December 2018 Volume 13 Number 2 ISSN 1557-6418

JAIF

Journal of Advances in Information Fusion

A semi-annual archival publication of the International Society of Information Fusion

Regular Papers Page

From the Editor-In-Chief

Guest Editorial: Evaluation of Uncertainty Representation and Reasoning Techniques

ISIF

INTERNATIONAL SOCIETY OF INFORMATION FUSION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied INFORMATION FUSION technologies. Technical areas of interest include target tracking, detection theory, applications for information fusion methods, image fusion, fusion systems architectures and management issues, classification, learning, data mining, Bayesian and reasoning methods.

Manuscripts are submitted at http://jaif.msubmit.net. If in doubt about the proper editorial area of a contribution, submit it under the unknown area.

Journal of Advances in Information Fusion (ISSN 1557-6418) is published semi-annually by the International Society of Information Fusion. The responsibility for the contents rests upon the authors and not upon ISIF, the Society, or its members. ISIF is a California Nonprofit Public Benefit Corporation at P.O. Box 4631, Mountain View, California 94040.

From the Editor-in-Chief:

December 2018

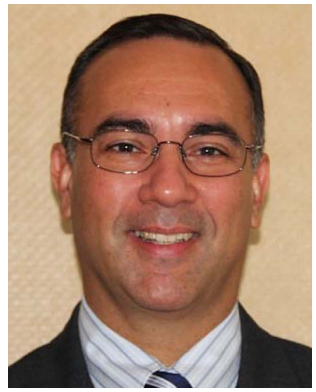

Paulo. C. G. Costa

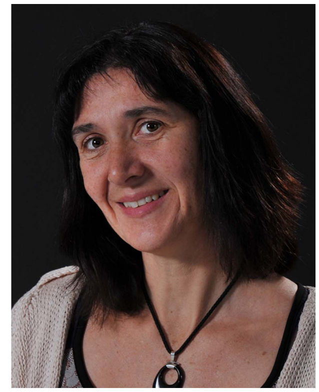

Anne-Laure Jousselme

Pieter de Villiers

Guest Editorial: Foreword to the Special Issue on *Evaluation of Uncertainty Representation and Reasoning Techniques*

In an era characterized by increasingly pervasive sensors, availability of large volumes of heterogeneous data and complex interactions between information systems, the problem of uncertainty representation and reasoning in high-level fusion information (HLIF) systems has attracted interest that extends beyond the Information Fusion (IF) community. For instance, fusing hard and soft information from diverse sensor or source types and the associated uncertainty is a task that still relies heavily on human intervention, creating a scalability conundrum that current technologies are incapable of solving. Despite the widespread acknowledgment that HLIF systems must support automated knowledge representation and reasoning in the presence of uncertainty, there is no consensus on the the appropriate approach to adopt (which theory, uncertainty function, fusion rule, etc), on the performance criteria that should guide the design of an HLIF system in terms of uncertainty handling, and on how to assess such criteria.

This special issue of JAIF aims at providing an overview of the most current efforts on evaluation of uncertainty representation and reasoning techniques in information fusion systems. In the opening paper of this issue, Costa et al. provide an overview of the Uncertainty Representation and Reasoning Evaluation Framework (URREF), which is currently in development by the ISIF Evaluation of Techniques for Uncertainty Representation Working Group (ETURWG). As an evaluation framework, the URREF is comprised of different components designed to provide support to researchers, developers, and other practitioners of high-level information fusion systems in the task of assessing and characterizing how choices on uncertainty representation and reasoning impact their performance. This is a multi-facet problem whose comprehensive and exhaustive coverage is challenging, but whose most basic and common facets are addressed here. This paper establishes the basic concepts and definitions, together with their links, as common grounds to be considered.

Jousselme and Pallota, in the second paper of this special issue, explore one of the most critical facets of the problem, which is how to identify performance criteria for uncertainty evaluation. They frame the comparison of six uncertainty representation and reasoning techniques in the URREF, with an illustrative example of HLIF on maritime anomaly detection. Next, Locher and Costa propose an overarching discussion on the difficulties in understanding where and how each criterion is applicable across a general fusion process environment, including a generic fusion system model. In the process, they provide some insight to the URREF ontology, a key component of the framework that offers a formal structure for representing the semantics of uncertainty evaluation.

In the fourth paper of this issue, De Villiers et al. discuss the role of uncertainty evaluation in the lifecycle of HLIF systems, while emphasizing how uncertainty impacts modeling and decision-making within these systems. In the discussion, the flow of abstraction in fusion system inception, design and implementation is contrasted to the flow of information and the flow of decisions/actions during the routine operation of a fusion system. This contrast is a good lead to the subject of the following paper, by Dragos et al., which explores an issue that pervades all the information flow: how to estimate trust in information received by and output from HLIF systems. The paper emphasizes how the URREF ontology can be used to characterize and track uncertainties arising within the development of HLIF systems, focusing on how trust can be estimated in the process.

The next papers emphasize key aspects of uncertainty representation and reasoning in HLIF systems, setting the stage for a discussion on the application of the URREF to different domains and techniques. An example of the latter is shown by Josang in the sixth paper of this special issue, which addresses the importance of selecting a belief fusion operator that adequately matches the situation to be modeled and analyzed. Moving the discussion from technique to applications, the last two papers of this special issue illustrate the role of uncertainty evaluation in two different application domains, Avionics (Insaurralde and Blasch) and Situational Assessment (Hintz and Darcy). The first proposes an Avionics Analytics Ontology (AAO) to bring together different types of uncertainties including semantic from operators, sensing from navigation, and situation from weather modeling updates. The approach is aligned with the URREF via its use of some of the URREF ontology concepts. Finally, Hintz and Darcy close the special issue by presenting the problem of measuring uncertainty over time to control the knowledge entropy in a situation awareness system.

As it can be inferred from the papers presented in this special issue and their associated reference lists, the problem of evaluating uncertainty representation and reasoning techniques in HLIF is still far from being solved. Yet, the IF community is clearly moving ahead towards that goal and its research on the key topics is starting to bear fruits.

Paulo C. G. Costa
George Mason University
Fairfax, VA, USA
E-mail: pcosta@gmu.edu

Anne-Laure Jousselme
NATO STO Centre for Maritime Research and Experimentation
La Spezia, IT
E-mail: anne-laure.jousselme@cmre.nato.int

Pieter de Villiers
University of Pretoria
Pretoria, South Africa
E-mail: pieter.devilliers@up.ac.za

Guest Associate Editors

URREF: Uncertainty representation and reasoning evaluation framework for information fusion

PAULO C. G. COSTA
ANNE-LAURE JOUSSELME
KATHRYN B. LASKEY
ERIK BLASCH
VALENTINA DRAGOS
JUERGEN ZIEGLER
PIETER DE VILLIERS
GREGOR PAVLIN

Uncertainty management is a key aspect of any information fusion (IF) system. Evaluation of how uncertainty is dealt with within a given IF system is distinct from, although closely related to, evaluation of the overall performance of the system. This paper presents the Uncertainty Representation and Reasoning Evaluation Framework (URREF), which is developed by the ISIF Evaluation of Techniques for Uncertainty Representation Working Group (ETURWG) for evaluating the uncertainty management aspects of IF systems. The paper describes the scope of the framework, its core element—the URREF ontology, the elementary fusion process it considers, and how these are related to the subjects being evaluated using the framework. Although material about the URREF has been previously published elsewhere, this work is the first to provide a comprehensive overview of the framework, establishing its scope, core elements, elementary fusion process considered, and relationship between these and the subjects they are designed to evaluate. We also briefly describe a few use cases of the framework, discussing how URREF can be applied in their evaluation.

Manuscript received October 22, 2018; revised May 5, 2019; released for publication May 13, 2019.

Refereeing of this contribution was handled by Sten Andler.

Authors' addresses: P. C. G. Costa, K. B. Laskey, George Mason University, Fairfax, VA, USA (E-mail: {pcosta,klaskey}@gmu.edu). A.-L. Jousselme, NATO STO Centre for Maritime Research and Experimentation, La Spezia, IT (E-mail: anne-laure.jousselme@cmre.nato.int). E. Blasch, Air Force Office of Scientific Research, Arlington, VA, USA (E-mail: erik.blasch.1@us.af.mil). V. Dragos, ONERA—The French Aerospace Lab, Palaiseau, France (E-mail: valentina.dragos@onera.fr). J. Ziegler, Competence Centers ISR, IABGmbH, Ottobrunn, Germany (E-mail: jziegler@iabg.de). P. de Villiers, University of Pretoria, Pretoria, South Africa, Council for Scientific and Industrial Research, Pretoria, South Africa (E-mail: pieter.devilliers@up.ac.za). G. Pavlin, D-CIS Lab, Thales Research and Technology, Delft, The Netherlands (E-mail: gregor.pavlin@d-cis.nl).

1557-6418/18/$17.00 © 2018 JAIF

I. INTRODUCTION

Evaluating how well an Information Fusion (IF) system performs requires defining the relevant criteria to be assessed and testing the IF system's fusion algorithm, data model, and architecture against that criteria. Empirical evaluation techniques are effective when assessing the latter two, but face a major limitation when addressing the former. More specifically, they often require embedding some uncertainty representation and its associated reasoning scheme within the fusion method, which serves as an enabler and becomes often the subject of evaluation itself. Inherently, it is not a trivial problem to isolate the uncertainty representation from either its reasoning scheme or the fusion algorithm, which prevents an effective assessment of the IF system since current methods cannot capture the impact of these in the overall IF system's performance. The work described in this paper focuses on addressing this limitation, providing a principled method for evaluating how the uncertainty representation and reasoning aspects of an Information Fusion impact its overall performance.

IF applications typically must deal with information that is incomplete, imprecise, inconsistent and otherwise in need of a sound methodology for representing and managing uncertainty. Complex and dynamic use cases make such tasks even more difficult, as apparently minor differences in how uncertainty is handled may drastically affect the output of the IF process. In short, it is fair to state that uncertainty management is a key aspect in most—if not all—IF systems. Despite this importance, the IF community still does not have a standardized framework for evaluating how uncertainty is represented and managed in IF systems. IF systems typically perform uncertainty reasoning to achieve their goals, which means they would benefit from a framework to evaluate how well they are performing on it.

The lack of an uncertainty evaluation framework for IF systems tends to be more widely acknowledged at higher levels of the Joint Directors of Laboratories (JDL) model [1]–[3]. More specifically, Low-Level Information Fusion (LLIF) systems (i.e., below JDL level 2) tend not to represent semantics explicitly. Semantics is commonly understood among theoreticians and algorithm developers, and is typically implicitly encoded in algorithms through devices such as variable naming conventions. LLIF systems tend to rely exclusively on probability theory as the paradigm for uncertainty representation and reasoning. This is justified by the typically large amount of available data, which justifies the use of statistical models to address the fusion problems at hand. Tools and techniques for evaluating probabilistic inference systems are well-understood. In contrast, because of the complexity and variety of semantic categories for High-Level Information Fusion (HLIF), applications usually require making semantics explicit and accessible to formal reasoning tools. Furthermore, HLIF systems make use of a variety of theories and methods to represent and reason with uncertainty. For example, deciding whether three different radars receiving echos from the same location are seeing one, two, or three tracks is a problem for which uncertainty is well understood and for which standard evaluation methods are well established. On the other hand, deciding whether the incoming fighter formation poses a danger to a radar installation may involve a multiplicity of sources of uncertainty, and may require consideration of complex semantic concepts such as enemy doctrine, the spatial configurations associated with hostile and innocuous formations, how danger should be defined, and the like.

Uncertainty analysis is even more critical for systems relying on multiple types of data and different uncertainty paradigms. Soft data is unstructured and intrinsically ambiguous [4], and tracking its uncertainty [5] often requires explicit semantics [6]. Heterogeneous fusion combines data of different natures, and uncertainty propagation for heterogeneous fusion still lacks a well-established and widely agreed upon theoretical foundation [7].

Clearly, a system that can reason about these and other HLIF problems must consider complex semantics, and may be required to employ multiple uncertainty formalisms (e.g., a fuzzy membership function might be used to transform verbal danger categories into a quantitative representation, which might be combined with a probability distribution on events leading to different levels of danger). The design of a HLIF system would definitely benefit from an uncertainty evaluation framework that would guide the selection of the most suitable uncertainty representation and reasoning technique. An ability to compare uncertainty handling approaches would enable exploitation of semantically rich representations to help assess its performance when facing an uncertain input. With the emergence of alternative uncertainty theories in addition to probabilities (see for instance [8] for a survey) came the question of which approach is the best suited for uncertainty handling in a specific problem setting. The question has been addressed both theoretically (e.g., [9]–[12]) and in practical implementation of fusion solutions (e.g., [13]–[15]). Handling uncertainty in fusion problems is indeed a major challenge for algorithm designers as it generates many questions, such as what "uncertainty" means, where it comes from, on what it bears, how to interpret the associated numerical values or measures, how to distinguish between its different varieties, etc. Acknowledging the existence of different types or facets of information quality provides partial answers (e.g., [16]–[18]). Nevertheless a deep understanding of the different uncertainty representation and reasoning techniques, their underlying mathematical frameworks, and associated hypotheses and semantics, is necessary to guide a fusion system's designer in making informed choices about the most suitable technique to the problem at hand. Such a deep understanding provides clearer explanations of the algorithms to the user for an improved synergy between the human and the machine [19].

The International Society of Information Fusion (ISIF) recognized this problem, and created a working group to address it. The ISIF Evaluation of Techniques for Uncertainty Representation Working Group (ETURWG) [20], [21] was created in the ISIF Board of Directors meeting just after the Fusion 2011 conference (Chicago, IL, USA) to specifically address this issue. The ETURWG's main goals are (1) to establish features required for any quantitative uncertainty representation to support the exchange of soft and hard information in a net-centric environment; (2) to develop a set of use cases involving information exchange and fusion requiring reasoning and inference under uncertainty; and (3) to define evaluation criteria supporting principled comparisons among different approaches applied to the use cases. As of this writing, the group has convened 104 general meetings spanning its 7 years of activities, and resulted in 43 peer-reviewed articles on the subject. The group's website[1] provides comprehensive information about its activities, including agendas and minutes of the meetings, datasets used, documentation on case studies and discussions, as well as a large amount of information related to the research efforts by the group.

This paper provides an overview of the Uncertainty Representation and Reasoning Evaluation Framework (URREF). It not only updates but also substantially enhances a similar paper published in the Proceedings of the Fusion 2012 conference [22]. After this brief introduction, Section II provides an overview of recent and current efforts in evaluating uncertainty in IF systems. Section III introduces the framework, which supports assessment of the impact of uncertainty representation on a fusion system. This is followed by a section covering the relationship between the framework elements and the subjects it is evaluating. Section V presents a brief description of case studies applying the framework. The final section contains discussion and conclusion.

II. EVALUATING UNCERTAINTY IN FUSION SYSTEMS

The evaluation of how uncertainty is dealt with within a given IF system is distinct from, although closely related to, the evaluation of the overall performance of the system [23], [24]. Figure 1 shows the elements of a generic IF model. The figure distinguishes between processes associated with low-level and high-level IF, a distinction dating to the seminal fusion model developed by the Joint Directors of Laboratories (JDL) [1]–[3]. Evaluation criteria and associated 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. These are more encompassing in scope than

[1] http://eturwg.c4i.gmu.edu, free registration required for full access

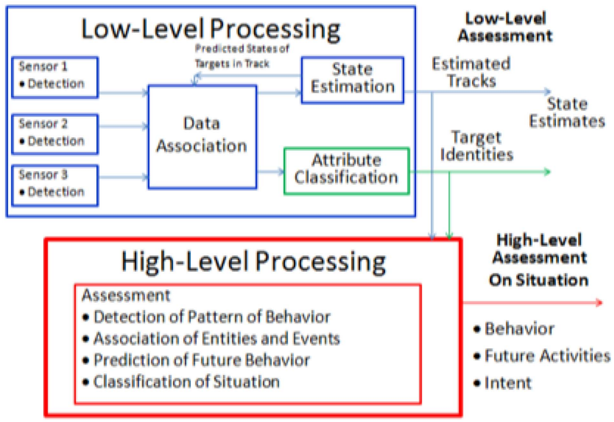

Fig. 1. The principal processing components of the IF process include both high level and low level processing components. Low level fusion processes include detection, association, state estimation and attribute classification, whereas high level fusion processes include behavioural pattern estimation, association, behaviour prediction and situation classification.

those focused on the uncertainty handling within the system. Metrics focused on uncertainty handling should address the contribution of uncertainty handling to the overall system performance.

For example, fusion-system-level metrics include timeliness, accuracy and confidence. Clearly, different choices in uncertainty representation approaches will affect the achievable timeliness, accuracy, and confidence of a system, and therefore must be considered when evaluating both the system's performance as a whole and the specific impact of the uncertainty handling approach. Yet, when evaluating timeliness (or any other system-level metrics), one will likely find some factors not directly related to the handling of uncertainty itself, such as object tracking and classification report updates (i.e., Level 1 fusion), situation and threat assessment relative to scenario constraints (i.e., Level 2/3 fusion), overall system architecture (e.g., centralized, distributed, etc.), data management processes and feedback/input control processes (i.e., Level 4 fusion considerations), and user-machine coordination based on operating systems (i.e., Level 5 fusion), and others.

The IF community envisions effortless interaction between humans and computers, seamless interoperability and information exchange among applications, and rapid and accurate identification and invocation of appropriate services. As the complexity of fusion solutions grows, we end up with a mixture of components handling different types of uncertainties, often by using different methods.

Here, the term "uncertainty" is intended to encompass a variety of aspects of imperfect knowledge, including incompleteness, inconclusiveness, vagueness, ambiguity, and others. The term "uncertainty reasoning" is meant to denote the full range of methods designed for representing and reasoning with knowledge when approaches based on Boolean algebra (e.g. propositional logic) are not applicable (e.g. when Boolean truth-values are unknown, unknowable, or inapplicable. Commonly applied approaches to uncertainty reasoning include probability theory, fuzzy logic, subjective logic, Dempster-Shafer theory, DSmT, and numerous other methodologies.

The problem of representing and reasoning with complex and heterogeneous data was addressed by a working group of the World Wide Web Consortium [25]. The working group's findings are relevant to the challenge considered in this paper. Information fusion under uncertainty is an intrinsic requirement for many of the problems in the World Wide Web domain. A full realization of the World Wide Web as a source of processable data and services demands formalisms capable of representing and reasoning under uncertainty.

- Automated agents are used to exchange Web information that in many cases is not perfect. Thus, a standardized format for representing uncertainty would allow agents receiving imperfect information to interpret it in the same way as were intended by the sending agents.
- Data often are intrinsically uncertainty-laden. Examples include weather forecasts or gambling odds. Canonical methods for representing and integrating such information are necessary for communicating it in a seamless fashion.
- Non-sensory collected information is also often incorrect or only partially correct, raising concerns related to trust or credibility. Uncertainty representation and reasoning helps to resolve tension amongst information sources having different confidence and trust levels.
- Dynamic composability of Web Services will require runtime identification of processing and data resources and resolution of policy objectives. Uncertainty reasoning techniques may be necessary to resolve situations in which existing information is not definitive.
- Information extracted from large information networks such as the World Wide Web is typically incomplete. The ability to exploit partial information is very useful for identifying sources of service or information. For example, that an online service deals with greeting cards may be evidence that it also sells stationery. It is clear that search effectiveness could be improved by appropriate use of technologies for handling uncertainty.

These problems all require IF, both low and high level. They bear an obvious relationship to the kinds of problems found in the sensor, data, and IF domain.

III. UNCERTAINTY REPRESENTATION AND REASONING FRAMEWORK

This section describes an evaluation framework to support assessment of how the choice of uncertainty

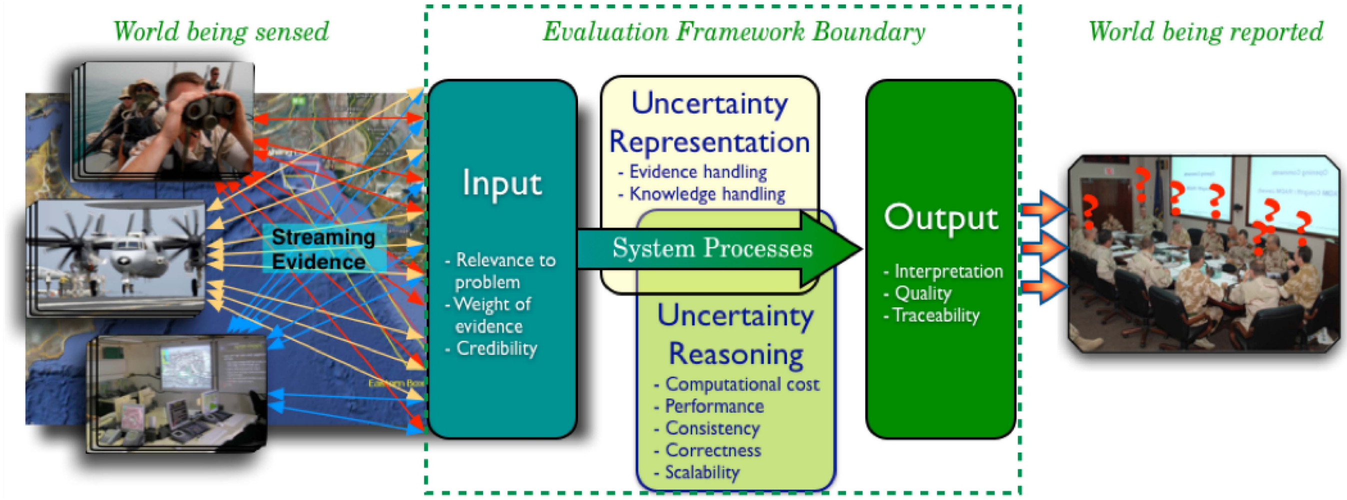

Fig. 2. URREF Boundary. This figure depicts the world being sensed on the left, the role of uncertainty representation and reasoning within a fusion system in the center, and the world being perceived on the right. The evaluation framework boundary encompasses the fusion system input, uncertainty representation, uncertainty reasoning and the fusion system output. Everything inside the evaluation framework boundary is known as the Uncertainty Representation and Reasoning Framework (URREF). The uncertainty representation and uncertainty reasoning are the primary subjects of evaluation, whereas the input and output are secondary subjects of evaluation.

representation and reasoning impacts the performance of an IF system. The scope of the framework is the main focus of the first sub-section, which is followed by an overview of its main component, the URREF ontology. Finally, an elementary fusion process is presented, as a means to identify the primary evaluation subjects of the evaluation methodology envisioned for the framework.

A. The URREF Scope

The basic idea behind the framework is to analyze an abstract fusion system and define its input data and output products. In a hypothetical IF system of the future, the uncertainty representation approach would be "plug-and-playable." That is, one might run the system with a Bayesian approach, then switch to a Dempster-Shafer approach, and then a Fuzzy Random Set approach. Alternatively, one might use a combination of uncertainty reasoning methods, as best suited for different aspects of the problem. The input data are the same in each case, as are the output products (but not necessarily the specific contents of the output products). Figure 2 depicts the uncertainty representation and reasoning evaluation framework (URREF) and its role in the overall fusion process.

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. The uncertainty representation and uncertainty reasoning are the primary subjects of evaluation, whereas the input and output are secondary subjects of evaluation.

The first external element refers to the events of interest to the IF system that happen in the world and are perceived by the system sources. Note that the implicit definition of sources in this case encompasses anything that can capture information and send it to the system. That is, both hard sources (e.g., imaging, radar, video, etc.) and soft sources (HUMINT reports, software alerts, etc.) are considered external to the evaluation system with respect to their associated sensorial capabilities, while the way they convey their information is within the scope of the system [24], [26], [27].

This reflects an important distinction between the evaluation of an IF system and the evaluation of its handling of uncertainty. To illustrate the distinction, consider the Input element in Figure 2. This element addresses the system's ability to represent uncertainty as an intrinsic part of the information being captured. As an example, information regarding trust of the input from a given sensor is important to evaluating how the overall system handles uncertainty, although it might not be as critical for its overall performance. A key question for evaluating uncertainty representation is what the uncertainty characteristics of the input data are, and how they affect the use of different uncertainty schemes. On the other hand, the format of the input might be important to evaluation of system interoperability, but is not included in Figure 2 because it does not relate to uncertainty handling. In general, the elements inside the evaluation framework boundary in the figure are important to evaluation of uncertainty handling, but not necessarily to evaluation of other aspects of fusion system performance. Likewise, elements that are critical to overall evaluation but not important to uncertainty handling are not included here.

In the ideal system model, having the appropriate data characteristics is critical. If the characteristics do not span the range of uncertainty techniques, then the model may not give meaningful results about the operationally significant differences between the techniques. Correctly identifying the desired input data characteristics will shape the future development of use cases and modeling data sets for those use case.

Once information is in the IF system, it will be processed to generate the system's deliverable that requires uncertainty characterization and reporting in the Output step. This process involves fusion techniques and algorithms that are directly affected by the uncertainty handling technique being used, as well as its impact on the system's inferential process. In this case, the URREF evaluation criteria focus on aspects that are specific to the way uncertainty is considered and handled within the fusion process. This is not an evaluation of the system's performance as a whole. We want to understand how the uncertainty representation affects system performance, and whether different uncertainty representation schemes are more or less robust to variations in the remaining parts of the IF system architecture. We want to focus specifically on the uncertainty representation aspects, and attempt, as best as possible, to separate those aspects from the overall system performance and architecture concerns.

After the information is fused and properly treated, then it is conveyed to the system's users. In Figure 2, these are represented by an image depicting decision-makers who would likely be supported by the IF system in their tasks. The URREF output step involves the assessment of how information on uncertainty is presented to the users and, therefore, how it impacts the quality of their decision-making process.

B. The URREF Ontology

The word "framework" in URREF's name reflects the conclusion we reached during the early ETURWG meetings, as we discussed how uncertainty in IF systems should be evaluated. From the very beginning, it became clear to us that we were not developing a tool to measure a set of metrics related to uncertainty in a given system. After all, because uncertainty is embedded in practically all aspects of the process, each application would have so many nuances that designing a "one-size-fits-all" evaluation tool would either be too specific for use in diverse IF systems, or too generic to be useful. In other words, we soon realized that what was needed to move the state-of-the-art in uncertainty evaluation was not a monolithic evaluation program or tool, but a set of standards, best practices, guidelines, and other development tools that provides coherent and consistent support for those tasked with evaluating uncertainty in information systems. We call this set an evaluation framework.

The reasons behind this view of URREF as a framework instead of a system, program, or tool, also implied that the diversity and complexity of the IF systems to be evaluated would require this framework to be flexible and adaptable enough to be used by developers with distinct backgrounds and requirements. We soon realized that defining common terminology was an enormous challenge, as a given term might have different meanings to different people, whereas a common idea might be given different names by different people. Designing a "mother of all evaluation taxonomies" was not an option, as it would be useless to various use cases, such as existing systems with already established semantics. Thus, when designing the framework we were naturally inclined to adopt ontology as a knowledge representation technique, as an ontology provides embedded support for reasoning and allows for explicit semantics that could be aligned, adapted, or reused when developing evaluation systems.

Designing an ontology for URREF proved to be a tall order though. Within the ETURWG we have people with distinct backgrounds, so it was natural to see some "semantic misalignment" regarding concepts such as data quality, accuracy, precision, etc. These differences in understanding proved to be challenging to deal with, but an accurate preview of the challenges that arise when using a framework that invariably includes concept definitions that may not fully match the views of different users. Not surprisingly, it took a considerable amount of time to arrive at a stable version of the URREF ontology, and while all in the group would prefer one or more specific concepts to be defined in a different way, the group agreed that the current version of the ontology is sufficient to support the evaluation of uncertainty in IF systems consistently and coherently. Most of the concepts used have been drawn from seminal work in related areas such as uncertainty representation (e.g., [27]–[35], evidential reasoning (e.g., [36]–[38]), and performance evaluation (e.g., [9], [39]–[41]). We now describe the main aspects of the URREF ontology, including its classes, properties, and key concepts. The reader would benefit from actually accessing the files, and even following the work of the ETURWG group. In addition to the information provided in the group's website, as indicated earlier in this paper, the ontology itself can be downloaded or opened directly from an ontology editor (e.g., Protégé [42]) via its official URL.[2] Alternatively, cloning the group's GitHub repository[3] would provide access to not only the current version of the ontology but also previous versions, references, and other related working documents.

Figure 3 depicts the main classes of the URREF ontology, which were identified by the ETURWG group as pertinent to the evaluation of uncertainty within an IF system. These classes represent concepts meant to be sufficient to support the design of evaluation processes that follow the same semantic constraints and

[2] http://eturwg.c4i.gmu.edu/files/ontologies/URREF.owl

[3] https://github.com/paulocosta-gmu/urref/tree/master

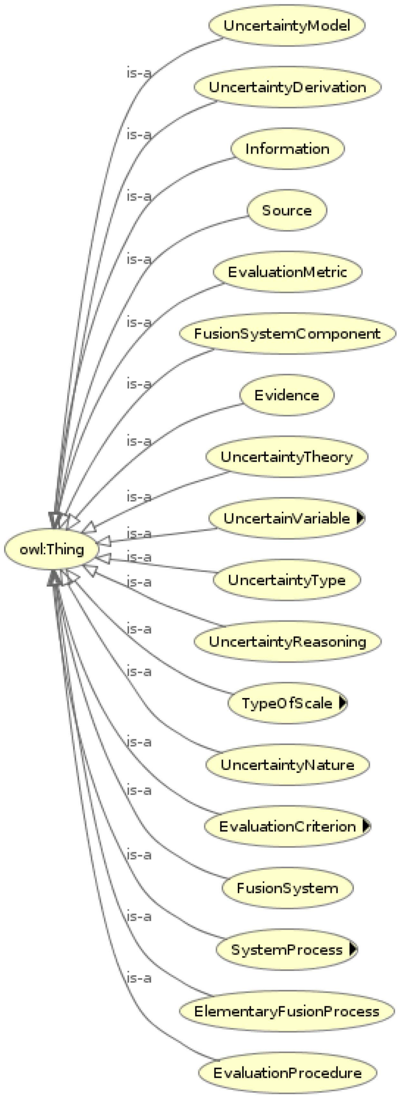

Fig. 3. Main classes of the URREF ontology

abide by the same principles of mathematical soundness. To emphasize the pragmatic aspect of the work of the ETURWG, it can be noted that these concepts capture the main aspects the group agreed upon when developing the use cases described in Section V. In fact, a brief comparison between these concepts and those of the first version of the ontology (cf. [22]) will show that many classes had to be added as a result of both the evolving discussions and the requirements elicited from the use cases.

The eighteen main classes of the URREF criteria focus on aspects that are specific to the way uncertainty is considered and handled within the fusion process. Figure 3 was built using the Protégé OWLviz plugin.[4] The classes are depicted as collapsed at the first level. Classes with a small black arrow head at the right have subclasses which can be shown in an expanded view. One example is the class *TypeOfScale*, which is depicted in its entirety in Figure 4. Its individuals correspond to

[4]https://github.com/protegeproject/owlviz

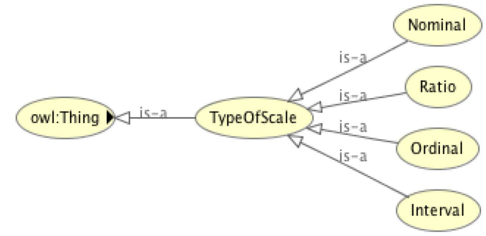

Fig. 4. URREF *TypeOfScale* class

specific scales used in quantifying the metrics employed when evaluating an IF system according to a given criteria, and its subclasses aggregate the types of quantification adopted. For instance, assume the precision of a given sensor (i.e., using the subclass *Precision* as evaluation criterion) would be evaluated using

$$u_{pre} = \sum_{t=1}^{n} L(r_t, a_t), \tag{1}$$

where n is the number of measurement trials, and L is a loss function with parameters r for reported value and a for actual value. In this case, the range of the loss function will dictate which type of scale should be used in that evaluation (e.g., a loss function returning a ratio between the two parameters would be classified under the associated type of scale). In the URREF framework, this class provides a way of mapping evaluation subjects and criteria chosen to the potential metrics and associated quantification types that can be used in a given evaluation.

While the type of scale defines how to quantify the metrics used to assess a given criterion in an evaluation, the *EvaluationMetrics* class defines what metric is being used (i.e., what is) the parameter being assessed. In the example of Eq. (1), the criterion being assessed is performance and the formula itself can be seen as the metric used to assess that criterion. Currently, the ontology only includes examples from NATO's Standardization Agreement 2511 (STANAG 2511) effort, which incorporates categories of reliability and credibility. Reliability has traditionally been assessed for physical machines to support failure analysis. Source reliability of a human can also be assessed. Credibility is associated with a machine process or human assessment of collected evidence for information content [43]. As the group work progresses, further standards are likely to be included as well.

