on which type of combatant it is. In a complete MTheory, every node of any MFrag has exactly one *home MFrag*, where its local distribution is defined. Input and context nodes (e.g., $SubType(Combatant, c)$ or $IsA(Combatant, c)$) influence the distribution of the resident nodes, but their distributions are defined in their own home MFrag. The context node $IsA(Combatant, c)$ for this MFrag represents whether or not the entity referred to by the variable c is an entity of type *Combatant*; and the input node $SubType(Combatant, c)$ represents the subtype (i.e., friendly or hostile) of the entity referred to by c . Notice that for the sake of conciseness, we are not displaying the probability distributions for the Aggressiveness MFrag.

Context nodes represent conditions that must be satisfied for the influences and local distributions of the fragment graph to apply. Context nodes may have value *True*, *False*, or

Absurd.² Context nodes having value *True* are said to be satisfied. As an example, if we substitute the unique identifier for *Ownership* (i.e., !C0) for the variable *c* in *IsOwnShip(c)*, the resulting hypothesis will be true. If, instead, we substitute a different vessel's unique identifier (say, !C1), then this hypothesis will be false. Finally, if we substitute the unique identifier of a non-vessel (say, !P1), then this statement is absurd (i.e., it is absurd to ask whether or not an entity that is a plan is one's ownership).

Typically, MFragS are small, because their main purpose is to model "small pieces" of domain knowledge that can be reused in any context that matches the context nodes. This is a very important feature of the logic for modeling complex, intricate situations and is one that can be seen as the knowledge representation version of the "divide and conquer" paradigm for decision-making. The correspondence is clear when we think that while the latter breaks a hard, complex decision problem in a set of smaller ones, the former uses a similar decomposition approach for representing intricate, complex military situations. This decomposition is accomplished by modeling a military situation as a collection of small MFragS, each representing some specific element of a tactical situation. The additional advantage of MEBN modeling is the ability to reuse these "small pieces" of tactical knowledge, combining them in many different ways in different scenarios.

Indeed, MFragS provide a flexible means to represent knowledge about specific subjects within the domain of discourse, but the true gain in expressive power is revealed when we aggregate these "knowledge patterns" to form a coherent model of the domain of discourse that can be instantiated to reason about specific situations and refined through learning. It is important to note that just collecting a set of MFragS that represent specific parts of a domain is not enough to ensure a coherent representation of that domain. For example, it would be easy to specify a set of MFragS with cyclic influences (i.e. a random variable which has its probability distribution influencing itself), or one having multiple conflicting distributions for a random variable in different MFragS (i.e. a random variable with more than one home MFrag, each defining a different distribution).

In order to build a coherent model we have to make sure that our set of MFragS collectively satisfies consistency constraints ensuring the existence of a unique joint probability distribution over instances of the random variables mentioned in the MFragS. Such a coherent collection of MFragS is called an MTheory, and it represents a joint probability distribution for an unbounded, possibly infinite number of instances of its random variables. This joint distribution is specified by the local and default distributions within each MFrag together with the conditional independence relationships implied by the fragment graphs.

A generative MTheory summarizes statistical regularities that characterize a domain. These regularities are captured and encoded in a knowledge base using some combination of expert judgment and learning from observation. To apply a generative MTheory to reason

² State names in this paper are alphanumeric strings beginning with a letter, including *True* and *False*. However, Laskey (2004) uses the symbols T for *True*, F for *False*, and \perp for *Absurd*, and requires other state names to begin with an exclamation point (because they are unique identifiers).

about particular scenarios, we need to provide the system with specific information about the individual entity instances involved in the scenario. On receipt of this information, we can use Bayesian inference both to answer specific questions of interest (e.g., how likely is it that the gunboat is executing an opportunistic attack?) and to refine the MTheory (e.g., each new tactical situation gives us additional statistical data about the likelihood of a given plan for that set of circumstances). Bayesian inference is used to perform both problem-specific inference and learning in a sound, logically coherent manner.

Findings are the basic mechanism for incorporating observations into MTheories, such as “!C0 (the Libyan gunboat in our model) is approaching !C0 (*Ownship*) at high speed”. A finding is represented as a special 2-node MFrag containing a node from the generative MTheory and a node declaring one of its states to have a given value. From a logical point of view, inserting a finding into an MTheory corresponds to asserting a new axiom in a first-order theory. In other words, MEBN logic is inherently open, having the ability to incorporate new axioms as evidence and update the probabilities of all random variables in a logically coherent manner.

