

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 02-12-2013		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 3-May-2011 - 2-Sep-2013	
4. TITLE AND SUBTITLE PROJECT FINAL REPORT Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion			5a. CONTRACT NUMBER W911NF-11-1-0176		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Paulo Costa, Kathryn Laskey			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES George Mason University Office of Sponsored Programs 4400 University Drive, MSN 4C6 Fairfax, VA 22030 -4422			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 59841-MA.12		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This is the final report on the research effort performed by George Mason University, under the agreement W911NF-11-1-0176, proposal number 59841MA, entitled "Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion" and focused on: (a) Establishing features required of any quantitative uncertainty representation for exchanging soft and hard information in a net-centric environment; (b) Developing a set of use cases involving information exchange and fusion requiring sophisticated processing and					
15. SUBJECT TERMS Uncertainty Representation, High-Level Information Fusion, Evaluation Framework					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Paulo Costa
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 703-993-9989

Report Title

PROJECT FINAL REPORT

Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion

ABSTRACT

This is the final report on the research effort performed by George Mason University, under the agreement W911NF-11-1-0176, proposal number 59841MA, entitled “Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion” and focused on:

- (a) Establishing features required of any quantitative uncertainty representation for exchanging soft and hard information in a net-centric environment;
- (b) Developing a set of use cases involving information exchange and fusion requiring sophisticated reasoning and inference under uncertainty;
- (c) Defining evaluation criteria supporting an unbiased comparison among different approaches applied to the use cases; and
- (d) Examining in detail how two popular formalisms, Bayesian and Dempster-Shafer, address the requirements in the context of the use cases.

The proposed research aimed to establish a commonly agreed understanding of the fundamental aspects of uncertainty representation and reasoning that a theory of hard and soft high-level fusion must encompass. Successful completion requires an unbiased, in-depth analysis of the associated enabling technologies, and a formalization of its fundamental principles.

Although this report covers the activities held between May 3rd, 2011, and September 2nd, 2013, the aspects already covered in the previous progress reports are not repeated here, unless some major changes were made.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

12/02/2013 5.00 Mark Locher, Paulo Costa. Ignorance and Uncertainty in Level Two Information Fusion, Information Fusion (05 2013)

TOTAL: 1

Number of Manuscripts:

Books

Received Paper

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Paulo Costa	0.25	
Kathryn Laskey	0.05	
FTE Equivalent:	0.30	
Total Number:	2	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

See Attachment

Technology Transfer

PROJECT FINAL REPORT

Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion

Period of Performance: May 3rd, 2011 to September 2nd, 2013
Agreement#: W911NF-11-1-0176

Prepared for: Dr. John A. Lavery
U.S. Army Research Office
P.O.Box 1221
Research Triangle Park, NC 27709-2211

GMU POCs: Dr. Paulo Costa and Dr. Kathryn B. Laskey
C⁴I Center and Systems Engineering and Operations Research Department
George Mason University
4400 University Dr., MS 4A6
Fairfax, VA 22030
Voice : (703) 993-1639, (703) 993-1644
Fax : (703) 993-1521
pcosta@gmu.edu, klaskey@gmu.edu

TABLE OF CONTENTS

1. Executive Summary	2
2. Overview of the Research Effort	2
2.1. Problem Statement	2
2.2. Project Goals and Tasks	3
2.3. Research Approach	4
2.4. Personnel Involved	5
2.5. Logistics and Administrative Issues	6
3. Synopsis of the Main Results	6
3.1. Task A: Perform In-Depth Analysis	6
3.1.1. Subtask A1: Identify requirements for quantitative uncertainty representation for exchanging hard and soft information	6
3.1.2. Subtask A2: Derive a balanced set of use cases encompassing a range of situations requiring HLIF in network-centric information exchange	8
3.1.3. Subtask A3: Design preliminary criteria for evaluating the performance of the techniques	9
3.1.4. Subtask A4: Present preliminary results and discuss project plan and objectives at Fusion 2012	10
3.1.5. Subtask A5: Develop test procedures and datasets to perform evaluation	10
3.1.6. Subtask A6: Perform preliminary evaluation	11
3.2. Task B: Develop Evaluation Framework	11
3.2.1. Subtask B1: Refine evaluation criteria from task A	11
3.2.2. Subtask B2: Design evaluation plan	11
3.2.3. Subtask B3: Convene session at Fusion 2012 to define final evaluation plan	12
3.2.4. Subtask B4: Define final criteria for evaluation contest at Fusion 2013	12
3.2.5. Subtask B5: Perform final evaluation	12
3.2.6. Subtask B6: Convene session at Fusion 2013 to present the results of the contest	13
4. Uncertainty in HLIF Applications	13
4.1. Scope	13
4.2. Nature of Level 2 HLIF	14
4.3. Defining a Level 2 HLIF Taxonomy	15
4.4. Uncertainty Types	17
4.5. Uncertainty Model	18
4.6. Uncertainty Considerations	21
4.6.1. Area 1: Uncertainty in the Input Data / Evidence	21
4.6.2. Area 2: Uncertainty in the HLIF Arguments	22
4.6.3. Area 3: Uncertainty in the Hypotheses	24
4.7. Integrating Uncertainty into the Process Model	24
4.8. Formal Construct	27
4.8.1. Uncertainty Modeling	27
4.8.2. Application to Level 2 HLIF	29
4.9. Discussion	30
5. Conclusion	31
Bibliography	32

1. Executive Summary

This is the final report on the research effort performed by George Mason University, under the agreement W911NF-11-1-0176, proposal number 59841MA, entitled “Mathematical Fundamentals of Probabilistic Semantics for High-Level Fusion” and focused on:

- (a) Establishing features required of any quantitative uncertainty representation for exchanging soft and hard information in a net-centric environment;
- (b) Developing a set of use cases involving information exchange and fusion requiring sophisticated reasoning and inference under uncertainty;
- (c) Defining evaluation criteria supporting an unbiased comparison among different approaches applied to the use cases; and
- (d) Examining in detail how two popular formalisms, Bayesian and Dempster-Shafer, address the requirements in the context of the use cases.

The proposed research aimed to establish a commonly agreed understanding of the fundamental aspects of uncertainty representation and reasoning that a theory of hard and soft high-level fusion must encompass. Successful completion requires an unbiased, in-depth analysis of the associated enabling technologies, and a formalization of its fundamental principles.

Although this report covers the activities held between May 3rd, 2011, and Sep 2nd, 2013, the aspects already covered in the previous progress reports are not repeated here, unless some major changes were made.

2. Overview of the Research Effort

2.1. Problem Statement

Military situations are inherently uncertain, and the available data are inevitably noisy and incomplete. The ability to represent and reason with uncertainty is fundamental to Information Fusion Systems at all levels of the JDL hierarchy. Tasks at higher levels of the JDL fusion framework, such as the Level 3 task of predicting threat behavior, require reasoning about complex situations in which entities of different types are related to each other in diverse ways. This is particularly true in asymmetric warfare where the threats are elusive, secretive, and decentralized. Related entities often appear unconnected and their stealthy behavior is very difficult to predict. Automated methods for reasoning about such complex situations require expressive representation languages that can represent and reason with uncertainty.

This research is concerned about the ability of a given uncertainty representation to provide principled support to information fusion systems in their handling of uncertainty. This is a key requirement for IFS in their support to commanders in the battlefield, as it allows for reliable information fusion and subsequent predictive analysis of the situation. In other words, making the best use of massive data coming from heterogeneous sources is essential for assessing the current situation and predicting its potential developments.

Real-time prediction requires sifting through a myriad of soft and hard information items to identify potential threats in modern asymmetric environments, such as a handful of malign actors attempting to blend in among millions of innocent people. The latter is an example of the typical “needle-in-a-haystack” problem, in which information clearly suggesting potential threats is only perceived in after-the-fact forensic analysis – too late for preventative action to be taken. High-level information fusion (HLIF) technologies for knowledge exchange require the use of expressive representational frameworks capable of capturing subtleties that current low-level information fusion (LLIF) techniques cannot handle. A key requirement for these HLIF frameworks is to provide principled support for representation and reasoning with uncertainty, which is ubiquitous to all applications involving knowledge exchange.

HLIF of hard and soft information from diverse sensor types still depends heavily on human cognition. This results in a scalability conundrum that current technologies are incapable of solving. The major issue preventing successful automated reasoning with incoming hard and soft information is the lack of a fundamental HLIF theory, backed by a consistent mathematical framework and supporting algorithms. Although there is no question about the fact that an HLIF framework must support automated knowledge representation and reasoning with uncertainty, there is no consensus on the most appropriate technology to satisfy this requirement. Further, the debate on the appropriateness of the prominent approaches is strewn with misconceptions and ill-supported assumptions, greatly jeopardizing attempts by the community to converge on a fundamental mathematical theory of HLF that: (1) supports representation of semantics and pragmatics, (2) provides a solid mathematical foundation underlying its algorithms, and (3) supports scalability of products such as common and user-defined operational pictures. This work aimed to address these misconceptions and to clarify the fundamental issues underlying the debate.

2.2. Project Goals and Tasks

As we have mentioned in the proposal phase of this effort, despite decades of long debate on the most appropriate automated uncertainty management methodology the matter is far from being settled. Unfortunately, the debate has been as much a source of misconception and wasted resources as it has been a source of insight to HLIF researchers and practitioners. *This state of affairs inhibits advancement of the community as a whole.* To work toward remediating this situation, this project was designed to achieve the following goals:

1. Perform an in-depth analysis of the major requirements for representing and reasoning with uncertainty from the HLIF perspective;
2. Develop a set of use cases with enough complexity to cover the identified requirements;
3. Define a comprehensive set of criteria to evaluate how well a given methodology addresses the representational and reasoning needs of each use case; and
4. Conduct an evaluation of two major uncertainty management approaches being used in HLIF research and development.