Another example is the *EvaluationCriterion* class, depicted in Figure 5 and is at the core of any evaluation procedure. Not surprisingly, it is the larger class of the URREF ontology and the one with more levels. When looking at its main sub-classes, the more detail-oriented readers would be able to establish a parallel between these subclasses and the items within the Evaluation Framework Boundary framework depicted in Figure 2. More specifically, the Uncertainty Representation and Uncertainty Reasoning boxes can be mapped directly

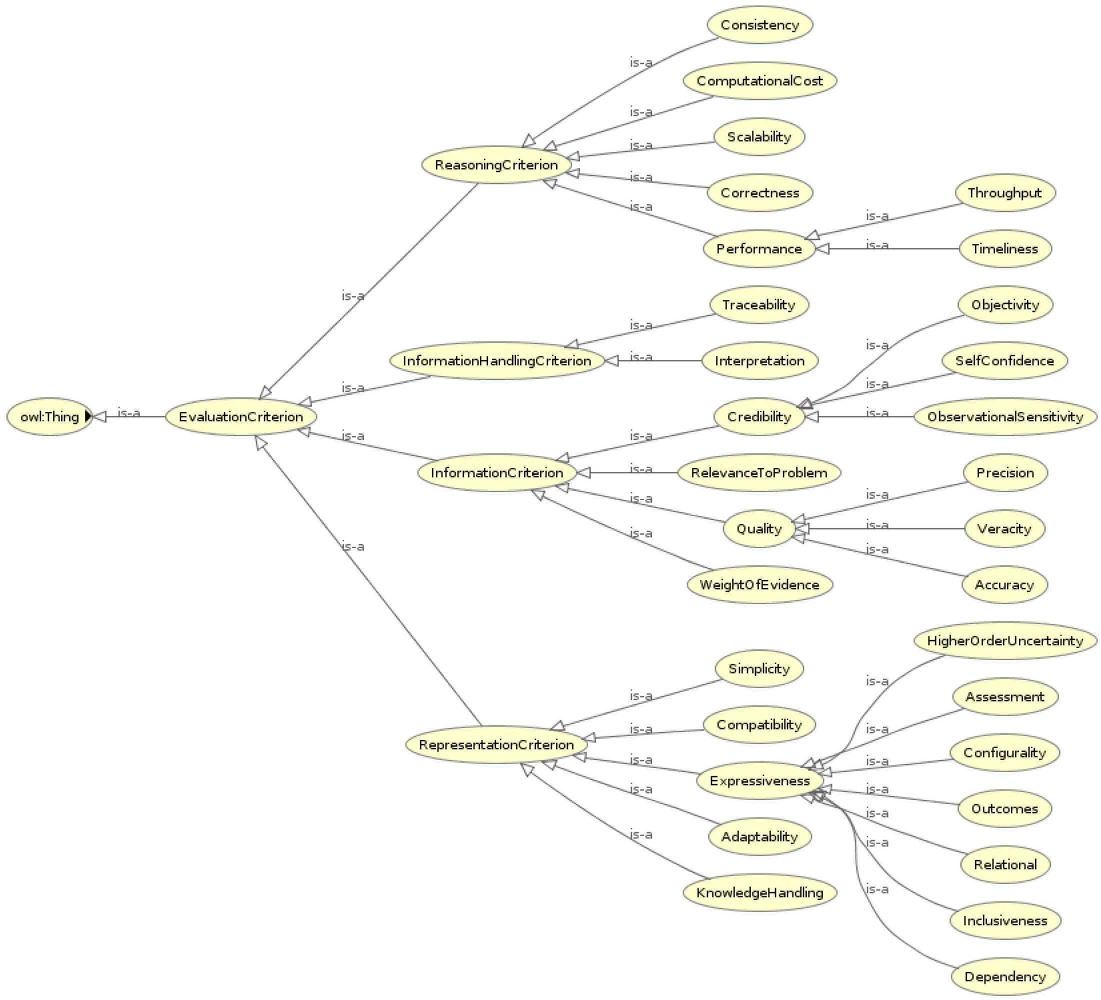

Fig. 5. *EvaluationCriterion* class

to the equally named sub-classes, while the classes *InformationHandlingCriterion* and *InformationCriterion* can be associated with the flow of information between Input and Output boxes.

The above classes form the structure of the URREF ontology, and were meant to collectively support the evaluation of uncertainty of an IF system. This is the third version of the URREF ontology, and at the time of this writing the group is now focusing on the case studies, which provide the necessary testbed for its ideas—and might force changes in the above classes. This approach privileges the pragmatism of having a good solution against having an "ideal" but unattainable solution. For instance, a definitive reference would involve having universally accepted definitions and usage for terms such as "Precision." This is unfeasible in any field of research that is not tightly controlled by a unique authoritative entity. The ETURWG approach also takes into consideration that more important than naming a concept is to ensure that it is represented clearly and distinctly within the ontology so to ensure the consistency of the latter.

Ontology reasoning requires axioms and properties to be defined, formally exposing the relationships between the above concepts that ultimately drive the logical reasoning that makes ontologies a very flexible and powerful technique. As an example, the object property *HasDerivationOfUncertainty* is used to map individuals of class *Evidence* (i.e., the domain of the property *hasDerivationOfUncertainty*) to individuals of class *UncertaintyDerivation* (i.e., the range of the property). The reasoner would use this relationship between these classes to support queries, automated classification, and other features the URREF could provide to its users.

A comprehensive description of the URREF ontology, with its classes, properties, and other elements is not within the scope of this paper. For a comprehensive overview of the URREF ontology, interested readers should refer to the ETURWG Github repository and the ETURWG website already mentioned in this paper.

C. The URREF Elementary Fusion Process

The elements of the Uncertainty Representation and Reasoning (URR) techniques to be assessed and compared will be referred within the URREF framework as

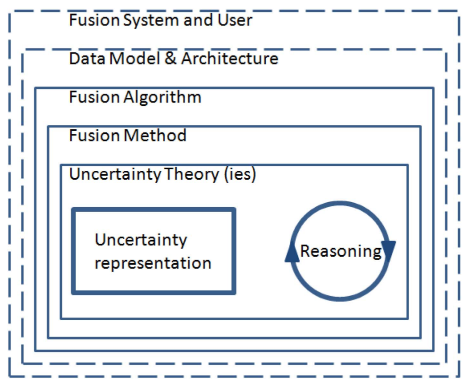

Fig. 6. An approximate hierarchy of fusion system components as possible evaluation subjects.

evaluation subjects. Owing to the complex and multiple connections between elements it seems difficult (if at all possible) to separate the uncertainty representation (e.g., an instantiated probability distribution) from its associated reasoning scheme (e.g., Bayes' rule), from its underlying uncertainty theory or mathematical framework (e.g., probability theory), from an underlying semantic representation (e.g., possible worlds, Ontology Web Language (OWL)), from the fusion method, from the fusion algorithm processing information (e.g., a specific implementation possibly involving some approximation), from a higher-level fusion system possibly including some human interaction.

Figure 6 illustrates some system components to assess and which interact to build a complete fusion system. As far as the URREF is concerned, the elements of an Uncertainty Representation and Reasoning scheme are the main evaluation subjects (thick lines in Figure 6), while the uncertainty theory, fusion method and fusion algorithm are of secondary focus. It is not the main purpose of the URREF to address the assessment of the fusion system nor the data model nor the architecture (dotted lines in the figure). Empirical evaluation techniques often require embedding some uncertainty representation and its associated reasoning scheme within the fusion method, which serves as an enabler and becomes often the subject of evaluation itself. Inherently, it is not a trivial problem to isolate the uncertainty representation from either its reasoning scheme or the fusion algorithm which may implement other contributing aspects, albeit minor.

For each evaluation subject, a series of evaluation criteria of interest is then defined in the URREF ontology [22] (see Section IV). It happens that the same criterion applies to different subjects with thus possible different associated metrics (or measures). For instance, *Accuracy* can be a quality criterion of information and of a source of information.

The fusion method is further detailed here by defining a generic procedure that highlights the main elementary constructs of uncertainty representation and reasoning that are the primary URREF evaluation subjects to be further defined in Section IV. The fusion method may be very complex, involving possibly several uncertainty representations, combination or inference rules, possibly framed in different uncertainty theories. Here, we abstract away complexities that are inessential to our purpose to obtain a simple, albeit quite general, fusion method aimed at clarifying the information flow. The result can be considered as an "atomic" fusion process.

The elementary constructs of a fusion process are shown in Figure 7, and illustrated with corresponding human intelligence fusion and multiple radar fusion examples in Table I:

① S is a source of information;

② ϕ is a piece of information provided by (or extracted from) S. It can be as simple as a measurement but could also be a natural language statement, a probability distribution, or in general a piece of information with some uncertainty already represented in a specific uncertainty theory;

③ h is the uncertainty representation process by which ϕ is transformed into a dedicated mathematical function conveying some notion of uncertainty. The process h is typically the choice of the solution designer who selects the way incoming information may be converted into a mathematical object. It can be learned from data when available or it can be general to all POIs, specified by type of source, by type of information, etc. Prior information on source's quality (e.g., reliability), source's self-confidence in statement, contextual information, comparison with other POIs, etc, may be captured by h;

④ $h(\phi)$ is the instantiated mathematical representation as built by h and expresses either the self-assessment of the source, an external assessment by the designer based on prior source's quality knowledge or an aggregation of both;

⑤ ρ is the inference process which transforms $h(\phi)$ into another $h_{\oplus}(\phi)$ within the same uncertainty theory. At this point, a series of POIs from other sources $\{h(\phi)\}_{i=1,\ldots,N}$ are combined, where other POIs are deduced, predicted, revised, etc;

⑥ $h_{\oplus}(\phi)$ is the resulting piece of information built from $h(\phi)$ and other related information;

⑦ l is the decision process which transforms $h_{\oplus}(\phi)$ to provide the decision, i.e., the output information ϕ';

⑧ ϕ' is the information output, to be possibly sent other systems. It can be a formal representation, i.e., an uncertainty function (such as a probability distribution), or a single measurement estimated after the decision process (soft versus hard decision). It can thus contain or not contain some uncertainty;

⑨ the reasoning process is $l \circ \rho$;

⑩ the Atomic Decision Procedure (ADP) is $l \circ \rho \circ h$.

TABLE I
Elementary fusion process constructs illustrated at the hand of a) a human intelligence fusion example and b) a multiple radar centralized fusion example

Element	Example
① S	a) Human observer b) Radar sensor
② ϕ	a) Human report, b) Radar range velocity measurement
③ h	a) Convert a natural language statement to a belief function over locations, b) convert a range and angle measurement and associated Root Mean Square Error (RMSE) error value to a Gaussian distribution with mean and variance
④ $h(\phi)$	a) Belief function b) Gaussian probability distribution
⑤ ρ	a) Dempster's combination rule (combine multiple reports) b) Bayes' rule (combine multiple measurements form different radars)
⑥ l	a) Maximum of plausibility rule b) Find expected value of posterior distribution
⑦ ϕ'	a) Element with maximum plausibility (or complete plausibility distribution over singletons) b) Expected value of the posterior distribution

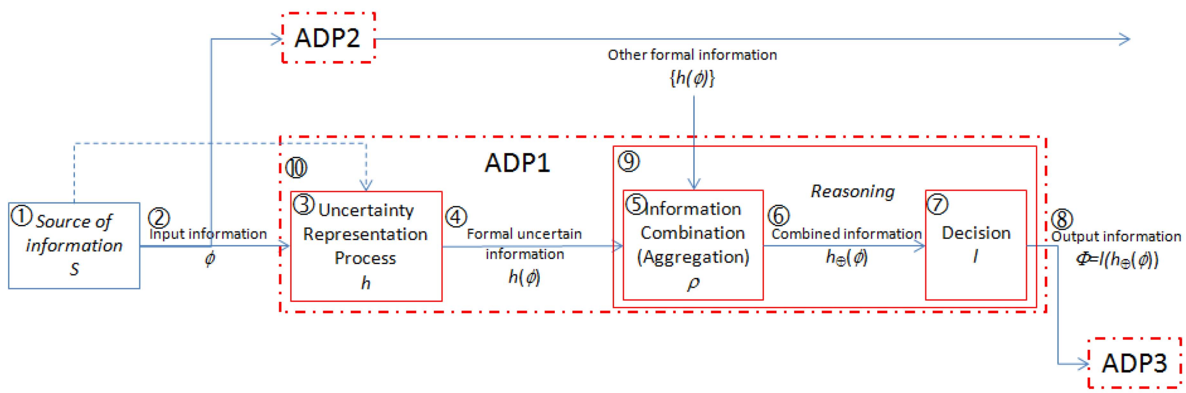

Fig. 7. Basic information flow and evaluation subjects.

Figure 7 illustrates this process and depicts each of the above 10 items in its appropriate place in the process.

As further detailed in [44], the method can distinguish between:

a) information processors (providing POIs): Elements ①, ③, ⑤, ⑦, ⑨;
b) the provided: Elements ②, ④, ⑥, ⑧;
c) the pairs (process; output information): (①,②); (③,④); (⑤,⑥); (⑦,⑧); (⑨,⑧); (⑩,⑧)=(①,②)

From an algorithmic standpoint, we may want to assess each of the 10 items above. However, based on the following observations some simplifications arise:

- Each information processor can be assessed through the information it provides, so it is natural to consider the pairs (processor; output information);
- The pair (①,②), (source; input information), is defined as a secondary evaluation subject and its previous characterization should be considered in the assessment of the primary subjects (see Section IV);
- In some cases, the reasoning process ($l \circ \rho$) may be considered as a whole, without separating the combination from the decision.

Thus the most important pairs (i.e., primary subjects) are:

- (③,④)—the uncertainty representation process h together with its output;
- (⑨,⑧)—the reasoning process together with its output;
- (⑩,⑧)—the pair (representation, reasoning) together with its output.

IV. URREF EVALUATION SUBJECTS

Following the previous detailed description of an elementary fusion process, this section defines the different evaluation subjects and identifies the corresponding criteria of the URREF ontology.

DEFINITION 1 (**Evaluation subject**) An *Evaluation Subject* is an item which can be assessed through the Uncertainty Representation and Reasoning Evaluation Framework according to the criteria defined in the URREF ontology.

Evaluation subjects correspond to design choices to assess for an enlightened solution design. The identification of the evaluation subjects helps to better specify and communicate the goal of the URREF ontology but also better focus the effort on the primary subjects that are uncertainty representations and reasoning schemes embedded in fusion algorithms. In the following we thus specify what is understood by "uncertainty representation" and by "reasoning."

The Joint Directors of the Laboratory (JDL) or updated version of the Data Fusion Information Group (DFIG) model fusion model (e.g., [45]) is a functional description of a series of fusion problems organized along levels. In order to solve these problems, a modeling step is required which isolates the real world entities and processes (RWEPs [46]) of interest, identifies the corresponding (uncertain) variables, possible sources of information, makes some assumption of the world's dynamics and states, represents the underlying uncertainty and finally designs the reasoning scheme by either merging, updating, revising information for an estimation (or prediction) of the variables states.

DEFINITION 2 (**Fusion problem**) A *fusion problem* corresponds to some unknown states or dynamics of the real world and for which several sources of information are available. Fusion problems typically correspond to the different levels of the JDL/DFIG model and encompass as subclasses for instance tracking, target classification, anomaly detection, threat assessment and resource management.

Note that the notion of source depends on the modeling and does not necessarily mean several sensors. Features in a classification problem could be considered as "sources." A fusion problem is solved by a fusion method.

DEFINITION 3 (**Fusion method**) A *fusion method* is a set of rules encoding a solution to the fusion problem at hand, involving several sources of information. It implements some uncertainty representations and reasoning schemes.

For instance, a Kalman filter is a fusion solution to a multi-sensor filtering problem in tracking applications. It implements an updating scheme involving a prediction step followed by a revision step within a probabilistic framework [47]. A naive Bayes classifier is a fusion solution to a classification problem, which is implemented as a naive Bayes (i.e. probabilistic) model where features (possibly provided by different sources) are assumed to be independent, followed by a maximum *a posteriori* (MAP) decision rule.

DEFINITION 4 (**Uncertain variable**) An *uncertain variable* represents a feature of the real world for which the state is unknown, partially known or uncertain. It describes the fusion problem and its state has to be estimated by the fusion method.

The concept of uncertain variable generalizes the one of random variable itself representing a random phenomenon (and generally expressed by a probability distribution), to encompass the cases of epistemic uncertainty where uncertainty is not due to the variability of the phenomenon, but to a lack of knowledge. We can define thus two types of variables relative to the nature of uncertainty (see class *UncertaintyNature* of the URREF ontology [22]): Random variable and epistemic variable.

For instance, in a Kalman filter the uncertain (random) variables correspond to the position and the speed of the target at time t and $t+1$, usually gathered into (random) state vectors $\mathbf{x}_t$ and $\mathbf{x}_{t+1}$, but also to the measurements received by the sensors represented by a state vector $\mathbf{y}_t$. In a vessel classification problem, the uncertain (epistemic) variable would be the class of the specific vessel observed.

The primary purpose of the URREF is to assess how uncertainty is handled in a given fusion method, with a specific focus on the uncertainty representation and the reasoning components. In a formal uncertainty handling, both components abide to rules and constraints defined by the *uncertainty theory* considered.

DEFINITION 5 (**Uncertainty theory**) An *uncertainty theory* is a set of axioms and rules describing uncertainty representation and reasoning. Two components can be distinguished, although possibly strongly connected:

1) The *representation* which defines *uncertainty relations (or functions)* through established sets of axioms;
2) The *reasoning* which defines *inference (or belief change) rules* to manipulate uncertainty functions and create new ones.

Uncertainty functions and inference rules can be assigned different semantics.

Examples of quantitative uncertainty theories are probability theory, evidence theory, fuzzy sets theory, random sets theory, possibility theory, and imprecise probability theory. Some qualitative theories are possibilistic logic, fuzzy logic or probabilistic logic.

A Kalman filter is framed into probability theory which itself defines probability functions to convey uncertainty notions. Probability functions must satisfy the three axioms of $P(\emptyset) = 0$ for the impossible event, $P(\Omega) = 1$ for the certain event and $P(A) + P(\bar{A}) = 1$ for any event (where $\emptyset$ denotes the empty set, Ω denotes the universe and $\bar{A}$ denotes the complement event of A). The most classical inference rule is Bayes' rule which defines the posterior probability of an event based on the occurrence of another one as $P(A \mid B) = P(B \mid A)P(A)/P(B)$. Several interpretations (or Uncertainty-Derivations [22]) can still be assigned to probability values, roughly either objective (e.g., frequentist) or subjective (e.g., degree of belief).

DEFINITION 6 (**Uncertainty relation**) An *uncertainty relation* is a mathematical or logical object conveying some notion of uncertainty. It can be an *uncertainty function* if each subset of the frame is related to a value between 0 and 1 or a binary relation such as an accessibility relation in modal logic.

The uncertainty relation covers uncertainty functions such as probability functions but also equivalence relations between states defining for instance rough sets. Uncertainty relations are the core representation of uncertainty, and express how much or how we or/and the sources are uncertain. They are defined over sets of variables, which themselves represent *what* we are uncertain about.

DEFINITION 7 (**Uncertainty Modeling Scheme**) An *Uncertainty Modeling Scheme* (UMS) is a theoretical concept that provides a mapping between (i) domain independent mathematical concepts and (ii) classes of fusion problems. A UMS

(1) introduces types of *uncertain variables* and the types of *relations* between these variables that are relevant for the modeling of a specific type of problem;
(2) provides *semantics* for a selection of uncertain relation types;
(3) formulates *assumptions* about the represented problem type;
(4) defines uncertainty functions over these variables.

For example, the UMS defining representations used by Kalman Filters introduce random variables representing the states of a dynamic process and observations. Moreover, it relates covariance matrices to the normally distributed process dynamics and observations, respectively. This model is based on the assumptions that the represented dynamic processes are linear and normally distributed. The UMS for causal Bayesian Networks associates basic conditional probabilities with uncertain causality. This model assumes Markov property, conditional independence theoretically captured by d-separation concepts and Markov Blankets. A UMS typically corresponds to a specific type of reasoning scheme. A UMS represents a theoretical basis for the solution of a specific use case (see Def. 8).

DEFINITION 8 (**Uncertain Domain Model**) An *Uncertain Domain Model (UDM)* is an artifact defined through (i) a set of uncertain variables and (ii) uncertainty relations which encode some assumptions about the real-world dynamics and states in a specific application. An UDM is a specific instantiation of a representation of the uncertainty associated with a specific real-world problem itself framed into an uncertainty theory and thus constrained by the rules and axioms. Such framing is provided by a suitable UMS (see Def. 7).

UMS defines the form of h and ρ, i.e. types of variables and functions in combination with a suitable uncertainty theory. The UDM defines the specific constellations of the variables and specific parameters used in h and ρ. The UMS supports theoretical analysis that facilitates (i) comparison of uncertainty representations and reasoning in a class of applications and (ii) an evaluation of the adequacy of a specific technique in a specific application (use case). The evaluation of a UDM supports the engineering process in the development of a specific fusion solution. An uncertain domain model could be the graphical part of a Bayesian network together with the instantiated joint probability distribution defining uncertainty over the set of variables. An uncertain domain model describes uncertainty about states of the variables and relations between variables and expresses thus some assumptions about either

(1) uncertain knowledge of possible states and dynamics of the world (generic knowledge/information/uncertainty);
(2) uncertain evidence about the current state of the world (singular information/uncertainty).

Although it is more common to associate singular evidence to a source of information, generic knowledge can also itself be derived from some source. For instance, a statistical model representing the maritime traffic and linking kinematic variables through some (possibly conditional) probability distributions (e.g. see [48]) can be interpreted as an uncertainty function derived from a specific AIS dataset covering a particular area during a given period of time, the source of this model.

DEFINITION 9 (**Uncertainty reasoning scheme**) An *uncertainty reasoning* scheme encodes some inference under uncertainty aiming at solving the fusion problem, by means of rules defined for several *uncertainty functions*.

For instance, Bayes' rule can be used "both for prediction from observations and revision of uncertain information" [49]. It can be used as a merging (fusion) rule performing a conjunction (product) of likelihoods provided by different sources. Dempster's rule itself encodes merging of (singular) testimonies for independent sources [50]. The combination rules have also different semantics and maybe thus dedicated to solve different types of problems (e.g., [49]).

DEFINITION 10 (**Source (of information)**) A *source of information* is any entity providing some piece of information.

A source of information is a relative notion and covers anything from where information can be extracted, i.e. a dataset, a database, an image, a video, a witness, etc, or the device providing it, i.e. a radar, a camera, an expert, etc. It can provide either generic or singular information.

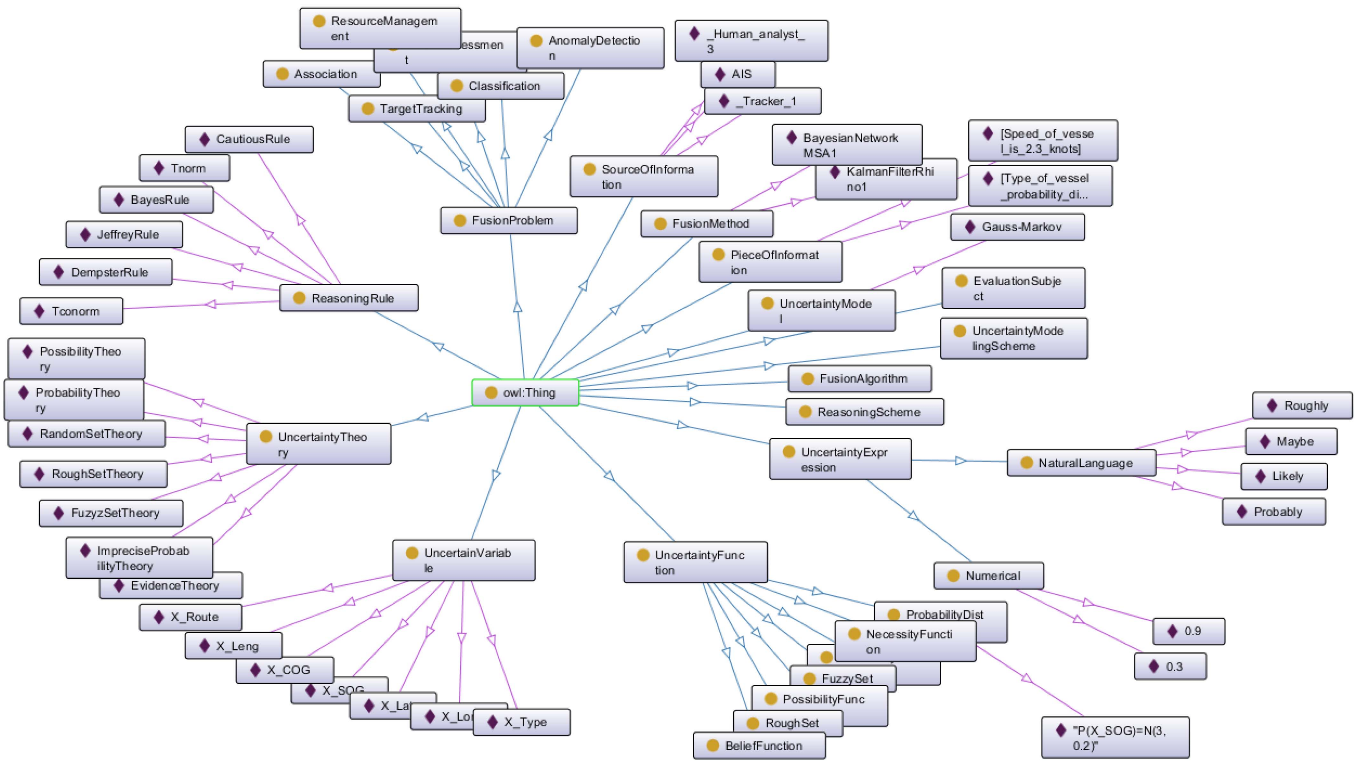

Fig. 8. URREF evaluation subjects, with sample instances (purple diamond bullets)

DEFINITION 11 (**Piece of information**) A *piece of information* is an item possibly conveying some information, and provided by a source.

The term "piece of information" is used in this paper in its most general meaning covering other notions such as evidence, knowledge and/or data. A piece of information can be as simple as a measurement (on the scale of real numbers) but could be a fact (i.e., an observation, known to be true), an uncertain statement already modeled into a given mathematical formalism (i.e., a probability distribution), an unstructured statement in natural language, etc.

Figure 8 lists the URREF evaluation subjects. Elements within rectangles with yellow circle bullets are classes. Examples of instance for each class are provided in rectangles with purple diamond bullets. The meaning of the relationship is displayed on arrows. N-ary relationships are displayed with blue arrows containing a triangle.

We identify the *primary evaluation subjects* of the URREF as:

- the **uncertainty representation**, which is either instantiated or theoretical: a particular probability distribution or probabilities in general; it may include instantiated uncertainty representations of processes in the real-world and how those processes are observed;
- the associated **reasoning** (or calculus) that comprises the combination, conditioning, updating, inference, decision, transformation rules. The calculus may be assessed while instantiated within a fusion method or theoretically, regardless any application or algorithm, focusing on the semantics for instance (e.g., Bayes' rule in general).

In URREF, the first is represented by the classes *UncertaintyTheory and UncertaintyModel*, while class *UncertaintyReasoning* represents the latter.

It is expected that a preliminary assessment of theoretical objects, either uncertainty representations or reasoning rules, is performed in the initial design phase (inception phase [51]), relying mainly on the literature and on the expertise of the fusion method designer. This pre-screening should provide guidance on the selection of appropriate models or reasoning schemes to be implemented which best suit the fusion problem at hand as far as uncertainty handling is concerned. In a second step, the assessment of instantiated representations and reasoning schemes should be assessed through a specific implementation of the fusion solution in a fusion algorithm, processing data. Then, output data analysis should provide some assessment on the implemented uncertainty handling method.

Secondary evaluation subjects of the URREF encompass other elements which either support or can be derived from the assessment of the primary subjects, but which are not the main concern of the URREF ontology:

- the **fusion method**, making use of instantiated uncertainty representations embedding **pieces of information** ϕ built according to a specific uncertainty representation process h and associated calculus $l \circ \rho$, and implemented by the **fusion algorithm**;

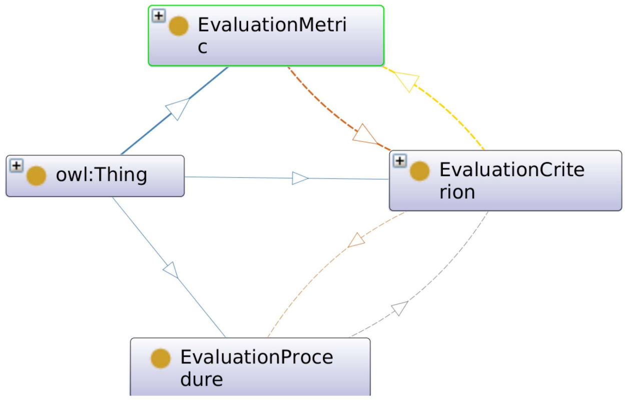

Fig. 9. In an EVALUATION PROCEDURE, EVALUATION SUBJECTS are assessed by EVALUATION CRITERIA, which are measured by EVALUATION METRICS.

- the **source** of information which provides the different and which quality may impact the whole fusion process. It can be expected that an uncertainty representation is able to properly capture and handle the meta-information about the source quality;
- the **pieces of information** input, processed and output throughout the process. Input and output information are only two special cases but others can be considered provided by internal steps such as for instance the aggregated information. The information assessment is at the core of the assessment of the uncertainty representation and reasoning. However, the development of such information quality criteria is not currently the main purpose of the ETURWG;
- the **uncertainty theory** (or framework) for uncertainty representation and reasoning (e.g., probability, fuzzy set, belief function theories). It can be assessed either theoretically, based on axioms, properties and original semantics as reported in the literature or through the assessment of the output provided by a specific **fusion algorithm** implementing the **fusion method** and specific instantiated uncertainty representations.

The fusion algorithm may be assessed either as a whole (assessing only the output) or through its different components that are the instantiated uncertainty representation (process and output information), and instantiated calculus (process and output information). Equivalently, the uncertainty theory can be assessed considering the theoretical uncertainty representation (i.e., general uncertainty function such as a probability or a belief function) on the one hand or/and the theoretical calculus apparatus (i.e., the set of reasoning tools available to this framework) on the other hand.

For each evaluation subject, there exists a corresponding set of evaluation criteria within the ontology, as illustrated in Figure 9. The quality of the source is assessed by *QualityCriterion*, the provided are assessed by *InformationCriterion*, the uncertainty representation part of the fusion method is assessed through *RepresentationCriterion* and the reasoning part is assessed through *ReasoningCriterion*.

A. Source criteria

Criteria about the source of information are necessary to characterize information input to the fusion process (other said, output by the source). The use of these criteria is rather informative than "judgmental." We assume that these initial assessments are known prior to processing the information and the question is *if* and *how* the fusion method, and especially the uncertainty representation and reasoning scheme are able to handle the different source quality dimensions. They are directly linked to the criteria on expressiveness (i.e., class *ExpressivenessCriterion*). As such, the source is a secondary evaluation subject and impacts the other subjects.

B. Information Criteria

Pieces of information (POIs) appear at different steps of the fusion process and include in particular, input data, measurement or declaration before any modeling of uncertainty (i.e., input information or dataset), the instantiated uncertainty representation (after uncertainty has been modeled), aggregated information (after the combination or inference process) and output information to be consumed by the user. Each of these POIs should be characterized according to the same subset of criteria although the expectations in their respect may differ. For instance, it is not expected that the input information be precise, nor true. Yet, it would be expected at the output. Also, comparing pieces of information at several steps of the process provides assessment of relevance (if one has an impact on the other one). Therefore, the same set of evaluation criteria should be used to assess input information, uncertain information (after h), combined information, and output information. If the same measure is used to capture this criterion, only the values (and the user's expectations) may change, not the criteria themselves. For input information, the assessment is rather a characterization, while for the other POIs during the process, the assessment criteria can be turned to optimization criteria to further tune the algorithm (e.g., maximize the Accuracy).

C. Representation Criteria

The Representation criteria (class *RepresentationCriterion*) are aimed at assessing the primary subject of evaluation within the URREF. Unsurprisingly, expressiveness is the main one. Indeed, at the inception phase [51], i.e. before any instantiation of an uncertainty representation, we are interested in the expressive power provided by its underlying uncertainty theory. This is a *prior* (theoretical) assessment driven by the problem at hand which mainly relies on analyzes of (1) the axiomatic constraints of the framework and (2) the current literature about the development of the approaches and tools to support the representation of concepts of interest as identified within the expressiveness list of criteria. The instantiated uncertainty representation should also be assessed along with the subset of criteria. An instantiated uncertainty representation is a piece of information and as such, will be assessed using the information criteria described above.

D. Reasoning criteria

This subset of criteria is so far not very detailed within the URREF ontology of criteria. Several interrelated elements must be considered:

a) the calculus and mathematical apparatus of the uncertainty theory, i.e., the set of reasoning tools available within this mathematical framework,
b) a particular instantiation of use of one of these rules, and
c) the fusion method making use of this apparatus.