In addition to the requirement that each random variable must have a unique home MFrag, a valid MTheory must ensure that all recursive definitions terminate in finitely many steps and contain no circular influences. Finally, as we saw above, random variable instances may have a large, and possibly unbounded number of parents. A valid MTheory must satisfy an additional condition to ensure that the local distributions have reasonable limiting behavior as more and more parents are added. Laskey [16] proved that when an MTheory satisfies these conditions (as well as other technical conditions that are unimportant to our example), then there exists a joint probability distribution on the set of instances of its random variables that is consistent with the local distributions assigned within its MFrag. Furthermore, any consistent, finitely axiomatizable FOL theory can be translated to an infinity of MTheories, all having the same purely logical consequences, that assign different probabilities to statements whose truth-value is not determined by the axioms of the FOL theory. MEBN logic contains a set of built-in logical MFrag (including quantifier, indirect reference, and Boolean connective MFrag) that provide the ability to represent any sentence in first-order logic. If the MTheory satisfies additional conditions, then a conditional distribution exists given any finite sequence of findings that does not logically contradict the logical constraints of the generative MTheory. MEBN logic thus provides a logical foundation for systems that reason in an open world and incorporate observed evidence in a mathematically sound, logically coherent manner.

Figure 3 shows an example of a generative MTheory for the Libyan gunboat scenario. For the sake of conciseness, the local distribution formulas and the default distributions are not shown here. In this model, each sub-type has only one parent type, but MEBN logic is flexible enough to accommodate all features of more complex typed systems, such as polymorphism (in a polymorphic MEBN, a given random variable could have different definitions for different types of entity), multiple-inheritance, etc.

First Order Logic (or one of its subsets) provides the theoretical foundation for the type systems used in popular object-oriented and relational languages. MEBN logic provides the basis for extending the capability of these systems by introducing a sound mathematical basis for representing and reasoning under uncertainty. Among the advantages of a MEBN-based typed system is the ability to represent *type uncertainty*. As an example, suppose we had two different types of “submerged entities”, submarines and whales, and we are not sure about the type of a given entity. In this case, the result of any query that depends on the entity type will be a weighted average of the result given that the entity is a submarine and the result given that it is a whale.