Given the multidisciplinary nature and broad spectrum of interests within the HFL community, our analysis addressed not only the fundamental issues of representing and reasoning with

uncertainty, but also pragmatic topics germane to knowledge exchange problems. These latter topics include consistency, accuracy and scalability.

In other to achieve the aforementioned goals, the project plan focused its efforts in executing the following tasks, as described in the project proposal:

A. Perform In-Depth Analysis

- A1 Identify requirements for quantitative uncertainty representation for exchanging hard and soft information
- A2 Derive a balanced set of use cases encompassing a range of situations requiring HLF in network-centric information exchange
- A3 Design preliminary criteria for evaluating the performance of the techniques
- A4 Present preliminary results and discuss project plan and objectives at Fusion 2012
- A5 Develop test procedures and datasets to perform evaluation
- A6 Perform preliminary evaluation

B. Develop Evaluation Framework

- B1 Refine evaluation criteria from task A
- B2 Design evaluation plan
- B3 Convene session at Fusion 2012 to define final evaluation plan
- B4 Define final criteria for evaluation contest at Fusion 2013
- B5 Perform final evaluation
- B6 Convene session at Fusion 2013 to present the results of the contest

C. Provide Documentation and Support to ARO

- C1 Support ARO technical exchanges and participate in technical review meetings
- C2 Prepare semi-annual and annual project technical progress reports

The results of tasks A and B are described in Section 3 below, while task C was conducted in close coordination with ARO and will not be covered in this report.

2.3. Research Approach

A key aspect that is implied in Tasks A and B, as well as in the overall nature of this research effort and its associated planning is the credibility issue. More specifically, we argue that years of debate about the best approaches for uncertainty representation and reasoning in IFS by some of the best minds of the International Society for Information Fusion (ISIF) community has failed to yield consensus not because of a lack of mathematical sophistication on the part of advocates of the different approaches. In fact, our ongoing work in this area leads us to conclude that the most commonly applied techniques (e.g. Bayesian, Dempster-Shafer, fuzzy) rest on mathematically correct foundations, although that does not mean all are equally suitable for IFS in military operations. On the contrary, we have argued [1]–[4] that a continuing stalemate in the long-standing debate on the most appropriate uncertainty management formalism for HLIF

applications can be traced at least in part to the lack of a comprehensive evaluation framework to support meaningful comparisons. Much of the debate has been at a philosophical level, focused on artificially simple problems. There has been little testing that explores behavior of the different formalisms on complex problems exhibiting characteristics fundamental to today's high-level fusion systems. A comprehensive evaluation framework accepted by the broad fusion community would facilitate meaningful comparisons on realistically complex problems

To address this deficit, we chose to lead a research effort aimed at developing a comprehensive evaluation framework. In collaboration with an international group of fusion researchers and practitioners, we founded the ISIF Working Group for Evaluation of Uncertainty Representation Techniques, ETURWG. This group has been led by the PIs since its inception, has received overwhelming support from the ISIF community, and has attracted a diverse multinational group of researchers to address the problem. The success of this initiative is evidenced by the high attendance at all ETUR workshops conveyed at the Fusion conferences.

In the one hand, leading an ISIF working group was our best option for achieving the community-wide results we wanted with the resources we had. As we describe below in this report, the evidence we had so far points to the validity of this approach. On the other hand, the complexity of leading volunteers from various countries on a major effort has exceeded our expectations in terms of the lead-time required to produce the desired results. As a result, we have been able to achieve full success on some tasks, but some tasks are still in progress. Details about the task progress are conveyed in Section 3 below.

2.4. Personnel Involved

For all the above reasons, ETURWG has been the primary vehicle for us to achieve the research goals of this effort. It is an official and active ISIF working group that relies on the leadership of the PIs, which requires a non trivial amount of time to bring the effort to fruition. Over the period of this research, the following personnel were directly involved:

Dr. Paulo Costa, PI – Partially supported by U.S. Army Research Office during the full period of this effort. Prof. Costa was responsible for the overall project management and technical progress, for the coordination of the evaluation process, as well as for the development of the evaluation criteria concerning the underlying semantic foundations for uncertainty representation and reasoning. He is the lead organizer of the ETURWG, defining the schedule of activities of the group and maintaining its website and collaboration tools. He also co-organized the ETUR special sessions held at the Fusion 2011, Fusion 2012, and Fusion 2013 conferences, as part of the ETURWG effort.

Dr. Kathryn Laskey, Co-PI – Partially supported by U.S. Army Research Office during the full period of this effort. Prof. Laskey supported the project management tasks, the coordination tasks of the evaluation process, and also shared responsibility for developing the evaluation criteria concerning the underlying mathematical foundations of uncertainty representation and reasoning. She is also a lead organizer of the ETURWG, responsible for the minutes and other tasks within the group. Together with Prof. Costa, she co-organized the ETUR special sessions held at the Fusion 2011, Fusion 2012, and Fusion 2013 conferences, as part of the ETURWG effort.

Mr. Mark Locher – GMU PhD Student, advised by Prof. Paulo Costa, not supported by ARO. Mr. Locher developed one of the use cases for this research effort. The use case he developed is now part of the ETURWG use case repository. He was also a key participant in the research on fundamental issues on uncertainty in HLIF, having published papers in conferences and journals devoted to fusion and uncertainty management.

ETURWG members – Constant commitment and participation was key for building up credibility within the community and ensuring the success of the evaluation framework proposed by the group. None of the various ETURWG members was supported by ARO. It is a relatively large group (50 registered members) and its membership list is publicly available via its website (registration required) – The URL is <http://eturwg.c4i.gmu.edu>. Participation in the meetings varied throughout the years, with a smaller set of active participants forming the core group. Attendance at the meetings and events can also be assessed from the website, but the average attendance at the 40 general meetings convened so far has been about 10 people; attendance at the ETUR special sessions has averaged 45 people.

2.5. Logistics and Administrative Issues

Ensuring a proper environment for collaboration was a key issue during the research, since the work relied heavily on a group of highly technical people with complex availability requirements. We have convened so far 40 biweekly meetings, using a combination of a paid teleconference service and the software Adobe Connect Pro.

Most of the results are posted in a members only section of the ETURWG website, which is freely available upon registration. The website was developed by the PI using the Drupal¹ content management system (CMS) and hosted by the C4I Center at George Mason University. The website includes various features aimed to provide not only a repository for the group's research and results, but also a working environment for collaboration among researchers.

3. Synopsis of the Main Results

The core subject of this research was to establish the requirements and develop an evaluation framework² for uncertainty representation and reasoning in the HLIF domain. We describe the main results for each of the tasks defined in the proposal document that led to this research agreement.

3.1. Task A: Perform In-Depth Analysis

3.1.1. Subtask A1: Identify requirements for quantitative uncertainty representation for exchanging hard and soft information

This Subtask was completed in full. The evaluation of how uncertainty is dealt with within a given IFS is distinct from, although closely related to, the evaluation of the overall performance

¹ <https://drupal.org/>

² In this research, we used the term “evaluation framework” to denote the set of requirements, use cases, and evaluation criteria that collectively form a comprehensive, unbiased means to evaluate how well a given approach addresses the needs of HLIF applications.

of the system. Metrics for evaluating the overall performance of IFSs are more encompassing in scope than those focused on the uncertainty handling within the system. The metrics for the overall system include the effects of the uncertainty representation, but there are also effects of other aspects of the fusion system that can affect the performance of the system.

For example, timeliness (how quickly the system can come to a conclusion within a specified precision level) is often given as an appropriate fusion-system-level metric. Clearly, the different choices in uncertainty representation approaches will affect the possible timeliness of a system. But non-uncertainty representation factors also play a role. Some are clearly not uncertainty related³, such as overall system architecture (centralized, distributed, etc.), data management processes, and feedback / input control processes (i.e. Level 4 fusion considerations). Other fusion system aspects may or may not be entangled with uncertainty representation considerations

During the course of this research, we have debated on the overall requirements for an IFS and defined the boundaries of the system as depicted in Figure 1 below, which was taken from [1].




Figure 1 - Boundaries of the Uncertainty Representation and Reasoning Evaluation Framework

There are two elements in the picture that are exogenous to the evaluation framework, named in the picture as “World being sensed” and “World being reported.” Between these two external elements, the boundary of the evaluation framework encompasses the way uncertainty is handled when data is input to the system, during the processes that occur within it, as well as when the final product is delivered to the IF system’s users.

Figure 1 was used as the external systems diagram for defining not only the boundaries of the system but also the main functions it has to perform. The derived requirements were key to define the criteria (item 3.1.3 below). More details can be found in the evaluation criteria session of the ETURWG website⁴.

³ At least not to the first order. There might be interactions between the uncertainty representation approach and these system factors, but we’ll begin with a presumption that they are not significant. This presumption, of course, should be explored by the working group as we address these considerations.

⁴ http://eturwg.c4i.gmu.edu/?q=URREF_Scope

3.1.2. Subtask A2: *Derive a balanced set of use cases encompassing a range of situations requiring HLIF in network-centric information exchange*

This Subtask was completed in full. The Use Case area⁵ of the ETURWG website includes six complete use cases that have data sets available for public use and one use case that does not have a data set due to the sensitive nature of its data. The use case area also contains additional information to support users of the use cases. Each use case follows a template that includes information that the group found to be necessary for leveraging the use case in IFS evaluations. Finally, a page containing an evaluation plan for each use case is linked to its main page. The use cases and their summary descriptions are listed below.

Use case 1 - Ship Locating and Tracking:

Within the available data on shipping worldwide or focused on an area of interest, either find a specific ship based on some indication that it is of interest or determine that a ship is of interest based on observed behavior. Then, resolve the uncertainty associated with the ship location and identification, reduce the uncertainty (i.e. finding the target location – position uncertainty), identify the target (classification uncertainty), and address the issue of data uncertainty.