For a more detailed analysis, these three subjects should be clearly distinguished, although the same criteria may be applicable and relevant to all of them. For instance, if we consider the *Consistency* criterion:

a) a particular rule of combination could be assessed according to its theoretical ability to provide consistent results,
b) a specific use of the rule which relies on other elements such as the universe of discourse selected or the type of uncertainty function to be combine, could be assessed according to the consistency criterion, and
c) a method embedding the rule with the uncertainty function and associated universe of discourse within a higher-level reasoning scheme (e.g., nearest neighbors approach, back-propagation) may also be assessed according to the same criterion of consistency.

V. CASE STUDIES

The URREF framework and its ontology component were developed through an iterative process, an essential part of which was to apply the framework to of a set of use cases. The use cases were selected to reflect a range of considerations relevant to evaluation of uncertainty representation within the context of an overall fusion application. Applying the framework to use cases grounds the ideas in concrete application areas, and helps to uncover requirements that emerge as the framework is applied to a concrete problem.

The requirements of the use cases in development are the main driver dictating what properties are needed within the URREF ontology. As such, the work on developing these use cases has been generating new insights and requirements for the URREF (e.g., [51]–[55]). The three use cases are described briefly below, with emphasis on how URREF was applied to the use case, what was learned through this process, and how the framework evolved in response to applying it to the use case.

A. Maritime Domain Awareness

We consider a use case of maritime surveillance where a harbor area is monitored by a set of sources mixing sensors and humans: *After being informed of the loss of the AIS contact with a particular fishing vessel one hour ago (at time 0), the Watch Officer (WO) now (at time t) needs to recover the track and locate the vessel. The locations of two unidentified tracks, called Vessel A and Vessel B, are provided as the only two possible locations for the missing vessel. The Watch Officer has to match the known features of the missing vessel, as reported by its last AIS contact, with the ones of the two unidentified tracks, as reported by the on-site sources. Hence, its name, MMSI, IMO, type, length, width, etc., must be known with a very high confidence to the Watch Officer.*

The sources of information available to the Watch Officer combine a variety of sensors both cooperative (e.g., Automatic Identification System (AIS)) and non-cooperative (e.g., radar, camera), whose measurement is processed either by automatic algorithms (e.g., tracker, Automatic Target Recognition (ATR) algorithm) or human analysts (e.g., camera analyst, cargo vessel's captain). The radar covers the whole area, the Infra-Red (IR) camera covers only the area around Track A, a cargo vessel is in the vicinity of Track B but too far from Track A for visual identification, and Synthetic Aperture Radar (SAR) imagery covering the whole area has been taken 30 minutes ago. Sources are imperfect and provide information which can be *uncertain* (the source itself is uncertain about its estimation or statement), *imprecise* (the source provides several possible values for the attribute estimated) and/or *false* (the value provided by the source does not correspond to the true value). Consequently, when combining the different POIs, the Watch Officer may face conflicting information.

In order to solve that fusion problem, several solutions can be designed. In [56], we illustrated how the URREF can support the designer in the decision of which uncertainty representation and reasoning method for fusion should be used. Two different fusion methods are compared: One framed into *probability theory* using Bayes' rule, and another one framed into *evidence theory* using Dempster's rule. The URREF criteria defined in classes *UncertaintyType*, *UncertaintyDerivation* and *UncertaintyNature* are used to categorize the input information highlighting the importance of the derivation of uncertainty values, as it has a direct impact on the interpretation of the output uncertainty. We stressed how the elements supporting uncertainty (e.g., variables, links between variables, uncertainty expression) crossed with the type of information (generic knowledge versus singular evidence) help in clarifying that Dempster's rule does not use generic knowledge but uncertain singular information (evidence), while Bayes' rule relies on generic information (knowledge).

B. Counter Rhinoceros Poaching Decision Support

The rhino poaching use case involves a decision support system that directs the patrol effort of the rangers to the areas with elevated risk of poaching [57], [58]. The central part of such a system is a set of Bayesian threat models, each with context evidence instantiated to correspond to a specific area or cell. A threat model is implemented as a Bayesian Network (BN) that captures the correlations between various context factors influencing the poaching (facilitators/inhibitors) as well as observable phenomena that might indicate an imminent threat. The system outputs a probability heat map that indicates the suitability for poaching at a specific point in space and time. The first attempt at applying the URREF ontology to the counter rhino poaching decision support system is presented in [59]. Given information in such a probability heat map, the rangers can position scarce resources distributed over large surface areas, such that the chance of preventing poaching is improved. Thus, the decision support system for counter rhino poaching operations covers all of the components of the OODA loop. The use of URREF concepts is demonstrated in [60] with reference to the OODA loop applied to the rhino decision support use case. Additional sources of information include human intelligence (HUMINT) reports of the field operations as well as the current status of the international rhino trafficking agencies.

Uncertainty may enter into a fusion system during both the design/modeling and routine operational phases. Selective application of the URREF to the anti-rhino poaching use case is demonstrated to characterize uncertainty during the design/modeling phase in [46] and during routine fusion system operation in [51], [60]. In particular, the URREF criteria are applied within the context of a fusion system development and deployment life cycle, as demonstrated on a high level context driven fusion approach to tracking poachers [51].

C. Cyber Threat Models

Systems for threat analysis enable users to understand the nature and behavior of threats and to undertake a deeper analysis for detailed exploration of threat profile and risk estimation. Models for threat analysis require significant resources to be developed and are often relevant to limited application tasks. In the Cyber Threat Use Case we presented and discussed a model for cyber threats which comprises an expert model and its translation into a Bayesian network (BN) as a tool for the development of practical scenarios for cyber threats analysis [61]. The BN for cyber threats is automatically generated from the expert model, highlighting vulnerabilities of systems along with threat-specific patterns, actors, actions and indicators [62]. For this use case, the goal of using the URREF ontology was to capture the quality of the knowledge. While the expert model was created manually by domain experts, by following a time consuming and expensive process, the BN was created thanks to an automatic procedure. Thus, the resulting models have different characteristics and granularity levels, and the question of their accuracy has to be addressed. For this purpose, the main URREF class considered for analysis was *RepresentationCriterion*, a general class regrouping several criteria explaining how uncertainty is characterized, captured and stored during modeling and representation stages, and introducing the most specific concepts of *Simplicity*, *Adaptability* and *Expressiveness* [52]. To analyze the model underlying the cyber threat application, *Simplicity* and *Expressiveness* criteria were considered. *Simplicity* is important since the expert model has to be created manually; *Expressiveness* is regarded to assess whether the knowledge encoded in the models is sufficient. Moreover, metrics were defined for those criteria, based on the characteristics of the models created (number of nodes in the model, density of connections). Several experiments carried out with different configurations of the model showed how the quality level of the knowledge representation, as captured by means of *Simplicity* and *Expressiveness*, is impacted by parameters of the model but also a complementary evolution of those criteria, as increasing the *Simplicity* goes hand in hand with decreases in *Expressiveness*. Future work is planned to carry out a complete assessment of knowledge representation using URREF criteria, to apply them to different BNs of different sizes and granularities, and to correlate the criteria for knowledge representation with other criteria of the URREF ontology.

VI. DISCUSSION AND CONCLUSION

Evaluation of IF systems presents intrinsic challenges due to the complexity of fusion systems and the sheer number of variables influencing their performance. In LLIF systems, the impact of uncertainty representation is well understood, and generally quantifiable. However, at higher levels of IF the approach chosen for representing uncertainty has an overall impact on system performance that is hard to quantify or even to assess from a qualitative viewpoint. This issue was recognized by the Fusion community when creating the ETURWG, with the main goal of providing an unbiased framework for evaluating the impact of uncertainty in IF systems. From the beginning, it became clear that the various approaches and technical considerations demand a common understanding that is only achievable by a formal specification of the contrasting semantics and pragmatics involved. As a result, the group developed the methodology for evaluation, the elements of the framework supporting it, a set of formal definitions of the distinct subjects under evaluation, as well as the linkage between these and the key aspects of the framework. As explained in this work, URREF is not a system or software application that can be "directly applied" to a use case. Yet, the use cases described here were essential for the group to achieve an understanding of all the nuances and idiosyncratic aspects of the process of evaluating techniques that are fundamentally different in their assumptions and views of the world. They provided the grounding for establishing the URREF concepts and mechanisms needed to mitigate the effects the underlying assumptions of each theory have in biasing the design of evaluations—each usually geared towards the strengths of one technique at the expense of the others. URREF does not completely remove the subjectivity and biases involved in evaluating uncertainty representation techniques, but is a strong step towards that direction.

ACKNOWLEDGMENT

Work in the ETURWG was commissioned by and had the full support from the International Society of Information Fusion, which vibrant community was the main driver of this volunteer effort by researchers from all over the world. Authors Kathryn Laskey and Paulo Costa would like to recognize the support of the US Army Research Office via Agreement #W911NF-11-1-0176 during the beginning of the ETURWG activities. During the last seven years, many people actively participated in the ETURWG effort, and undoubtedly left their contribution to the framework. Although we do not have the space to cite all of them, we would like to mention some who have devoted a reasonable amount of their time in the group. The authors are specially grateful for the contributions from Sten Andler, Mark Locher, Matt Roberts, Amandine Belenger, Dafni Stampouli, Gavin Powell, Max Kruger, Brian Ulcliny, Audun Josang, Alta de Waal, Claire Laudy, Simon Maskell, Kellyn Rein, Joe Steinhauer, David Hall, and Spandana Jagtap.

REFERENCES

[1] A. N. Steinberg, C. L. Bowman, and F. E. White "Revisions to the JDL data fusion model," in *Sensor Fusion: Architectures, Algorithms, and Applications III*, B. V. Dasarathy, Ed., vol. 3719. Orlando, FL, USA: SPIE, Mar. 1999, pp. 430–441. [Online]. Available: http://link.aip.org/link/?PSI/3719/430/1.

[2] J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White "Revisiting the JDL Data Fusion Model II," in *In P. Svensson and J. Schubert (Eds.), Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, Stockholm, Sweden, Jul. 2004.

[3] E. Blasch and P. Hanselman "Information fusion for information superiority," in *Proceedings of the IEEE 2000 National Aerospace and Electronics Conference. NAECON 2000. Engineering Tomorrow (Cat. No.00CH37093)*, Oct. 2000, pp. 290–297.

[4] V. Dragos "Shallow semantic analysis to estimate humint correlation," in *Information Fusion (FUSION), 2012 15th International Conference on*. IEEE, 2012, pp. 2293–2300.

[5] ——— "Assessment of uncertainty in soft data: a case study," in *Information Fusion (FUSION), 2014 17th International Conference on*. IEEE, 2015, pp. 1–8.

[6] ——— "An ontological analysis of uncertainty in soft data," in *Information Fusion (FUSION), 2013 16th International Conference on*. IEEE, 2013, pp. 1566–1573.

[7] V. Dragos, X. Lerouvreur, and S. Gatepaille "A critical assessment of two methods for heterogeneous information fusion," in *Information Fusion (Fusion), 2015 18th International Conference on*. IEEE, 2015, pp. 42–49.

[8] D. Dubois and H. Prade *Formal representations of uncertainty*. ISTE, London, UK & Wiley, Hoboken, N.J. USA, 2009, vol. Decision-making—Concepts and Methods, ch. 3, pp. 85–156, invited paper.

[9] P. Walley "Measures of uncertainty in expert systems," *Artificial Intelligence*, vol. 83, pp. 1–58, 1996.

[10] J. Pearl "Reasoning with belief functions: an analysis of compatibility," *Int. Journal of Approximate Reasoning*, vol. 4, pp. 363–389, 1990.

[11] D. Dubois and H. Prade "Evidence, knowledge, and belief functions," *Int. Journal of Approximate Reasoning*, vol. 6, pp. 295–319, 1992.

[12] G. J. Klir "Probabilistic versus possibilistic conceptualization of uncertainty," in *Proceeding of the first international Symposium on Uncertainty modeling and analysis*, 1990, pp. 38–41.

[13] H. Leung and J. Wu "Bayesian and Dempster-Shafer target identification for radar surveillance." *IEEE Transactions on Aerospace and Electronic Systems*, vol. 36, no. 2, pp. 432–447, 2000.

[14] B. Ristic and P. Smets "Target classification approach based on the belief function theory," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 2, pp. 574–583, 2005.

[15] S. Challa and D. Koks "Bayesian and Dempster-Shafer fusion," *Sādhanā*, vol. 29, pp. 145–176, April 2004.

[16] P. Smets "Imperfect information: Imprecision—uncertainty," in *Uncertainty Management in Information Systems. From Needs to Solutions*, A. Motro and P. Smets, Eds. Kluwer Academic Publishers, 1997, pp. 225–254.

[17] R. Y. Wang and D. M. Strong "Beyond accuracy: What data quality means to data consumers," *Journal of Management Information Systems*, vol. 12, no. 4, pp. 5–34, 1996.

[18] A.-L. Jousselme, P. Maupin, and E. Bossé "Uncertainty in a situation analysis perspective," in *Proceedings of the 6th Annual Conference on Information Fusion*, Cairns, Australia, July 2003, pp. 1207–1214.

[19] D. Ongaro and J. Ousterhout
"In Search of an Understandable Consensus Algorithm,"
in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC'14. Berkeley, CA, USA: USENIX Association, 2014, pp. 305–320. [Online]. Available: http://dl.acm.org/citation.cfm?id=2643634.2643666.

[20] "ISIF.org/Working Groups—Evaluation of Techniques for Uncertainty Representation Working Group."
[Online]. Available: http://isif.org/evaluation-techniques-uncertainty-representation-working-group/.

[21] "Evaluation of Techniques for Uncertainty Representation Working Group Website."
[Online]. Available: http://eturwg.c4i.gmu.edu/.

[22] P. C. G. 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: International Society for Information Fusion, 2012.

[23] A.-L. Joussselme and P. Maupin
"A brief survey of comparative elements for *uncertainty calculi* and decision procedures assessment,"
in *Proc. of the 15th Int. conf. on Information Fusion*, 2012, panel *Uncertainty Evaluation: Current Status and Major Challenges*.

[24] P. C. G. Costa, R. N. Carvalho, K. B. Laskey, and C. Y. Park
"Evaluating uncertainty representation and reasoning in HLF systems,"
in *14th International Conference on Information Fusion*, Jul. 2011, pp. 1–8.

[25] K. J. Laskey and K. B. Laskey
"Uncertainty Reasoning for the World Wide Web: Report on the URW3-XG Incubator Group,"
in *Proceedings of the Fourth International Conference on Uncertainty Reasoning for the Semantic Web—Volume 423*, ser. URSW'08. Aachen, Germany, Germany: CEUR-WS.org, 2008, pp. 108–116. [Online]. Available: http://dl.acm.org/citation.cfm?id=2889804.2889814.

[26] P. C. G. Costa, K. Chang, K. Laskey, T. Levitt, and W. Sun
"High-level fusion: Issues in developing a formal theory,"
in *2010 13th International Conference on Information Fusion*, Jul. 2010, pp. 1–8.

[27] A.-L. Jousselme, P. Maupin, and E. Bossé
"Quantitative Approaches,"
in *Concepts, Models and Tools for Information Fusion*. Artech House, Inc., Mar. 2007, vol. Chapter 8, p. 376.

[28] A. P. Dempster
"A generalization of Bayesian inference,"
Journal of the Royal Statistical Society, vol. 30, pp. 205–247, 1968.

[29] L. A. Zadeh
"Fuzzy logic and approximate reasoning,"
Synthese, vol. 30, pp. 407–28, 1975.

[30] ———
"Fuzzy Sets as a Basis for a Theory of Possibility,"
Fuzzy Sets and Systems, vol. 1, pp. 3–28, 1978.

[31] J. Pearl
Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference,
1st ed. Morgan Kaufmann, Sep. 1988.

[32] G. Shafer
"The Dempster-Shafer theory,"
New York, NY, USA, pp. 330–331, 1992.

[33] J. Pearl
Causality: models, reasoning, and inference.
Cambridge University Press, Mar. 2000.

[34] J. Dezert and F. Smarandache
"DSmT: A New Paradigm Shift for Information Fusion,"
in *Proceedings of Cogis'06 Conference*, Paris, France, Mar. 2006.

[35] A. Jøsang
Subjective Logic: A Formalism for Reasoning Under Uncertainty,
ser. Artificial Intelligence: Foundations, Theory, and Algorithms. Springer International Publishing, 2016. [Online]. Available: //www.springer.com/gp/book/9783319423357.

[36] D. Schum
Evidential Foundations of Probabilistic Reasoning.
New York: Wiley, 1994.

[37] J. B. Kadane and D. A. Schum
A Probabilistic Analysis of the Sacco and Vanzetti Evidence,
1st ed. Wiley, May 1996.

[38] D. A. Schum
"Inference Networks and the Evaluation of Evidence: Alternative Analyses,"
arXiv:1301.6737 [cs, stat], Jan. 2013, arXiv: 1301.6737. [Online]. Available: http://arxiv.org/abs/1301.6737.

[39] K. C. Chang, Z. Tian, S. Mori, and C. Chong
"MAP track fusion performance evaluation,"
in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002. (IEEE Cat.No.02EX5997)*, vol. 1, Jul. 2002, pp. 512–519 vol.1.

[40] N. UzZaman and J. F. Allen
"Temporal evaluation,"
in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers—Volume 2*, ser. HLT '11. Stroudsburg, PA, USA: Association for Computational Linguistics, 2011, pp. 351–356. [Online]. Available: http://dl.acm.org/citation.cfm?id=2002736.2002809.

[41] E. Blasch, P. Valin, and E. Bosse
"Measures of effectiveness for high-level fusion,"
in *Information Fusion (FUSION), 2010 13th Conference on*, Jul. 2010, pp. 1–8.

[42] M. A. Musen
"The Protégé Project: A Look Back and a Look Forward,"
AI matters, vol. 1, no. 4, pp. 4–12, June 2015. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4883684/.

[43] E. Blasch, K. B. Laskey, A.-L. Jousselme, V. Dragos, P. C. Costa, and J. Dezert
"URREF reliability versus credibility in information fusion (STANAG 2511),"
2013.

[44] A.-L. Jousselme, A.-C. Boury-Brisset, B. Debaque, and D. Prévost
"Characterization of hard and soft sources of information: A practical illustration,"
in *Proceedings of the International Conference of Information Fusion*, Salamanca, Spain, 2014.

[45] A. N. Steinberg and C. L. Bowman
"Rethinking the JDL data fusion levels,"
in *National Symposium on Sensor and Data Fusion*, 2004.

[46] J. P. de Villiers, K. Laskey, A.-L. Jousselme, E. Blasch, A. de Waal, G. Pavlin, and P. Costa
"Uncertainty representation, quantification and evaluation for data and information fusion,"
in *Information Fusion (Fusion), 2015 18th International Conference on*. IEEE, 2015, pp. 50–57.

[47] S. Benferat, D. Dubois, and H. Prade
"Kalman-like filtering and updating in a possibilistic setting,"
in *Proc. 14th European Conf. on Artificial Intelligence (ECAI 2000)*. Berlin, Germany: IOS Press, 2000, pp. 8–12.

[48] B. Ristic, B. La Scala, M. Morelande, and N. Gordon
"Statistical analysis of motion patterns in AIS Data: Anomaly detection and motion prediction,"
in *Proc. of the 11th Conference on Information Fusion*, Cologne, Germany, 2008.

[49] D. Dubois and T. Denœux
"Conditioning in Dempster-Shafer theory: prediction vs. revision,"
in *Proc. of the Conference on Belief Functions (BELIEF 2012)*, ser. Advances in Intelligent and Soft Computing, T. Denoeux and M.-H. Masson, Eds., vol. Belief Functions: Theory and Applications 164, Compiègne, France, 2012, pp. 385–392.

[50] G. Shafer
A Mathematical Theory of Evidence.
Princeton University Press, 1976.

[51] G. Pavlin, A.-L. Jousselme, J. P. De Villiers, P. C. G. Costa, and P. de Oude
"Towards the Rational Development and Evaluation of Complex Fusion Systems: a URREF-Driven Approach,"
in *Proceedings of the Twenty-First International Conference on Information Fusion (FUSION 2018)*, Cambridge, UK, Jul. 2018.

[52] V. Dragos, J. Ziegler, and J. P. de Villiers
"Application of URREF criteria to assess knowledge representation in cyber threat models,"
in *Information Fusion (Fusion), 2018 21th International Conference on*. IEEE.

[53] M. Kruger
"Experimental Comparison of Ad Hoc Methods for Classification of Maritime Vessels Based on Real-life AIS Data,"
in *Proceedings of the Twenty-First International Conference on Information Fusion (FUSION 2018)*, Cambridge, UK, Jul. 2018.

[54] A. de Waal and K. Yoo
"Latent Variable Bayesian Networks Constructed Using Structural Equation Modelling,"
in *Proceedings of the Twenty-First International Conference on Information Fusion (FUSION 2018)*, Cambridge, UK, Jul. 2018.

[55] E. Camossi and A.-L. Jousselme
"Information and Source Quality Ontology in Support to Maritime Situational Awareness,"
in *Proceedings of the Twenty-First International Conference on Information Fusion (FUSION 2018)*, Cambridge, UK, Jul. 2018.

[56] A.-L. Jousselme
"Semantic criteria for the assessment of uncertainty handling fusion models,"
in *Proc. of the 19th Int. Conf. on Information Fusion*, Heidelberg, GE, July 2016.

[57] H. Koen, J. P. De Villiers, G. Pavlin, A. de Waal, P. de Oude, and F. Mignet
"A framework for inferring predictive distributions of rhino poaching events through causal modelling,"
in *Information Fusion (FUSION), 2014 17th International Conference on*. IEEE, 2014, pp. 1–7. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6916121.

[58] H. Koen, J. P. De Villiers, H. Roodt, and A. de Waal
"An expert-driven causal model of the rhino poaching problem,"
Ecological Modelling, vol. 347, pp. 29–39, 2017.

[59] J. P. de Villiers, G. Pavlin, P. Costa, K. Laskey, and A.-L. Jousselme
"A URREF interpretation of Bayesian network information fusion,"
in *Information Fusion (FUSION), 2014 17th International Conference on*. IEEE, 2014, pp. 1–8.

[60] J. P. de Villiers, A.-L. Jousselme, A. de Waal, G. Pavlin, K. Laskey, E. Blasch, and P. Costa
"Uncertainty evaluation of data and information fusion within the context of the decision loop,"
in *Information Fusion (FUSION), 2016 19th International Conference on*. IEEE, 2016, pp. 766–773.

[61] T. Kiesling, M. Krempel, J. Niederl, and J. Ziegler
"A model-based approach for aviation cyber security risk assessment,"
in *Availability, Reliability and Security (ARES), 2016 11th International Conference on*. IEEE, 2016, pp. 517–525.

[62] J. Ziegler and B. Haarmann
"Automatic generation of large causal bayesian networks from user oriented models,"
in *Proceedings of the 6th Workshop on Sensor Data Fusion (INFORMATIK 2011: Informatik schafft Communities)*. Citeseer, 2011.

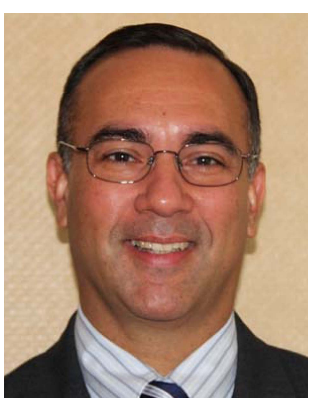

Paulo C. G. Costa is an Associate Professor of Systems Engineering and Operations Research at George Mason University, and Associate Director of the C4I & Cyber Center's Radio and Radar Engineering Laboratory.

His teaching and research interests comprise the areas of probabilistic ontologies, multi-sensor information fusion, Bayesian reasoning, predictive analysis, cybersecurity and decision theory.

He is a former fighter pilot with extensive experience in tactical and operational planning, and an expert in requirements engineering for complex systems, such as intelligent transportation and health-care support systems. Dr. Costa is an IEEE senior member, a member of the International Council of Systems Engineering, and currently serves as President of the International Society for Information Fusion (ISIF—tenures 2019 and 2020).

Anne-Laure Jousselme received her PhD degree from the Electrical Engineering Department of Laval University in Quebec City (Canada) and the Institut National Polytechnique de Grenoble (France) in 1997. Formely with Defense Research and Development Canada (DRDC), she is now with the NATO STO Centre for Maritime Research and Experimentation (CMRE) in La Spezia (Italy), where she conducts research activities on reasoning under uncertainty, high-level and hard & soft information fusion, information quality assessment and serious gaming applied to maritime situational awareness and anomaly detection. She is area editor of the International Journal of Approximate Reasoning and associate editor of the Perspectives on Information Fusion magazine. She is a member of the Boards of Directors of the International Society of Information Fusion (ISIF) where she serves as VP membership and of the Belief Functions and Applications Society (BFAS) where she serves as Secretary. She serves on program committees of the International Conference of Information Fusion and the International Conference on Belief Functions. She was Tutorial Chair of FUSION 2007 in Quebec City (CA), International Co-chair of FUSION 2015 in Washington and Technical Co-chair of FUSION 2019 in Ottawa (CA). She was general Chair of the Canadian Tracking and Fusion Conference (CTFG) in 2014 in Ottawa (CA) and Local Organizer of the International Conference of Scalable Uncertainty Management (SUM) in 2015 in Quebec City (CA).

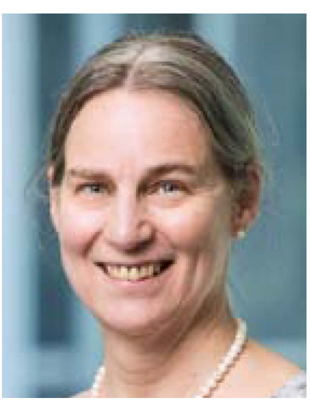

Kathryn Blackmond Laskey, Ph.D., is Professor of Systems Engineering and Operations Research at George Mason University and Associate Director of the Center of Excellence in Command, Control, Communications, Computing and Intelligence (C4I Center). She teaches and performs research on multisource information fusion, decision theoretic knowledge representation and reasoning methodology, data analytics, and decision support. A major focus of her research has been knowledge representation and reasoning for higher level multi-source fusion to support situation awareness and decision support. She has performed research in diverse application areas, including modeling the emplacement of improvised explosive devices, detecting insider threats, predicting aircraft delays, managing terrorist risk at public facilities, and planning military engagements. Dr. Laskey developed multi-entity Bayesian networks (MEBN), a language and logic that extends classical first-order logic to support probability. She was a key contributor to the development of the PR-OWL language for representing uncertainty in OWL ontologies. She serves on the ISIF Board of Directors and has is co-founder and active participant in the ISIF Evaluation of Techniques for Uncertainty Management Working Group (ETURWG). She serves on the Board of Directors of the Washington Metropolitan Area chapter of INCOSE and is past board chair of the Association for Uncertainty in Artificial Intelligence. Dr. Laskey served on several boards and committees of the United States National Academy of Sciences.

Erik Blasch is a program officer at the Air Force Research Laboratory (AFRL)—Air Force Office of Scientific Research (AFOSR) in Arlington, VA. Previously he was he was a principal scientist at the AFRL Information Directorate in Rome, NY, USA (2012–2017), exchange scientist to the Defence Research and Development Canada (DRDC) in Valcartier, Quebec (2010–2012), and Information Fusion Evaluation Tech Lead for the AFRL Sensors Directorate—COMprehensive Performance Assessment of Sensor Exploitation (COMPASE) center in Dayton, OH (2000–2009). Additional assignments include USAF Reserve Officer Col supporting intelligence, acquisition, and space technology. He was an adjunct associate professor in Electrical and Biomedical Engineering (2000–2010) at Wright State University and the Air Force Institute of Technology (AFIT) teaching classes in signal processing, electronics, and information fusion as well as research adjunct appointments at the Univ. of Dayton (2001–2014), Binghamton University (2012–2017), and Rochester Institute of Technology (2015–2017).

Dr. Blasch was a founding member of the International Society of Information Fusion (ISIF), (www.isif.org), 2007 President, and Board of Governors (2000–2010). He served on the IEEE Aerospace and Electronics Systems Society (AESS) Board of Governors (2011–2016), distinguished lecturer (2012–2018), co-chair of 5 conferences, and associate editor of 3 academic journals. He has focused on information fusion, target tracking, robotics, and pattern recognition research compiling 800+ scientific papers and book chapters. He holds 25 patents, received 33 team-robotics awards, presented 60+ tutorials, and provided 9 plenary talks. His co-authored books include High-Level Information Fusion Management and Systems Design (Artech House, 2012), Context-enhanced Information Fusion (Springer, 2016), Multispectral Image Fusion and Colorization (SPIE, 2018), and Handbook of Dynamic Data Driven Applications Systems (Springer 2018).

Dr. Blasch received his B.S. in Mechanical Engineering from the Massachusetts Institute of Technology ('92) and Masters' Degrees in Mechanical ('94), Health Science ('95) and Industrial Engineering (Human Factors) ('95) from Georgia Tech and attended the University of Wisconsin for a MD/PhD Neuroscience/Mechanical Engineering until being call to military service in 1996 to the United States Air Force. He completed an MBA ('98), MS Econ ('99), and PhD ('99) in Electrical Engineering from Wright State University and is a graduate of Air War College ('08). He is the recipient of the IEEE Bioengineering Award (Russ-2008), IEEE AESS magazine best paper Award (Mimno-2012), Military Sensing Symposium leadership in Data Fusion Award (Mignogna-2014), Fulbright scholar selection (2017), and 15 research/technical and team awards from AFRL. He is an American Institute of Aeronautics and Astronautics (AIAA) Associate Fellow, Society of Photonics and Industrial Engineers (SPIE) Fellow, and Institute of Electrical and Electronics Engineers (IEEE) Fellow.

Dr. Valentina Dragos is a research scientist, member of the Department of Information Modeling and Systems at ONERA, The French Aerospace Lab in Palaiseau, France. Valentina received Master and PhD degrees in Computer Science from Paris V University and area of research interest is in artificial intelligence, with emphasis on natural language processing, semantics technologies and automated reasoning. Since joining ONERA in 2010, Valentina has been active in several national and EU projects focused on crisis management, maritime surveillance and cyber terrorism. For those projects, her contributions addressed various topics such as: semantic interoperability for command and control systems, heterogeneous information fusion, exploitation of semantic data (HUMINT, OSINT) for situation assessment and analysis and exploration of social media. Valentina is currently involved in NATO Research Task Groups, focusing on social media exploitation for operations in the information environment.

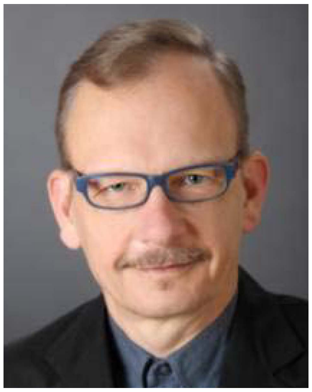

Juergen Ziegler is a Senior Technical Manager for Information Fusion at Industrieanlagenbetriebsgesellschaft mbH. He is a member of the department Competence Centers ISR.

His interests comprise the areas of Situational Awareness using methods of higher-level information fusion with applications in identification, Cyber Situational Awareness, medical diagnostics and reconnaissance. He is an expert in applications of Bayesian networks. One of the main issues of his work is model building with a focus on automatic generation of models, automatic generation of Bayesian Networks, assessment of the quality of knowledge models and ergonomic aspects of model generation. Another focus of his work is interoperability of data exchange and data fusion for situational pictures. He was one of the main authors of a STANAG about identification (STANAG 4162 Edition 3).

Pieter de Villiers is an associate professor at the University of Pretoria, South Africa and was a principal researcher at the Council for Scientific and Industrial Research (CSIR) until October 2017. He obtained his Bachelors and Masters degrees at the University of Pretoria, South Africa, and a PhD in 2008 at the University of Cambridge, UK, in statistical signal processing (particle filtering). From 2010 until 2018 he was performing research into data fusion at the Radar and Electronic Warfare competency at the CSIR. His research interests include data fusion, target tracking, Bayesian inference, nonlinear filtering, pattern recognition, graphical models and machine learning. Pieter has been regularly attending the International Conference of Information Fusion since 2010 and his ISIF activities include membership of the technical program committees, a tutorial selection committee and acting as session chairs over the years. He is the general co-chair for the 23rd International Conference on Information Fusion to be held in 2020 in South Africa. Pieter is a member of the official ISIF Evaluation Techniques for Uncertainty Representation and Reasoning Working Group (ETURWG). He is also a guest editor for a special issue at the Journal of Advances in Information Fusion (JAIF).