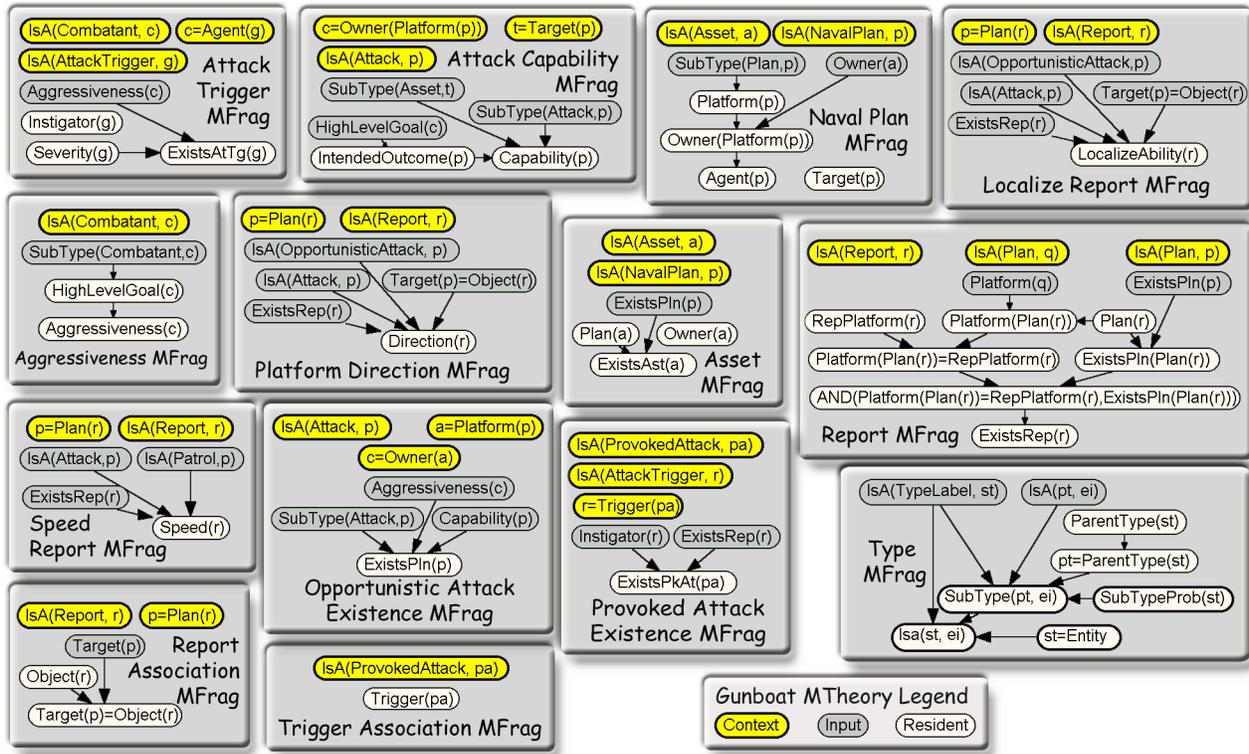


Figure 3 – The Gunboat MTheory

Further advantages of a MEBN-based type system include the ability to refine type-specific probability distributions using Bayesian learning, assign probabilities to possible values of unknown attributes, reason coherently at multiple levels of resolution, and other features related to representing and reasoning with incomplete and/or uncertain information.

MEBN logic supports *finite or countably infinite recursion*, which includes temporal and other kinds of recursion, and can represent and reason about hypothetical entities. Uncertainty about whether a hypothesized entity actually exists is called *existence uncertainty*. In the Asset MFragment of our MTheory, the random variable *ExistsAst(a)* is used to reason about whether its argument is an actual asset. For example, we might be unsure whether a sensor report corresponds to one of the vessels we already know about, a vessel of which we were previously unaware, or a spurious sensor report. In this case, we can create an asset instance, say !A2, and assign a probability of less than 1.0 that *ExistsAst(!A2)* has

value *True*. Then, any queries involving !A2 will return results weighted appropriately by our belief in the existence of !A2. Furthermore, our belief in *ExistsAst(!A2)* is updated by Bayesian conditioning as we obtain more evidence relevant to whether !A2 denotes a previously unknown asset. Representing existence uncertainty is particularly useful for counterfactual reasoning and reasoning about causality [18; 19].

A very common problem in multi-sensor data fusion systems is *association uncertainty*, which means uncertainty about the source of a given report (e.g. whether a given report refers to vessels !C1, !G0 or !G1). Many weakly discriminatory reports coming from possibly many vessels produces an exponential set of combinations that require special *hypothesis management* methods [c.f. 20]. In the Gunboat model this problem can be seen in the Report MFrag, which captures the complex reasoning involved in inferring the existence of a platform given the available sensor reports. As we will see below, MEBN logic can represent and reason with association uncertainty, and thus provides a sound logical foundation for hypothesis management in multi-source fusion.

Finally, another important aspect of MEBN logic is its flexibility. The generative MTheory in Figure 3 is just one of the many possible (consistent) sets of MFrag that can be used to represent the same joint distribution. These MFrag were designed to mimic the recognition/metacognition model presented by Cohen, et. al. [17] for time-stressed decision-making, since we wanted to show that their cognitive approach can be modeled as a Bayesian process³. However, the approach to be taken when building an MTheory will depend on many factors, including the model's purpose, the background and preferences of the model's stakeholders, the need to interface with external systems, etc.

5. Implementing and running the model

In order to use MEBN logic to draw inferences in a given scenario, we have to have an initial generative MTheory, a Finding set (which conveys the new information we have) and a Target set (which indicates the nodes of interest to us). We implemented the MTheory in Figure 3 using Quiddity*Suite™. Figure 4 shows an example of an MFrag (Aggressiveness) and its respective implementation in Quiddity*Suite's frame-based syntax.

For our experiments, we implemented the MTheory in figure 3, a set of findings that matched the parameters defined in Cohen, et. al. [17] (e.g. information on the two U.S. cruisers, the Libyan Gunboat, the reports, etc.), and a target set consisting of nodes that would allow us to assess the hypotheses considered in their article.