Use case 2 - Situation Assessment:

This is a "Determine what the story line is" type of scenario. That is, some event occurs that triggers an inquiry to determine what the situation is. In this specific scenario, there is neither *a priori* identification of the specific class of situation (which distinguishes this use case from use case 1), nor a reliable source describing the situation. These must be inferred from the data.

Use case 3 - Vehicle-Borne IED:

This scenario is focused on a VBIED (Vehicle-Borne Improvised Explosive Device) attack on an administrative building. It assumes prior information about an individual under surveillance due to previous unstable behavior who drives customized white Toyota (WT) vehicle, and includes a series of related observations and its associated sources. The base data set is the freely available SYNCOIN, and the emphasis of the use case is on the ability to the theoretical framework to represent and integrate information from physics-based sources (hard) and from human-based sources (soft); as well as to deal with conflicting pieces of information.

Use case 4 - Image Fusion and Tracking:

This scenario focuses on the monitoring of activity from EO/IR cameras. It is assumed that the data are continuous, registered, and ready for analysis. Numerous data sets are available from such sources as www.imagefusion.org as well as data sets from various programs listed in the ETURWG page for this use case. Uncertainty in this use case mostly comes from the target identification (classification) and the target location accuracy. Numerous issues are needed in the analysis as to the “completeness” of the situation awareness of the objects, their activities, and predicted outcomes.

Use case 5 - Vehicle Identification:

⁵ <http://eturwg.c4i.gmu.edu/?q=usecases>

The goal of this use case is to classify each target in three different scenarios. Each scenario consists of a single track of a target moving through an area of interest. The dataset supplied is entirely synthetic. The target trajectories consist of straight-line segments and each target's kinematic behavior is based on openly available information combined with assumptions about its mobility. As such, the target trajectories provided within this dataset are not assumed to be realistic. The original purpose of this data was not to create realistic scenarios — it was purely to create scenarios that were more realistic than existing ones that consisted of simple linear target trajectories with no concept of terrain.

Use case 6 - Asymmetric Threat Detection:

This is a “Needle-in-a-Haystack” problem that also includes the possibility that several asymmetric threats are developing in parallel. The possible types of threat are very diverging but must nevertheless be regarded in parallel. Goals include finding the data relevant for assessment of asymmetric threats, making a forecast about developing threats, assessing possible course of actions and which types of threats are most likely to occur, and alert the user about relevant information on the threats.

Use case 7 - CBRN:

This use case deals with environmental monitoring in the Port of Rotterdam, the Netherlands, a densely populated industrial area with a high density of chemical plants. The main goal in this use case is to detect invisible gaseous substances and localize the origin of the pollution (i.e. localize the leak). Usually such pollution is caused by routine operations (e.g. cleaning the systems) or incidents of different scales. In most cases the released substances are not fatal, but frequent exposure is harmful for the residents. In order to be able to enforce the regulations in this area, the local agency needs the capability to detect the pollution and localize the source as quickly as possible. The use case belongs to a class of problems that can be characterized by dynamic phenomena spreading over large areas. Other examples are water pollution, spreading diseases and flooding. Due to the classified nature of the data, experiments with this use case will be only performed with synthetic data sets, which are currently not available at the ETURWG webpage for this use case.

3.1.3. Subtask A3: Design preliminary criteria for evaluating the performance of the techniques

This task was completed in full. The preliminary criteria are depicted in Figure 2 below, which shows the criteria class of the ontology of Uncertainty Representation and Reasoning Evaluation Framework (URREF). More details, including the complete description of each class and subclass, the associated OWL file, and the discussions regarding each item can be found in [1] and in the ETURWG website page on the first version of the URREF ontology. The associated URL is http://eturwg.c4i.gmu.edu/?q=URREF_Ontology.




Figure 2 – The Criteria Class of the URREF Ontology

The Criteria class of the URREF ontology was the first complete set of criteria to be used in the evaluation framework, and was the basis of most of the ETURWG work until Fusion 2013. It is the main class of the URREF ontology, and it is meant to encompass all the different aspects that must be considered when evaluating uncertainty handling in multi-sensor fusion systems

3.1.4. Subtask A4: Present preliminary results and discuss project plan and objectives at Fusion 2012

This Subtask was completed in full. In Fusion 2012 we convened an ETURWG face-to-face meeting, a special session, and two panels - “Uncertainty Evaluation: Current Status and Major Challenges” [2], and “Issues of Uncertainty Analysis in High-Level Information Fusion” [3]. Detailed material on all these activities can be accessed via the URL <http://eturwg.c4i.gmu.edu/?q=fusion2012eturSS>.

3.1.5. Subtask A5: Develop test procedures and datasets to perform evaluation

This task was completed in full. Test procedures and the datasets are available at the ETURWG website. The associated URL for the datasets is <http://eturwg.c4i.gmu.edu/?q=datasets>.

3.1.6. Subtask A6: Perform preliminary evaluation

This task has been partially executed. The original plan for the preliminary evaluation was to apply the procedures and data sets developed under Subtask 3.1.5 above to the use cases and criteria developed under Subtasks 3.1.2 and 3.1.3, respectively. This evaluation was scheduled for presentation at Fusion 2012. However, the group deemed the preliminary version of the criteria as too immature for conducting serious evaluations of uncertainty (cf. ETURWG general meeting minutes at <http://eturwg.c4i.gmu.edu/?q=minutes>). As a result, it was decided that a new round of discussions was needed. Section 4 presents some of the aspects of the aforementioned discussion.

3.2. Task B: Develop Evaluation Framework

3.2.1. Subtask B1: Refine evaluation criteria from task A

This task was completed in full. The ETURWG held discussions over the period of time between September 2012 (after Fusion 2012) and February 2013 (when the updated version of the URREF was defined). The completion target of early 2013 was met, allowing use of the URREF in papers and evaluations submitted to Fusion 2013. During this period, various terms were adjusted and compatibility with standards was sought.

However, it is important to note that the evaluation criteria are *still being modified* through the ETURWG meetings at the time of this writing. Two factors originated this delay in reaching a commonly agreed version of the evaluation criteria. The first was the complexity of the problem addressed resulting in a need for extra work to achieve its goals. After the refined version of the URREF ontology was used to support the various papers submitted to the ETUR section, it became clear that there were issues that we did not oversaw and that would jeopardize the overall evaluation framework's goal of getting community acceptance. In short, to address these issues we had to rethink the original planning and allocate an extra round of discussions and testing, which extended the final version of the framework to Fusion 2014.

The second factor was the complex scheduling, involving various researchers from different countries, located at diverse time zones, and with different constraints in terms of availability. These scheduling limitations, when added to the fact that all collaborators are volunteers, resulted in a few meeting cancellations that impacted the overall schedule.

These two factors, when taken together, impose a limitation in terms of how fast the group can make progress. Still, the group is clearly committed in reaching the goals of the evaluation framework and moving forward to a commonly agreed final version of the evaluation criteria.

3.2.2. Subtask B2: Design evaluation plan

This task was completed in full. The plan was presented in Fusion 2012 and updated in Fusion 2013. The latest version of the evaluation plan is the basis of the current discussions being held at the ETURWG meetings. Basically, the work was divided into nine different subtasks that together were intended to result in a process for using the evaluation framework. These tasks are listed at <http://eturwg.c4i.gmu.edu/?q=tasks>.

3.2.3. Subtask B3: Convene session at Fusion 2012 to define final evaluation plan

This task was completed in full. Refer to the item above.

3.2.4. Subtask B4: Define final criteria for evaluation contest at Fusion 2013

This task has been revised and is currently in execution. An evaluation contest at the current stage of the research was considered inappropriate by the group, as it could generate unwanted competition among people when a collaborative effort should be the focus. Therefore, we devoted the events at Fusion 2013 to evaluating the framework itself and to addressing the potential limitations in terms of applying it to a software platform. This is an ongoing discussion, and Figure 3 depicts the changes proposed so far to the original URREF.




Figure 3: Fundamental Reasoning Process

3.2.5. Subtask B5: Perform final evaluation

This task has been partially executed. The original plan for the final evaluation included the refinement of the work performed in the preliminary evaluation, the development of a formal framework to support the comparisons between approaches for uncertainty representation and reasoning, and a direct comparison between two of these approaches (Bayesian Theory and Dempster-Shafer Theory). We performed the framework refinements and developed the formal support for the modeling. However, as mentioned above, a formal comparison was not conducted since it relies on the ongoing collaborative work from Subtask B4, which is the current focus of the ETURWG general meetings.

3.2.6. Subtask B6: Convene session at Fusion 2013 to present the results of the contest

As mentioned above, the work at Fusion 2013 focused on evaluating the framework and planning its refinement for the final version, which will be presented at Fusion 2014.

4. Uncertainty in HLIF Applications

This Section discusses the main aspects of the framework and its use in HLIF applications. This material supports and expands on Section 3.1.6. Subtask A6: Perform preliminary evaluation (item 3.1.6. above in this document). It also includes the partial results obtained in Subtask B5: Perform final evaluation (item 3.2.5. above in this document). As mentioned in Section 3 above, the preliminary evaluation focused on assessing the various aspects of the evaluation framework.

Work on the final evaluation (Subtask B5) refined the evaluation framework in various respects, including the formal aspects of the framework. Results of this latter work are included in the present Section. The planning for Subtask B5 also included a direct comparison between two uncertainty representation approaches using the framework. As mentioned in Section 3 above, this latter part has been postponed pending the completion of the final version of the evaluation criteria.