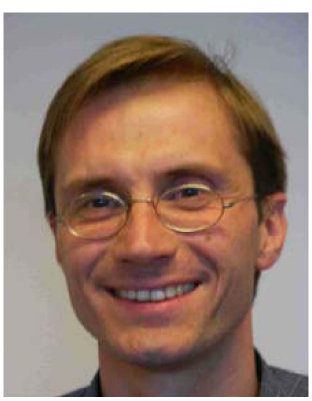

Gregor Pavlin received the M.Sc. degree in theoretical engineering and the Ph.D. degree in computer science from Graz University of Technology, Austria in 1995 and 2001, respectively. He has extensive industrial experience in safety critical software systems as well as complex AI-driven solutions. His current research interests are (i) robust algorithms and architectures supporting distributed probabilistic AI, (ii) machine learning and (iii) interoperability in complex service oriented processing systems. Since 2006 he has been a senior researcher and project manager at a corporate research lab of the Thales Group in Delft, the Netherlands. Between 2006 and 2015, he was also a part-time visiting researcher at the Intelligent Autonomous Systems lab, University of Amsterdam. He also has an extensive experience with the coordination of European and national collaborative projects. He served in the organizing committee of the Fifth International Symposium on Intelligent Distributed Computing in Delft (IDC 2011) and is also a member of the organizing committee of the 23rd International Conference on Information Fusion to be held in South Africa in 2020. He is also a member of the official ISIF Evaluation Techniques for Uncertainty Representation and Reasoning Working Group (ETURWG).

Dissecting uncertainty handling techniques: Illustration on maritime anomaly detection

ANNE-LAURE JOUSSELME
GIULIANA PALLOTTA

Detecting and classifying anomalies for Maritime Situation Awareness highly benefits from the combination of multiple sources, correlating their output for detecting inconsistencies in vessels' behaviour. Adequate uncertainty representation and processing are crucial for this higher-level task where the operator analyses information in conjunction with background knowledge and context. This paper addresses the problem of performance criteria identification and definition for information fusion systems in their ability to handle uncertainty. In addition to the classical algorithmic performances such as accuracy, computational cost or timeliness, other aspects such as the interpretation, simplicity or expressiveness need to be considered in the design of the technique for uncertainty management for an improved synergy between the human and the system. The Uncertainty Representation and Reasoning Evaluation Framework (URREF) ontology aims at connecting these criteria to other uncertainty-related concepts. In this paper, we dissect six classical Uncertainty Representation and Reasoning Techniques (URRTs) in their basic form framed into three uncertainty models of probability, belief functions and fuzzy sets, and addressing a fusion problem for maritime anomaly detection. We introduce the *Uncertainty Supports* as a means to capture what is the carrier of uncertainty and distinguish between three types of supports, that are single variables, sets of variables and uncertainty representations. The latter type indeed captures second-order uncertainty. The different URRTs are qualitatively evaluated according to their expressiveness along the uncertainty supports, and quantitatively evaluated according their accuracy and conclusiveness (uncertainty and imprecision) when processing real AIS data with pseudo-synthetic anomalies. This study illustrates a possible use of the URREF for the assessment and comparison of uncertainty handling methods in fusion systems. The framework provides solid basic foundations for a formal assessment to guide further development and implementation of fusion schemes, as well as for the definition of associated criteria and measures of performance.

Manuscript received April 27, 2018; revised March 5, 2019; released for publication April 17, 2019.

Refereeing of this contribution was handled by Pieter de Villiers.

Authors' addresses: A-L. Jousselme, NATO STO Centre for Maritime Research and Experimentation, La Spezia, Italy (E-mail: anne-laure.jousselme@cmre.nato.int). G. Pallotta, Lawrence Livermore National Laboratory, Livermore, California, USA (E-mail: pallotta2@llnl.gov).

The authors wish to thank the NATO Allied Command Transformation (NATO-ACT) for supporting the CMRE project on Data Knowledge and Operational Effectiveness (DKOE).

This paper is an extension of a preliminary version presented in [24].

1557-6418/18/$17.00 © 2018 JAIF

I. INTRODUCTION

In the field of Maritime Situation Awareness (MSA), detecting and classifying vessels' abnormal behaviour is a challenging and crucial task at the core of the compilation of the maritime picture [32, 31]. It requires not only the extraction of relevant contextual patterns-of-life information shaped for instance as maritime routes or loitering areas [42], but also the real time monitoring of the maritime traffic by a set of sensors mixing co-operative self-identification systems (such as the Automatic Identification System (AIS)) and non-cooperative systems such as coastal radars or satellite imagery, to overcome the possible spoofing of the AIS signal [44]. In many cases, intelligence information is of great help to refine and guide the search in the huge amount of data to be processed, filtered and analysed.

In order to take informed decisions, the operator needs to get good quality information. Furthermore, he/she needs to understand additional characteristics of the provided information, including for instance, how that information has been obtained, processed, or what was the context of its creation. In particular, understanding how an anomaly detector came up with an alert is of great importance to the Vessel Traffic System (VTS) operator. More specifically, the operator would benefit from knowing which were the reference data used, which were the sources processed, if the information and associated uncertainty were obtained in objective or subjective manner, whether the decision process considered the sources' quality and how, if the contextual information was considered in the decision, what was the meaning of numerical output values expressing uncertainty, and what was the underlying logical reasoning providing the answer. Second-order information quality may also be highly valuable. For example, probability maps about possible threats could be supplemented by uncertainty assessments about the validity of the probability values, represented as intervals or error estimations on algorithms performance. The benefit of including these different information quality dimensions is twofold: on the one hand, they increase the operator's situation awareness and, on the other hand, they improve trust in the use of the system.

To characterise the outputs provided to operators by some information system, the standard performance criteria of algorithms such as precision, accuracy, False Alarm Rate (FAR), Area Under the Receiver Operating Characteristic (ROC) curve (AUC), timeliness or computational cost [1, 33, 11] may not be sufficient and should be complemented by others to cover the interaction of humans and systems. For instance, some criteria such as explanation, adaptability, simplicity, expressiveness could be considered as well. The Evaluation of Techniques for Uncertainty Representation (ETUR) working group of the International Society of Information Fusion (ISIF) addresses since

2011 the definition and articulation of assessment criteria for uncertainty models and frameworks, uncertainty types, uncertainty derivation, uncertainty nature [8]. Outcomes of this work provide guidance for the selection and design of adequate tools for reasoning support, uncertainty traceability and understandability (e.g., [5, 10, 43]). It is also a first step towards some standardisation of the characterisation and assessment of uncertainty management techniques and, by extent, of fusion schemes.

Evaluating or comparing uncertainty *calculi* in the absolute is not trivial task because these make different fundamental assumptions about the nature and interpretation of uncertainty they aim at representing or processing (see for instance [29, 19]). Fundamental and global formal evaluation and analysis have been in particular presented in [48, 55, 56, 17, 29, 52] and more recently in [19]. For instance, probability, possibility and fuzzy set theory are non comparable since they are appropriate to deal with different types of uncertainty. Rather than competitors, they appear to be "complementary theories of uncertainty that utilise distinct types of uncertainty for expressing deficient information" [29]. Belief functions [47, 54] are "[...] aimed directly at modeling incomplete evidence, but certainly not incomplete knowledge," and designed to handle singular uncertainty [47, 19]. Fusion rules have their own meaning and application constraints as well. While being an updating rule, Bayes' rule is also widely used for fusion purposes (e.g., [34, 7]). However, Bayes' rule is not applicable in case of probable knowledge, unanticipated knowledge and introspective knowledge [15]. In Shafer's view, Dempster's rule is specifically dedicated to combine uncertain and imprecise singular information, such as testimonies. Dempster's rule should also be applied only to independent and reliable sources [47, 53]. It appears thus that rather than competitors the different models for uncertainty representation and associated reasoning schemes are dedicated to different problems and different types of information. As a step toward a formal analysis of uncertainty representation and reasoning techniques, the work presented in this paper aims at bringing the comparisons and descriptions of the classical uncertainty models under the Uncertainty Representation and Reasoning Evaluation Framework (URREF).

In this paper, we compare six (6) different approaches (hereafter called Uncertainty Representation and Reasoning Techniques, URRTs) to combine pieces of information from a set of heterogeneous sources (hard and soft) as the core of a maritime anomaly detector for route deviation. In complement to comparative analyses (e.g., [26, 2]), this paper identifies additional comparison elements which may have an impact on the behaviour (and performances) of the fusion schemes. The maritime anomaly detection problem is first introduced in Section II, covering route extraction and route association problems together with some associated uncertainty-related challenges. In Section III, we briefly review the current state of the URREF ontology and introduce the *uncertainty support* as part of possible refinement of the EXPRESSIVENESS criterion. Six URRTs are introduced in Section IV as alternative basic fusion schemes to solve the above defined problem, with an emphasis on the uncertainty representation. The six URRTs are compared in Section V in a qualitative way regarding their expressiveness (relatively to their uncertainty support and imperfection type captured) but also in a quantitative manner through more classical but complementary quality criteria, processing a real AIS dataset augmented with pseudo-synthetic anomalies. We conclude in Section VI on future work and further challenges to be addressed in the coming years by the Evaluation of Techniques for Uncertainty Representation working group.

II. MARITIME ANOMALY DETECTION

We illustrate the discussed methods via a real-world example of maritime anomaly detection. Although a unique definition of anomalies in the maritime domain is not available yet, we here use the term "maritime anomaly" to indicate a *deviating* behaviour from traffic normalcy, which we learn from spatio-temporal data of ships at sea. More specifically, the analysis of traffic spatio-temporal data streams provided by the AIS, a cooperative self-reporting system allows detecting and characterising inconsistencies or ambiguities, which can be ultimately transformed into usable and actionable knowledge [40]. We briefly introduce in this section the problem of route extraction, which builds the normalcy models, in our case, traffic normalcy, in such a way that these models can be further exploited for anomaly detection. We then introduce the problem of associating a vessel to a pre-defined route selected from the extracted system of routes which represents the traffic normalcy functional to the anomaly detection. We conclude this section with some uncertainty challenges related to the way we represent the maritime routes which affects the maritime anomaly detection.

A. Route extraction

The Traffic Route Extraction and Anomaly Detection (TREAD) tool presented in [42] implements an unsupervised classification approach which we here use to derive a dictionary of the maritime traffic routes by processing spatio-temporal data streams from terrestrial and satellite AIS receivers. The analysis and synthesis of the activity at sea as patterns of life is referred to as *maritime routes* and summarises the normal maritime traffic over a given period of time, a given area and a specified set of employed sensors (or sources). The AAP-6-2014 NATO glossary of terms defines a route as "*The prescribed course to be travelled from a specific point of origin to a specific destination*." A

TABLE I
Examples of source statements expressed by ϕ about different vessel attributes.

Attribute i	ϕ	Type of statement
SOG (knots)	10.3	Single measurement (precise and certain)
Type [Type1 Type2 Type3 Type4 Others]	[0.2 0.1 0 0.7 0]	Probability vector
Size	"Big vessel"	Natural language

TREAD *route* is then defined by a starting point and an ending point, together with a subset of intermediate waypoints, describing a physical path on a portion of the sea. If the area under surveillance is captured by a big enough bounding box, the route starting and ending points are the centroids of stationary areas, either coastal areas such as ports, either offshore areas such as islands, either offshore platforms, or open-sea areas such as fishing areas. The TREAD algorithm first reconstructs the single-vessel trajectories by linking the vessels' contacts and then clusters the trajectories followed by vessels into groups having the same starting and ending points. Each of these clusters represents a maritime route. The average path along a route is called *synthetic route*. The basic uncertainty around this path is computed using the trajectories of all the vessels which transited along that route in the given time window.

While only temporal streams of positional information is processed to extract the set of maritime routes, they can be further characterised by additional attributes representing the traffic of vessels composing it, such as speed, type or heading distributions. The associated uncertainty characterisation of the route along these attributes can be more or less complex, ranging from simple average values, to added variance parameters, to histograms, to estimated complete probability distributions, to sets of distributions (see Section II-C). The maritime traffic, and thus the set of routes, may be influenced by meteorological conditions (some areas may be avoided), season, economical context (ships may decide their destination based on the current stock market linked to their cargo) or areas of conflict. Also, in order to derive the average path (i.e., synthetic route) from the route cluster an extent parameter is included in the TREAD algorithm which allows adjusting the search range radius dynamically, thus enhancing the computation of intermediate waypoints while still avoiding issues such as land crossing.

The set of routes summarises thus some kinematic patterns of life of vessels over a given period of time and region, possibly layered by specific vessel types (e.g. fishing vessels, tankers, passenger vessels). This synthetic information and associated uncertainty characterises part of the context or background knowledge for the problem of route association and detection of anomalies at sea. It provides a *reference* or *normalcy* against which the current vessel contacts will be compared, and the anomalies detected.

B. The route association problem

A route deviation detector is to be designed to help the Vessel Traffic System (VTS) operator to (1) associate vessels to existing routes (and possibly predict their destination), and (2) detect abnormal behaviours to be further investigated.

We consider a vessel V observed by a series of heterogeneous sources $\mathcal{S} = \{s_1, \dots, s_N\}$ such as a coastal radar and its associated tracker, a SAR (Synthetic Aperture Radar) image with associated either ATR (Automatic Target Recognition) algorithm or a human analyst, a visible camera operated by a human analyst, the AIS information sent by the vessel itself or some intelligence source. Let $\mathcal{A}$ be the set of features of interest, either observed and thus about which information is either provided by or extracted by some sources of $\mathcal{S}$, or to be inferred. For solving our problem of route association, we consider attributes such as the position (latitude, longitude), Course Over Ground (COG), Speed Over Ground (SOG), Type, Length, and also the maritime route followed by the vessel. Let denote by $\mathcal{A}$ the set of features of interest, by $\mathcal{X}$ the set of uncertain variables corresponding to features of $\mathcal{A}$, by X_i the variable of $\mathcal{X}$ corresponding to feature $i \in \mathcal{A}$ and by U a subset of variables of $\mathcal{X}$. We further denote by Ω_i the domain of definition of X_i containing the set of its possible values, by $x_i \in \Omega_i$ a singleton of Ω_i and by $A_i \subseteq \Omega_i$ a subset of Ω_i. Let Ω be the corresponding space, defined as the Cartesian product of the Ω_i corresponding to vessel features of interest at a given timestamp t. Also, $\mathbf{x}_t = \{\phi(X_i, s, t)\}_{(i \in \mathcal{A}, s \in \mathcal{S})}$ denotes a set of information items jointly provided by some sources from $\mathcal{S}$ about some features in $\mathcal{A}$ at a specific instant in time t. This notation of information item encompasses the general case where sources provide some uncertainty about their statement and thus ϕ denotes a source statement either as a single measurement (precise and certain), either as a probability vector (expressing some uncertainty interpreted as provided by the source itself), either as a natural language expression (possibly vague), etc. Table I lists some examples. In the specific case of precise and certain measurements defined over a scale of real numbers, $\mathbf{x}_t$ would simply be a vector of real values of Ω. For the purpose of the discussion in this paper, we consider that each feature estimation is provided by a single source (while in general several sources may provide information about the same feature). Moreover, we focus on the fusion of all (singular) observations obtained at the same instant in time t. Thus for the sake of simplicity of the exposure, the index t and s will be omitted, and information items will be denoted simply by $\phi(X_i)$ or ϕ_i. Uncertainty about state transitions $\mathbf{x}_t \rightarrow \mathbf{x}_{t+1}$ will be considered in further extension of this work.

Let $\Omega_R = \{R_0, R_1, \ldots, R_K\}$ be the finite set of possible routes followed by the vessel V for the given area of interest, where R_k for $k = 1, \ldots, K$ is a pre-computed route and R_0 represents "none of the K routes": R_k, for $k = 1, \ldots, K$, is the label to be output by the fusion process corresponding to the event "The vessel V follows route R_k" and R_0 is a rejection class corresponding to the vessel following no specific pre-computed route. This class gathers the events of "The vessel is physically off-route," "The vessel is in the reverse traffic on the route," "The speed is not compatible with the route followed," "The type of the vessel is not compatible with the route followed," representing some Maritime Situational Indicators of possible interest to the VTS operator. In the following, we consider a quite simple reasoning scheme according to which an anomaly is detected based on a joint assessment (fusion) of the 5 features of Position, COG, SOG, Length, Type provided by the AIS report of the vessel and describing the route. Other said, the behaviour of a vessel V is detected as being abnormal if *the set* of its estimated features is *not compatible* with any existing route. Compared to [40], the nature of the anomaly will not be specified. However, identifying the features which contribute the most to the disbelief toward any of the routes would provide information about the nature of the anomaly.

For convenience, we partition Ω into the observation space, say Ω_o and decision space Ω_R. The fusion scheme to be designed aims thus at establishing a mapping $\Psi : \Omega_o \rightarrow \Omega_R$ such that $\hat{R} = \Psi(\mathbf{x})$ is the route label assigned to V represented by $\mathbf{x}$ (at time t). The underlying reasoning is that any observed feature at t combined with possible background knowledge contributes to a global belief (disbelief) that V is following a pre-established route from Ω_R. Indeed, if all the observed (measured) features match the corresponding feature values of a specific existing route, then the corresponding route label is assigned to the vessel. If some "inconsistency" or "conflict" exists between the set of observed features and the routes features (e.g. if the distance between $\mathbf{x}$ and each of the R_k is too high, or if the set of compatible routes according to the speed does not match the equivalent set according to the type) then V is assigned to no route and an anomaly is reported (label R_0).

The same set of pieces of information would then be used for two purposes:

(1) Associating a vessel to route, under the assumption that the sources are reliable and
(2) detecting anomalies, under the assumption that an inconsistency among the set of estimated features would reveal a possible behaviour of interest.

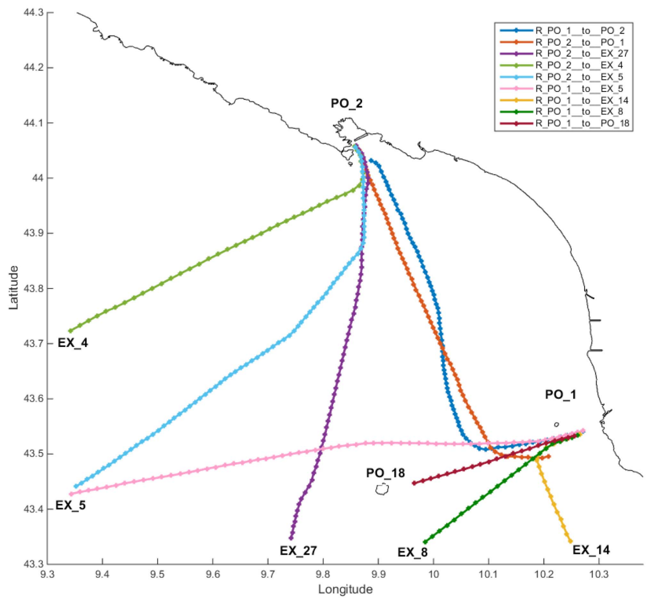

Fig. 1. Historical route prototypes extracted via the TREAD algorithm [42] in the area between La Spezia and Livorno, Italy, from AIS data (Jan 1–Feb 20 2013).

However, on the one hand, information is inherently imperfect (incomplete or imprecise, uncertain, gradual, granular [19]—See Section III-B) and on the other hand inconsistencies may arise either from sources limitations (e.g. gaps in or weak coverage of sensors, limited reasoning abilities, storage limitations, false detections or identifications), and lack of reliability in general) or malevolent behaviour of the vessel such as deception. The appropriate detection and identification of anomalies highly relies on the technique for fusing the different pieces of information and detecting inconsistencies, which include the handling of uncertainty.

C. Uncertainty in Maritime Anomaly Detection

Figure 1 illustrates the set of maritime routes previously extracted with TREAD algorithm [42] from a large number of AIS contacts for the area between La Spezia and Livorno in Italy. The used AIS data are part of a reference dataset published at CMRE [41].

As computed by TREAD, a maritime route is a cluster of vessel detections (positions from AIS contacts) with label R_k and identifies the geographical area where vessels have been observed travelling between a predefined entry point and exit point in the past temporal window. From this set of contacts (cluster) several synthetic representations can be extracted more or less complex, more or less rich, more or less precise. As an example, each route is represented in a synthetic way by a series of *intermediate* waypoints with associated average headings. Additional features characterising the traffic can be further extracted such as the distribution of speed, length and type of vessels traveling on this route. As a matter of fact, routes are, by nature, uncertain objects and the characterisation and representation of their uncertainty is of primary importance for

TABLE II
Dictionary of routes and examples of simple associated uncertainty representations

Route	Label	Synthetic route $\mathbf{r}^{(k)}$	Traffic information statistics			
	$R - Origin - to - Destin$	Position ± width [km]	$\overline{COG}$±std	SOG [knots]	Length [m]	Type [frequency]
R_1	R_PO_1__to__PO_2	$\{WP\}^{(1)} \pm 2.77$	$297 \pm 85°$	$\mathcal{N}(10,4)$	[80 : 100];[260 : 300]	[0.8 0.1 0 0 0.1]
R_2	R_PO_2__to__PO_1	$\{WP\}^{(2)} \pm 3.01$	$140 \pm 49°$	$\mathcal{N}(11,4)$	[0 : 130];[250 : 300]	[0.37 0.13 0 0 0.5]
R_3	R_PO_2__to__EX_27	$\{WP\}^{(3)} \pm 1.19$	$185 \pm 11°$	$\mathcal{N}(12,2)$	[120 : 250]	[0.75 0 0 0 0.25]
R_4	R_PO_2__to__EX_4	$\{WP\}^{(4)} \pm 2.86$	$221 \pm 31°$	$\mathcal{N}(15,4)$	[100 : 200];[260 : 350]	[0.97 0 0 0 0.03]
R_5	R_PO_2__to__EX_5	$\{WP\}^{(5)} \pm 5.08$	$209 \pm 19°$	$\mathcal{N}(11,1)$	[100 : 300];[200 : 210]	[1 0 0 0 0]
R_6	R_PO_1__to__EX_5	$\{WP\}^{(6)} \pm 1.91$	$255 \pm 18°$	$\mathcal{N}(13,4)$	[0 : 25];[110 : 300]	[0.82 0.09 0 0.09 0]
R_7	R_PO_1__to__EX_14	$\{WP\}^{(7)} \pm 1.25$	$210 \pm 90°$	$(\mathcal{N}(10,2);\mathcal{N}(18,2))$	[50 : 100];[120 : 200]	[0.93 0 0 0.07 0]
R_8	R_PO_1__to__EX_8	$\{WP\}^{(8)} \pm 0.86$	$225 \pm 14°$	$(\mathcal{N}(11,2);\mathcal{N}(19,2))$	[100 : 150];[190 : 240]	[0.38 0.24 0 0.38 0]
R_9	R_PO_1__to__PO_18	$\{WP\}^{(9)} \pm 0.98$	$244 \pm 21°$	$\mathcal{N}(11,3)$	*not reported*	[0 0 0 1 0]

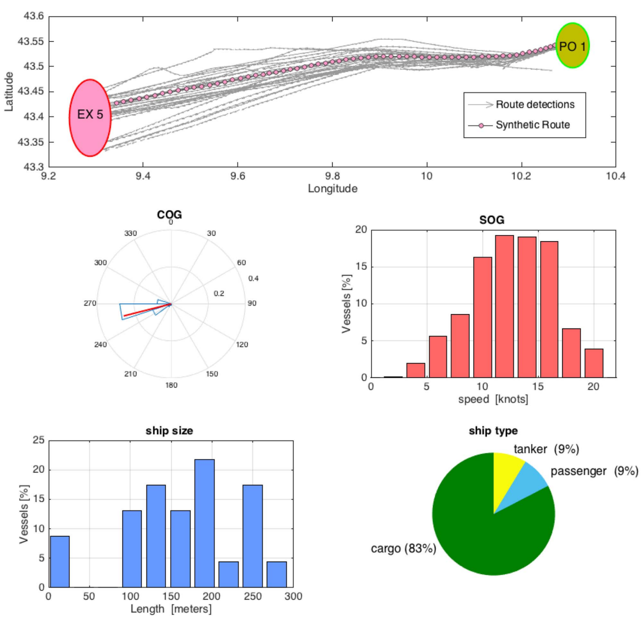

Fig. 2. An example of multi-dimensional uncertainty representation for Route R_6 with label R_PO_1__to__EX_5 reported in Table II.

a proper use of this information for the anomaly detection task. Figure 2 gives an example of how some dimensions of uncertainty for a specific route can be represented: top panel—the geographical displacements of vessel positions with respect to the synthetic (average) route; middle left panel—distribution of COGs; middle right panel—distribution of SOGs; bottom left panel—distribution of ship length; middle right panel—frequency of types of the ships which transited along that route in the given time window).

Table II lists several examples of simple uncertain representations for the different routes in the derived dictionary and illustrates how this multi-dimensional uncertainty of the routes can be encoded in a compact way.[1]

For instance, each route R_k may be represented by a prototype $\mathbf{r}^{(k)}$ corresponding to the mean or most frequent trajectory of the cluster. Those features are *precise* and *certain* values to which some *imprecision* or *uncertainty* can be added for a richer representation, based on the statistical information from the raw AIS messages which contain many additional fields of interest. The route width $w^{(k)}$ is defined as the maximum of the distances of each route point (i.e., vessel positions associated to the route) to the closest waypoint on the synthetic route. It defines an area where the transited vessels have been observed in the past.

The statistics extracted from the raw AIS dataset may serve two purposes: On the one hand, they can be used as the basic ingredient for the generic uncertainty representation captured by the route objects and, on the other hand, they are possibly transferred to express some uncertainty about new singular measurements. The histograms of the different features (SOG, COG, Length, Type) can be further interpreted as likelihood functions $p(X_i = x \mid R_k)$ (see Section IV-C) and approximated by different models. For instance, the distribution of the speed variable X_S for Route R_1 can defined by the couple $(\bar{s}_1; \sigma_1^{(s)})$ representing the mean and standard deviation of speed values estimated on the training dataset used to build R_1. With the additional assumption of a Gaussian (normal) model, these two parameters would completely define one estimation of a probability distribution for X_S. A Mixture-of-Gaussian (MoG) model could be used for the conditional likelihoods of the Speed and Length for instance, as well as more sophisticated techniques of joint density estimation, or models of dynamics of vessels, considering as well the interaction between speed and position (e.g., [46, 38]).

However in some cases, the amount of data (e.g., number of trajectories) building the cluster may not be large enough to estimate reliable distributions and considering second-order uncertainty could be appropriate (see Sections IV-E and IV-F). Also some AIS fields, especially the ones entered manually, are often missing or miss-spelled. For instance, the destination may not be specified or may not be valid, the Estimated Time of Arrival (ETA) may not be updated. The positional and

[1]The field 'TYPE' in Table II corresponds to the following encoding: [T1 T2 T3 T4 T5]=[Cargo Tanker Fishing Passenger Others]

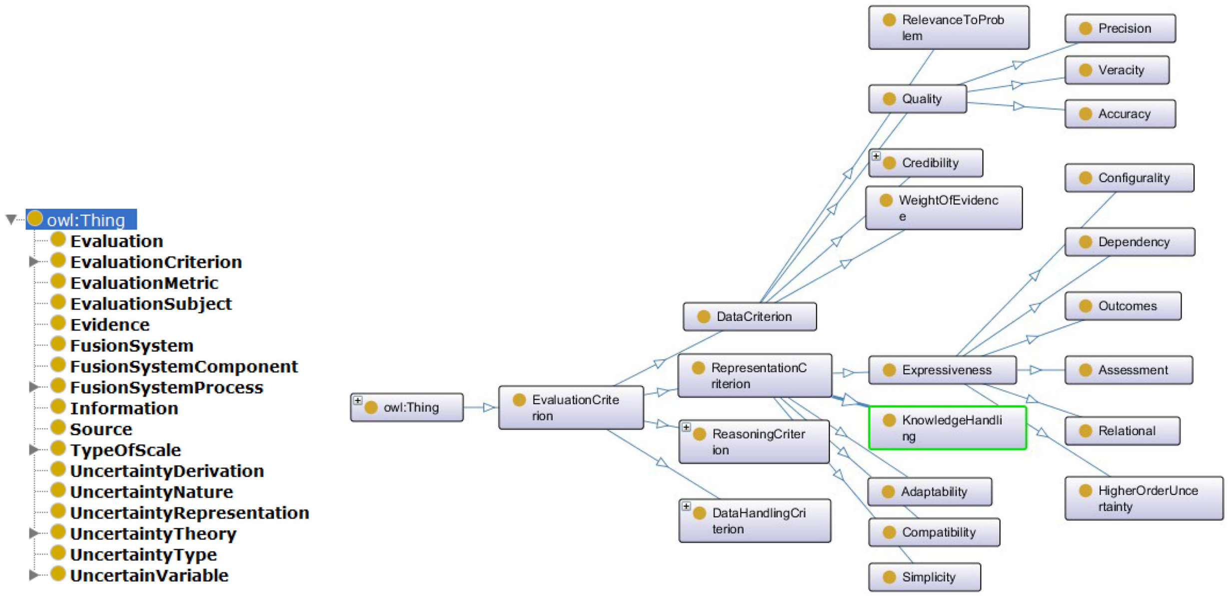

(a) (b)

Fig. 3. Excerpt of the URREF ontology for the EvaluationCriterion classes. UncertaintyType can be used to refine the Expressiveness criterion. Displayed with the Protégé software [39]. The full and last version of the ontology is available at eturwg.c4i.gmu.edu/files/ontologies/URREF.owl. (a) Top-level concepts of the URREF ontology. (b) URREF EvaluationCriterion class with subclasses. Expressiveness is a subclass of UncertaintyRepresentationCriterion. Precision and Accuracy are subclasses of DataQualityCriterion.

kinematic information being automatically sent is more reliable but can suffer from incompleteness to due a lack of coverage of the AIS receivers resulting in missing reports for a certain period of time. The non-reception of the AIS signal may arise as well from an intentional manipulation, either simply to conceal some activity either legal (e.g., fishing) or illegal (e.g., smuggling), or to keep hidden from pirates. Finally, the AIS signal can be spoofed for instance shifting the positional information to another area, or by modifying the MMSI or IMO identifier of the vessel for instance [44]. Previous studies have demonstrated that roughly 5% of AIS data is generally inconsistent (see e.g. [35]).

The consideration of these different imperfections of information is crucial in the design of maritime anomaly detection solutions. However, it requires a prior proper understanding of the origins of uncertainty, of the kinds of imperfection, of the type of information (be it relevant to a population of situations or to a single one—generic or singular) to provide a meaningful solution and to properly interpret the estimates output by the algorithms and made available to the user. In the following we provide a brief overview of the URREF ontology which aims at capturing assessment criteria on the one hand, and relevant uncertainty-related concepts that impact the solution assessment on the other hand.

III. THE UNCERTAINTY REPRESENTATION AND REASONING EVALUATION FRAMEWORK (URREF)

The URREF ontology [8] identifies, defines and links uncertainty-related concepts which come into play when evaluating the uncertainty representation and reasoning approaches underlying information fusion schemes. As the work is still on-going and some elements are currently under discussion within the ETUR group, this section only provides a partial description of the ontology focusing on the concepts relevant to this paper. The reader is referred to the ETUR working group collaboration website for an up-to-date description of the URREF ontology.[2]

The top level concept Thing (see Figure 3(a)) contains concepts such as UncertaintyNature (epistemic vs aleatory), UncertaintyType, UncertaintyTheory (mathematical framework), UncertaintyDerivation (objective vs subjective), Source (of information), EvaluationSubject and associated EvaluationCriterion. The EvaluationCriterion class is further split into DataCriteria, DataHandlingCriterion, RepresentationCriterion and ReasoningCriterion classes (see Figure 3(b)).

A. Evaluation subjects

Evaluation subjects are the elements composing the URRT which assessment through the URREF is meaningful [25]. An evaluation subject is any item which can be compared and evaluated through the URREF ontology according to a series of corresponding criteria. The uncertainty representation process (that we denote here by h) corresponds to the abstraction process of modelling [9] and aims at capturing the uncertainty (i.e. imperfection) arising from (in particular but not only):

[2]eturwg.c4i.gmu.edu/files/ontologies/URREF.owl

- the measurements, including the links between the variables, the mapping from the measurement space to the decision space, and finally the uncertainty over the decision space, including the route definition (i.e., the normalcy definition);
- the source quality, either provided by the source itself (i.e., self-confidence) which expresses some doubt about the estimated value or testimony provided, or estimated by the algorithm designer (or user) based on past experience with the source (i.e., reliability). If we relate reliability to the ability of the source to consistently provide correct outputs, then self-confidence and reliability differ in the sense that the source may have a low confidence in its declaration (singular information) while being still highly reliable (generic information), or being highly confident while being always wrong (low reliability).

The uncertainty representation is assessed by the REPRESENTATIONCRITERION of the URREF ontology.

The fusion method builds a series of uncertainty functions over the space Ω, that we split for convenience between the measurement and decision spaces, i.e. $\Omega = \Omega_o \times \Omega_R$, and involves at least one instance of the following elements: (1) a combination function ρ acting over (possibly some subsets of) Ω, (2) a mapping function g from Ω_o to Ω_R, and (3) a decision mapping l from an uncertainty function over Ω_R to a singleton of Ω_R.