The inference process begins when a query is posed to assess the degree of belief in a target random variable given a set of evidence random variables. The first step in MEBN inference is to construct a *situation-specific Bayesian network* (SSBN). This is an ordinary Bayesian network constructed by creating and combining instances of the MFrag in the

³ It is interesting to note that the authors in [17] explicitly stated that the CO's decision process was not Bayesian, since the U.S. officers "considered different interpretations of each cue in the context of alternative situation pictures". Indeed, this kind of situated, adaptive reasoning cannot be captured by standard Bayesian networks, but it can be modeled quite naturally by MEBN logic and SSBN construction.

generative MTheory. Next, a standard Bayesian network inference algorithm is applied. Finally, the answer to the query is obtained by inspecting the posterior probabilities of the target nodes. A MEBN inference algorithm is provided in Laskey [16].

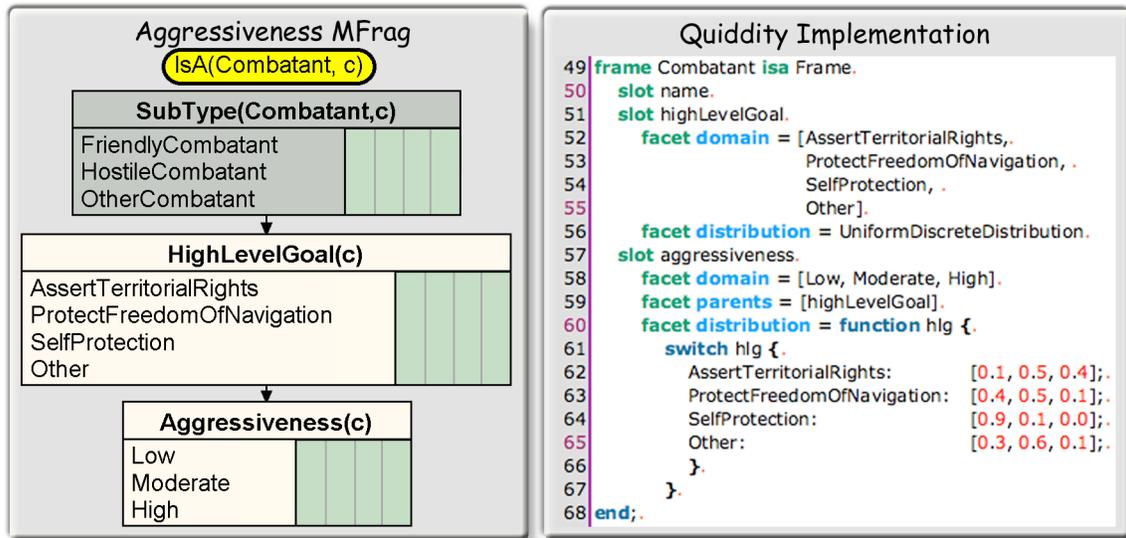


Figure 4 – Aggressiveness MFrag and its Respective Quiddity Implementation

In some cases the SSBN can be infinite, but under conditions given in Laskey [16], the algorithm produces a sequence of approximate SSBNs for which the posterior distribution of the target nodes converges to their posterior distribution given the findings. Mahoney and Laskey [21] define a SSBN as a minimal Bayesian network sufficient to compute the response to a query. A SSBN may contain any number of instances of each MFrag, depending on the number of entities and their interrelationships.

Figure 5 represents one of the SSBNs that resulted when we applied this process to the MTheory in Figure 3 with evidence from table 1 below. For a detailed account of the SSBN construction algorithm, the interested reader should refer to Laskey [16]. There, it is possible to find the mathematical explanation and respective logical proof for the many intricate possibilities when instantiating MFrag, such as nodes with an infinite number of states, situations where we face the prospect of large finite or countably infinite recursions, what happens when the algorithm is started with an inconsistent MTheory, etc.

In addition, the text provides a detailed account of how to represent any First Order Logic sentence as an MFrag using Skolem variables and quantifiers, and an overview of Bayesian learning, which is treated in MEBN logic as a sequence of MTheories. These issues go beyond the scope of the present paper.

For the present work, we dealt with the hypothesis management problem using the PLASMA architecture described by Fung, et al. [22]. In short, we address the possibly infinite number of hypotheses by applying logical reasoning to identify what MFrag to retrieve and instantiate, and then how to construct the SSBN using probabilistic generalizations of forward-chaining and backward-chaining of rules, encoded in *first-order logic suggestors*. In other words, our use of suggestors can be viewed as a way to control the construction of

SSBNs, thus ensuring that the most relevant hypotheses are included while avoiding the incorporation of large numbers of irrelevant hypotheses.