4.1. Scope

Uncertainty propagation in HLIF applications is affected by a number of considerations, which must be understood as a key aspect in the evaluation of uncertainty management methods in HLIF. We specifically addressed JDL level 2 and leveraged the discussions held by the ETURWG prior to Fusion 2013, as well as the initial versions of the URREF ontology. Six different HLIF types are identified, ranging from simple entity attribute refinement using situation status data to the development of a complete situation assessment assembled from applicable situational fragment data. Additional considerations include uncertainty handling in the input data, uncertainty representation, the effects of the reasoning technique used in the fusion process, and output considerations. Input data considerations include the data's relevance to the situation, its credibility, and its force or weight. Uncertainty representation concerns follow the uncertainty ontology developed by the W3C Incubator Group on Uncertainty Reasoning. For uncertainty effects of the fusion process, a basic fusion process model is presented, showing the impacts of uncertainty in four areas. Finally, for output uncertainty, the significance of a closed-world versus open-world assumption is discussed. A more comprehensive account is currently under review by the Information Fusion journal⁶. The paper includes the above considerations and the proposal of a formal framework for analyses of processes occurring at JDL fusion level 2. This formal framework provides a basis for correlating such processes with the uncertainty considerations.

⁶ Locher, Mark; and Costa, Paulo C. G. Ignorance and Uncertainty in Level Two Information Fusion. Submitted to Information Fusion in Sept 2013.

4.2. Nature of Level 2 HLIF

Any level 2 HLIF involves both ontological structures and logical reasoning, whether formally recognized or not. An ontology specifies the classes of entities that may exist, their attributes or properties, and relationships between class members. It is an abstraction of data models, with a focus on modeling knowledge about individuals, their attributes, and their relationships to other individuals [5]. Ontologies provide a logical, structured approach to organize the information content of a problem domain. In doing so, they also provide a basis for the semantics (the meaning of the information) necessary for effective reasoning within a HLIF process.

Although ontologies can be developed and defined using data modeling tools such as UML or entity-relationship diagrams, an ontological approach requires that semantic content be explicitly preserved in the development approach, and that semantic information not be hardcoded into the data management processes. In data modeling, the semantic content is normally discarded once the conceptual schema is translated into the physical schema [6]. Most ontology tools and processes used today are deterministic – they assume a binary true / false logic. As such, they do not explicitly incorporate various forms of uncertainty in their structure. There has been significant work recently in using ontological approaches in developing fusion techniques; some have taken uncertainty considerations into account (e.g. [7]–[11]).

In addition to being ontological, HLIF involves extensive reasoning. In any fusion process, one follows a fundamental reasoning process, which logically uses a series of reasoning steps, often of an “if, then” form. Beginning with a set of events, one forms a chain of reasoning to come to one or more conclusions. Figure 4a models a simple case, while Figure 4b gives an example of that case. More complex structures can be easily created [12].




Figure 4: Fundamental Reasoning Process

This reasoning is often expressed symbolically, using a variety of logic-based approaches [13]. The reasoning involved most often exceeds the range of logics normally considered (usually a fragment of first order logic, one of several variants of description logic) in an ontological approach, and can take any form of deductive, inductive, and abductive reasoning.

4.3. Defining a Level 2 HLIF Taxonomy

In examining various level 2 HLIF use cases and problem sets, we were struck by the significant diversity among them. We suspected that this diversity potentially masked significant differences in uncertainty considerations, differences that went beyond mathematical technique considerations. We did not find a significant discussion of this in the literature, so we developed a taxonomy of level 2 HLIF that categorizes how uncertainty differences could be captured and presented. In doing so, we were greatly aided by a top-level entity categorization developed by John Sowa. He defined twelve ontological categories for all entities, which together comprise a very attractive framework for analyzing fusion processes at level 2. He suggests that one way of categorizing entities in the world is to consider them from three orthogonal aspects. For more details, the interested reader is referred to [14].

In developing our taxonomy, we reviewed a number of use cases and scenarios to see what was considered to be a level 2 HLIF within the community. We examined the case studies provided for the ETURWG. These case studies covered a wide variety of areas, including suspicious ship detection, image fusion, counterinsurgency support, multi-sensor vehicle tracking, asymmetric threat detection, industrial area chemical plume detection and characterization, and various criminal situation assessments. We also reviewed fusion systems documented in the literature, especially [15], and system work such as [16] and [17] performed at George Mason University.

The use cases and examples range widely in complexity and the analytic inferences they require. Four findings were evident. One, there appears to be a tendency to equate multi-sensor fusion with high level fusion, even if the purpose of the fusion is solely to enhance an understanding of a single entity, without consideration of other entities or relationships (e.g. improve a track estimate). Second, the role of time, situation-knowledge and context knowledge varied significantly between cases, affecting the levels and types of uncertainties they brought into the case. Third, the outputs (the supported hypotheses) of the fusion process were wide-ranging, from an elementary assessment of the existence of an additional entity based on high-likelihood relationships, to the complete development of an entire case history (a fully fleshed out situation, such as a forensic reconstruction of a criminal case). Fourth, not all level 2 information fusion scenarios had to infer a situation state. Rather, in some cases, the overarching situation state is given in the data, but inferences are required to provide specific details of interest about the situations. This insight became a core factor in the proposed taxonomy

Shown in Figure 5, the taxonomy begins with an understanding that all situations involve a purpose or reason. But this does not necessarily mean that the purpose or reason is unknown. In a number of cases, the overarching purpose or reason is known, but one or more key aspects about the existence of entities, their attributes, or their specific relationships with other entities are unknown and of interest. After this first division, there are three subdivisions in each half, making a total of six cases.




Figure 5: Taxonomy of Situation Assessment Cases

In general, the least complex case is for known situations where one is determining the existence of an entity, based on its relationship to another. This case straddles the level 1 / 2 line. It is object / process identification where the relationship between entities of interest may vary. An example is the radar / vehicle case above. The defined situation is that a Tin Shield radar has been detected at a particular location. The question is whether an SA-10 battery is at that location, or whether the radar is operating in a stand-alone mode (whether operationally, for system testing, or for system maintenance). The inferences generally are based on schema-based evidential reasoning (e.g. “there is a 95% chance that this radar will be associated with an SA-10 battery in its immediate vicinity”).

When the reason or purpose is known, but multiple entities are involved (more than a few), there is a significant step up in complexity. One is moving from a strict focus on the relative level to the mediating level. The second case is where the situation is well defined but the objective is to identify a specific object of interest within the situation. For example, one might have very credible evidence that a terrorist group will attempt to smuggle a radiological bomb into the United States via a freighter. In this case, the situation itself is known (one knows the purpose / intention), but the actors may be hidden. Inferring which freighter (an object identification) is a likely carrier of the bomb is the question of interest. Another example would be to determine who committed a robbery of a bank, when one has a video of the act itself (the situation is a robbery). In this case, the evidence is extracted from a variety of sources, which can be classified as being junctures, participations, histories or descriptions.

The third case is where the overall purpose or reason is known, but the details of the associated entities must be identified and their interrelationships developed. Building a military order of battle is an example. In developing an enemy order of battle for a nation-state’s standing military, one has a basic understanding of the objects and relationships that constitute a modern military force. One has a series of templates that can be used in developing aspects of the structure. A country may not have all of the elements, and the organizational structure will vary. Yet, it is very likely that the structure and deployment will follow patterns similar to those used by other countries.

The inferential process generally becomes more complex when the specific situation itself is not known, but must be inferred. The taxonomy outlines three such cases, each with an increasing level of complexity. The first is when the specific situation is not known, but there is a set of well-defined situation choices to select from. This case is a situation version of a state transition. A classic example is the military indications and warning question, which can be raised when an increase in activity at military locations in a country is detected. The question then becomes “what is the purpose of the activity?” Four major choices exist: a major military exercise, suppression of domestic unrest, a coup d’état, or preparing to go to war. Each is a relatively well-defined situation with known entities, attributes and relationships. Another case is where one is continually monitoring an unfolding situation to determine when major transitions are occurring. The selection among them becomes a pattern-matching exercise.

The next level of complexity occurs when not only is the situation itself unknown, the situation itself must be developed. Unlike the case above, the issue now is not choosing among a set of possible situations but to build the situation or structure from the data. Structure or situation development can be divided into two subcases. In the first subcase, one has to refine the situation presented. For instance, assume that one decides that the purpose of the unusual military activity is to go to war. Immediately, a series of questions arise. What is the overall purpose of the war? What are the expected military objectives? What is the specific war plan? In this case, refining the situation is necessary in order to support a decision-maker’s specific needs.

The final case is the most complex one. Here, one must develop a situation where the basic purpose itself must be determined. For example, consider the case when a government agency is notified that something is significantly amiss at a particular location, with enough information to spark interest, but not enough to understand what is happening. In that case, the evidence must be assembled without a common template to guide the fusion. Rather, the evidence must be fused using fragmentary templates, that themselves must be integrated to provide the overall situation. Such forensic constructions require significant reasoning skills, and involve major issues of uncertainties. This case also includes much of the predictive intelligence workload. Integrating the data to “connect the dots” that could have predicted the September 11, 2001 commercial airliner strikes on the World Trade Center and the Pentagon falls into this category.

4.4. Uncertainty Types

Uncertainty varies in its forms and manifestations, and these different forms of uncertainty affect the performance of a fusion process. But not all uncertainty representation schemes can accommodate all possible forms of uncertainty. Understanding what specific forms of uncertainty might exist in a particular fusion problem is necessary to select the appropriate uncertainty representation scheme. Several authors have summarized views on uncertainty (e.g. [18]–[21]). Although we have considered all of them in the present discussion, our main reference was the Uncertainty Ontology [22], depicted in Figure 6, which was used in the discussions at the ETURWG meetings and incorporated into the URREF ontology.




Figure 6 – Uncertainty Ontology

4.5. Uncertainty Model

Aligning the various types of ignorance and uncertainties with a level 2 fusion process model requires having a satisfactory model with sufficient process detail to account for all of the identified uncertainties. According to Salerno, over thirty fusion process models had been proposed by 2002 [23]. Several teams have reviewed selected subsets, including Esteban et al. [24], Bedworth and O'Brien [25], Whitney, Posse and Lei [26] and Roy et al. [27]. Kokar, Bedworth and Frankel categorized many of the models by whether they were information centered or process centered. Information-centered models focus on differences in the level of abstraction of the data. They identified the JDL model, Dasarathy model and the Waterfall model as being information-centered. Process centered models are organized from a functional point of view, and included the intelligence cycle model, Boyd control loop and Omnibus model [28]. Carvalho et al.'s General Data Fusion Architecture [29] can also be classified as an information-centered architecture.