The Atomic Decision Procedure (ADP) underlying the fusion method Ψ is thus composed of the elements $\{h, g, \rho, l\}$. The scheme Ψ before decision (l), outputs an uncertainty function over Ω_R representing some belief degrees we may have at time t regarding the different hypotheses of Ω_R, based on a set of pieces of information (either singular measurements received by the sources at t or generic information extracted from historical data or background knowledge and can formally be denoted as:

$$\Psi(\phi(U,S)) = \phi(X_R, \Psi) \tag{1}$$

where $U \subseteq \mathcal{X}$, $S \subseteq \mathcal{S}$ and ϕ is an information item provided by S over U. Equation (1) expresses that Ψ processes some pieces of information defined over a subset U of variables from $\mathcal{X}$, provided by a subset S of sources $\mathcal{S}$ and including some uncertainty, and outputs another piece of information defined over Ω_R, then provided by ψ as a source. As we will illustrate in Section IV, the order of the elements of Ψ is not fixed, since the fusion operation ρ can be performed within different subsets of Ω or Ω_o (e.g., URRT#1, URRT#2, URRT#3) or solely within Ω_R (e.g., URRT#4, URRT#6), the fusion can occur after the decision step (e.g., majority vote in classifier combination), etc. The reasoning scheme is assessed by the REASONINGCRITERIA of the URREF ontology.

B. Information deficiencies

In the current state of the URREF ontology, some information quality dimensions are covered by the UNCERTAINTYTYPE class (Ambiguity, Incompleteness, Vagueness, Randomness, Inconsistency). Alternative categorisations of information deficiencies could be considered instead, such as either Smets' structured thesaurus of imperfection of information [51], either Klir and Yuan's typology [30], or the typology of defects of information of Dubois and Prade [19]. In this paper, we will refer to the later one, and following the authors we will distinguish between the four information defects of *incompleteness* (or yet imprecision), *uncertainty*, *graduality*, and *granularity*.

- *Imprecision*—Refers a set of possible values, regardless how they have been obtained: The bigger the size of the set, the higher the imprecision. It represents the inability of the source to provide a single value or to discriminate between several values. Imprecision is interpreted as a type of incompleteness as it arises from a lack of information. For instance, the statement "The vessel is following either route R_1 or R_2" is imprecise and provides only incomplete information not allowing to answer the question "What is the route followed by the vessel?".
- *Uncertainty*—Arises when an agent does not know (or partially knows) if a proposition is true or false. It can be expressed by a degree (or a set of degrees) of confidence assigned to a specific (or set of) value(s) to be "true." Its nature can either correspond to a lack of knowledge (epistemic uncertainty) or to the variability of an underlying process (aleatory uncertainty). When assigned by the source itself it corresponds to "self-confidence." Uncertainty can also be expressed at the output of the fusion process itself with an equivalent interpretation, meaning that the fusion process does not provide a maximal confidence toward its output. For instance, the probability distribution over the set of possible types of vessels C_i as output by some classifier can be interpreted as an uncertainty expression, i.e., expressing a set of (normalised) degrees of confidence in the truth of the proposition "The vessel is of class C_i."
- *Graduality (or gradualness)*—Arises usually from linguistic expressions and induces propositions with some possible degrees of truth (i.e., non Boolean). That kind of imperfection allows a proposition to be more or less true or false. For instance, "The vessel is fast" is a gradual information item, using the *gradual predicate* "fast," and is typically represented by fuzzy sets. As we will illustrate later, "on-route" can be considered as a gradual predicate making maritime routes ill-defined objects.
- *Granularity*—Refers to the support over which the proposition is defined, i.e. to the set of pre-established possible values. Granularity refers to the partition granules used in the definition of a set. For instance,

the set $\Omega_1 = \{$FISHING VESSEL; NOT FISHING VESSEL$\}$ describing exhaustively the types of vessels has a rougher granularity than the set $\Omega_2 = \{$FISHING VESSEL; CARGOS; TANKERS; OTHERS$\}$ covering also exhaustively the possible types of vessels. The change of granularity is done through the operations of refinement or coarsening (see for instance [47]).

Not a defect *per se*, we also consider the dimension of "trueness" (vs falseness):

- *Trueness*—It is considered here as the criterion relating a piece of information (either input or output) to truth or to some reference value. It is defined in [22] as the "closeness of agreement between the expectation of a test result or a measurement result and a true value." The notion of trueness covers two different aspects that are *how close* the results are to truth, especially in case on measurements on continuous scales or *how frequently* the results correspond to truth, especially in case of nominal scales such as output of classifiers.

Usually, on the one hand, imprecision (or precision) and uncertainty (or certainty) are opposed [51]: "I'm certain that the speed of the vessel is between 3 and 6 knots" (Imprecise but certain statement) versus "I'm not certain that the speed of the vessel is 5 knots" (Precise but uncertain statement). On the other hand, precision and trueness are often associated in performance assessment of systems, and are gathered under the term *accuracy* in ISO 5725 [22], referring to a series of independent tests. The way these information deficiencies relate to the concepts of UNCERTAINTYTYPE, UNCERTAINTYDERIVATION and DATACRITERIA is still under discussion within the ETUR working group and is not addressed in this paper.

These five deficiencies (or imperfections) introduced above will be used in the following to characterise both input and output information of the fusion method. We will denote in the following by η the imperfection to be captured by the uncertainty representation process h.

C. Type of information

Following [19], we distinguish between *generic* and *singular* information. Generic information refers to a population of situations such as statistical models, physical rules, logical rules or commonsense knowledge. It is a synthesis of previous knowledge. Singular information is about the current state of the world such as an observation, a testimony or a sensor measurement. This distinction is similar to the one sometimes made between knowledge and evidence: According to Pearl (as cited in [16]) knowledge is understood as "judgments about the general tendency of things to happen," whereas evidence refers to the description of a specific situation."

Therefore, as a matter of convention in this paper, the notions of data, knowledge, evidence and information are all covered by the single term *information*. This is driven only by the need to avoid confusion between the terms and by no means to deny any existence of distinction between these notions. Consequently, "incomplete knowledge," "uncertain evidence," "erroneous data," etc, are all covered by the general term "*imperfect information*."

Moreover, we reserve the term *uncertainty* to the definition introduced in Section III-B. Indeed, *uncertainty* may be used sometimes abusively to cover the different types of imperfection (or information defects) as they all induce some uncertainty in the decision maker's mind. Note that *uncertainty* is also considered as the dual of *information* as classically understood in the field of Generalised Information Theory (GIT) [28]: To some increase of information corresponds equivalent reduction of uncertainty, as captured for instance by Shannon entropy measure. Hence, instead of uncertainty, we rather use the general terms of *imperfect information*, *imperfection*, *information defects*, *information deficiencies*, uncertainty being one of them.

D. Uncertainty theory

The UNCERTAINTYTHEORY class contains the mathematical theories for the representing and reasoning with uncertainty. It typically includes, but is not limited to, probability theory, fuzzy set theory, possibility theory, belief function theory, rough set theory, imprecise probability theory (see [19] for a survey). In the following, we will consider the three mathematical frameworks of probabilities, belief functions and fuzzy sets. Although geometry is not traditionally considered as an uncertainty theory, contrary to probabilities, belief functions or fuzzy sets, we also provide in this section the description of distance measures together with some justifications for its consideration.

Let us denote by Ω a set of hypotheses which could correspond either to the joint space $(\Omega_o \times \Omega_R)$, either to the measurement space only Ω_o, either to the decision space only Ω_R or to any other subset of it.

A Probability Mass Function (PMF) p satisfies the following properties:

(p.1) $p : \Omega \to [0;1]$
(p.2) $\sum_{x\in\Omega} p(x) = 1$

A probability measure P satisfies the following properties:

(P.1) $P : 2^\Omega \to [0;1]$
(P.2) $P(\emptyset) = 0$ and $P(\Omega) = 1$
(P.3) $P(A) = \sum_{x\in A} p(x), \quad \forall A \subseteq \Omega$ and p the PMF
(P.4) $P(A\cup B) = P(A) + P(B)$ if $A \cap B = \emptyset$

We have that $P(\{x\}) = p(x)$. The additivity property (P.4) constrains in particular $P(A) + P(\bar{A}) = 1$, if $\bar{A}$ denotes the negation (or complement) of A, i.e., $\bar{A} = \Omega\backslash A$.

A state of complete ignorance about the value of x is usually represented by a uniform distribution over Ω such that $P(x) = 1/|\Omega|$, for all $x \in \Omega$, where $|.|$ denotes the cardinality. The additivity property is what distinguishes probability measures from other non-additive measures such as belief functions.

Dempster-Shafer theory, or evidence theory, or belief function theory [12, 47], is often described as an extension of probability theory in which the axiom of additivity is relaxed for an axiom of sub-additivity on belief functions. In other words, the underlying distribution of a belief function is no longer defined over the singletons of Ω but rather over its powerset 2^{Ω}.

A Basic Probability (or Belief) Assignment (BPA or BBA) is a function m such that

(m.1) $m : 2^{\Omega} \to [0;1]$
(m.2) $\sum_{A \subseteq \Omega} m(A) = 1$
(m.3) Closed-world assumption: $m(\emptyset) = 0$ OR Open-world assumption: $m(\emptyset) \neq 0$

A belief function is a function Bel such that:

(Bel.1) $\mathrm{Bel} : 2^{\Omega} \to [0;1]$
(Bel.2) $\mathrm{Bel}(\emptyset) = 0$ and $\mathrm{Bel}(\Omega) = 1$
(Bel.3) $\mathrm{Bel}(A) = \sum_{B \subseteq A} m(B), \quad \forall A \subseteq \Omega.$
(Bel.4) $\mathrm{Bel}(A \cup B) \leq \mathrm{Bel}(A) + \mathrm{Bel}(B)$ for all $A, B \subseteq \Omega$ such that $A \cap B = \emptyset$

A plausibility function is a function Pl such that:

(Pl.1) $\mathrm{Pl} : 2^{\Omega} \to [0;1]$
(Pl.2) $\mathrm{Pl}(\emptyset) = 0$ and $\mathrm{Bel}(\Omega) = 1$
(Pl.3) $\mathrm{Pl}(A) = \sum_{B \cap A \neq \emptyset} m(B), \quad \forall A \subseteq \Omega.$
(Pl.4) $\mathrm{Pl}(A \cup B) \geq \mathrm{Pl}(A) + \mathrm{Pl}(B)$ for all $A, B \subseteq \Omega$ such that $A \cap B = \emptyset$

The belief function Bel and plausibility function Pl are thus respectively sub-additive ($\mathrm{Bel}(A) + \mathrm{Bel}(\bar{A}) \leq 1$) and super-additive ($\mathrm{Pl}(A) + \mathrm{Pl}(\bar{A}) \geq 1$). The uncertainty functions Bel and Pl are dual of each others ($\mathrm{Bel}(A) = 1 - \mathrm{Pl}(\bar{A})$) and can be interpreted (under Dempster's statistical view [12]) as respectively lower and upper bounds of an (unknown) probability of A: $\mathrm{Bel}(A) \leq P(A) \leq \mathrm{Pl}(A)$, $\forall A \subseteq \Omega$. The open-world assumption [54] relaxes the exhaustivity of the original Dempster-Shafer model, allowing the empty set to have a non-null mass. That means that other hypotheses than the ones initially considered in Ω can actually be true. It is interesting in our practical case of route association as this empty set would then act as a rejection class for "off-route" vessels (see Section IV-F). Evidence theory "includes extensions of probabilistic notions (conditioning, marginalisation) and set-theoretic notions (intersection, union, inclusion, etc.)" [13]. The conjunctive rule is based on the intersection between sets (see (10)). A non-null mass to the empty set denotes thus a conflict (or inconsistency) between the two belief functions combined and may be interpreted as an indicator to an anomaly. A state of complete ignorance is represented by the vacuous BPA $m(\Omega) = 1$ (or equivalently by $[\mathrm{Bel}(A); \mathrm{Pl}(A)] = [0;1]$ for all $A \subseteq \Omega$, $A \neq \emptyset$ and $A \neq \Omega$), which is distinct from the uniform distribution.

A fuzzy set μ satisfies the following properties [57]:

(f.1) $\mu : \Omega \to [0;1]$
(f.2) $\max_x \mu(x) = 1$
(f.3) $\mu(A \cup B) = \max(\mu(A), \mu(B))$
(f.4) $\mu(A \cap B) = \min(\mu(A), \mu(B))$

Compared to probabilities and belief functions which define degrees of belief regarding the occurrence (or truth) of an event, being itself either true or false, fuzzy sets define degrees of truth for events which are thus allowed to be more or less true.

Geometric distances are not an uncertainty model *per se*. However, they are at the basis of the computation of trueness, precision or accuracy in measurement data (e.g. [22]) which all convey notions of uncertainty. Moreover, pattern matching techniques (see Sections IV-A and IV-B) rely on distances computation. Finally, uncertainty may be derived from distance measures as the farther to a route the vessel, the higher our uncertainty that it follows that route (see Section IV-D). For these reasons we include here the basic properties of distance measures.

A (metric) distance function d satisfies the following properties:

(d.1) $d : \Omega \times \Omega \to [0;1]$
(d.2) $0 \leq d(x_1, x_2) \leq 1$
(d.3) $d(x_1, x_2) = d(x_2, x_1)$
(d.4) $d(x, x) = 0$
(d.5) $d(x_1, x_2) = 0 \Rightarrow x_1 = x_2$
(d.6) $d(x_1, x_2) \leq d(x_1, x_3) + d(x_3, x_2)$

All these properties define *metric distances*, but relaxing some of them lead to weaker forms of distances such as pseudo-metrics or semi-metrics. The properties of the functions introduced here correspond to some desirable behaviours of the uncertainty handling models within the fusion method to be designed. One of the tasks of the designer is to identify and select the uncertainty representation together with the associated mathematical framework in order to meet the requirements of the expected underlying logic of the method. To sum up, and referring to the basic information quality dimensions identified in Section III-B, probabilities convey the notion uncertainty only, belief functions convey both uncertainty and imprecision, while fuzzy sets convey the notion of graduality which can be assessed using distance measures.

E. Uncertainty supports

In order to refine the assessment of uncertainty representations, we introduce the concept of *uncertainty support* as an item about which some uncertainty (or

TABLE III
Examples of pieces of information for different uncertainty supports and generic and singular information

Uncertainty Support		Generic uncertainty ("0")	Singular uncertainty ("t")
(us.1)	X_i	$\eta(X_i,s,0)$ *Uncertainty about the type of vessels deduced from past AIS records*	$\eta(X_i,s,t)$ *Uncertainty about the type of a specific vessel, as provided by an ATR classifying a SAR imagery*
	X_R	$\eta(X_R,s,0)$ *Prior uncertainty about the routes followed*	$\eta(X_R,s,t)$ *Uncertainty about the route followed by the vessel at t*
(us.2)	(X_i,X_j)	$\eta((X_i,X_j),s,0)$ *Uncertainty linking type and speed of vessels, in general*	$\eta((X_i,X_j),s,t)$ *Uncertainty about the type of the vessel given the current speed*
	(X_i,X_R)	$\eta((X_i,X_R),s,0)$ *Uncertainty linking the speed and the route*	$\eta((X_i,X_R),s,t)$ *Uncertainty about the route followed by a specific vessel at t given its speed*
(us.3)	$\eta(.,0)$	$\eta(\eta(.,s_2,0),s_1,0)$ *Uncertainty about the prior distribution over the routes as lower and upper bounds*	$\eta(\eta(.,s_2,0),s_1,t)$ *Uncertainty at t about the routes previously extracted*
	$\eta(.,t)$	$\eta(\eta(.,s_2,t),s_1,0)$ *Uncertainty about the source s_2 declaration provided at t (e.g., prior reliability)*	$\eta(\eta(.,s_2,t),s_1,t)$ *Uncertainty about the current source s_2 statement itself including some uncertainty*

imperfection in general) needs to be captured and represented (in other words, what "we are uncertain about") and distinguish between:

(us.1) *Individual states of the world* as represented by any single variable of $\mathcal{X}$

(us.2) *links between states* as represented by subsets of variables from $\mathcal{X}$

(us.3) *uncertainty expression* η over the above supports (us.1) or (us.2).

Supports (us.1) are a special case of (us.2). The supports of type (us.3) correspond to abstract states covering for instance uncertainty or imprecision about a probability distribution, about a probabilistic model linking several variables, etc. The joint distribution of length and types of vessels can be itself the support of some uncertainty or imprecision since its estimation may not reflect the real distribution (due to a lack of data for instance). This is a HIGHERORDERUNCERTAINTY (i.e. second-order uncertainty), which a subclass of EXPRESSIVENESS criteria captured in the URREF ontology under the REPRESENTATIONCRITERION class (see Figure 3(b)). Advantages of considering second-order uncertainty are for instance discussed in [45, 34, 7, 2]. One of the purposes of the URREF is to analyse and capture these features of second-order uncertainty.

Table III lists examples of uncertainty supports (for both generic and singular information) together with the notation and meaning.

The examples of uncertainty supports provided in Table III are for two variables only, although these cover any subsets of variables. To distinguish between generic and singular information, we will use the indexes 0 and t respectively to the corresponding uncertainty supports. Moreover, we assign the symbol of the information source s from which the imperfection has to be captured. For instance:

$$\eta(X_T, \text{AIS dataset}, 0)$$

denotes the imperfection of the type of vessels observed in the past pertaining to the AIS dataset of interest.

In Section V-A, the URRTs will be compared according to their ability to capture the different imperfection types of our problem at hand as exemplified by Table III.

F. Evaluation criteria

We will focus in Section V-A on the EXPRESSIVENESS criterion of the REPRESENTATIONCRITERION class of the URREF ontology. Expressiveness is defined as the power of an uncertainty representation technique to convey relevant aspects of a given fusion problem [8]. The uncertainty supports are a "relevant aspect" of the problem as they are able to convey the idea of DEPENDENCY (between variables), HIGHER-ORDER UNCERTAINTY, (source) SELF-CONFIDENCE and extend to the source's reliability. Note that this assessment along the expressiveness criterion is not be ordinal in the sense that the methods are not be ordered according to their expressiveness. It is rather a comparative assessment where the methods are characterised according to their expressiveness. Instead of establishing some ranking of the URRTs, the expressiveness assessment is aimed at improving the understanding of the semantics of the different approaches. We further expand the EXPRESSIVENESS criterion to cover the ability of the URRTs to capture the different types of imperfection as defined in Section III-B.[3] Figure 3(b) displays the EVALUATIONCRITERION class split into REPRESENTATIONCRITERION, REASONINGCRITERION, DATACRITERION and DATAHANDLINGCRITERION. EXPRESSIVENESS is a subclass of REPRESENTATIONCRITERION, having itself other subclasses such as HIGHERORDERUNCERTAINTY or DEPENDENCY.

Additionally in Section V-B, we also assess the URRTs globally on their outputs through the DATACRITERION-QUALITY. Notions of TRUENESS (or Falseness), IMPRECISION and UNCERTAINTY are quantitatively evaluated when the fusion scheme solving the route association problem is implemented processing real AIS data.

IV. UNCERTAINTY REPRESENTATION AND REASONING TECHNIQUES

Six different uncertainty representation and reasoning techniques (URRTs) are presented below, as six instantiations of the fusion scheme Ψ, to be further assessed through the URREF. The URRTs presented here are very basic and simple schemes far less complete than the ones reported in the literature addressing the problems of maritime anomaly detection or route association. However, this deliberate simple exposure is aimed at "*dissecting*" the underlying uncertainty representation and reasoning, as a first step for comparison and improved understanding.

A. URRT#1: Pattern matching—Euclidean

Intuitively, the closer the observed vessel under consideration is to the centroid of the routes, the more likely it is to belong to the route. A pattern matching approach captures this basic reasoning. Prototype matching differs from template matching (such as 1-nearest neighbour) in a way that a perfect match is not expected. It provides better flexibility and allows some tolerance to handle uncertainty. A standard pattern (prototype) matching approach computes the Euclidean distances between $\mathbf{x}$ and each of the routes of Ω_R as:

$$d^{(E)}(\mathbf{x}, R_k) = \sqrt{(\mathbf{x}-\mathbf{r}^{(k)})'(\mathbf{x}-\mathbf{r}^{(k)})} \quad (2)$$
$$= \sqrt{\sum_{i\in\mathcal{A}}(x_i - r_i^{(k)})^2}$$

where $\mathbf{r}^{(k)}$ is the prototype corresponding to route R_k (see Table II), defined in the feature space Ω and $\mathbf{x}'$ is the transpose vector of $\mathbf{x}$. The ith components of $\mathbf{x}$ and $\mathbf{r}^{(k)}$ are denoted by x_i and $r_i^{(k)}$ respectively. The quantity $(d_i^{(k)})^2 = (x_i - r_i^{(k)})^2$ can be interpreted as an inverse *degree of match* of the observation x_i to the equivalent prototypical element of R_k, that we denote as $r_i^{(k)}$: The lower the square distance, the higher the degree of membership of the vessel to that route. Let us define by $\mu_i^{(k)}$ the degree of membership of $\mathbf{x}$ to R_k according its feature x_i. Then, adopting a similarity view of fuzzy sets [3, 18], $\mu_i^{(k)}$ can be defined through $d_i^{(k)}$ as, for instance:

$$\mu_i^{(k)} = \exp(-(d_i^{(k)})^2)$$

which tends toward 0 whenever the distance tends toward infinity and equals to 1 if the distance is null. Equation (2) can then be written as:

$$d^{(E)}(\mathbf{x}, R_k) = \sqrt{-\sum_{i\in\mathcal{A}}\log(\mu_i^{(k)})} \quad (3)$$

where $\mu_i^{(k)} \in]0;1]$ is a normalised degree of membership. Eq. (3) is a bisymmetrical continuous strictly monotonous mean [6]. The fusion operator in (2) is a sum (disjunction) which averages local dissimilarities with R_k along the different features. It acts as a compromise between min (conjunctive) and max (disjunctive) operators.

We then consider the following decision rule:

$$\hat{R} = \begin{cases} \arg\min_k d^{(E)}(\mathbf{x}, R_k) & \text{if } d^{(E)}(\mathbf{x}, R_k) < \epsilon_1 \\ R_0 & \text{otherwise} \end{cases} \quad (4)$$

where ϵ_1 is a threshold to be set according to the operator's needs or expectations, representing some tolerance over the global distance over the 5 features. In practice, ϵ_1 can be deduced from some aggregation of the individual thresholds ϵ_1^i for each feature. This decision rule allows some imprecision in the decision space as it can lead to a set of possible routes, without identifying a single one. An anomaly is detected if it does not match any route. Many anomaly detection approaches are based on distances computation as an implementation of the notion of "closeness to normalcy" (e.g. [11]). Semantic distances can also be used to assess the different meanings between attributes (e.g. [4]).

B. URRT#2: Pattern matching—Mahalanobis

A modified version of the Euclidean pattern (prototype) matching scheme is obtained by using the Mahalanobis distance:

$$d^{(M)}(\mathbf{x}, R_k) = \sqrt{(\mathbf{x}-\mathbf{r}^{(k)})'\boldsymbol{\Sigma}^{-1}(\mathbf{x}-\mathbf{r}^{(k)})} \quad (5)$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of the random vector $\mathbf{X}$ associated to $\mathbf{x}$, whose coordinates are r.v. X_is. The superscript $^{-1}$ denotes the inverse matrix. The element $\sigma_{i,j}$ of $\boldsymbol{\Sigma}$ is the covariance of X_i and X_j defined as $E(X_i, X_j) - E(X_i)E(X_j)$ where E is the expectation operator such that $E(X) = \sum xp(X = x)$ for a discrete random variable X. The same decision rule (4) than for the Euclidean pattern matching is used. However, another threshold ϵ_2 must be used instead of ϵ_1, based on the covariance matrix.

As in (2), the fusion operator in (5) is a disjunction but including weights which would discount the local individual dissimilarities relatively to the variance of their corresponding feature, and pairs of errors relatively to their covariance.

[3]Note that this link between EXPRESSIVENESS and UNCERTAINTYTYPE is not currently implemented in the URREF ontology and is at a stage of proposal for inclusion.

The Euclidean and Mahalanobis distances in (2) and (5) are well suited to features defined over numerical and continuous scales while they reduce to logical AND for nominal variables such as the type. Better suitable distance measures are usually used based on the aggregation of individual for each feature, possibly using different definitions than the square difference (e.g. [21]). Other distances such as the log-normal probability density (e.g. [1]) would account for the routes statistics as well. Mahalanobis distance is used in [36] to associate vessel tracks to maritime routes.

C. URRT#3: Probability-based—Bayesian

In the standard Bayesian approach to fusion, the function $p(\mathbf{X} = \mathbf{x} \mid R_k)$ represents the likelihood of observing a specific set of values $\mathbf{x}$ on a given route R_k, and is usually derived from past observations used to compute the routes. The different observations are combined following Bayes' rule:

$$P(R_k \mid \mathbf{x}) \propto p(R_k) \prod_{i \in \mathcal{A}} p(x_i \mid R_k), \quad \forall R_k \in \Omega_R \tag{6}$$

under the assumption of *independent and identically distributed observations*. $p(R_k)$ is some prior probability that the vessel follows a specific route. The resulting posterior probability $P(R_k \mid \mathbf{x})$ represents some belief that the route followed by the vessel of interest is R_k given that we currently observe $\mathbf{x}$. A normalisation factor ensures that a probability distribution is obtained. Equation (6) is known as Naïve Bayes model in classification. This combination rule (6) can be written using the individual posterior probabilities as $P(R_k \mid \mathbf{x}) \propto p(R_k)^{-(|\Omega_R|-1)} \prod_i p(R_k \mid x_i) p(x_i)$. The decision rule is the Maximum A Posteriori (MAP) probability:

$$\hat{R} = \begin{cases} \arg\max_k p(R_k \mid \mathbf{x}) & \text{if } p(R_k \mid \mathbf{x}) > \epsilon_3 \\ R_0 & \text{otherwise} \end{cases} \tag{7}$$

where ϵ_3 is a threshold: if the posterior probability is too uniformally distributed among the routes, then no clear matching is detected and an anomaly is returned. The Bayesian reasoning scheme is at the basis of the Bayesian network approach proposed for instance in [23].

The fusion operator is a conjunctive operator, i.e. the product of individual likelihoods. It has the property of decreasing very fast to 0 as the number of features to be combined increases. Also, the result is exactly 0 if only one likelihood is null.

D. URRT#4: Probability-based—Non-Bayesian

In a still probabilistic but non-Bayesian approach, each measured feature is considered providing some evidence about the membership of $\mathbf{x}$ to a given route R_k. For instance, $p_s(R_k) = p(R_k \mid x_s)$ is the contribution of the speed observation to the membership of the vessel V to R_k and is interpreted as the probability that V belongs to R_k given (or according to) the estimated speed. Then, the observations are aggregated by a weighted sum as:

$$p(R_k \mid \mathbf{x}) = \sum_{i \in \mathcal{A}} \alpha_i p(R_k \mid x_i), \quad \forall R_k \in \Omega_R \tag{8}$$

where $\alpha_i \in [0,1]$ is a weight reflecting either the confidence in the soft decision values computed by the individual sources, and possibly be deduced from $p(x_i)$, or the relevance of the features to the fusion problem (for instance, the position and heading may be given a higher weight than the type). This rule is derived in [27] from (6) under the assumption of uniform $p(R_k)$. Contrary to the Bayesian approach, the posteriors are combined. The decision rule is then (7).

The fusion operator is a disjunctive operator, as in (2) and (5), but probabilities are combined rather than distances.

E. URRT#5: Transferable Belief Model (TBM) model-based

The reasoning scheme considered here is the one proposed in [45, 14] within the Transferable Belief Model (TBM) framework and making use of the Generalised Bayes Theorem (GBT) [50] as the combination rule, given by the following plausibility measure for a subset of routes A:

$$\mathrm{Pl}(A \mid \mathbf{x}) = 1 - \prod_{R_k \in A} (1 - \mathrm{Pl}(\mathbf{x} \mid R_k)), \quad \forall A \subseteq \Omega_R \tag{9}$$

where $\mathrm{Pl}(A) = \sum_{A \cap B \neq \emptyset} m(B)$ is the plausibility of $A \subseteq \Omega_R$, with m being a Basic Belief Assignment (see Section III-D). $\mathrm{Pl}(A \mid \mathbf{x})$ is the conditional plausibility of A and is interpreted as the maximum confidence that can be assigned to A (i.e., that the route followed belongs to the subset A) given that $\mathbf{x}$ has been observed. As proposed in [45], $\mathrm{Pl}(\mathbf{x} \mid R_k)$ is the least committed plausibility function corresponding to the probabilistic likelihood function considered as the pignistic probability. For a BBA m, the pignistic probability [49] is defined for any singleton of Ω_R as $\mathrm{BetP}(R_k) = \sum_{R_k \in A} m(A)/|A|$. As introduced in Section III-D, pairs of plausibility and belief values can be interpreted as intervals over the probability of any subset of routes $A \subseteq \Omega_R$. However, if we restrict to singletons only, (9) reduces to the product of plausibility under the independence assumption. The decision rule requires then two steps: (a) the transformation of the Pl measure into a probability distribution over Ω_R (e.g. the pignistic probability) such that (b) the MAP rule (7) can be applied (with the appropriate threshold).

The fusion operator is again a conjunctive operator with similar properties than the ones described in Section IV-C.

F. URRT#6: Belief functions—Database query

Similarly to the probabilistic non-Bayesian URRT#4, each observed feature x_i of $\mathbf{x}$ is assumed to provide some evidence about route R_k being followed by V. The uncertainty is modeled by belief functions rather than probabilities. Each observation x_i is regarded as a query to Ω_R such that only the items (i.e. routes) satisfying the associated criterion are retrieved, to form a set of possible routes A_i according to x_i. A_i is the subset of routes satisfying the query x_i:

$$A_i = \{R \in \Omega_R \mid x_i \in \Omega_i\}$$

For instance, A_1 is the set of routes compliant with a measured speed of 5 knots. The multivalued mapping between the observation space Ω and decision space Ω_R assigns to any singleton of Ω a subset of Ω_R. Let us consider that some singular information about the observation x_i under the form of a probability, provided for instance by a classifier: $p^{(T)}(X_T = \text{Cargo}) = 0.4$ is the probability that the observed vessel is a Cargo type, as estimated based on current observations. Let $\mathbf{p}_T = [0.4\ 0.3\ 0\ 0.3\ 0]'$ be the uncertainty of the classifier (source) expressed as a probability distribution about the type of the vessel. This uncertainty is transferred to the corresponding subsets of Ω_R previously defined by the multivalued mapping, defining thus a BBA m_i over Ω_R, where the numerical weight $m_i(A_i) = p_i(x_i)$ is interpreted as the degree of belief that can be assigned to A_i and to none other subset of A_i. Then, $A_{\text{Cargo}} = (R_1, R_2, R_3, R_4, R_5)$ is the set of routes possibly followed by cargo vessels and is assigned a weight of 0.4. Equivalently, $A_{\text{Tanker}} = (R_2, R_3, R_5)$ and $m(A_{\text{Tanker}}) = 0.3$ and $A_{\text{Passenger}} = (R_2, R_5)$ and $m(A_{\text{Passenger}}) = 0.3$. This multivalued mapping does not induce a probability distribution over Ω_R but a BBA.

The resulting BBA m over Ω_R is obtained by combining the individual contributions of each feature by the conjunctive rule, where weights are assigned to conjunctions of sets of routes A_i and A_j:

$$m(A) = \sum_{A_i \cap A_j = A} m_i(A_i) m_j(A_j), \quad \forall A \subseteq \Omega_R \tag{10}$$

The rule (10) defines a conjunctive fusion based on the intersection between sets. The decision rule is similar to (7) but considers the conflict measure as a criterion for anomaly:

$$\hat{R} = \begin{cases} \arg\max_k \text{BetP}(R_k) & \text{if } m(\emptyset) < \beta \\ R_0 & \text{otherwise} \end{cases} \tag{11}$$

where BetP is the pignistic transformation of m. The quantity $m(\emptyset)$ is the BBA of the empty set after combination and represents the global weight of conflict between all the sources (or features).

V. ASSESSMENT OF URRTS

We now characterise the different approaches previously described through the URREF and its associated ontology, Expressiveness, in Section V-A and output Quality criteria in Section V-B.

A. Expressiveness assessment

Table IV summarises the comparative description of the 6 URRTs presented in Section IV as candidate solutions to the same problem of maritime route detection. The expressiveness of the URRTs relatively to different uncertainty supports identified in Section III-E is first assessed in a binary way, so that an empty cell means that the technique (as actually defined in the previous section) does not account for the uncertainty on the corresponding support. The types of imperfection (*graduality*, *uncertainty*, *imprecision*) are mentioned in case the URR technique captures them, together with the corresponding notation. The *granularity* is kept constant for all the methods and is just reported as the list of possible values for all variables in the first rows. In the third part of the table the reasoning schemes are compared along their respective uncertainty representation, marginalisation, decision elements.