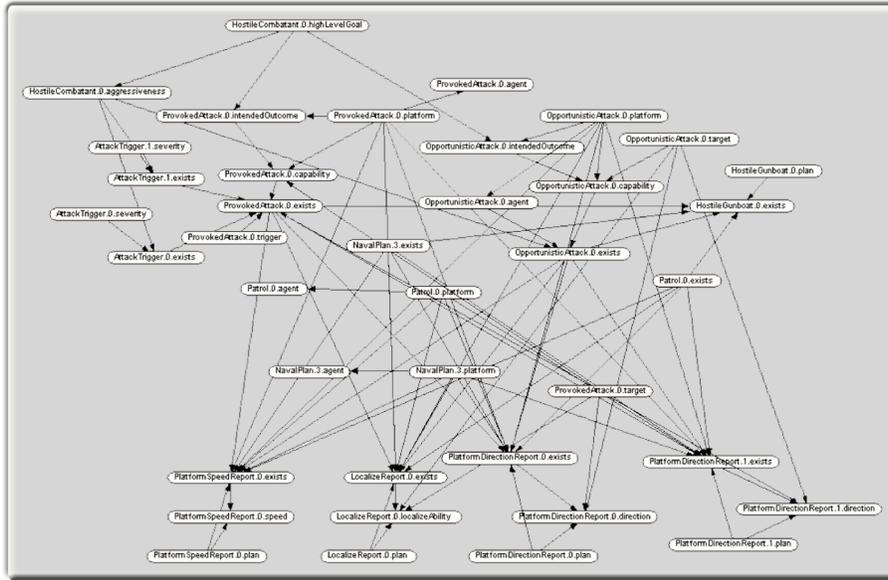


Figure 5 – An Example SSBN based on the Gunboat MTheory

In our implementation, we hypothesized plans of action for the gunboat based on available evidence. We developed a proof of concept knowledge base for this example consisting of:

- A set of frames representing generic knowledge about naval conflict. Frame types include:
 - *Combatants*: There is a generic combatant frame, as well as subtype frames for friendly and hostile combatants. These frames represent background knowledge about the high-level goals and the level of aggressiveness of combatants.
 - *Plans*: There are plans for provoked attacks, opportunistic attacks, and patrols, as well as a generic, non-specific plan representing the types not included among the hypotheses being considered.
 - *Attack Triggers*: An attack trigger is a kind of behavior that might provoke an attack. We modeled the U.S. presence below the “Line of Death” as an instance of an attack trigger.
 - *Reports*: Reports are used to model observable evidence that bears on the hypotheses under consideration.
- A set of suggestors representing knowledge about which hypotheses should be explicitly considered in specific situations. Suggestors are represented as a set of rules expressed in first-order logic. The rules apply observed evidence (e.g., the cruiser is situated below the “Line of Death;” the gunboat is on a rapid direct approach),

together with information about the current status of the Bayesian network (e.g., no attack hypothesis has yet been enumerated for the gunboat), to nominate suggested network construction actions. Although not implemented in this proof-of-concept, suggestors can also nominate pruning actions to help make inference more efficient.

- Particular facts germane to the scenario, including background knowledge and observed evidence.

The suggestors were based on the following general rules which, according to the description of Cohen, et al. [17], are reflective of the kinds of reasoning the TAO applied in this scenario.

- IF a ship is sailing in its own territorial waters, THEN the ship may be on patrol.
- IF a hostile asset is approaching on a bearing directly toward own ship, THEN the asset may be attacking.
- IF an attack by a hostile is hypothesized AND the asset's approach is rapid and direct, THEN the strength of the attack hypothesis is increased.

All the abovementioned components were implemented using Quiddity*Suite's probabilistic and logical toolbox. We comment and analyze the results of our experiments in the next section.

6. Results and Discussion

Our results are displayed in Table 1 below, and indicate that the MEBN model presented in this work is consistent with the qualitative reasoning of the original study by Cohen, et. al. [17].

Initially, the provoked attack hypothesis dominates, but as the evidence is processed, its incongruity with the provoked attack hypothesis becomes more apparent. This increases the probability of the "other" hypothesis. A natural alternative hypothesis to consider is the patrol hypothesis, but it too is incongruent with the available evidence. When the opportunistic attack hypothesis is nominated in response to the failure of other plans to account for the evidence, it becomes the dominant hypothesis.