A number of authors weighed in on the architectural consideration of a fusion processing system. The papers listed above each included some architectural considerations. One way to specify a fusion architecture is along four independent dimensions:

- Centralized versus decentralized structure
- Local versus global component interaction
- Modular versus monolithic component design
- Single versus multiple control centers [30]

Authors who explicitly addressed HLIF functions in significant detail include Roy et al. [27], Salerno [23], Steinberg [31], and Schum [32]. Roy and his associates examined in some detail four different fusion models. They provided a significant history of the JDL model, from its origination in 1985 until 2006. They then reviewed the Visual Data-Fusion model. This is an extension of the JDL model that emphasizes the visual presentation of information in a manner that maximizes the human's understanding, without overwhelming the human with detail. This allows the human to solve the problem at hand. From there, they explored Lambert's Unified Data Fusion Model, which uses a two-dimensional categorization of situation assessment. One dimension uses Endsley's situation awareness decomposition, into Perception (data ingest and understanding), Comprehension (awareness) and Projection (predictive understanding). The other dimension looks at the human (psychological), machine (technology) and human / machine integration aspects of a fusion process. This structure allows for the mapping of specific fusion

capabilities to their function and performers. The final area they reviewed was Salerno's conceptual functional flow model [26].

Based on these models and our understanding of the various uncertainties potentially involved, we developed the comprehensive fusion process model shown in Figure 7, which was expanded from the original idea in [33]. The first thing to observe is that the raw data can come in at any level, as evidenced by the incoming arrows at the right side of the figure. The model does not require that all data be signal or feature (Level 0) data, which is then aggregated into higher-level conclusions. For instance, object identification data (level 1) could come from an on-scene observer or from an image analyst reporting on an image. Communications intercepts or human reporting could provide evidence on relationships (level 2) or future intentions (level 3). Note that if a level 3 fusion process is active, its outputs could affect the level 2 process in two places. It can either be a controlling variable in the fusion process itself, or it can affect the interpretation and extraction of evidence. However, a level 3 process will have an effect only if it has separate evidence that is not being used in the level 2 fusion process (otherwise one has circular reporting).




Figure 7 – Level 2 Fusion Process Model

There are six basic processes in this model. The structuralization and data alignment process structures the incoming data and brings it to a common reference base, appropriate for follow-on use. Level 2 processes, unless executed solely by a human, requires structured data, organized by some ontology. However, many level 1 and 2 systems today transmit free text reports, not structured text. They are meant to be human readable, not machine-readable. Some form of entity and relationship extraction is required to transform those reports into machine-understandable data. If data are already structured according to an understood ontology, then this process is unnecessary. Data may also come in with different reference bases, and need to be aligned to a common baseline in order to be used in the extraction and fusion processes.

The core process is evidence marshaling, the process by which the evidence is organized and evaluated. This process can use multiple schemes to arrange the data to best provide insights into potential arguments and hypothesis. Unlike the model proposed by Schum [32], we moved hypothesis management and argument management into their own sections. Evidence marshaling does include a function that identifies what additional data are needed to support or disconfirm a hypothesis, and requests those data.

Another important aspect of this model is that not all of the evidence that goes into the model-based process is assumed to be in an immediately usable form. There is an evidence extraction process as part of the fusion process, which some data undergo to have the appropriate evidence extracted from them. For example, the raw level 2 data may be a series of people association data, which must be combined into a social network analysis to reveal the full extent of the relationships. Observe that the uncertainty representational scheme used plays a significant role in establishing the kinds of uncertainties that can be assessed in the evidence extraction process. For example, if the incoming data is heavily ambiguous, but the evidence extraction process has no mechanism for representing that ambiguity, the evidence output may be specified as being more definitive than the data warrants.

Another example may be that one is interested in whether two ships met and transferred cargo in the open ocean. Suppose there is a track file on each ship, which has long revisit rates between collections. This does not provide an obvious indication that the ships met and stopped for a while. But the track files show that both ships were on tracks that did put them at a common location at a given period, and that the average speed dropped significantly during the time a meeting could have occurred (implying that the ships may have stopped for a while). Given these data, one could conclude with some level of certainty that they did meet and stopped to transfer something. The level of certainty is driven by at least two factors: the quality of the track file data (establishing how certain one is in concluding that the tracks allowed them to meet), and how likely is it that two ships showing these track characteristics actually would have met and stopped.

A significant part of the evidence extraction process could be comparison to historical or reference data. For example, a vehicle may be moving outside of a normal shipping lane / airway or off-road. This requires a reference to a map base. For this reason, the process model includes a data store, for both reference information and for previous data.

The fifth process is hypothesis management. This process maintains the active hypotheses under consideration. It also is involved in generating hypotheses and in the pruning of hypotheses.

The final process is the argument management process. This process can take one of two forms. In less complex situation assessments, it takes some form of direct symbolic reasoning, often a model-based process. To account for the uncertainty in the data and process, current models often take the form of Bayesian networks [31], [34], [35], although alternative approaches have been proposed using graphical belief models [36] and general purpose graphical modeling using a variety of uncertainty techniques [37]. For more complex situation assessments, such as forensic reconstruction, the argument management process is a meta process, responsible for constructing the model used to provide the response. As such, there is a close interaction

between argument management and hypothesis management. The argument management process is the process that provides the level 2 data output.

Finally, note that the level 2 process includes the possibility of a direct use of level 0 data. An area of active research is the multi-source integration of level 0 data that is not of sufficient quality, or that does not have enough quantity, to allow a high quality single-source conclusion.

4.6. Uncertainty Considerations

Within HLIF level 2 there are at least three areas for uncertainty considerations: the uncertainty in the input data / evidence, the uncertainty effects of the argument management process, and the uncertainty in the hypothesis generation and management process. All of these affect the uncertainty in the fusion process output [1].

4.6.1. Area 1: Uncertainty in the Input Data / Evidence

All conclusions are ultimately grounded on evidence, drawn from a variety of data sources. But often evidence is “inconclusive, ambiguous, incomplete, unreliable, and dissonant” [12]. Any conclusions drawn from a body of evidence are necessarily uncertain. In Figure 6, uncertainties in the data / evidence apply to the relevance and believability of the input data, and the errors introduced by the structuralization / data normalization and the evidence extraction processes. There is a distinction often made between hard and soft data, where hard data is received from machine sources, while soft data comes from human sources. The latter often adds uncertainty due to linguistic imprecision, as well as well-documented errors and biases in human measurement and judgment. Yet, as M. Pravia et al. point out, very few machine sources provide data without human involvement. Much machine-generated sources of data require human interpretation in order to be of use in a fusion process. For example, interpretation of imagery data and translation / transcription of voice communications today require humans to perform [38].

Regardless of the source of the evidence, Schum [12] found that one must establish the credentials of any evidence used in a reasoning process. These credentials are its relevance to the question / issue at hand, its believability and its weight or force. The force (or weight) of the event establishes how important the existence of that event is to the conclusion one is trying to establish, which is related to the Area 2 explained below.

Data becomes evidence only when it is relevant. Relevance assesses whether the evidence at hand is germane to the question(s) being considered. Irrelevant information makes no contribution to the conclusion drawn, and potentially confuses the fusion process by introducing extra noise. Relevance is specific to a hypothesis, and can vary (even to the joint irrelevance) between hypotheses. Evidence can be either positively (supportive) or negatively (disconfirmatory) relevant to a particular hypothesis. Any analytic effort is obliged to seek and evaluate all relevant data.

Once data are shown to be relevant to a particular problem (i.e., it becomes evidence), Schum points out that there is an important but often overlooked distinction between an entity and the evidence about that entity. Joe’s claim that “I saw Bob hit Bill with a golf club” does not mean that such an event actually happened. The claim should be seen only as evidence of the event.

Believability establishes the considerations as to how trustworthy a piece of evidence is about the entity it reports on. Table 1 gives the five elements of believability [17].

Believability and relevance uncertainty exist in the data that are fed into the fusion process. Entering data are subject to uncertainty arising from confusion or inaccuracies in the structuralization and data normalization processes, and from the evidence extraction process. Context mismatches and term indeterminacy issues are potentially significant contributors, as is uncertainties that arise from incompatibilities in the world context between evidence items.

Table 1: Elements of Evidential Believability

Veracity: Source is telling what it believes to be true (note that the source may be deceived)
Objectivity: Source has received the evidence on which it based its reporting. This includes consideration of system biases and false alarms
Observational Sensitivity Source has the ability to actually observe what it report: (e.g. Observer actually has the visual acuity needed to see what was going on, or an electronic intercept was of such low quality the operator guessed part of the conversation)
Competence: Source has the ability to correctly understand what it reports (e.g. the source was in the right location and had the proper capabilities to make the report).

Believability includes a number of data acquisition considerations. Lapinski identified several sources of error that can occur during the data gathering process (see [39]).

4.6.2. Area 2: Uncertainty in the HLIF Arguments

As noted earlier, the argument management process is where the evidence is fused via a logical reasoning process. Uncertainty in the argument management process depends on the force of the evidence, the structure of the argument process and the marshaling process used to construct the arguments.

The force of the evidence establishes how important the existence of an entity is to the conclusion one is trying to establish. Note that force here assumes that the entity exists. The earlier believability considerations assess how well the evidence supports the existence of the claimed entity. Force measures the strength of the relationship between an entity A and an outcome B (B can be another entity, an intermediate hypothesis, or the final hypothesis). This varies significantly between the event and the outcome. For example, the event “Bob hit Bill with a golf club” would have a significant force in establishing a conclusion that Bill was seriously injured. It would have somewhat less force in establishing that Bill was committing an intentionally violent act, and even less force in concluding that Bob was angry with Bill.