1) URRTs analysis:

The uncertainty supports introduced in Section III-E are mentioned for each method in Table IV. We thus refer the reader to Table IV for details on the uncertainty supports about the URRTs analysis.

URRT#1—We observe that the standard pattern matching approach (URRT#1) does not account for many uncertainty supports: The route representation is considered as precise and certain since the prototypes are defined by single values (either the mean, or the mode for the type); the dependency between variables is not considered, nor is the possible links between routes; sources' uncertainty (or self-confidence) about their singular declaration at t is not considered; sources' reliability is not represented, nor is any second-order uncertainty. URRT#1 captures a single imperfection type as a notion of *graduality* through a distance measure, the route prototype being considered as a reference: the distance to route can be interpreted as a degree of membership of $\mathbf{x}$ to R_k. From this generic information, a singular imperfection is further derived as $\eta(X_R, t)$ combining with the observation of the vessel at t. The fusion is performed through the distance definition by a sum operator acting as an average of inverse of similarities along the different features of $\mathcal{A}$: The higher the local similarities, the lower the global distance and the higher the membership of $\mathbf{x}$ to R_k.

URRT#2—The extension of URRT#1 using the Mahalanobis distance as described by URRT#2, accounts for both the spread of the routes along the different features (through the individual standard deviations σ_is) and the dependency between variables (through the covariances $\sigma_{i,j}$s). The variance can be interpreted as a measure of *imprecision* regarding X_i. The covariance describes how the variables vary with each other, measures the dependency between them, and expresses then some statistical *uncertainty* on the link between X_i and X_j. Compared to URRT#1, URRT#2 considers some imperfection about the reference objects (the routes). Still,

TABLE VI
Expressiveness comparison of Uncertainty Representation and Reasoning Techniques based on Uncertainty Support.

	URRT Name	URRT#1	URRT#2	URRT#3	URRT#4	URRT#5	URRT#6
	Mathematical framework	Geometry	Geometry, Statistics	Probability, Bayesian	Probability, non-Bayesian	Evidence theory, TBM	Evidence theory
Domains	Ω_1-Latitude	$\{[43.2; 43.4], \dots,]44; 44.2]\}$					
	Ω_2-Longitude	$\{[9.3; 9.5], \dots,]10.1; 10.3]\}$					
	Ω_3-Speed	$\{[2.0; 3.0], \dots,]33; 34]\}$					
	Ω_4-Heading	$\{[0; 20], \dots,]340; 360]\}$					
	Ω_5-Type	{Cargo; Tanker; Fishing; Passenger; Other}					
	Ω_R-Route	$\{R_0, R_1, \dots, R_8\}$					
UNCERTAINTY REPRESENTATION (EXPRESSIVENESS)							
Uncertainty supports (see Table III)	X_i		Imprecision $\eta(X_i^0) = E(\{X_i^0\}^2)$				Uncertainty $\eta(X_i^t) = p_i(X_i^t)$
	X_R			Uncertainty $\eta(X_R^0) = p(R_k)$			Uncertainty + Imprecision $\eta(X_R^t) = m_i(A)$
	(X_i, X_j)		Uncertainty $\eta(X_i^0, X_j^0) = E(X_i^0, X_j^0)$				
	(X_i, X_R)	Graduality $\eta(X_i^0, X_R^0) = \mu_i^{(k)}$	Graduality $\eta(X_i^0, X_R^0) = \mu_i^{(k)}$	Uncertainty $\eta(X_i^0, X_R^0) = p(x_i \vert R_k)$	Uncertainty $\eta(X_i^0, X_R^0) = p(x_i \vert R_k)$	Uncertainty $\eta(X_i^0, X_R^0) = \mathrm{Pl}(x_i \vert R_k)$	Imprecision $\eta(X_i^0, X_R^0) = A \subseteq \Omega_R$
	η^t				Uncertainty $\eta(\eta^t) = \omega_i$		
	η^0					Imprecision $\eta(\eta^0) = [\mathrm{Bel}_i(A); \mathrm{Pl}_i(A)]$	Imprecision $\eta(\eta^0) = [\mathrm{Bel}_i(A); \mathrm{Pl}_i(A)]$
REASONING							
ADP elements	h	$\eta(X_i^0, X_R^0)$	$\eta(X_i^0, X_R^0)$ $\eta(X_i^0)$ $\eta(X_i^0, X_j^0)$	$\eta(X_i^0, X_R^0)$ $\eta(X_R^0)$	$\eta(X_i^0, X_R^0)$ $\eta(\eta^t)$	$\eta(X_i^0, X_R^0)$ $\eta(\eta^0)$	$\eta(X_i^0, X_R^0)$ $\eta(X_i^t)$ $\eta(X_R^t)$ $\eta(\eta^0)$
	ρ	Ω, $\sqrt{\sum_i (.)^2}$	Ω, $\sqrt{\sum_{i,j} \sigma_{i,j} (.)^2}$	Ω, $\prod_i (.)$	Ω_R, $\sum_i w_i(.)$	Ω_R, $\prod_i (.)$	Ω_R, $(\cap, \sum, \prod_i)$
	$\downarrow \eta^t(X_R)$	Graduality $d^{(E)}(\mathbf{x}, \mathbf{r}^{(k)})$	Graduality $d^{(M)}(\mathbf{x}, \mathbf{r}^{(k)})$	Uncertainty $p(R_k \vert \mathbf{x})$	Uncertainty $p_{\mathcal{A}}(R_k)$	Uncertainty + Imprecision $\mathrm{Pl}(R_k \vert \mathbf{x})$	Uncertainty + Imprecision $m_{\mathcal{A}}(A)$, $A \subseteq \Theta$
	l	$\min d + \epsilon_1$	$\min d + \epsilon_2$	$\max p(R_k \vert x_i) + \epsilon_3$	$\max p_{\mathcal{A}}(R_k) + \epsilon_4$	$\max \mathrm{BetP}(R_k) + \epsilon_5$	$\max \mathrm{BetP}(R_k) + \epsilon_6$ on $m(\varnothing)$

there is no consideration of singular uncertainty about the observations at t, excepted the *graduality* measured by the distance to the prototype route.

URRT#3—As in URRT#1, the independence assumption between variables applies to the Bayesian approach presented in URRT#3. No consideration for either the source's reliability nor self-confidence and the measurement itself is assumed both certain and precise by the source. Rather *uncertainty* is considered over the mapping between Ω and Ω_R where the likelihoods $p(x_i \mid R_k)$ describe how likely it is to obtain some specific measurement given that the vessel follows route R_k. Prior *uncertainty* about routes is explicitly considered by $p(R_k)$ which could be based on other contextual information such as meteorological or seasonal. The fusion is done through a product operator which has the drawback of decreasing very rapidly to 0 once one of the likelihoods is very low. This rule is named "severe" for that reason [27], since it is very sensible to one source's negative opinion. The product is a conjunctive operator (corresponding to a logical AND) making the underlying assumption either that all the measurements are correct, or that all the sources are reliable. Although the independence assumption between features is in our case wrong, this naive Bayesian fusion rule is however shown to provide good (accurate) results. This can be explained by the randomness of likelihood estimates, the low variance mitigating the obvious bias [20]. Including the source's reliability about measurements is a direct extension of URRT#3 (see for instance [34]), as well as considering the dependencies between variables. The final assessment $\eta(X_R,t)$ expresses some *uncertainty* degree that the vessel is actually following route R_k.

URRT#4—In the probabilistic non-Bayesian approach of URRT#4, the individual probabilities are assumed to provide local belief degrees toward each route. They are summed up to give a global belief so that the higher the belief degree according to each feature, the higher the global belief. URRT#4 does not consider the dependency between features. However, some notion of source's *reliability* can be captured by the weights ω_i that can be derived from some likelihood measures extracted from a confusion matrix. This expresses some *second-order uncertainty* about the source's declaration at t. The combination rule is a disjunction (logical OR) and is known to be less sensitive to estimation errors (unreliable sources), and to single source's opinion [27] making the approach more robust. This is a more cautious rule to be used in case of less reliable sources.

URRT#5—URRT#5 may be seen as an extension of URRT#3 within the TBM model, where non-additive functions (i.e., plausibility functions) are used rather than probabilities. The plausibility function $\mathrm{Pl}(\mathbf{x} \mid R_k)$ models some *imprecision* about the (assumed precise but unknown) likelihood function $p(\mathbf{x} \mid R_k)$ (itself capturing some *uncertainty*) used in URRT#3. Equation (9) is obtained under the assumption of a vacuous prior on Ω_R, meaning that no prior uncertainty on routes is considered. The output of the GTB expressed by $\eta(X_R,t)$ being also a plausibility function, assigns plausibility values to *subsets* of routes and captures thus some *imprecision* over Ω_R. $\eta(X_R, t)$ defines then *second-order uncertainty* by means of a couple belief-plausibility measure expressing some uncertainty about the posterior event $(R_k \mid \mathbf{x})$. This second-order uncertainty is not considered in the traditional Bayesian approach where the probability estimations are considered certain. Other equivalent approaches exist framed into imprecise probability or robust Bayesian frameworks.

URRT#6—In URRT#6, the *uncertainty* output by the sources about the measurement provided at t is considered. Rather than a single (precise and certain) measure, each source outputs a probability distribution over the set of values of their respective feature which induces as many multivalued mappings over Ω_R when querying the dictionary of routes. The multivalued mappings define some *imprecision* over the set of routes, since to a single value in Ω_i corresponds a subset A of Ω_R. The prior imperfection on the links within Ω or between Ω_i and Ω_R is characterised as sets of routes (imprecision) satisfying some criteria about the features. This *imprecision* is further combined with the singular uncertainty of the source at time t defining the resulting BBA $\eta_i(X_R, t)$. The main characteristic of this scheme is to deal with subsets of routes, in a qualitative way, with an additional quantification. The explicit notion of *conflict* is a way to detect inconsistencies between the subsets of routes compatible with each feature. The fusion is performed through a conjunctive rule, assuming the independence between sources as well as totally reliable sources.

2) Interpretation:

The type of imperfection handled by URRT#1 and URRT#2 is graduality meaning that the route is considered as an "ill-defined object," with fuzzy boundaries, to which vessels belong more or less. The distance measure provides an aggregated inverse degree of membership of the vessel to a given route: If the distance is low then the vessel belongs to the route with a high degree of membership. Contrarily, the other methods (URRT#3 to URRT#6) express a "degree of belief" that the vessel is following the route. This is a difference between a binary event (URRT#3 to URRT#6) and a fuzzy event (URRT#1 and URRT#2). This semantic aspect highlights the need for a clear semantics for the concept of *maritime route*, whether it means either "following a specific path and thus ending in a specific destination" (binary event) or "being positioned on a portion of the sea with ill-defined boundaries" (fuzzy event).

3) Enrichment of basic URRTs:

Each of the URRT above could be enriched to account for more uncertainty supports. As examples only, the reliability of the sources is classically considered in URRT#6 by introducing discounting (or reinforcement) operations for belief functions such as described in [37]. Also, the reasoning scheme of Equation (6) in URRT#3 can be enriched by considering the reliability of the sensors in providing accurate measurements, and introducing factors $p(Z_i \mid X_i)$ where Z_i is the measurement provided by the source while the true value was X_i, as proposed in [34] for instance. URRT#3 can be easily implemented as a Bayesian network (e.g., [23]) where the dependency between variables is considered. A Bayesian network has the advantage of a better *transparency* in the reasoning for the user, which could also be an interesting assessment criterion to be considered in the URREF ontology. Moreover, the *computational cost* is improved by local computations.

B. Output quality assessment

The qualitative analysis above is now complemented by a quantitative analysis based on more standard criteria. We provide below a series of possible criteria for quantitative assessment of the six URRTs discussed in this paper, that we implement to discover abnormal behaviours of vessels within a real AIS dataset complemented by pseudo-synthetic anomalies.

1) Output criteria:

We consider the output quality criteria of Trueness (or falseness), Precision (or imprecision) and Certainty (or uncertainty). The Trueness notion captures how correct the results are after decision. To measure this criterion we use the standard F_β-score (or measure), classically defined as:

$$\mathrm{Tru}(\Psi) = F_\beta(\Psi) = \frac{(1+\beta^2)TP}{(1+\beta^2)TP + \beta^2 FN + FP} \quad (12)$$

where $\beta \in [0; 1]$ is a parameter weighting the two types of errors, TP, FN and FP are the number of true positives, false negatives and false positives respectively, N is the number of negative samples and P is the number of positive samples.

The Imprecision and Uncertainty are assessed *before* the final decision (labelling to a single route) is taken, and quantify how much the URRT is non-specific and uncertain before the labelling, respectively. They are assessed through the Hartley measure and Shannon entropy[4]:

$$\mathrm{Imp}(\Psi) = \frac{1}{\log_2(|\Omega_R|)} \log_2(|A|) \quad (13)$$

$$\mathrm{Unc}(\Psi) = -\frac{1}{\log_2(|\Omega_R|)} \sum_{R \in \Omega_R} p(R) \log_2(p(R)) \quad (14)$$

where $|.|$ denotes the cardinality of sets and p is the probability distribution over the set of routes before decision is taken. The equations above are normalised versions of the measures. In (13), A is the set of compatible routes according to the corresponding decision criteria.

[4] In (14), the distances in URRT#1 and URRT#2 are transformed into probability distributions over the set of routes, with thus a different meaning. Shannon entropy may not be an adequate measure in this case.

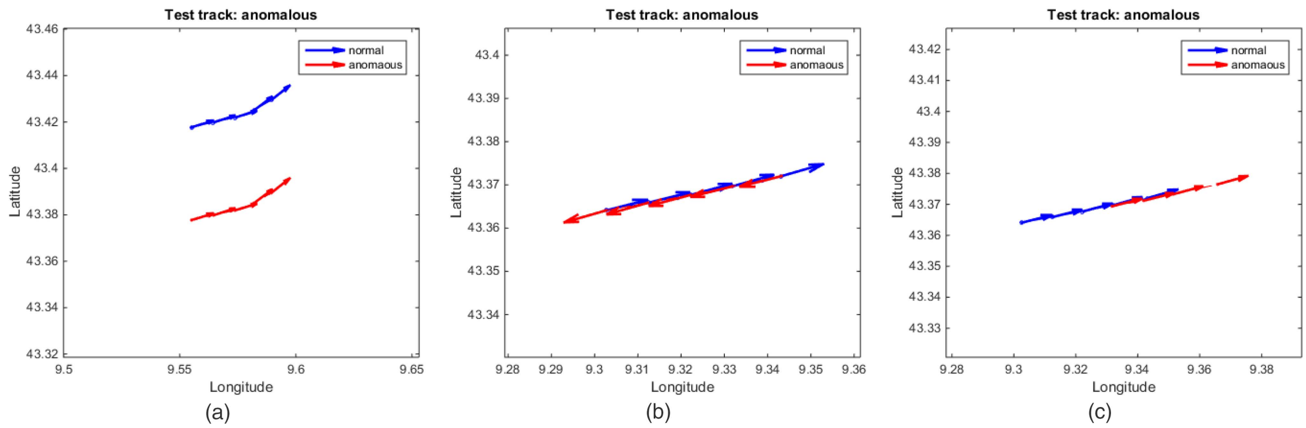

Fig. 4. Examples of simulated anomalies starting from the real data in [41]: the blue track is the normal track derived from the trajectories belonging to the subset of 8 routes, the red one is the synthetic anomalous track, reproducing a specific anomalous behaviour. (a) Positional anomaly: shifted track. (b) Directional anomaly: reverse flow track. (c) Kinematic anomaly: high speed track.

2) Dataset of anomalous tracklets:

The six URRTs are tested on a reference data set of AIS data developed at CMRE. The tracklet dataset consists of raw positional data collected for research purposes via the ground-based Automatic Identification System (AIS) receiver located in Castellana (La Spezia–Italy) owned by CMRE. The dataset contains the reports of the vessels equipped with AIS transponders, which were transiting over a section of the Northern Tyrrhenian Sea framing La Spezia harbour during the time period which goes from January 1st through February 20th 2013. The dataset contains both real tracklet data (labelled as "normal tracklets") and pseudo-synthetic tracklet data (labelled as "anomalous tracklets"). The original Castellana dataset [41] which has to be considered the source of the current dataset, is in the form of terrestrial AIS (T-AIS).

Two classes are considered: Class R_1 corresponds to normal trajectory segments and Class R_0 corresponds to anomalous trajectory segments. The normal trajectory segment of each evaluation trajectory is constructed by first selecting a random tracklet from the set of normal evaluation trajectories of a given length of 5 consecutive points: 95 tracklets are extracted from the system of pre-computed routes. Each route is decomposed into single-vessel trajectories and then further divided into tracklets of 5 consecutive points. The anomalous trajectory segment of each evaluation trajectory is constructed by first selecting a random tracklet from the set of normal evaluation trajectories of a given length of 5 points, replicating it and then altering its features. More specifically, a total of 275 anomalous tracklets were generated as follows:

- *Positional anomalies:* 80 Off-route tracks were created by shifting either the LONGITUDE or LATITUDE sequence (of a given magnitude);
- *Directional anomalies:* 108 high-speed tracks were created by increasing the initial instant speed of the track and by using a Near-Constant-Velocity Model to derive the new coordinates (LONGITUDE, LATITUDE), given the observed reported course;
- *Kinematic anomalies:* 87 opposite-flow tracks were created by changing the initial heading of the track and by using a Near-Constant-Velocity Model to derive the new coordinates (LONGITUDE, LATITUDE), given the observed reported speed SOG.

Figures 4 shows examples of the three types of simulated anomalous tracklets. As the traffic normalcy, we considered a subset of 8 routes as displayed in Figure 1.

3) Results and discussion:

We present here results of anomaly detection, thus considering two classes only, R_0 the class of anomalous tracklets containing three kinds of anomalies as described above and R_1, the class of normal tracklets belonging to the subset of 8 routes. Figure 5 displays the output quality results on a spider (radar) graph, with the three criteria of TRUENESS ($F1$-score), UNCERTAINTY (reverse entropy) and IMPRECISION (reverse non-specificity). The best method is the one covering the widest area in the graph. The ranges of the criteria are indicated in brackets. The TRUENESS criterion as measured by F_1 aggregates the TP and FN and hides thus the contribution of each corresponding type of errors. Table V expands the criterion of TRUENESS by displaying additional measures to the TPR, as the TNR, the F_1, F_2 and $F_{0.5}$ measures. While F_1 assigns equal weights to false negatives and true positives, F_2 gives more emphasis on false negatives and $F_{0.5}$ attenuates the influence of false negatives.

Through these criteria and associated performance measures, we observe that URRT#1 provides excellent results in terms of TRUENESS and PRECISION. That means that URRT#1 was able to correctly detect the anomalies and the on-route vessels. Moreover, before decision the set of compatible routes was minimum (a singleton). However, the entropy was quite high meaning some UNCERTAINTY before decision. The extension of

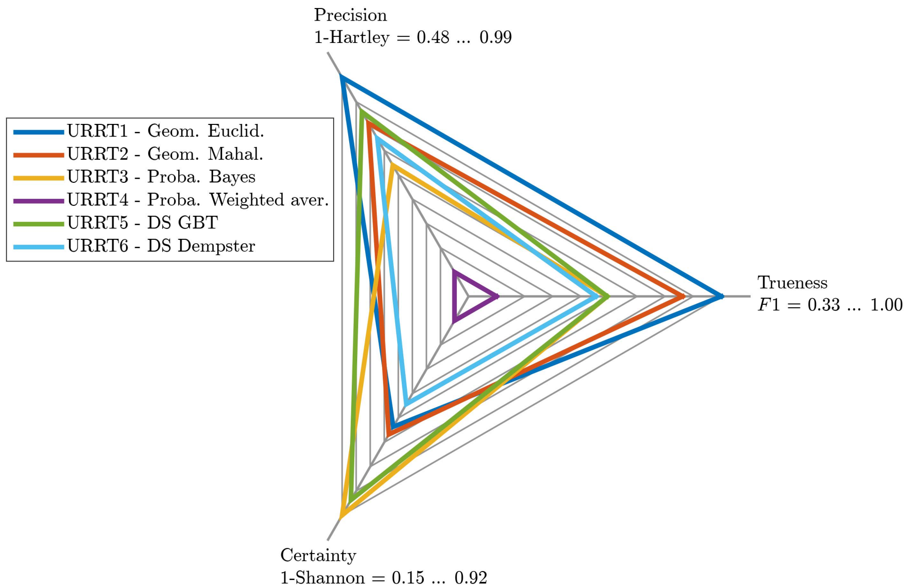

Fig. 5. Spider graph of three output quality criteria for the six URRTs.

TABLE V
Trueness measures for the six URRTs.

	TPR	TNR	F_1	F_2	$F_{0.5}$
URRT#1	1.00	1.00	1.00	1.00	1.00
URRT#2	1.00	0.91	0.88	0.95	0.83
URRT#3	1.00	0.64	0.66	0.83	0.55
URRT#4	0.21	0.97	0.33	0.25	0.48
URRT#5	1.00	0.64	0.66	0.83	0.55
URRT#6	0.77	0.76	0.62	0.70	0.56

URRT#1 to the Mahalanobis distance provides slightly lower results in terms of TRUENESS though, especially regarding the TNR (some anomalies have been missed) as we can see in Table V. Indeed, it appears that considering the dependency between the attributes in the observation space, although more correct than the naive independence assumption under URRT#1, leads to a slight decrease in the performances. In both URRT#1 and URRT#2, the uncertainty representation is based on the distance of the tracklet to the routes, computed by a Hausdorff distance. If the set of points of the tracklet belongs to the set of points of the routes, then the distance will be very low, or null.

The Bayesian approach URRT#3 and its evidential extension URRT#5 provide similar performance results. Compared to the pattern matching approaches, the TPR is still maximum while the TNR is only 60%. However the UNCERTAINTY is lower meaning that the decisions could be taken with a quite high confidence. However, combined with the low TNR, this is not a desirable behaviour as this apparent confidence of the algorithm may be miss-interpreted by the decision maker. These two approaches use the likelihoods extracted from the routes' statistics as a basis for uncertainty representation. No probability distribution estimation method was applied and the likelihoods were simply extracted from the histograms. The evidential approach based on the Generalised Bayes Theorem (URRT#5) uses plausibility functions instead of probabilistic likelihoods and allows by that to account for some IMPRECISION on the probability distributions. It is particularly interesting when the amount of data available does not guarantee a reliable estimation of the probability distribution. Indeed, as illustrated in Table II, some routes are built upon only a few trajectories and their uncertainty may be better represented by lower and upper bounds of unknown probability distributions (as provided by belief and plausibility measures respectively) or simply by crisp intervals.

The weighted average of probabilities (URRT#4) provides the worse results along the three criteria, while the TNR is actually better than most of the other approaches. From Table V, it appears that the bad performance of URRT#4 is mainly due to a very low TPR (around 20%). That means that on-route vessels are seldom detected and wrongly detected as anomalies instead. The disjunctive operator (+) averages the posterior probabilities and a very low probability along one feature (denoting an anomaly) would be diluted among other higher probabilities. It would thus be more difficult to detect the directional and kinematic anomalies. As mentioned previously, the disjunctive operator is a rather cautious fusion operator, more suited to a consensus. We should not however conclude that URRT#4 is not a good approach, as its strength is to be robust to errors and unreliable sources, something that was not reflected in our dataset.

The evidential approach using the conjunctive rule (URRT#6) provides mediocre TPR and TNR while this pair of values is actually better than all the approaches expected the pattern matching ones. The UNCERTAINTY and IMPRECISION are both quite high meaning that the decision was taken with still a high hesitation. URRT#6 is the only method which rejection criterion is based on a measure of conflict (here Dempster's conflict). The conflict is represented by the empty set between subsets of routes. The core of the reasoning relies thus on the intersection of the subsets of routes compatible with the features. In case this intersection is empty, no route is actually detected as compatible and the tracklet is classified as abnormal. The BBA was set to represent the uncertainty originating from the source's quality, which acts as a discounting over the categorical BBA of the set of compatible routes. However, in case the source expresses some (lack of) SELF-CONFIDENCE about its declaration, this singular uncertainty could be considered as well with this approach.

Finally, note that all the URRTs but the URRT#6 rely on *generic imperfection* only. URRT#6 is the only approach (again, as currently implemented) which accounts for the uncertainty expressed at the current instant in time t. All the other approaches rely on uncertainty, imprecision or graduality derived from past observations.

REMARK The results presented here should be read as an instantiation of the exploitation of the URREF mainly, as the application of such techniques to maritime anomaly detection requires deeper work. In particular, the synthetic anomaly generation technique may have a high impact on the results. The fact that the technique essentially shifts tracklets from their original position in the feature space may explain why the pattern matching approaches provide better results.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we dissected six (6) uncertainty and reasoning techniques (URRTs) to information fusion and proposed detailed description and comparison in their ability to handle uncertainty, in representation and fusion. We selected a variety of classical and simple schemes from (or adapted from) the literature which are all good candidates to solve the two problems of maritime route association and anomaly detection. We introduced the *uncertainty support* as an element conveying uncertainty, which allowed to make clearer which uncertainty is actually captured in the different reasoning schemes. We distinguished between uncertainty over individual variables or links between them, as well as second-order uncertainty. We framed our discussion within the Uncertainty Representation and Reasoning Evaluation Framework (URREF) and illustrated that considered jointly with the type of information either generic (from historical data or prior knowledge) or singular (at the time of the observation), the uncertainty support concept covers some elements of EXPRESSIVENESS of the URREF ontology (DEPENDENCY, HIGH-ORDER UNCERTAINTY, SELF-CONFIDENCE) and could expand to other criteria such as RELIABILITY.

The implementation of the URRTs to detect anomalies of a real AIS dataset allowed us to illustrate that the expressiveness criterion should not be assessed in isolation and that it is the joint assessment of the various criteria that makes the URREF powerful. Indeed for instance, a lack of expressiveness about the dependency between variables may still provide a good overall accuracy of the algorithm through some natural balance process.

Rather than identifying a "winner" approach, the comparison between the URRTs presented herein aimed at highlighting the *differences* and possible *complementarity* in uncertainty representation and reasoning. The approaches have been kept simple for a clearer understanding and in future works we will build upon this thin characterisation of the basic techniques together with the quality of the data available, taking advantage of the diversity of the different approaches, to design an efficient algorithm with easily interpretable results for detecting the anomalies at sea.

REFERENCES

[1] B. Auslander, K. M. Gupta, and D. W. Aha
A comparative evaluation of anomaly detection algorithms for maritime video surveillance.
In *Proc. SPIE, Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense X*, volume 8019, Orlando, Florida, USA, May 2011.

[2] A. Benavoli and B. Ristic
Classification with imprecise likelihoods: A comparison of TBM, random set and imprecise probability approach.
In *Proc. of the 14th Int. Conf. on Information Fusion*, 2011.

[3] T. Bilgiç and I. B. Türkşen
Measurement of membership functions: Theoretical and empirical work.
Handbook of Fuzzy Sets and Systems, 1 Fundamentals of Fuzzy Sets, 1995.

[4] E. Blasch, E. Dorion, P. Valin, and E. Bossé
Ontology alignment using relative entropy for semantic uncertainty analysis.
In *Proceedings of the IEEE National Aerospace and Electronics Conference (NAECON)*, pages 140–148, Fairborn, OH, 14–16 July 2010.

[5] E. Blasch, A. Josang, J. Dezert, P. C. G. Costa, K. Laskey, and A.-L. Jousselme
URREF self-confidence in information fusion trust.
In *Proc. of the International conference on Information Fusion*, Salamanca, Spain, July 2014.

[6] I. Bloch
Information combination operators for data fusion: A comparative review with classification.
IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans, 26(1):52–67, 1996.

[7] S. Challa and D. Koks
Bayesian and Dempster-Shafer fusion.
Sādhanā, 29:145–176, April 2004.

[8] P. C. G. Costa, K. Laskey, E. Blasch, and A.-L. Jousselme
Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology.
In *Proc. of the 15th Int. Conf. on Information Fusion*, Singapore, 2012.

[9] J. P. de Villiers, A.-L. Jousselme, A. de Waal, G. Pavlin, K. Laskey, E. Blasch, and P. Costa
Uncertainty evaluation of data and information fusion within the context of the decision loop.
In *Proc. of the 19th Int. Conf. on Information Fusion*, Heidelberg, Germany, July 2016.

[10] P. de Villiers, G. Pavlin, P. Costa, K. Laskey, and A.-L. Jousselme
A URREF interpretation of Bayesian network information fusion.
In *Proc. of the International Conference on Information Fusion*, Salamanca, Spain, July 2014.

[11] G. K. D. de Vries and M. van Someren
Machine learning for vessel trajectories using compression, alignments and domain knowledge.
Expert Systems with Applications, 39(18):13426–13439, Dec. 2012.

[12] A. Dempster
Upper and lower probabilities induced by multivalued mapping.
The Annals of Mathematical Statistics, 38:325–339, April 1967.

[13] T. Denœux
Introduction to belief functions.
Fourth School on Belief Functions and their Applications, 2017. URL https://www.hds.utc.fr/ tdenoeux/dokuwiki/_media/en/lecture1.pdf.

[14] T. Denoeux and P. Smets
Classification using belief functions: Relationship between case-based and model-based approaches.
IEEE Transactions on Systems, Man and Cybernetics—Part B: Cybernetics, 36(6):1395–1406, 2006.

[15] P. Diaconis and S. Zabell
Some alternative to Bayes's rule.
In B. Grofman and G. Owen, editors, *Proc. of the Second University of California, Irvine, Conference on Political Economy*, pages 25–38, 1986.

[16] D. Dubois and H. Prade
Evidence, knowledge, and belief functions.
Int. Journal of Approximate Reasoning, 6:295–319, 1992.

[17] D. Dubois and H. Prade
Fuzzy sets and probability: Misunderstandings, bridges and gaps.
In *Proc. of the 2nd IEEE Int. Conf. on Fuzzy Systems*, pages 1059–1068, 1993.

[18] D. Dubois and H. Prade
The three semantics of fuzzy sets.
Fuzzy Sets and Systems, 90:141–150, 1997.

[19] D. Dubois and H. Prade
Formal representations of uncertainty,
volume Decision-making—Concepts and Methods, chapter 3, pages 85–156. ISTE, London, UK & Wiley, Hoboken, N.J. USA, 2009. Invited paper.

[20] J. H. Friedman
On bias, variance, 0/1-loss, and the curse-of-dimensionality.
Data Mining and Knowledge Discovery, 1:55–77, 1997.

[21] K. M. Gupta, D. W. Aha, and P. Moore
Case-based collective inference for maritime object classification.
In *Proceedings of the Eighth International Conference on Case-Based Reasoning*, pages 443–449, Seattle, WA, 2009. Springer.

[22] ISO 5725
Accuracy (trueness and precision) of measurement methods and results—part 1: Introduction and basic principles.
Technical report, ISO International Standardization, 2011. Available at standardsproposals.bsigroup.com/Home/getPDF/830.

[23] F. Johansson and G. Falkman
Detection of vessel anomalies—a Bayesian network approach.
In *Proc. of the 3rd International Conference on Intelligent Sensors, Sensor Networks and Information*, pages 395–400, Melbourne, Qld., Dec. 2007.

[24] A.-L. Jousselme and G. Pallotta
Dissecting uncertainty-based fusion techniques for maritime anomaly detection.
In *Proc. of the 18th Int. Conf. on Information Fusion*, Washington, D.C., USA, July 2015.

[25] A.-L. Joussselme and P. Maupin
A brief survey of comparative elements for *uncertainty calculi* and decision procedures assessment.
In *Proc. of the 15th Int. conf. on Information Fusion*, 2012. Panel *Uncertainty Evaluation: Current Status and Major Challenges*.

[26] A. Karlsson, R. Johansson, and S. F. Andler
Characterization and empirical evaluation of Bayesian and credal combination operators.
Journal of Advances in Information Fusion, 2011.

[27] J. Kittler
Combining classifiers: A theoretical framework.
Pattern Analysis and Applications, 1(18–27), 1998.

[28] G. J. Klir
Generalized information theory.
Fuzzy Sets and Systems, 40:127–142, 1991.

[29] G. J. Klir
Probability-possibility transformations: A comparison.
International Journal of General Systems, 21:291–310, 1992.

[30] G. J. Klir and B. Yuan
Fuzzy Sets and Fuzzy Logic: Theory and Applications.
Prentice Hall International, Upper Saddle River, NJ, 1995.

[31] R. O. Lane, D. A. Nevell, S. D. Hayward, and T. W. Beaney
Maritime anomaly detection and threat assessment.
In *Proc. of the 13th Int. Conference on Information Fusion*, Edinburgh, UK, 2010.

[32] R. Laxhammar
Anomaly detection for sea surveillance.
In *Proc. of the Int. Conference on Information Fusion*, Firenze, Italy, July 2008.