The results suggest that MEBN logic with SSBN construction provides a natural model for the reasoning process engaged in by the CO and the TAO. In addition, given the nature and flexibility of our approach, the model has the potential to be applied "as is" to scenarios involving arbitrary numbers of combatants and assets, and with varied patterns of incoming evidence.

MEBN's modular design also allows our model to be extended to include additional complexity, such as subtypes of existing entity types or other asset types (e.g., carrier, air squadron), additional features of existing entity types, and additional entity types as well. It was also clear to us that the model's modularity facilitates maintenance, modification, and iterative improvement, so we can change part of the model without affecting unrelated parts, allowing different experts to focus on different parts of a model.

The modularity of the MEBN approach and its ability to handle the complexity of a real world tactical environment provides a very flexible situation assessment framework for tactical decision systems. Thus, we extend our initial statement to claim that the technique has the potential to be applied in a wide variety of real-life tactical decision problems that involve uncertain, ambiguous, incomplete information asynchronously arriving from many, diverse sensors.

Table 1 – Results of MEBN Inference given different evidence

Evidence (ordered as input into the model)	Hypotheses	Probabilities	Target of Provoked Attack
Cruiser 1 instigates attack - trigger moderate severity Cruiser 2 instigates attack - trigger high severity	Provoked attack Other	69.2% 30.8%	Cruiser 1: 31.5% Cruiser 2: 68.5%
Gunboat approaching Cruiser 1	Provoked attack Other	81.4% 18.6%	Cruiser 1: 77.9% Cruiser 2: 22.1%
Gunboat not approaching Cruiser 2	Provoked attack Other	78.8% 21.2%	Cruiser 1: 89.0% Cruiser 2: 11.0%
Gunboat approaching fast	Provoked attack Other	93.2% 6.8%	Cruiser 1: 96.3% Cruiser 2: 3.7%
Gunboat probably cannot localize Cruiser 1	Provoked attack Other	62.2% 37.8%	Cruiser 1: 79.5% Cruiser 2: 20.5%
	Provoked attack Patrol Other	42.7% 31.5% 25.7%	Cruiser 1: 70.0% Cruiser 2: 30.0%
	Provoked attack Patrol Opportunistic attack Other	4.7% 2.5% 90.8% 2.1%	Cruiser 1: 51.5% Cruiser 2: 48.5%

From a technical standpoint, the formulation of MEBN logic provided in Laskey [16] is to our knowledge the first probabilistic logic to possess all of the following properties: (1) the ability to express a globally consistent joint distribution over interpretations of any consistent, finitely axiomatizable FOL theory; (2) a proof theory capable of identifying inconsistent theories in finitely many steps and converging to correct responses to probabilistic queries; and (3) a built in theory for refining theories in the light of observations.

As such, MEBN should be seen not as a competitor, but as a logical foundation for the many emerging languages that extend the expressive power of standard Bayesian networks and/or extend a subset of first-order logic to incorporate probability. That said, it is important to emphasize the potential application of MEBN logic to command and control systems, given its ability to, among others:

- Store domain knowledge in “small pieces” that can be reused in future occasions.
- Be extended to richer and more complex situations as needed.
- Deal with finite or countably infinite recursion, providing a basis for temporal reasoning.
- Use Bayesian learning to infer possible pattern correlations given a corpus of data.
- Deal with type, association, and existence uncertainty.
- Treat the hypothesis management problem in real time.
- Represent and reason with uncertainty in a mathematically principled way, while keeping the computational time compatible with command and control applications.

Up to the last decade, command and control systems have been designed under the information paradigm, a situation that led Creveld to conclude in 1985 that “The history of command can thus be understood in terms of a race between the demand for information and the ability of command systems to meet it” ([23], page 265). Today, we live under a different paradigm, and while the above assertion can be updated simply by changing the word “information” by the word “knowledge”, updating command systems is far less trivial. In this case, a paradigm shift is required, for which old techniques are not suitable for that, no matter how successful they were a mere decade ago and how much time and resources are spent to update it.

To be successful, any implementation of the NCW concept in today’s battlespace must have at least the above feature list. We have found MEBN logic to be a very promising technology to meet these challenges from the knowledge engineering standpoint, and we hope this work may help to articulate and clarify that potential.

Acknowledgments

Grateful acknowledgement is due to the Brazilian Air Force for supporting Paulo Costa during his PhD studies at George Mason University. The other authors were partially supported under a contract with the Office of Naval Research, number N00014-04-M-0277.