The strength of the force is captured by a conditional association. Numerical values are often established by an inductive reasoning process (such as a data mining or an expert elicitation process), although it can also be derived using abductive or deductive reasoning. In a fusion process, forces are typically established in an off-line, or separate, process, rather than occurring during the evidence fusion process. For example, if a Bayesian network model is used, force of evidence is usually measured by likelihood ratios, which are established during the initial fusion system development process, and are stored in the system's database. But in some cases, such as some forensic reconstruction cases, selected forces can be developed as part of the fusion process. In either case, developing the specific value(s) for force of evidence requires careful consideration of the conditions under which the value is applied, the uncertainty representation scheme being used, the evidence base from which the value is defined, and the analytic process by which it is developed. Most commonly, evidential force is measured by a single number, although there has been some exploration both of using linguistic values and of using interval schemes.

The argument structure plays a key role in the argument uncertainty assessment. The best-known and probably most widely used fusion structure is the combination of independent sources of like data. Averaging with associated variance reduction is a prime example. The conjunctive combination of unlike evidence also reduces uncertainty (e.g. John was seen leaving the crime scene immediately after the shooting AND the bullet markings on the fatal bullet match those produced by the gun he owns AND.....). Schum investigated a number of different ways in argument structures affected the uncertainty in the final conclusions. He observed that while having independence in the data simplifies computation, cases involving conditional nonindependence often provide surprising results. He especially focused on cases where there was a relationship between the believability of the evidence and the final hypothesis. He demonstrated that, under certain circumstances, having strong data on the believability of a data source can have a more significant force on the conclusion than the force of the event reported in the data (e.g. knowing that a witness is lying can have a stronger effect on a jury's assessment of a defendant's guilt than the topic about which the witness is lying). For more on argument structures for fusion processing, see [12].

A critical item in uncertainty propagation is the proper fit between the types of uncertainty in the input data and the representational model(s) used in the fusion process. Failure to account for all of the uncertainty types in the input data can result in an erroneous process output. Various approaches can be used to model different types of uncertainty in a reasoning process. These include (but are not limited to):

- Bayesian Probability
- Dempster-Shafer Theory
- Possibility Theory
- Imprecise Probability
- Random Set Theory
- Rough Sets
- Certainty Factors

A classic survey of uncertainty models, with a discussion on applicable uncertainty types, is given in [40], with a recent review state-of-the-art in [41].

4.6.3. Area 3: Uncertainty in the Hypotheses

The output of the fusion process is one or more hypotheses that are best supported by the evidence, together with associated argumentation connecting evidence to hypotheses. A key concern is whether the set of considered hypotheses include the hypothesis that best conforms to the reality being assessed. If that hypothesis is not included, the quality of the fusion process can be significantly degraded. Depending on the particular problem being addressed, hypothesis management can be a crucial process. There is, in general, no way to guarantee that the hypothesis management process will correctly generate or maintain the correct hypothesis. For some problems, generating hypotheses can be relatively straightforward, but the numbers can grow beyond the practical computational capacity of real-world systems. For those problems, a viable hypothesis pruning process is required, which has a low possibility of eliminating the correct hypothesis. For other problems, generating and refining the set of hypotheses is a significant issue.

There are several indicators available that can signal that the hypothesis set is incomplete. A key problem is that the signaling is ambiguous. One indicator is that there is a conflict (contradiction or contrariness) in the evidence set. Evidentiary contradiction often results from believability issues with one or more pieces of evidence, caused either by deception in the source, or some error in the collection / transformation process. But contradiction can also occur if the correct hypothesis that unifies the apparently contradicting evidence is missing. A missing hypothesis is more likely if the evidence base is contrary, where different evidence items strongly support different hypotheses, but do not directly contradict the other hypotheses. Such conflict should be noted, and used as a guide to seek additional hypotheses.

4.7. Integrating Uncertainty into the Process Model

Figure 8 maps the various uncertainties identified in Subsection 4.4 to the process model laid out in Subsection 4.5. The lower half of the figure organizes the various uncertainties in line with the structure suggested by the W3C Incubator Group on Web Uncertainty.




Figure 8 – Level 2 Fusion Process Model

We begin with a set of sentences, both evidentiary sentences and hypothesis sentences. In examining the uncertainty, we consider the sentences both individually and as a collective. Individually, we observe that the sentences are stated by an agent, they possess uncertainty, and they are about a world. The agent (whether human, machine or both) has uncertainty considerations as to the truthfulness of the sentence. Following Schum, these considerations can be divided into veracity (the agent reports what it believes to be true), bias (the agent reports what it observes / senses without an offset), observational sensitivity (the agent has the necessary capabilities to observe what it reports) and competence (the agent has the expertise and position to be able to make the report). Observational sensitivity encompasses a number of uncertainty sources, such as reported in [39].

The agent is reporting on something observed in the world. That observation has a context. It is about one or more entities, and it has a setting. The setting is normally a geospatial-temporal one that defines when and where the observed entity(ies) are located. But the setting can be virtual, or undetermined (e.g. a telecommunications intercept may say very little about where the communicating entities are actually located). Finally, the agent's perspective may introduce uncertainties in the sentence.

Each sentence also has a degree of uncertainty. It is important to assess the types, nature, and derivation of the uncertainties in a sentence. A key observation here is that a sentence may have multiple types of uncertainties in it. It could be vague, either because of an inherent fuzziness in the reported item (“John is near the plaza”) or because of nonspecificity (“The car was seen somewhere in the city”). The uncertainty could be empirical, that is, its truth status is unknown, but is theoretically determinable in specific instances. Entities whose status follows statistical laws (e.g. random events) or natural variations (e.g. predator-prey cycles) have empirical uncertainty. There could be ambiguity in the sentence, in which one is uncertain to which specific entity the sentence is referring to (e.g. “The victim had either a heart attack or a stroke”). Lastly, the process by which the agent determined the sentence may have introduced confusion, which results in an incorrect statement.

This is an especial source of uncertainty when an agent is forced to prune possibilities during an estimating process, such as in track hypothesis pruning in a multiple hypotheses tracking process, or when the agent states only the most likely outcome, rather than providing the range of viable outcome along with their estimated uncertainty. Context mismatch occurs when one fails to specify the context in which the sentence is to be understood, and an inappropriate context is used. The final area is term indeterminacy. Meanings of terms can change over time, and terms can lose their original meanings in translations from one language to another.

But sentences also have uncertainty considerations when they are viewed collectively as a set. Implicit in this assertion is that there is some underlying factor that makes a set of sentences, that the sentences all have some relevance to a particular problem of interest. Given this, the set most often is incomplete; one or more sentences that could assist in solving the problem are missing.

A second area of uncertainty is that there is conflict in the set. This conflict can come from three different sources. The first is outright contradiction, resulting in a logical inconsistency in the set. A second area is if sentences are contrary. Contrary sentences are ones that cannot both be correct, but both can be wrong (unlike contradiction, in which one must be right). The third area of conflict is incompatibility in the context of the sentences. An implicit assumption is contradiction or contrariness is that the conflicting sentences are both stated in the same context; they refer to the same entities and situations. But since sentences in a fusion system can come from a variety of agents, they may not have the same context. This can lead to conflict if not recognized and accounted for. As stated above, context can be misinterpreted, distorted or lost.

The upper half of the figure summarizes the uncertainty generation in the fusion process. A fusion process is the bringing together of evidence and hypotheses in order to address a problem for some user. We can split these uncertainties into three main topics: uncertainty in the evidence, in the argument, and in the hypothesis.

The uncertainty in the evidence is derived from believability considerations, which tie back to assessments of the believability of the agent, and can also arise when one has conflict in the evidence set. Conflicts raise doubt about the general believability of all evidence sentences involved in the conflict. Uncertainty also arises in the relevance of the data. This is an interesting uncertainty, because it is specific to a hypothesis sentence, and can vary between hypotheses. Finally, a specific fusion process can introduce errors (especially confusion or inaccuracies, but also context mismatches or term indeterminacies) in the process of aligning evidence,

structuralizing the entities and relationships in a report, or in the internal evidence extraction process.

There is also uncertainty in the set of hypothesis sentences, which primarily revolve around the completeness of the hypotheses set, and in the accurate representation of the uncertainties that are expressed in the hypotheses. For example, expressing specific hypotheses in circumstances where a degree of ambiguity is appropriate affects the goodness of the estimate of the uncertainty that exists in the output of the fusion process.

In the course of developing and exercising arguments, uncertainty both occurs in, and can be reduced by, the evidentiary force and the argument structure. The evidentiary force is the strength by which a particular piece of evidence supports a particular hypothesis. The evidentiary force is affected both by the uncertainties in the individual evidence sentences, and in the collective uncertainty in the set of relevant sentences. The collective uncertainty is mediated through the argument structure. That is, how the evidence is marshaled together and organized into chains of reasoning will affect the force of the evidence. Two considerations dominate here. The first is the appropriateness of the argument structure. How well does it match the underlying reality that it is trying to model? An argument structure that fails to take into account significant relationships between the evidence will misrepresent what the evidence supports. A second consideration is the limits of the representational system that is used. For example, if it fails to accurately account for the degree of fuzziness in the evidence, then that failure will add to the uncertainty inherent in the output of the fusion process. Note that this failure may not be reflected in the system's assessment of the degree of uncertainty remaining in the output.