[33] R. Laxhammar
Anomaly Detection in Trajectory Data for Surveillance Applications.
PhD thesis, School of Science and Technology at Örebro University, 2011.

[34] H. Leung and J. Wu
Bayesian and Dempster-Shafer target identification for radar surveillance.
IEEE Transactions on Aerospace and Electronic Systems, 36(2):432–447, 2000.

[35] B. Liu, E. N. de Souza, S. Matwin, and M. Sydow
Knowledge-based clustering of ship trajectories using density-based approach.
In *proc. of the IEEE International Conference on Big Data (Big Data)*, pages 603–608, Oct 2014. doi: 10.1109/BigData.2014.7004281.

[36] F. Mazzarella, M. Vespe, and C. Santamaria
SAR ship detection and self-reporting data fusion based on traffic knowledge.
12:1685–1689, 2015.

[37] D. Mercier, B. Quost, and T. Denœux
Refined modeling of sensor reliability in the belief function framework using contextual discounting.
Information Fusion, 9:246–258, 2008.

[38] L. Millefiori, P. Braca, K. Bryan, and P. Willett
Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction.
IEEE Transactions on Aerospace and Electronic Systems, 52(5):2313–2330, October 2016.

[39] M. Musen
The Protégé project: A look back and a look forward.
AI Matters. Association of Computing Machinery Specific Interest Group in Artificial Intelligence, 1(4):4–12, June 2015. URL DOI: 10.1145/2557001.25757003.

[40] G. Pallotta and A.-L. Jousselme
Data-driven detection and context-based classification of maritime anomalies.
In *Proceedings of the 18th International Conference on Information Fusion*, Washington, D. C. (USA), July 2015.

[41] G. Pallotta and M. Vespe
Vessel traffic dataset from ground-based AIS receiver Castellana: description and use.
Technical Report CMRE-DA-2014-001, NATO STO CMRE, 2014. NATO UNCLASSIFIED.

[42] G. Pallotta, M. Vespe, and K. Bryan
Vessel pattern knowledge discovery from AIS data—A framework for anomaly detection and route prediction.
Entropy, 5(6):2218–2245, 2013.

[43] G. Pavlin, A.-L. Jousselme, P. J. de Villiers, P. C. G. Costa, and P. de Oude
Towards the Rational Development and Evaluation of Complex Fusion Systems: A URREF-Driven Approach.
In *Proc. of the 21st Int. Conf. on Information Fusion*, pages 679–687, Cambridge, UK, 2018.

[44] C. Ray, R. Gallen, C. Iphar, A. Napoli, and A. Bouju
DeAIS project: Detection of AIS spoofing and resulting risks.
In *OCEANS 2015—Genova*, pages 1–6, May 2015. doi: 10.1109/OCEANS-Genova.2015.7271729.

[45] B. Ristic and P. Smets
Target classification approach based on the belief function theory.
IEEE Transactions on Aerospace and Electronic Systems, 41(2):574–583, 2005.

[46] B. Ristic, B. La Scala, M. Morelande, and N. Gordon
Statistical analysis of motion patterns in AIS Data: Anomaly detection and motion prediction.
In *Proc. of the 11th Conference on Information Fusion (FUSION)*, Cologne, Germany, 2008.

[47] G. Shafer
A Mathematical Theory of Evidence.
Princeton University Press, 1976.

[48] G. Shafer and A. Tversky
Languages and designs for probability judgment.
Cognitive Science, 9:309–339, 1985.

[49] P. Smets
Constructing the pignistic probability function in a context of uncertainty.
Uncertainty in Artificial Intelligence, 5:29–39, 1990. Elsevier Science Publishers.

[50] P. Smets
Belief functions: The disjunctive rule of combination and the generalized Bayesian theorem.
International Journal of Approximate Reasoning, 9:1–35, 1993.

[51] P. Smets
Imperfect information: Imprecision—uncertainty.
In A. Motro and P. Smets, editors, *Uncertainty Management in Information Systems. From Needs to Solutions*, pages 225–254. Kluwer Academic Publishers, 1997.

[52] P. Smets
Probability, possibility, belief: Which and where?
In P. Smets, editor, *Quantified Representation of Uncertainty & Imprecision*, volume 1, pages 1–24. Kluwer, Doordrecht, 1998.

[53] P. Smets
Analyzing the combination of conflicting belief functions.
Information Fusion, 8:387–412, 2007.

[54] P. Smets and R. Kennes
The transferable belief model.
Artificial Intelligence, 66:191–234, 1994.

[55] L. Sombé
Reasoning under Incomplete Information in Artificial Intelligence,
volume 5 of *Int. J. of Intelligent Systems (Special Issue)*. John Wiley and Sons, INC., 1990.

[56] P. Walley
Measures of uncertainty in expert systems.
Artificial Intelligence, 83:1–58, 1996.

[57] L. A. Zadeh
Fuzzy sets.
Information and Control, 8:338–353, 1965.

Anne-Laure Jousselme received her PhD degree from the Electrical Engineering Department of Laval University in Quebec City (Canada) and the Institut National Polytechnique de Grenoble (France) in 1997. Formely with Defense Research and Development Canada (DRDC), she is now with the NATO STO Centre for Maritime Research and Experimentation (CMRE) in La Spezia (Italy), where she conducts research activities on reasoning under uncertainty, high-level and hard & soft information fusion, information quality assessment and serious gaming applied to maritime situational awareness and anomaly detection. She is area editor of the International Journal of Approximate Reasoning and associate editor of the Perspectives on Information Fusion magazine. She is a member of the Boards of Directors of the International Society of Information Fusion (ISIF) where she serves as VP membership and of the Belief Functions and Applications Society (BFAS) where she serves as Secretary. She serves on program committees of the International Conference of Information Fusion (FUSION) and the International Conference on Belief Functions (BELIEF). She was Tutorial Chair of FUSION 2007 in Quebec City (CA), International Co-chair of FUSION 2015 in Washington and Technical Co-chair of FUSION 2019 in Ottawa (CA). She was general Chair of the Canadian Tracking and Fusion Conference (CTFG) in 2014 in Ottawa (CA) and Local Organizer of the International Conference of Scalable Uncertainty Management (SUM) in 2015 in Quebec City (CA).

Giuliana Pallotta is a data scientist at Lawrence Livermore National Laboratory (LLNL), Computational Engineering Division, where she is working on diverse projects relevant to national security (e.g., climate, global security, evolving networks). Prior to joining LLNL as a staff member, she was a Research Scientist in maritime big data analytics at NATO Science and Technology Organization CMRE during 2012–2015. She is a Fulbright Scholar, recipient of a 2013 IEEE JPLC Information Fusion Best Paper Award for her work on pattern recognition and anomaly detection from vessel motion data. In 2011 she was a post-doctoral Research Associate at the University of Washington, Seattle–US. She received her PhD in Applied Statistics in 2009. Her research interests are in the areas of Bayesian inference, Statistical modeling, Big Data, Machine Learning, Anomaly Detection, Situational Awareness and Information Fusion for surveillance and security applications.

Assessing uncertainty handling representations of HLIF systems with URREF

MARK LOCHER
PAULO COSTA

Researchers have extensively explored uncertainty issues in Low Level Information Fusion (DFIG L0/L1 process levels) systems, and predominately use probabilistic uncertainty representations. However, this prominence does not happen in High-Level Information Fusion (HLIF) systems. One reason for this discrepancy is that HLIF systems ingest a wider range of evidence, with its associated uncertainties, and execute a broader scope of inferential reasoning than LLIF systems. Researchers developed multiple techniques to address these uncertainties and reasoning needs, but it is not clear when and where in a specific fusion system a particular technique should be applied. ISIF established the Evaluation of Technologies for Uncertainty Reasoning Working Group (ETURWG) to provide some clarity on this issue. As a first step, the ETURWG created the Uncertainty Representation and Reasoning Evaluation Framework (URREF). The framework formally represents concepts and criteria needed to evaluate the uncertainty management capabilities of HLIF systems. It provides 26 criteria for evaluating the effectiveness and resource efficiency of a fusion system's uncertainty management capabilities. However, given the recency of the framework and the complexity of the issues it addresses, practitioners face difficulties in understanding where and how each criterion is applicable across a general fusion process environment, including a generic fusion system model. This paper's primary contribution is to address this gap by providing a discussion of the significant application factors and considerations regarding the usage of the framework, while providing examples of such usage in the process.

Manuscript received February 20, 2018; revised October 31, 2018; released for publication December 10, 2018.

Refereeing of this contribution was handled by Johan de Villiers

Authors' addresses: Department of Systems Engineering and Operations Research and the Center of Excellence in C4I & Cyber, George Mason University, Fairfax, VA, 22020 USA (E-mail: {mlocher, pcosta}@gmu.edu).

Both authors state that they have no conflicts of interest or financial interests affected by this work.

Available as an OWL file at http://eturwg.c4i.gmu.edu/files/ontologies/URREF.owl.

1557-6418/18/$17.00 © 2018 JAIF

1. INTRODUCTION

This paper describes the use of the Uncertainty Representation and Reasoning Evaluation Framework (URREF) in evaluating an information fusion system's ability to appropriately handle the various uncertainties that arise in the fusion process. Information fusion transforms information from different sources and different points in time into a unified representation that supports human or automated decision-making [8]. This decision-making focus demands that information fusion results are sound. Unfortunately, data sources used are often "inconclusive, ambiguous, incomplete, unreliable and dissonant" [59]. It is important to evaluate the different forms of uncertainty a fusion system has to deal with, where and how they occur, and the impact they have on the fusion processes and system outputs. The URREF provides a set of uncertainty definitions and evaluation criteria to support such an evaluation.

High Level Information Fusion (HLIF) is defined as the situation (L2) and impact (L3) levels of the Data Fusion and Information Group (DFIG) model [72], [5]. It is distinguished from L0/1, which is called Low Level Information Fusion (LLIF). LLIF has been widely explored and issues of uncertainty determination and propagation are extensively documented. It typically uses crisp data from homogenous, credible sources. Classical probabilistic uncertainty representations with fixed probabilities, rather than belief functions or imprecise probabilities, predominate in LLIF [31]. HLIF involves more complex environments, reasoning about complex situations, with a diversity of entities and multiple relationships between those entities. HLIF uses more diverse information sources, with significant evidential vagueness or ambiguity, and incompleteness and inconsistencies between evidence items. The credibility of individual sources may vary significantly. The community has developed a range of techniques and models to address these issues, but there is no consensus on how to compare their effectiveness and system impacts.

The International Society for Information Fusion (ISIF) chartered the Evaluation of Technologies for Uncertainty Reasoning Working Group (ETURWG) to provide a forum to collectively address this common need in the ISIF community, coordinate with researchers in the area, and evaluate techniques for assessing, managing, and reducing uncertainty [25]. The group developed the Uncertainty Reasoning and Representation Evaluation Framework (URREF) as a first step towards sound evaluation of uncertainty representations in HLIF systems. First documented in [13], the current version and associated documentation can be found at the ETURWG website.[1] These criteria focus on evaluating the effectiveness and resource efficiencies of the uncertainty representation(s) within a fusion system. The ETURWG does not expect URREF to identify a "silver bullet" technique that will adequately address all the

[1] Use of an ontology editor such as Protégé suggested.

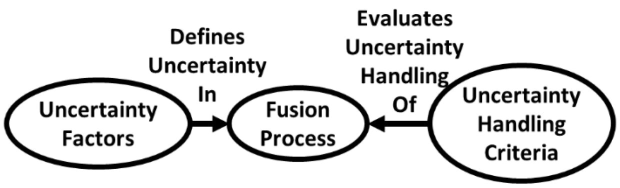

Fig. 1. URREF Top-Level Model

significant relevant uncertainties in a fusion system's environment but will assist designers in incorporating the appropriate range of techniques to meet their specific requirements.

This paper's primary contribution is to provide a discussion of the significant application factors and considerations regarding the usage of the framework, while providing examples of such usage in the process. Section 2 highlights the URREF and provides the criteria. Section 3 defines the key characteristics of the overall fusion process environment, including fusion system model, that affect uncertainty representation. Section 4 maps the URREF evaluation criteria to this environment and discusses how they are used to understand the uncertainty representation capabilities of a fusion system.

2. THE URREF

Figure 1 shows the URREF's top-level model. Uncertainty Factor provide a core description of the type, nature, derivation and models of the uncertainties that can be found in the fusion process. The Fusion Process includes the source, fusion system (in both a component and process view) and evidence/information.[2] These will be the subjects of an uncertainty handling evaluation. The Uncertainty Handling Criteria are measures useful for evaluating how well a specific fusion process handles its uncertainties. The ETURWG grounded the URREF on earlier work done by the W3C Incubator Group for Uncertainty Reasoning [47]. This work provides a basic framework of world/agent/sentence where an agent makes a statement about some aspect of the world using a logical sentence format. A logical sentence is a statement stated precisely enough that it can be assigned a truth value. This truth value may be binary, qualitative or numerical. The ETURWG identified three basic uncertainty characteristics: the nature, derivation and type of uncertainty, described in Table 1.

Although uncertainty has been understood qualitatively since the Greek philosophers of the early 5th Century BCE, an understanding of the different types of uncertainty began with the development of quantitative probability, addressing randomness, started by Fermat, Pascal and Huygens in the 17th century [3]. In 1921, Knight distinguished between problems with known probabilities (which he called risk) from those with unknown probabilities (called uncertainty—also a form of ignorance) [41]. The concepts of vagueness and ambiguity were given formal form by Black in 1937 [4]. Since that time, numerous taxonomies of uncertainty have been developed, both for general use and for specific fields. Jousselme et al. reviewed six taxonomies for potential application in fusion systems [35]. The two most comprehensive characterizations they identified were by Smithson [68] and by Krause and Clark [43]. Both use the classic randomness (probability)/vagueness/ambiguity classification. Smithson also included knowledge incompleteness and distortion as types of uncertainty. Distortion occurs when biases/inaccuracies in one's knowledge or when the knowledge transformation process introduces confusion in the knowledge [68]. Krause and Clark's taxonomy made two important distinctions. The first was between uncertainty induced by the classic sources and uncertainty induced by conflict. Second was the need to distinguish between uncertainty in a single information item and uncertainty in a set of information items. Conflict (also called inconsistency) most often occurs in an information set, although equivocation is identified as an internal conflict in a single item. Incompleteness is also primarily a characteristic of a set, although a single item may have missing information as well [43].

TABLE 1
URREF Uncertainty Factors

Uncertainty Nature	**Uncertainty is either inherent in the phenomenon expressed by the sentence or is result of lack of knowledge about that phenomenon.**
Aleatory	Uncertainty is inherent property of the world.
Epistemic	Uncertainty from lack of complete knowledge
Uncertainty Derivation	**Uncertainty derivation refers to the way it can be assessed. That is, how the uncertainty metrics can be derived.**
Objective	Assessed in a formal way, e.g., via a repeatable derivation process.
Subjective	Assessed via a subjective judgment. Even if one uses formal methods for this assessment, if the assessment involves subjective judgment, the Uncertainty Derivation is subjective.
Uncertainty Type	**Underlying characteristics of the information that make it uncertain.**
Ambiguity	Sentence has multiple possible interpretations
Vagueness	No precise correspondence between terms in the sentence and referents in the world
Randomness	The information comes from a process whose outcomes are non-deterministic.
Inconsistency	No world exists that satisfies the sentence.
Incompleteness	Occurs when information is missing.

[2]This paper will use the terms evidence and information interchangeably.

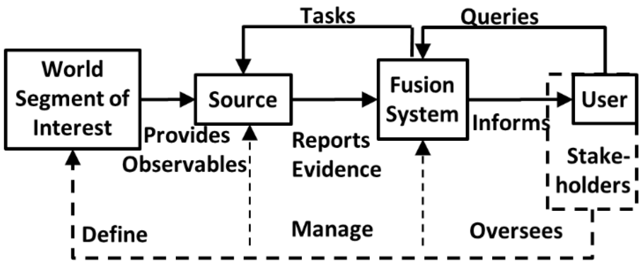

Fig. 2. Fusion Process Environment Model

Finally, the ETUWG identified the most common uncertainty representations (models):

- Belief functions[3]
- Fuzzy methods
- Probabilistic methods
- Random set
- Rough set

Additional choices can be found in Khaleghi et al. [40] or Castanedo [12]. The uncertainty handling criteria and their definitions are given in Table 2 below. The criteria are in four categories. Data criteria assess how a fusion process's design, including its uncertainty model(s), address aspects of uncertainty in data, both for individual items and for the collective set. Data Handling criteria focus on the effect of the uncertainty representation on explaining the reasoning used to create the output, and to maintain a record of what data was used in the process. Reasoning Criteria assess the overall approach to uncertainty handling in two areas:

- The correctness and consistency criteria assess the effects on the system outputs.
- The remaining criteria assess the effects on the overall system performance. These highlight the resource demands made by an uncertainty handling approach.

Representation Criteria assess internal characteristics of the uncertainty handling representation(s) and its integration with the fusion process.

It is an irony that the literature on uncertainty has a significant amount of ambiguity, redundant or overlapping terms, and conflicting definitions to describe aspects of uncertainty. In identifying these criteria, the ETURWG often had to select one term out of a range of choices for that aspect of uncertainty. In this paper, we generally do not attempt to identify synonymous terms or conflicting meanings.

3. FUSION PROCESS ENVIRONMENT

To apply the URREF criteria, one needs a model of the overall fusion process environment. We derived the model in Figure 2 from the DFIG model [5]. The main extension was to subsume the user in a larger group we call stakeholders, for reasons discussed below. This section describes each component, providing the context and key considerations for applying the URREF criteria.

3.1. Stakeholders/User

Any fusion system has a group of stakeholders, who collectively have an influence on the design and operation of the fusion system. The focus, scope and extent of a fusion system is driven by stakeholders' objectives, values and plans (collectively "stakeholders' interests"). A key subset of this group are the system users. These are the decision-makers, operators, and analysts who are the primary interactors with the system. Other stakeholders manage or influence aspects of the fusion process. For example, many fusion system users do not control the sources that provide evidence to their system. They submit information requests to one or more centralized management groups. Other stakeholders may require that the fusion process maintain records on how it created its outputs and the uncertainties associated with it. For example, the law of armed conflict requires a military commander to gather a reasonable amount of information to determine whether the target was a military objective and whether incidental damages to non-military targets are proportionate [48]. Uncertainties in the gathered information are a consideration in judging whether a commander acted properly. For such a system, the military legal community (as a stakeholder) may require that a fusion system be able to identify and trace the uncertainties in the evidence and how they were addressed in the fusion process to support a judgment of the legality of a commander's planned actions.

3.2. World Segment of Interest

The world segment is those aspects of a "real" world that stakeholders of the fusion system are interested in. Their points of view define the world segment. A world segment is defined as an area in the real world or cyber domain and possibly a time frame of interest (Figure 3).

The stakeholders' information needs define the world segments aspects of interest, including boundaries, key characteristics and entities of interest along with their attributes and relationships with other entities. In the same ocean area, a fusion system supporting a naval commander will focus on different entities than one supporting biologists studying marine mammals. The entities in the world segment generate observables, features detectable and reported by some source. Some entities may have a very limited set of observables, which may require a very specific approach to detect and collect the observable.

This is distilled into a world segment model using an ontological structure [9]. Entities should be categorized broadly, such as using Sowa's ontological categories. This allows for both concrete and abstract entities, with either time-stable (objects) or time varying (events) characteristics. It also allows for modeling

[3]Belief functions encompass approaches derived from Evidential Reasoning (Dempster-Shafer [62]). It includes Transferable Belief Model [66], Dezert-Smarandache Theory (DSmT) [18], and Subjective Logic [33].

TABLE 2
Uncertainty representation handling criteria

Data Criteria			Criteria on aspects of data and their relationship to source, reported object, and objectives of the fusion process.
	Credibility		Degree to which an evaluation subject can be believed or accepted as true, real, or honest
		Objectivity	Whether an evaluation subject reports in an unbiased manner
		Observational Sensitivity	Effect of a change in a result of evaluation subject and the corresponding change in a value of a quality being observed
		Self Confidence	Evaluation subject's assessment of its own credibility
	Quality		Assessment of the informational quality of the data
		Accuracy	Closeness of agreement between reported value and true value of reported entity
		Precision	Closeness of agreement between reported values obtained by replicate measurements on the same or similar evaluation subjects under specified conditions
		Veracity	Extent to which a source reports what it assesses to be the case
	Relevance To Problem		Degree to which information has direct bearing on the objectives of the fusion process
	Weight Of Evidence		Ability to assess the degree of impact of an evaluation subject on the result of fusion
Data Handling Criteria			Measure of the way data is managed by the fusion system
	Interpretation		Fusion system's ability to provide a coherent explanation that can be used to guide assessment, to understand the system's conclusions, and to provide a basis for reasoning and action
	Traceability		Ability to provide an accurate and unbroken historical record of its inputs and the chain of operations that led to its conclusions
Reasoning Criteria			How an information fusion system transforms its input data into knowledge. Focus on uncertainty model effects
	Computational Cost		Amount of system's computational resources required by a given representational technique to produce its results
	Consistency		Ability to produce the same results when provided with the same data under the same conditions
	Correctness		Ability to produce correct results, as measured against ground truth or an accepted gold standard
	Performance		Assess how well the fusion system and its representational model handle the functional requirements of an information fusion system
		Throughput	Measure the average and peak rate of conversion of inputs to outputs
		Timeliness	Ability to produce results within a required timeframe
	Scalability		Ability to handle an expanded work load
Representation Criteria			Encompasses criteria related to how uncertainty is characterized, captured and stored in a manner that can be processed by the fusion system
	Adaptability		Ability of the representational model to allow for different configurations of the model.
	Compatibility		Degree to which a given knowledge representation complies with data standards, and is related to the degree of flexibility it has in being coded with various standards
	Expressiveness		Ability to convey all relevant aspects of a given fusion problem
		Assessment	Ability to capture the types of uncertainty present in the evaluation subject
		Outcomes	Ability to represent appropriate scale for the outcomes
		Dependency	Ability to capture dependency among propositions (e.g., cause and effect, relevance, statistical association)
		Relational	Ability to represent uncertainty about domains with relational structure, i.e., domains in which there are types of objects with type-specific attributes and structure, having relationships to other types of objects
		Higher order uncertainty	Ability to capture uncertainty about the uncertainty model, including parameters, structure, and/or type of model
		Configurality	Ability to combine different types of uncertainty in multiple entities / relationships
	Knowledge Handling		Ability of a given uncertainty representation technique to convey knowledge
	Simplicity		User 's ability to execute common operations without requiring deep knowledge about its inner workings

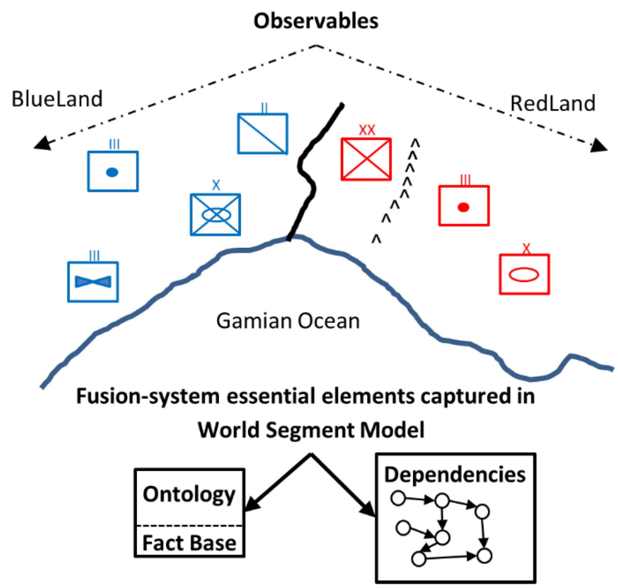

Fig. 3. World Segment of Interest defined by stakeholders/users needs and interest. Defines relevant observables, entities/attributes/relationships (captured in an ontology) and key dependencies that can infer new information

complex structures or situations, along with assigning attributes like purpose to them [70]. The world segment model generally becomes part of the fusion system, and mismatches between the world and model can result in significant errors and uncertainties. A key part of the world segment model is the dependencies. These are linkages between the attributes and relationships of entities, both within an entity and between entities. They have an "If A, then B" structure. The dependency between A and B is established from prior knowledge (include expert elicitation) or learned from collected evidence. The core of HLIF reasoning hinges on dependencies; when we have good reasons to believe A exists, then our understanding of B's existence, attributes or relationships change. Dependencies are expressed as rules, clauses (for logic programs) or graphical models (e.g. Bayesian networks, Markov networks).

A system may have multiple world segments within it (e.g. a global health epidemic system may be divided into regions or countries) or it may have multiple system copies, each with a different world segment. A system may also be deployable and load different world segments models as needed.

3.3. Source and Evidence

A source gathers observables and transforms them into evidence on some aspect of a world segment, through new observation or analysis of previously collected data (Figure 4). Source here means a specific mode of accessing data (e.g. panchromatic imagery, communications intercept, seismic detection, human reporting, database searches, etc.). When humans are part of the source process, at least some of the functions in Figure 4 are done mentally. Some source systems are multi-mode (e.g. radar with both Synthetic Aperture Radar and Surface Moving Target Indication modes) or multi-sensor (e.g. imaging and signals intercept on the same platform). Uncertainty should be assessed for each mode. A source may be dedicated to a specific fusion system or provide data to multiple fusion systems. A source may perform L0 fusion of observable samples (e.g. SAR change detection) using either internally generated data or integrating externally provided data. A source system may also conduct Level 1 fusion, using either self-generated or externally provided evidence. When data from those different sources are fused, the overall fusion process must be aware of this to avoid multiple counting of the same evidence.

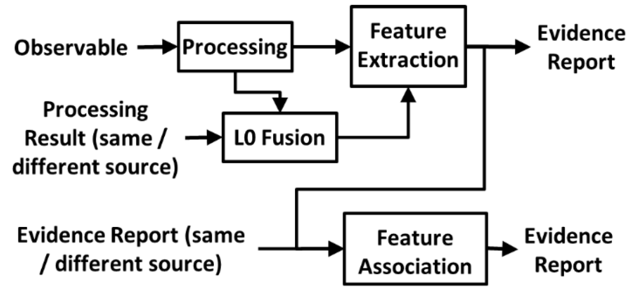

Fig. 4. Source process model

A common source differentiator is the hard/soft distinction, which aligns with the URREF Uncertainty Derivation criterion of objectively or subjectively derived evidence. Technical sensors are considered to provide hard or objective evidence, based on a repeatable derivation process. They generally provide consistent data with little possibility of source-generated untruthfulness, bias or deception. Evidence developed from human reporting is considered soft or subjective, with issues of source credibility, including deception; significant use of vague or ambiguous terms, or inconsistent application of terms between individual human sources [37]. The distinction is useful but benefits from being refined. Many sources have a machine/human partnership, where the extraction of useful information is done by humans. Imagery and communications intercepts sources are two examples. Such sources are generally classified as hard sources. In classifying a source as hard or soft, there are at least four considerations:

- Degree of calibration. Almost all technical sources undergo some type of calibration prior to employment, to ensure a level of accuracy and consistency. For some sources, human data exploiters undergo training to provide a level of consistency across different individuals. This consistency may not be tight as for a technical source.
- Use of source quality standards and reporting reviews prior to evidence release.
- Source recording. If the source maintains a record of the data that generated the evidence, it can be

TABLE 3
Classes of Evidence

Unequivocal testimony	Statement from a source (written, verbal)
Equivocal testimony	Hedged source statement ("I think I saw....")
Tangible	Evidence that may be physically examined: e.g. objects, documents, images, recordings
Missing evidence	Evidence one expects to find but does not.
Accepted facts	Statements whose truthfulness as evidence is not questioned (e.g. gold has a higher density than iron).

reviewed in cases where there are questions about the evidence.
- Source quality improvement efforts to identify and correct deficiencies, adjusting their accuracy and credibility over time.

Each source has its own characteristics that define how it gathers and processes its data. The source model describes, to some level of detail, how the source gathers and processes its data. An accurate model for each type of source is necessary for doing an uncertainty assessment on that source.

Sources generate evidence that is used in the fusion process. Evidence can be expressed using logical sentences with an uncertain truth value (which include "100% true" and "0% true"). Evidence can take a variety of forms. Table 3 provides a classification scheme [61]. Testimony is a statement made by a source. The statement may be based on direct observation, or on secondhand sourcing/hearsay. The statement may be either unequivocal ("It is the case that...") or equivocal ("I think that....", "I'm not positive, but..."). An equivocal assertion may include a reason for the equivocation ("It was dark, but I'm pretty sure I saw..."). Opinion is a form of equivocal testimony. It is defined as "A view or judgement formed about something, not necessarily based on fact or knowledge" or "A statement of advice by an expert on a professional matter." [53]. The key here is whether an opinion statement comes from a competent and knowledgeable source, able to support that statement. Expert judgment is a form of opinion that is a valid form of evidence. Missing evidence is not negative evidence, which is evidence that something does not exist at a point in time one is interested in. In some cases, missing evidence can be significant. For example, evidence intentionally destroyed can have a negative connotation for the destroyer.

Evidence may be at any level of the DFIG model, and it does not have to come from a process that moves sequentially through the levels. While sensor-derived data goes through L0 processing, human derived data often does not (although some may go through a form of preprocessing, such as summation or statistical processing). Evidence, especially from human or communications intercept sources, can also be about relationships between entities, situation or structure identification, or intentions (specific plans and objectives).

3.4. Fusion System

Understanding how to apply the URREF criteria to a HLIF process benefits from a generic system fusion model allows aligning the criteria with fusion system processes/components. After initially exploring the literature, we established these model requirements:

- Identifies key functions within a fusion process.
- Maps the flow between the functions, including feedback and reevaluation requests.
- Allows varying human/machine divisions of effort.
- Is not bound to a specific uncertainty representation or fusion methodology.
- Uses general domain-independent terminology.

According to Salerno, over 30 fusion process models had been proposed by 2002 [58]. Several teams have reviewed selected subsets, including Esteban et al. [24], Bedworth and O'Brien [2], Whitney, Posse and Lei [78] and Roy et al. [57]. Foo and Ng published an updated review in 2013 [26]. We found most of the models before 2005 very limited in their functional description. These included Pau's Model [55], Intelligence Cycle model [2], Thomopoulos' model [74], JDL model [30], Dasarathy model [16], Waterfall model [2], Extended OODA loop [63], Omnibus model [2] and the General Data Fusion Architecture [11]. Although they also had limited functional details, models that incorporated humans as part of the fusion process included the Visual Data-Fusion model [39], JDL level 5 [7], Endsley's situation awareness model [22], [23] and Lambert's Unified Data Fusion Model [44]. Four models published between 2002 and 2016 included significant details about their functions, shown in Figure 5. They were by Salerno [58], Steinberg [71], Lambert [45] and García, Snidaro, and Llinas [27]. There is a high degree of commonality in the functions described. All have some form of data ingest function that performs reference base alignment and semantic (ontological) registration. Some models explicitly depicted entity extraction from unstructured information sources (e.g. free text reporting). García et al.'s model was the only one to explicitly depict an uncertainty characterization process, while Steinberg's model discussed it in the text describing the model. Salerno's model explicitly depicted a number of information development activities to support the overall fusion process, including

- Data mining activities, including link analysis, pattern learning and pattern matching.
- Model development support, including pattern identification and model generation. Models may be built ahead of time, or created from the data stream.