Bibliography

- 1 Alberts, D. S., Garstka, J. J., & Stein, F. P. (1999). *Network Centric Warfare: Developing and Leveraging Information Superiority*: National Defense University Press.
- 2 Sowa, J. F. (2000). *Knowledge representation: logical, philosophical, and computational foundations*. Pacific Grove: Brooks/Cole.

- 3 Costa, P. C. G. (1999). The Fighter Aircraft's Autodefense Management Problem: A Dynamic Decision Network Approach. Master of Science Thesis, George Mason University, Fairfax, VA.
- 4 Laskey, G., & Laskey, K. B. (2002). Combat Identification with Bayesian Networks, 7th International Command and Control Research and Technology Symposium (pp. 13). Quebec City, Canada: CCRP.
- 5 Pearl, J. (1988). Probabilistic reasoning in intelligent systems: networks of plausible inference. San Mateo, CA: Morgan Kaufmann Publishers.
- 6 Lauritzen, S., & Spiegelhalter, D. J. (1988). Local computation and probabilities on graphical structures and their applications to expert systems. *Journal of Royal Statistical Society*, 50(2), 157-224.
- 7 Charniak, E., & Goldman, R. P. (1989, August 1989). A semantics for probabilistic quantifier-free first-order languages with particular application to story understanding. Paper presented at the Eleventh International Joint Conference on Artificial Intelligence, Detroit, Michigan.
- 8 Charniak, E., & Goldman, R. P. (1989). Plan recognition in stories and in life. Paper presented at the Fifth Workshop on Uncertainty in Artificial Intelligence, Mountain View, California.
- 9 Binford, T., Levitt, T. S., & Mann, W. B. (1987). Bayesian Inference in Model-Based Machine Vision. In a. P. C. C. T. S. Levitt (Ed.), *Uncertainty in Artificial Intelligence: Proceedings of the Third Workshop*. Seattle, WA.
- 10 Heckerman, D. (1990). Probabilistic Similarity Networks. Ph.D. Thesis, Stanford University, Stanford, CA.
- 11 Hansson, O., & Mayer, A. (1989, August, 1989). Heuristic search as evidential reasoning. Paper presented at the Fifth Workshop on Uncertainty in Artificial Intelligence, Windsor, Ontario.
- 12 Heckerman, D., Mamdani, A., & Wellman, M. P. (1995). Real-world application of Bayesian networks. *Communications of the ACM*, 38(3), 24-68.
- 13 Koller, D., & Pfeffer, A. (1997). Object-Oriented Bayesian Networks. Paper presented at the Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference, San Francisco, CA.
- 14 Heckerman, D., Meek, C., & Koller, D. (2004). Probabilistic models for relational data. Redmond, WA: Microsoft Corporation.
- 15 Nielsen, M. A., & Chuang, I. L. (2000). *Quantum Computation and Quantum Information*. Cambridge, UK: Cambridge University Press.
- 16 Laskey, K. B. (2004, 2004/10/16). MEBN: Bayesian Logic for Open-World Reasoning. Retrieved Dec 8, 2004, from <http://ite.gmu.edu/~klaskey/publications.html>
- 17 Cohen, M. S., Freeman, J. T., & Wolf, S. (1996). Metarecognition in Time-Stressed Decision Making: Recognizing, Critiquing, and Correcting. *Human Factors and Ergonomics Society*, 38(2), 206-219.
- 18 Druzdzel, M. J., & Simon, H., A. (1993). Causality in Bayesian belief networks. Paper presented at the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93), San Francisco, CA.
- 19 Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge, U.K.: Cambridge University Press.
- 20 Stone, L. D., Barlow, C. A., & Corwin, T. L. (1999). *Bayesian Multiple Target Tracking*. Boston, MA: Artech House.

- 21 Mahoney, S. M., & Laskey, K. B. (1998). Constructing Situation Specific Networks. Paper presented at the Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference, San Mateo, CA.
- 22 Fung, F., Laskey, K. B., Pool, M., Takikawa, M., & Wright, E. J. (2004). PLASMA: Combining Predicate Logic and Probability for Information Fusion and Decision Support. Paper presented at the AAAI Spring Symposium, Stanford, CA.
- 23 Creveld, M. V. (1985). *Command in War* (Reprint edition (March 1, 1987) ed.). Cambridge, MA: Harvard University Press.