4.8. Formal Construct

4.8.1. Uncertainty Modeling

Several authors have developed mathematical constructs for use in assessing the uncertainty of a situation assessment [12], [22]. The model below is a modification of one put forth by Karlsson [42], expanded in scope and recast using the terminology by Franconi [43]. Karlsson's version focuses primarily on relationships, and does not explicitly include entities and attributes. While one can model entities and attributes using relationships, it aids understanding to separate the relationship descriptions from the entity and attribute spaces. In addition, the construct formed in this paper acknowledges level 2 HLIF as explicitly including entities and their attributes in the fusion process as well as relationships between entities. Including attributes as separate from entity relationships, rather than defining relationships to include attribute states, makes this clearer. A second change from Karlsson's approach is to more clearly articulate the existence and completeness concepts discussed above. As done throughout this paper, entities can include both enduring objects and transient processes / events. Per [43], the language consists of:

- E_i , the entities, stated as 1-ary predicates
- A_k , the attributes assigned to entities, stated as 2-ary predicates
- R_p , the relationships, stated as n-ary predicates

There is an interpretation function $I = \langle D, \cdot^I \rangle$ where domain D is a non-empty set $= \Omega \cup B$, Ω is the set of all entities instances, B is the set of all attribute values, and $\Omega \cap B = \emptyset$. Then

- $E_i^I \subseteq \Omega$
- $A_k^I \subseteq \Omega \times B$
- $R_p^I \subseteq \Omega \times \Omega \times \dots \times \Omega = \Omega^n$

x_{ij} are the specific instances (i indexes the entity type, j the individual count of that entity type) and $x_{ij} \in \Omega$

The attributes and relationships are defined as follows:

$A_k \stackrel{\text{def}}{=} ((x_{ij}, b_{ky}): P(x_{ij}, b_{ky}), x_{ij} \in \Omega, b_{ky} \in B)$, where $P(x_{ij}, b_{ky})$ is the predicate defining the attribute assignment; i, j are defined as above; and k, y are the attribute type index and specific attribute type value index.

$R_p \stackrel{\text{def}}{=} ((x_{11}, \dots, x_{mn}): P(x_{11}, \dots, x_{mn}), x_{ij} \in \Omega)$, where $P(x_{11}, \dots, x_{mn})$ is the predicate defining the relationship assignment.

One can define a situation as:

$S_q \stackrel{\text{def}}{=} ((R_1, \dots, R_j, A_1, \dots, A_j): P_s(R_1, \dots, R_j, A_1, \dots, A_j))$, where P_s is a situation purpose predicate, and j may vary between tuples.

The point that j may vary between tuples and still be the same kind of situation is significant, and points to a key distinction between relationships and situations. A relationship has a fixed number of elements, called the arity of the relationship. But situations do not have a fixed arity in terms of the number of relationships and attributes they encompass. Consider, for example, a birthday party as a situation. The canonical birthday celebration in the United States includes a cake with a number of lit candles on it. If there are no candles on the cake, does this mean it is not a birthday celebration? What if the cake is replaced with cupcakes or muffins? This variability in what constitutes the same kind of situation is a feature that adds uncertainty to situation assessment.

We can make at least three existence uncertainty assessments. First, there is a fundamental existence uncertainty for each entity.

$$u_{Ei}(x_{ij} \mid E_B^*, S, I) \quad (1)$$

where E_B^* is the body of evidence available for making the assessment, and S, I are any already known situation or impact states. This existence uncertainty combines all identifiable sources of uncertainty, including believability considerations and errors /inaccuracies, to give an estimate as to whether the entity really exists (or existed).

There is a corresponding attribute existence uncertainty assessment

$$u_{Tk}((x_{ij}, b_{ky}) \in A_k \mid E_B^*, S, I) \quad (2)$$

For any specific entity tuple $(x_{11}, \dots, x_{mn})$, we have a level of uncertainty as to whether that tuple is a member of a specific relationship. For a generic uncertainty measure u_T , the basic equation for whether a tuple is correctly associated with a defined relationship is

$$u_{Tj}((x_{11}, \dots, x_{mn}) \in R_j \mid E_B^*, S, I), 1 \leq j \leq p \quad (3)$$

Given a set of tuples one believes is part of an attribute or relationship, we can assess the completeness of that set. Let $\{A_k^h\}$ be the set of all binary pairs which the evidence E_B^* supports being members of the attribute A_k , and $\{R_p^h\}$ be all tuples that are a member of R_p (e.g. $\{A_k^h\} = \{(x_{11}, b_{11}), \dots, (x_{ij}, b_{ky})\}$

$$u_{Ak}(\{A_k^h\} = A_k \mid E_B^*, S, I) \quad (4)$$

$$u_{Rp}(\{R_p^h\} = R_p \mid E_B^*, S, I) \quad (5)$$

Next, we have an uncertainty measure u_s . Given a set of q possible situations and a body of evidence E_B^* from which one has extracted a set of entities, attribute assignments and relationships that one has assigned to a tentative situation

$$S_{\text{current}} = ((x_{11}, \dots, x_{mn})_1, \dots, (x_{ij}, \dots, x_{rs})_k, (x_{11}, b_{11}), \dots, (x_{ij}, b_{ky})).$$

Given this, we can assess the following uncertainty:

$$u_s(S_{\text{current}} = S_m \mid E_B^*, S, I) \quad 1 \leq m \leq q \quad (6)$$

Finally, one can compound the uncertainty measures. For instance, combining equation (1) and (3) gives us

$$u_{Tj}((u_{Ei}(x_{11})), \dots, (u_{Ei}(x_{xm}))) \in R_j \mid E_B^*, S, I), 1 \leq j \leq p \quad (7)$$

4.8.2. Application to Level 2 HLIF

We can use this model to better understand the varying complexities of the different situation assessment cases given in section 3. For the simplest case, entity existence, equation (1) is appropriate. Using the example given there, from the existence of one entity (the Tin Shield radar), we are inferring the existence of a second entity, based on the relationship between the two. For the second case, entity selection, we again have a defined situation, but now are seeking a specific object within multiple choices of objects. We are operating with equations (2) through (5) – we are seeking a specific set of attributes and relationships that ship i is the ship of interest, but we want confidence that we identified all the appropriate members of the relationship. Based on the evidence, we will create multiple tuples for the different attributes and relationships that could lead us to the ship (using equations (2) and (3)) and then combine

the results to get to equations (4) and (5). For the third case, entities development, we are operating with equation (6), with a focus on (4) and (5). We know what the basic situation is in case 3, but we are looking for completeness in terms of the entities, their attributes and relationships.

For the fourth case, structure / situation selection, we invoke equation (6) as the basic equation. We are choosing between multiple choices as to what the situation is. We use equations (2) and (3) to determine if various attributes and relationships exist, and based on those findings, determine which situation model is the correct one for this body of evidence. For the fifth case, structure / situation refinement, we again use equations (2), (3) and (6). But we also use equations (2) and (3) to determine what the exact set of attributes and relationships is. Case 5 differs from case 4 in that we are trying to determine what the relationships are that are appropriate for this situation (or structure). One can consider case 5 to be what case 3 would be if one was not certain about the situation type.

For the sixth case, structure / situation creation, we have all of the uncertainties addressed above, and we add an uncertainty not immediately obvious in the generic equations. Relook equation (6). One of the stated requirements is that we are selecting among a set of defined situations. This essentially is a closed world assumption. However, in case 6 we are building the situation, rather than determining which situation among a choice of situations is the applicable one. We still have a number of models to choose from, but they are more fragmentary than in previous cases. The previous cases represent more of a “pieces of the puzzle” approach, where one is assembling the puzzle according to one or more available pictures to help guide you. Case 6 represents the case where we one is assembling the puzzle without a picture or set of pictures to guide one. Rather, you are assembling the puzzle guided by basic puzzle rules about matching shapes and picture colors. So, in case 6, we are also determining what the applicable S_{ks} are.

4.9. Discussion

In our research we observed the existence of a number of uncertainty considerations when analyzing a level 2 HLIF. Most of these are not necessarily obvious at a first glance, suggesting the importance of a framework that supports the analytical process. The framework aims to support the analysis of processes occurring at JDL fusion level 2, providing the ability to correlate such processes with the uncertainty considerations raised so far. Figure 9 summarizes these considerations as they relate to the heart of the basic process model shown in Figure 7.




Figure 9 – Mapping Uncertainties to Fusion Model

The taxonomy of level 2 HLIF types discussed in Subsection 3.2.3 defines the complexity of the uncertainty considerations that must be accounted for. Six different types are identified, ranging from simple entity attribute refinement using situation status data to the development of a complete situation assessment assembled from applicable situational fragment data. The uncertainty in the input data / evidence must be assessed for relevance, credibility, and force / weight, per the ontology of evidence presented in Laskey et al. [44]. The representation uncertainties that drive the modeling methodologies can be classified per the uncertainty ontology developed by the W3C Incubator Group for Uncertainty Reasoning [22]. A variety of different models can be used to properly capture the aspects of uncertainty in the data [44]. Finally, the output uncertainty strongly depends on the a priori identification of possible situation choices, or upon having a fusion process that allows for an effective open world assumption. Understanding these uncertainty considerations is a key step towards a principled evaluation of the effectiveness of various uncertainty management methods in HLIF.

5. Conclusion

Our planning included constant interaction with the HLIF community during the development of the evaluation framework. This interactive approach provided and is still providing an external (to ETURWG) means to further ensure that our work is perceived as part of the efforts from the ISIF community on evaluation of uncertainty in information fusion. This recognition is an important asset in avoiding the framework to be deemed as biased towards a given approach.

The use cases exemplified complex situations involving difficult aspects of uncertainty representation and reasoning that a comprehensive theory of HLIF must be capable of handling. The use cases were derived from, and support identification of, requirements that any comprehensive framework for HLIF must address. The evaluation plan included the development of metrics to evaluate performance of the uncertainty management methods on the use case scenarios.

To ensure we leveraged the ongoing debate within the ISIF community, we synchronized deliverables with the Fusion conference series. Preliminary results of the requirements definition and an initial draft of the evaluation framework were presented at Fusion 2011 conference. Then, building on feedback from this preliminary effort, we further developed the framework and used it as the basis of a special session at Fusion 2012. Finally, we conducted a comprehensive refinement of the requirements, use cases, and evaluation criteria. The current version of the evaluation framework (as of December 2013) was presented at Fusion 2013.

Due to the circumstances related to the complexity of the problem and the limitations of a collaborative work by volunteers from different countries, the final version of the framework and its associated results were rescheduled to Fusion 2014. Yet, it is fair to state that the main goals of this research effort were either achieved (goals 1-3) or are on track for being achieved (goal 4) as we continue with the voluntary work on ETURWG. More specifically, throughout the research period we performed an in-depth analysis of the major requirements for representing and reasoning with uncertainty from the HLIF perspective (GOAL 1), we developed a set of use cases with enough complexity to cover the identified requirements (GOAL 2); and we have

defined a comprehensive set of criteria to evaluate how well a given methodology addresses the representational and reasoning needs of each use case (GOAL 3). Finally, we did not conduct a full evaluation of two major uncertainty management approaches being used in HLIF research and development (GOAL 4). This latter effort is contingent to completing the final version of the evaluation criteria, which is currently being refined at ETURWG and should be ready in time for presentation at Fusion 2014.

Bibliography

- [1] P. C. Costa, K. B. Laskey, E. Blasch, and A.-L. Jousselme, "Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology," in *Proceedings of the Fifteenth International Conference on Information Fusion*, Singapore, 2012.
- [2] P. C. G. Costa, E. P. Blasch, K. B. Laskey, S. F. Andler, J. Dezert, A.-L. Jousselme, and G. Powell, "Uncertainty evaluation: current status and major challenges," in *Proceedings of the Fifteenth International Conference on Information Fusion*, Singapore, 2012.
- [3] E. P. Blasch, P. C. G. Costa, K. B. Laskey, D. Stampouli, G. Ng, J. Schubert, R. Nagi, and P. Valin, "Issues of uncertainty analysis in high-level information fusion: panel discussion," in *Proceedings of the Fifteenth International Conference on Information Fusion*, Singapore, 2012.
- [4] 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, 2011.
- [5] T. Gruber, "Ontology (Computer Science) - definition in Encyclopedia of Database Systems," 16-Sep-2009. [Online]. Available: <http://tomgruber.org/writing/ontology-definition-2007.htm>. [Accessed: 16-Sep-2009].
- [6] M. Uschold, "Ontologies and Database Schema: What's the Difference?," presented at the Semantic Technology Conference, San Francisco, CA, 25-Jun-2010.
- [7] M. M. Kokar, C. J. Matheus, K. Baclawski, J. A. Letkowski, M. Hinman, and J. Salerno, "Use cases for ontologies in information fusion," in *Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, Stockholm, Sweden, 2004, pp. 415–421.
- [8] E. G. Little and G. L. Rogova, "Designing ontologies for higher level fusion," *Information Fusion*, vol. 10, no. 1, pp. 70–82, Jan. 2009.
- [9] P. C. G. Costa, "Bayesian semantics for the Semantic Web," PhD Dissertation, George Mason University, Fairfax, VA, USA, 2005.
- [10] R. N. Carvalho, "Probabilistic ontology: Representation and modeling methodology," Doctoral Dissertation, George Mason University, School of Information Technology and Engineering, Fairfax, VA, USA, 2011.
- [11] C. J. Matheus, M. M. Kokar, and K. Baclawski, "A core ontology for situation awareness," presented at the Information Fusion, 2003. Proceedings of the Sixth International Conference of, 2003, vol. 1, pp. 545–552.
- [12] D. Schum, *Evidential Foundations of Probabilistic Reasoning*. New York: Wiley, 1994.
- [13] D. A. Lambert, "A blueprint for higher level fusion systems," *Information Fusion*, vol. 10, no. 1, pp. 6–24, Jan. 2009.

- [14]J. F. Sowa, “Knowledge Representation: Logical, Philosophical and Computational Foundations,” Brooks-Cole Publishers, 2000.
- [15]M. L. Nichols, “A survey of multisensor data fusion systems,” in *Handbook of Multisensor Data Fusion: Theory and Practice*, 2nd ed., Boca Raton, Fla: CRC Press, 2008, pp. 691–699.
- [16]R. N. Carvalho, P. C. G. Costa, K. B. Laskey, and K.-C. Chang, “PROGNOS: Predictive Situational Awareness with Probabilistic Ontologies,” presented at the 13th International Conference on Information Fusion, Edinburgh, UK, 2010.
- [17]G. Tecuci, M. Boicu, D. Marcu, D. Schum, and B. Hamilton, “TIACRITIS system and textbook: Learning intelligence analysis through practice,” in *Proceedings of the Fifth International Conference on Semantic Technologies for Intelligence, Defense, and Security*, Fairfax, VA, USA, 2010, pp. 108–115.
- [18]H. M. Regan, M. Colyvan, and M. A. Burgman, “A taxonomy and treatment of uncertainty for ecology and conservation biology,” *Ecological Applications*, vol. 12, no. 2, pp. 618–628, Apr. 2002.
- [19]A.-L. Jousselme, P. Maupin, and E. Bossé, “Knowledge, belief and uncertainty,” in *Concepts, Models and Tools for Information Fusion*, Boston, MA: Artech House, Inc., 2007, pp. 141–150.
- [20]M. Smithson, *Ignorance and uncertainty: emerging paradigms*. Springer-Verlag, 1989.
- [21]P. Krause and D. Clark, *Representing Uncertain Knowledge: An Artificial Intelligence Approach*. Springer-Verlag New York Incorporated, 2012.
- [22]K. J. Laskey, K. B. Laskey, P. C. G. Costa, M. M. Kokar, T. Martin, and T. Lukasiewicz, “W3C Uncertainty Reasoning for the World Wide Web Incubator Group Report,” World Wide Web Consortium, Final Report XGR-urw3-20080331, Mar. 2008.
- [23]J. Salermo, “Information fusion: a high-level architecture overview,” in *Proceedings of the Fifth International Conference on Information Fusion*, Annapolis, MD, USA, 2002, vol. 1, pp. 680–686.
- [24]J. Esteban, A. Starr, R. Willetts, P. Hannah, and P. Bryanston-Cross, “A Review of data fusion models and architectures: towards engineering guidelines,” *Neural Computing & Applications*, vol. 14, no. 4, pp. 273–281, Dec. 2005.
- [25]M. Bedworth and J. O’Brien, “The Omnibus model: a new model of data fusion?,” *Aerospace and Electronic Systems Magazine, IEEE*, vol. 15, no. 4, pp. 30–36, 2000.
- [26]P. D. Whitney, C. Posse, and X. C. Lei, “Towards a unified approach to information integration a review paper on data/information fusion,” Pacific Northwest National Lab., Richland, WA, Technical Report PNNL-15428, 2005.
- [27]J. Roy, D. A. Lambert, F. E. White, J. Karakowsky, J. Salermo, and M. Hinman, “Data- and information-fusion models,” in *Concepts, Models and Tools for Information Fusion*, Boston, MA: Artech House, Inc., 2007, pp. 69–118.
- [28]M. M. Kokar, M. D. Bedworth, and C. B. Frankel, “Reference model for data fusion systems,” in *Sensor Fusion: Architectures, Algorithms, and Applications IV*, 2000, vol. 4051, pp. 191–202.
- [29]H. S. Carvalho, W. B. Heinzelman, A. L. Murphy, and C. J. N. Coelho, “A general data fusion architecture,” in *Information Fusion, 2003. Proceedings of the Sixth International Conference of*, 2003, vol. 2, pp. 1465–1472.
- [30]H. Durrant-Whyte and T. Henderson, “Multisensor Data Fusion,” in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer Berlin Heidelberg, 2008, pp. 585–610.

- [31] Alan N. Steinberg, "Foundations of Situation and Threat Assessment," in *Handbook of Multisensor Data Fusion*, 0 vols., CRC Press, 2008, pp. 437–501.
- [32] D. Schum, "Evidence marshaling for imaginative fact investigation," *Artificial Intelligence and Law*, vol. 9, no. 2–3, pp. 165–188, 2001.
- [33] M. Locher and P. C. G. Costa, "Ontological considerations for uncertainty propagation in high level information fusion," in *Proceedings of the Seventh International Conference on Semantic Technologies for Intelligence, Defense, and Security*, Fairfax, VA, USA, 2012, vol. 966, pp. 12–19.
- [34] S. K. Das, *High-Level Data Fusion*. Artech House, 2008.
- [35] K. B. Laskey, P. C. G. Costa, and T. Janssen, "Probabilistic ontologies for knowledge fusion," in *Proceedings of the Eleventh International Conference on Information Fusion (Fusion 2008)*, Cologne, Germany, 2008.
- [36] R. G. Almond, *Graphical Belief Modeling*. Taylor & Francis, 1995.
- [37] P. Shenoy, "Valuation-based systems for Bayesian decision analysis," *Operations Research*, vol. 40, no. 3, p. 463, May 1992.
- [38] M. A. Pravia, R. K. Prasanth, P. O. Arambel, C. Sidner, and C.-Y. Chong, "Generation of a fundamental data set for hard/soft information fusion," in *Information Fusion, 2008 11th International Conference on*, 2008, pp. 1–8.
- [39] A.-L. S. Lapinski, "A Strategy for Uncertainty Visualization Design," Defence Research and Development Canada, 2009.
- [40] P. Walley, "Measures of uncertainty in expert systems," *Artificial Intelligence*, vol. 83, no. 1, pp. 1 – 58, 1996.
- [41] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi, "Multisensor data fusion: A review of the state-of-the-art," *Information Fusion*, vol. 14, no. 1, pp. 28 – 44, 2013.
- [42] A. Karlsson, "Dependable and generic high-level information fusion : methods and algorithms for uncertainty management," Institutionen för kommunikation och information, HS- IKI -TR-07-003, 2007.
- [43] E. Franconi, "Description Logic Tutorial Course." Available online. Last accessed on 3/3/2013. URL: <http://www.inf.unibz.it/~franconi/dl/course/>.
- [44] K. B. Laskey, D. A. Schum, P. C. G. Costa, and T. Janssen, "Ontology of Evidence," presented at the OIC 2008, Fairfax, VA, USA, 2008, vol. 440, pp. 20–24.