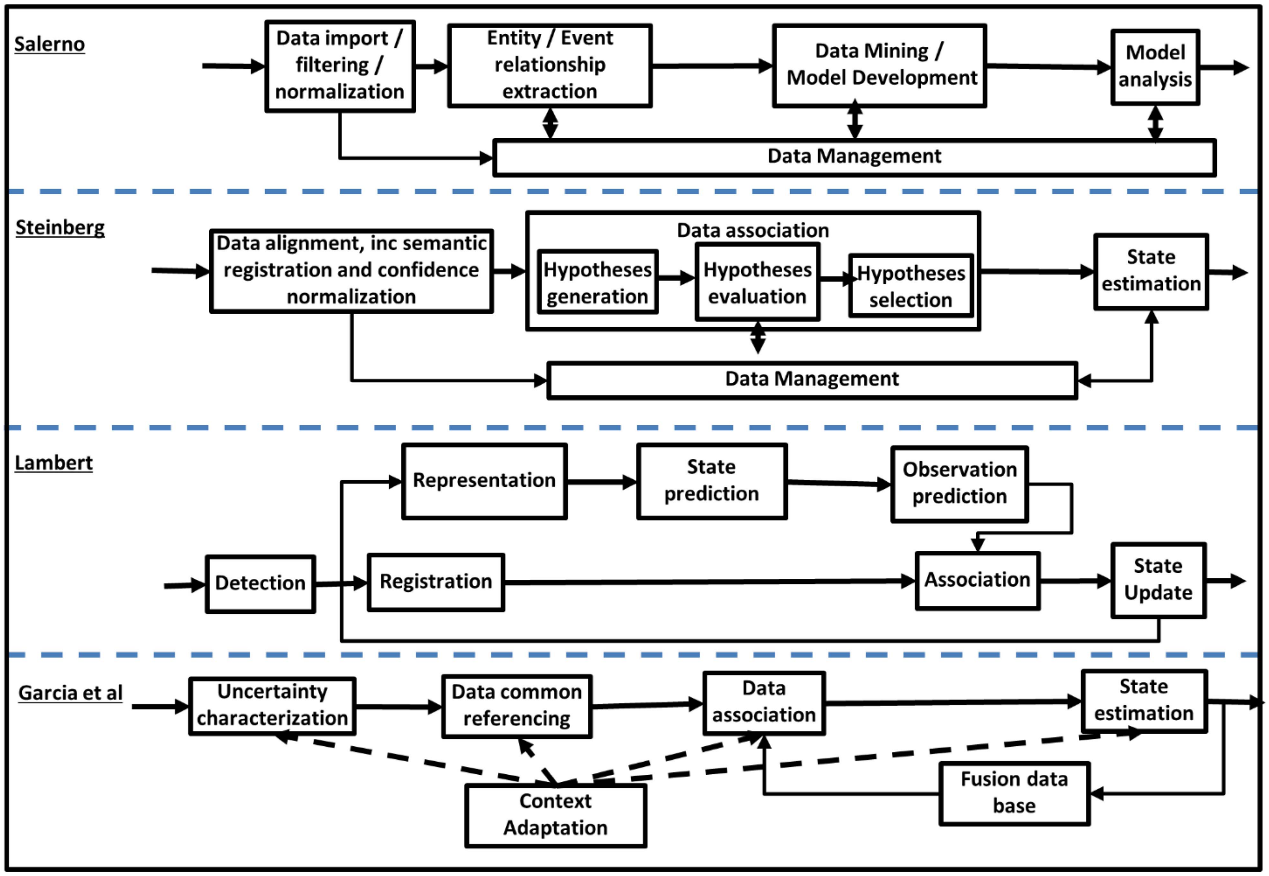

Fig. 5. Functional process/component elements of four major fusion system models

The other three models call out these functions as data association. For example, at level 2 HLIF, Steinberg model focused on finding and estimating relationships in the data, expressed as possible hypotheses. This is done by three subfunctions: hypotheses generation, hypotheses evaluation and hypotheses selection.

All four models had state estimation or state modeling. For HLIF, this process can use a variety of techniques, including link analysis, graph matching, templating methods, belief networks, compositional methods for model detection and development, and various algorithmic techniques [71].

Lambert's model differed from the others in using state transitions as a focusing element. This concept extends the idea of a Kalman filter to observing, predicting and updating state data, including tracking which scenario is being executed (L3 fusion) [45].

Because of differences between soft and hard sensors García et al.'s model have data from each type flow through a distinct path designed for the characteristics of that data [27]. They also explicitly include the use of context information. In the last five years, there has been significant work done on incorporating contextual information such as map data, weather, and procedural data (e.g. traffic rules, doctrinal concepts, patterns of life, hierarchies) for HLIF. Such non-sensor information can be used to both constrain and explain behaviors seen in sensor data [27] [67] [69].

To identify where to apply the various criteria, we merged these four models together to create the generic fusion system model shown in Figure 6. Based on our criteria, we realized that we needed to explicitly include several processes that one or more models discussed in their text but did not include in their visual model. The model assumes that input data may be L1, L2 or L3 data, including contextual data. The model has eight basic processes. Many source systems transmit free text reports, not structured text. Some form of entity and relationship extraction is required to transform those reports into machine-understandable data. The *Data Extraction/Alignment/Registration* process does this, including named entity recognition, coreference resolution, relationship extraction, and event extraction [56]. It also aligns the incoming data to a common reference base and ontological structure, appropriate for follow-on use. If the data is already structured according to an understood ontology, then this process is unnecessary.

For incoming evidence, *Source Uncertainty* manages all aspects of source uncertainty, as described in section 4.3. The *Data Store* captures all incoming evidence for access by the various processes. This includes both current and historical source evidence and reference information such as maps and equipment capability records.

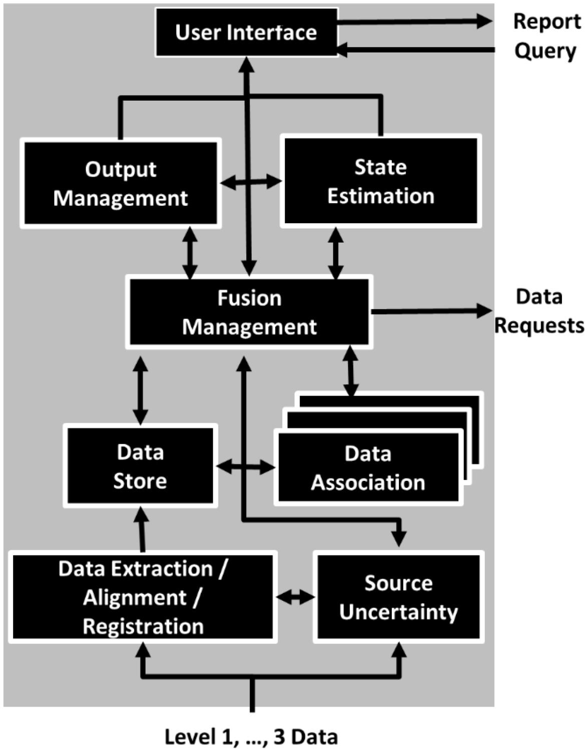

Fig. 6. HLIF fusion system model

An important aspect of this model is that not all the information is assumed to be in an immediately usable form for high level fusion processing in the State Estimation module. *Data Association* provides one or more services in which some or all of evidence, include context information, undergo to have the appropriate information extracted from them. For example, a fusion reasoning process may require relationship information. But 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. A key distinction between LLIF and HLIF is the significantly broader range of information in a HLIF, requiring a diverse set of data association processes to create that information [58], [24], [14], [71]. These processes can be implemented via middleware services [69].

Fusion Management involves all activities necessary to marshal information for the various fusion processing components and to sequence the fusion processes. This function can use multiple schemes to arrange the information to best provide insights into potential reasoning arguments and output hypotheses. It also identifies what additional information is needed to in the fusion process, and requests it [60].

The *State Estimation* process is the core of the fusion process. This process can take one or both of two forms. In less complex HLIF systems, 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 [71], [15], [46], although alternative approaches have been proposed using graphical belief models [1] and general-purpose graphical modeling using a variety of uncertainty techniques [64]. For more complex situation assessments, such as forensic reconstruction, the reasoning management process is a meta process, responsible for constructing the model used to provide the response. As such, there is a close interaction between reasoning management and output management.

The seventh process is *Output Management*. This process maintains the active hypotheses under consideration. It provides the output interpretation process (how did system arrive at this conclusion) and the traceability function (what evidence and functions did it use). It also is involved in generating hypotheses and in the pruning of hypotheses [32], [49].

The final process is the *User Interface*, which provides the information output and accepts user queries.

4. UNCERTAINTY ASSESSMENT

This section describes where and how the URREF criteria in Section 2 are applied to the process described in Section 3. The focus is on HLIF systems, but the criteria can also be applied to Level 0/1 systems as well. They do not cover the fusion management process levels (L4/5/6). These criteria guide fusion system developers and assessors through a comprehensive assessment of how well their uncertainty representations addresses the uncertainties both embedded in the evidence and generated by the fusion system's processes. Of the 26 criteria, thirteen can be specified as quantitative uncertainty measures, while the other thirteen are qualitative measures.

4.1. Stakeholder/User Uncertainty Tolerance Assessment

Identifying the stakeholders' concerns should drive the overall system uncertainty assessment. The first need is to understand their sensitivities to different kinds of uncertainties in the system. This focuses the main areas of evaluation, including the relative importance of different types of uncertainty. A second consideration is the uncertainty—system effects trade-off of addressing the various uncertainties via different uncertainty handling representations. Collectively, this information will focus and scope the uncertainty handling assessment.

4.2. World Segment Uncertainty Assessment

A fusion system uncertainty evaluation assesses the world segment to understand two important items:

- The uncertainties inherent in the observables.
- The uncertainties in the world segment model.

Uncertainties exist as variability in the world segment's observables and can propagate to the accuracy and precision of the collecting source. One needs to know the types and nature of these uncertainties. For any fusion system assessment, one must assume the world segment has a factual state. It is possible that the ground truth of that state may never be completely known, but it must be estimable well-enough to conduct meaningful assessments on the overall performance of a fusion system. The key component here is the world segment model. This model is a central part of fusion system, used both in data association and state estimation. Any model is an abstraction of a reality, and the fit with reality is imperfect. The key question is whether the fit is good enough. This is part of an overall assessment of the suitability and acceptability of a fusion system. For the uncertainty assessment, the primary question is whether the world segment model incorporates the key uncertainties inherent in the world segment. These will propagate through the source and into the fusion model, affecting both the correctness of the output and the demands placed on the fusion system's resources to address those uncertainties [17].

Second, epistemic uncertainties exist in world segment model and affect both the fusion system's output's correctness and consistency criteria and the data input's relevance criterion. In addition, limits on the expressiveness of the world segment model can induce uncertainty. The three characteristics are dependency uncertainty, higher order uncertainty and relational uncertainty. Dependency uncertainty occurs when there is significant doubt about the existence of or strength of the dependency between two or more world segment elements. This is a problem encountered during the model building effort. While the exact degree of dependency is often uncertain, the issue here is when is the uncertainty significant enough to affect the outcome (often detected by a sensitivity analysis). This leads to epistemic uncertainty because one does not know whether the model should include the dependency, or what strength value should be assigned to dependencies that are possible but not required (e.g. a probabilistic dependency). Higher order uncertainty is when one has significant doubt about the quantification values assigned in the model. All uncertainty representations require some form of quantification (e.g. basic probability assignments, membership functions). It is very possible to have uncertainties about the specific quantification scheme. This also leads to an epistemic uncertainty about the outcomes. Relational uncertainties occur in world segment models that allow for a varying number of entities and relationships. If so, then sources may make mistakes in assigning observables to entities. The evidence, including extracted information, will then have relational uncertainties. This can also occur in the fusion system when associating multiple evidence from different sources, or from the same source at different time periods. These are also a significant form of epistemic uncertainty in HLIF systems. There are five types of relational uncertainty:

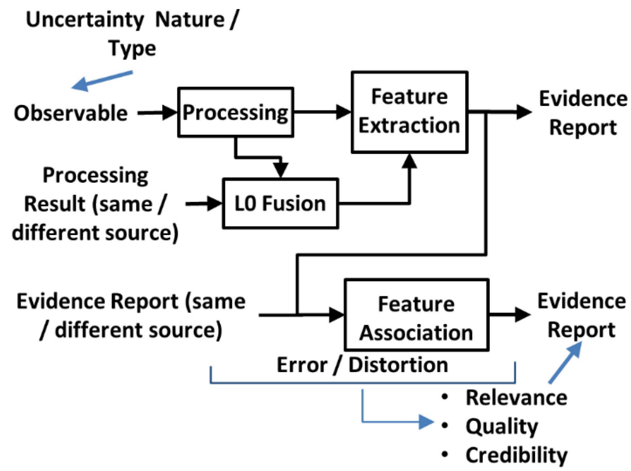

Fig. 7. Source errors and distortions combine with the uncertainty in the observables to create relevance, quality and credibility uncertainty

- Existence uncertainty for a key relationship or entity [28].
- Reference uncertainty is a dependency between two entities, but which specific entity has the dependency is uncertain (from a choice of several possible entities) [28].
- Type uncertainty is when one has determined the existence of an entity, but its reference class is uncertain [42].
- Identity uncertainty occurs when one is not certain if an entity is a new instance or one that has been previously identified [54].
- Number uncertainty occurs when the number of possible entities varies in a specific situation [52].

The primary effects of these uncertainties are seen when comparing the outputs of the fusion system to ground truth estimates in the world segment. This will be taken up in Section 4.4.2.

4.3. Source Uncertainty Assessment

Source uncertainty assessment focuses on the uncertainty in the evidence. The source ingests the variability, vagueness and ambiguity inherent in the observable. In the process, it often reduces the effects of variability, but can add uncertainty via process errors/distortions/limitations, especially for human-involved sources (Figure 7). For example, vagueness occurs when the source cannot apply a quantitative value to the observable. The discussion below follows Schum's classic work on evidence analysis and effects in probabilistic reasoning [59]. There are two basic questions when assessing uncertainties regarding evidence from a source:

- Is it relevant to the issues of interest to the fusion system's users?
- Is the evidence right?

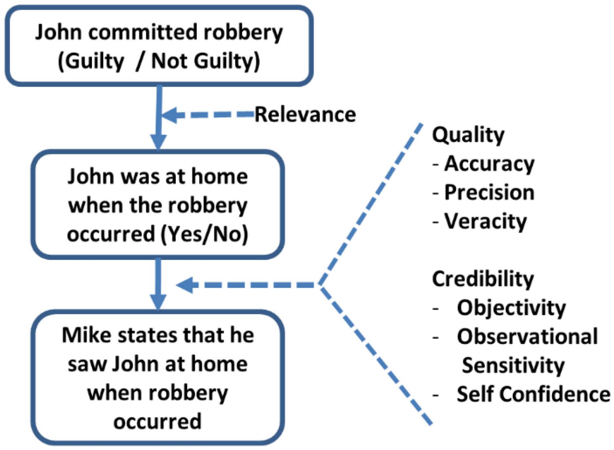

Fig. 8. Robbery Scenario

We use the example in Figure 8 to illustrate applying the criteria. John is accused of committing a robbery. If he did so, he would not have been at home when it occurred. If he did not do it, then he may or may not have been at home. This makes knowledge of John's whereabouts relevant to whether he committed the crime. A useful definition of relevance comes from the US Rules of Evidence [76]:

"Evidence is relevant if:

(a) it has any tendency to make a fact more or less probable than it would be without the evidence; and

(b) the fact is of consequence in determining the action."

Relevance measures the force of an item of evidence on some intermediate or final output of reasoning process. Probabilistically, relevance means that for a specific hypothesis H and any information E that could affect the belief in that hypothesis:

$$\text{Relevance} \stackrel{\text{def}}{=} P(H) <> P(H \mid E) \tag{1}$$

Relevance, as force of evidence, is always conditional on a particular hypothesis. It is not an inherent source characteristic. But we introduce it here because source uncertainty can modify the force of the evidence, sometimes in surprising ways. Relevance assumes a piece of evidence is true. There are several relevance measures in the literature [21]. The Bayes factor is one measure of the force of evidence:

$$\text{Relevance} = \frac{P(E \mid H)}{P(E \mid \bar{H})}, \quad (\bar{H}) \text{is the complement of } H \tag{2}$$

In Figure 8, we have a testimonial statement from Mike that he saw John at home at the robbery. Is his statement right? This is assessed by the Credibility and Quality criteria.

Credibility assesses the source's ability to understand the information in the observables. Although Credibility is most applicable to human sources, there are elements that may occur with technical sources. It

TABLE 4
Credibility measures

Credibility	
Objectivity	u_{o}(Source Understood State | Competence, Bias)
Observational Sensitivity	u_{os}(Source Understood State | Environment, sensor factors)
Self Confidence	u_{sc}(Source Understood State | Source Equivocation)

has three subcriteria: Objectivity, Observational Sensitivity, and Self Confidence. Table 4 provides mathematical measures for each, where u is a general uncertainty measure which assigns a value between 0 to 1. This measure represents common measures of uncertainty (probability, belief, fuzzy or possibility measures). These measures represent a dependency, where "|" is "Given" or "If", modeling "If B, then A." If u_x is a probability measure where A and B have discrete states, "|" becomes the conditioning operator, and u_x is measured via a conditional probability table on A and B's states (e.g. a confusion matrix). Observe that these measures focus on what the source understands from the observable, not what it reports. Objectivity has two elements: competence and bias. Competence addresses two areas. One, did the source have the access and ability to observe what the source reported? Ability in this case refers to the source's general capabilities. Two, in the case where the source is providing an opinion, does the source has the competence and data necessary to make the judgment expressed in the opinion. Incompetent sources cannot make an objective statement. Bias is any source characteristic that affects the source's ability to objectively understand the received data and influences them to ignore or misinterpret the data. Both human and technical bias are well-documented in the literature. Both can be hard to detect, especially if one is not looking for them. Bias can also be dependent on what is being reported on.

Observational sensitivity complements objectivity by noting when adjustments need to be made for situation-specific differences. For example, descending darkness near the time of the robbery could impair Mike's ability to correctly identify John. Technical sensors can also suffer from transient environmental effects that impair but not eliminate the ability to detect an observable. Self-confidence is the criterion that assesses equivocal evidence. This is evidence where the source specifically casts doubt on the accuracy of what it is reporting. Human sources may use vague or nonspecific phrases such as "It was getting dark, so I'm not sure…" or "I think it was him." Technical equivocation occurs when a source reports using abnormal sensor settings, system limitations or releasing below normal quality standards. Both human and technical equivocation affect the fusion system's understanding of the

TABLE 5
Quality measures

Quality	
Accuracy	u_{acc}(Reported State | Actual State)
Precision	$u_{\text{pre}} = \sum_{\text{trials}}$ Loss Function (Reported, Actual
Veracity	u_{ver}(Reported State | Source Understood State)

source's accuracy and requires an adjustment for that specific evidence.

Quality recognizes that even trustworthy sources makes mistakes. Quality has three subcriteria: Accuracy, Precision and Veracity. Quality measures are in Table 5. In these measures, the focus is on what the source reports.[4]

Accuracy assesses how close the reported information is to what is true in the world segment. It recognizes that no source is infallible. Whether technical or human, there is always the possibility that a source makes a mistake, with no intention to do so. Confusion matrices, Receiver Operating Characteristics, or Precision/Recall are all measures of accuracy. The Precision criterion complements Accuracy by assessing the degree of measurement variability between repeated observations of the same or similar entities under similar conditions. It is a measure of the consistency of the observation.[5] Precision is related to variability in the sensing environment, which can change a sensing measurement over time. A source with low precision will vary significantly more than a high precision source, decreasing the confidence one may have in the evidence. Veracity measures whether the source believes it is telling the truth (even if the evidence statement itself is not true). As such, Veracity is applicable to sources that have humans in a significant judgment role.

The evidential force of a source report as a standalone item depends on a function of its relevance, credibility and quality. The predominant understanding of credibility and quality is that they reduce (discount) the evidential force. But not always. Schum's explorations of the effects of veracity and credibility show that under some circumstances, knowledge about credibility and veracity factors can give more evidential force than the evidence contents themselves [59].

Figure 9 extends the model in Figure 8 to demonstrate this, giving two approaches to modeling veracity effects. In both, the prior probability of John's guilt is 10%. If John is guilty, he could not have been home at the time of the robbery. If not guilty, there is still a 70% chance he was not at home at the time of the robbery. Finally, if Mike is a truthful witness, his accuracy is 95%.

TABLE 6
Results of two different credibility models

Common data	**Guilty**	**Not Guilty**
Initial Belief (priors)	0.10	0.90
At Home—Yes	0	1
At Home—No	0.14	0.86
Single Thread		
Truthful—Source "Seen"	0.02	0.98
Truthful—Source "Not Seen"	0.13	0.87
Liar—Source "Seen"	0.09	0.91
Liar—Source "Not Seen"	0.11	0.89
Multi-Thread (Mike may know John's role in robbery)		
Truthful—Source "Seen"	0.02	0.98
Truthful—Source "Not Seen"	0.13	0.87
Liar—Source "Seen"	0.17	0.83
Liar—Source "Not Seen"	0.06	0.94

Figure 9A gives a classical discounting approach, using a single thread model. Here, the source Mike is a suspected liar, and the probability of his evidence being true in either case is assessed at 60%. In Figure 9B, one suspects that Mike has some knowledge about whether John committed the robbery, and that he is willing to lie to protect John if John is guilty. If he has some knowledge that John is not guilty, he will tell the truth about what he observed (he will not risk perjury in this case). If John did commit it, Mike has only a 60% chance of telling the truth (we are not certain he will lie). Because Mike's statement has a dependency on whether John is guilty or not, as well on whether John was at home, this is a multi-thread model. Table 6 gives the results of the two models. First, see the effect of knowing for certain whether John was at home or not. If he was, then he is not guilty. If he wasn't, then the probability that he is guilty increases from 10% to 14%.

Second, in both models, a truthful source has the same result: 2% guilty if Mike says he saw John at home, 13% if he says he did not. This is a dilution of the 0%/14% result of John's actual state and results from the 95%/5% accuracy distribution. Now look at the liar results. In the single thread model (Mike has 60% of telling the truth in any case), one sees a further dilution of the relevance. It stays closer to the 10% prior probability than the case where the source is credible. But there is a surprise in the multi-thread case. If Mike lies when he knows John committed the robbery and says that he saw John at home, the probability of being guilty climbs to 17%. This is opposite of what happens in the truthful case. This is because

[4] Which is why Veracity was classified as a quality criterion, not a credibility criterion

[5] The term Precision has at least three different uses in uncertainty discussions. The one given is the most common. Other uses include the proportion of true positives out of the total items classified as true in a confusion matrix (precision/recall), and the value of the least significant digit in a measurement.

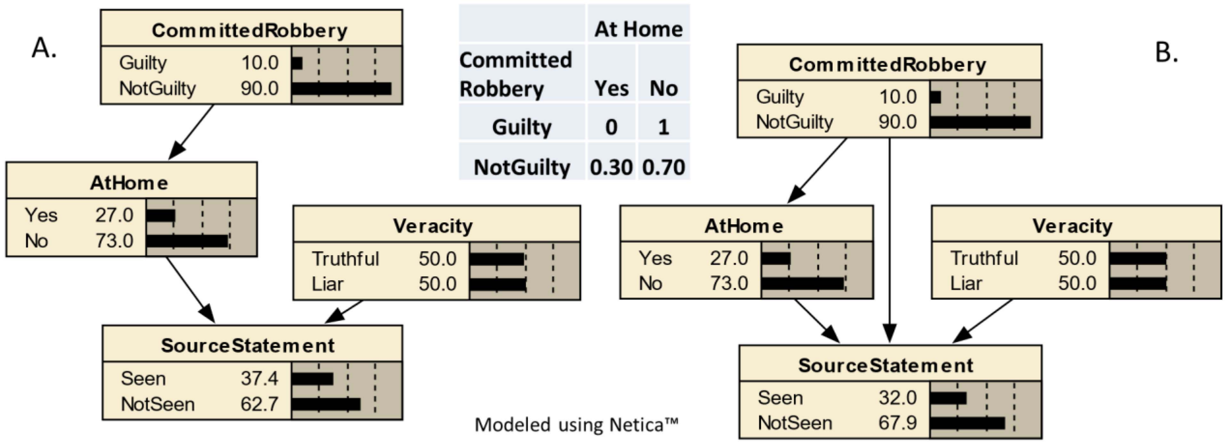

	At Home	
Committed Robbery	Yes	No
Guilty	0	1
NotGuilty	0.30	0.70

		Source Statement	
Veracity	At Home	Seen	Not Seen
Truthful	Yes	0.95	0.05
Truthful	No	0.05	0.95
Liar	Yes	0.60	0.40
Liar	No	0.40	0.60

			Source Statement	
Veracity	Committed Robbery	At Home	Seen	Not Seen
Truthful	Guilty	Yes	0.95	0.05
Truthful	Guilty	No	0.05	0.95
Truthful	Not Guilty	Yes	0.95	0.05
Truthful	Not Guilty	No	0.05	0.95
Liar	Guilty	Yes	0.60	0.40
Liar	Guilty	No	0.60	0.40
Liar	Not Guilty	Yes	0.95	0.05
Liar	Not Guilty	No	0.05	0.95

Fig. 9. Simplified robbery scenario with a suspected lying witness

if we think that someone who has knowledge about the ultimate hypothesis we are seeking will lie under certain circumstances, then telling the lie increases our probability of the ultimate hypothesis being true if the lie is told. There are many subtleties like this in doing source evidence assessments. See Schum [59] for an in-depth discussion on this issue.

The bottom line is that all source uncertainty assessment must determine to what these quality and credibility issues exist in their sources and select an uncertainty model and associated representation that address all the significant issues.

4.4. Fusion Model Uncertainty Assessment

Fusion model assessments focus on uncertainty representation in three areas: input evidence, output information, and the components of the fusion system. Figure 10 maps the URREF criteria in Section 2 to the fusion model in section 3.

4.4.1. Input Uncertainty Assessment Criteria

The input uncertainty assessment criteria can be divided into two categories: criteria applicable to individual evidence items and those for the collective set of evidence. For individual evidence items, the criteria are Credibility, Quality, and Assessment (an Expressiveness criterion). Credibility and Quality were discussed in the previous section. Assessment evaluates whether the fusion system can appropriately address the range of uncertainty types in the evidence. Uncertainty types identifies the basic uncertainty introduced by the world segment uncertainties and the specific characteristics of the source's process. In the fusion model, the characteristics of source evidence establish the uncertainty models needed for the source uncertainty, data association and state estimation modules. For individual evidence items, the source uncertainty module has the primary responsibility, since it establishes the credibility and veracity of each item.

In addition to uncertainty in the individual items of evidence, there is also uncertainty associated with the collective set of evidence. There are three criteria that apply: Assessment, Relevance, and Weight of Evidence. Assessment evaluates the fusion's system's ability to address the uncertainty types of incompleteness and inconsistency. Incompleteness is missing data, either partial (missing fields in a piece of evidence) or entirely. The most likely causes often are lack of source resources to obtain the evidence, observational problems in collection (e.g. cloud obscured image) and failure to request the evidence. When missing data is not available in time, the fusion process needs to be robust enough to provide its best estimate without the data, and able to identify what data was missing and its effects on the

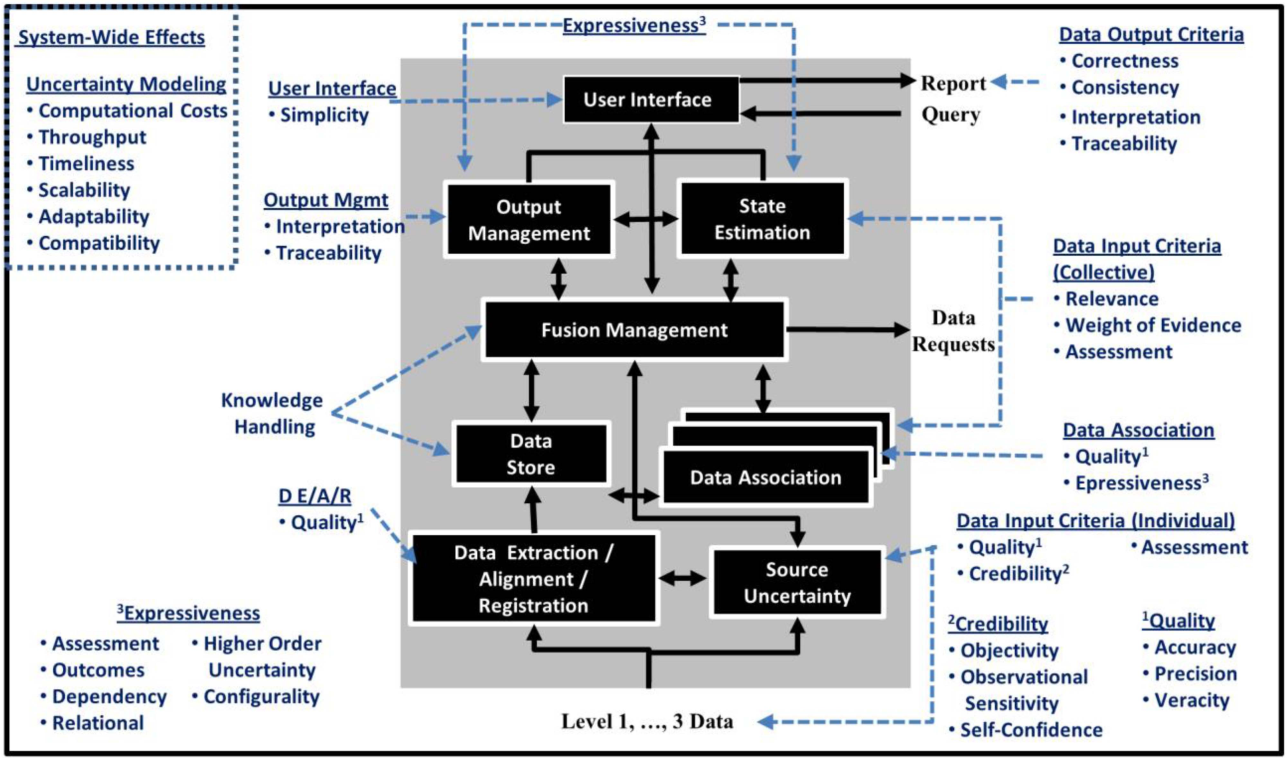

Fig. 10. Application of URREF evaluation criteria to different components of the generic fusion system model

output (see Traceability in section 4.4.2). Understanding how the system provides default values is important in these cases.

Inconsistency occurs when two or more inputs are contrary (they support different outputs in the fusion process) or contradictory. This is also called conflict and is a common issue in fusion systems. Conflict generally decreases the overall evidential force, as the conflicting items favor different outcomes. Conflict also increases uncertainty about source credibility and veracity, especially when one item of evidence favors an outcome significantly different that the remaining relevant evidence from different sources. Conflict has multiple causes, including non-source-initiated deception, source credibility/veracity issues, world segment model mismatches, and incomplete or uncertain model specification. Subject to available time, the desired approach is for the fusion process to alert the users to conflicts and allow them to conduct the necessary investigations to identify and resolve the root cause of the conflict. If resolution is not possible, then the system must be able to form a judgment based on the credibility of the evidence. Conflict can result in a significant amount of uncertainty that hinders decision making. Conflict modeling is usually addressed via probabilistic [59] or a belief function-type approach [62] [33] [36] [65].

Relevance, as an assessment of the force of an individual piece of true (from a credible and truth-telling source) evidence, is often dependent upon the related pieces of evidence. In many cases, evidence can be synergistic, either positively or negatively; its force is greater or lesser than its force when considered individually. Evidential relevance for additional like evidence tends to decrease if the multiple items provide limited or no additional new information. The synergy needs to be accounted for in the modeling.

Weight of Evidence (WOE) is an assessment of the totality of the available data and its effects on the output of the fusion. It is a holistic measure. It assesses the completeness of the evidence in supporting the fusion system output. It involves both the input evidence and the reasoning processes within the fusion system. There are multiple approaches to establishing the weight of evidence [77] [6]. Consider a physics analogy—weight is a function of the force of gravity and the mass of an object. Here, we will use effective force of evidence. This force results from the collective effects of Credibility and Quality on Relevance for each piece of evidence.

$$\text{WOE} = f(\text{Credibility, Quality, Relevance, Mass}) \quad (3)$$

The first three have already been discussed. Mass as used here is a measure of the comprehensiveness of the evidence—how many possible outputs are ruled out by the data. This makes Mass more than a simple count of how many items of evidence the system has. Rather, it focuses on the reasoning process in the fusion. A fusion process making a situation or impact assessment works as much by eliminating possible outputs as by supporting a specific output. Outputs that are neither positively or negatively supported remain as doubt in the system. WOE is also a useful tool in explaining the fusion system's output results.

4.4.2. Fusion Outputs Uncertainty Assessment Criteria

The Reasoning criteria of Correctness and Consistency are the core criteria for assessing uncertainty in the fusion system's output. How well the output mirrors the reality of the world segment it models is the primary measure of goodness of a fusion system. This makes output Correctness the central URREF criterion. However, this criterion is different than the Accuracy criterion for input information. Fusion system outputs normally come with an uncertainty hedge. Most often, this hedge is presented probabilistically—"There is a 90% chance this ship is the ship of interest." If a Correctness measures does not account for the probabilistic nature of the output, it will provide an incorrect view of the system's performance. Correctness can be assessed quantitatively using scoring rules [51], [29]. The original scoring rule is the Brier score. There are several versions; the most common applies to cases where the predicted outcome occurred or did not occur.

$$\mathrm{BS} = \frac{1}{N}\sum_{i=1}^{n}(f_i - a_i)^2 \tag{4}$$

Where N is the total number of outputs for which both a forecast probability (f_i) and an actual outcome (a_i) are available [10].

Closely following is the Consistency criterion. There are two considerations in this criterion:

- How repeatable are the results, when the same kind of evidence is provided?
- How sensitive is the output to minor changes in the input conditions?

Within the Brier score, there is a measure of the consistency of the forecasts. This assesses whether something predicted to be true 80% time actually occurs 80% of the time. It is also called reliability or calibration in the literature. It is

$$\frac{1}{N}\sum_{k=1}^{J} n_k(\mathbf{f}_k - \bar{\mathbf{o}}_k)^2 \tag{5}$$

Where N is the total number of outputs for which both a forecast probability (f_k) and an actual outcome are available, J is the number of forecast probabilities (assumed finite), n_k is the number of forecast probabilities in bin k, $\mathbf{f}_k$ is the forecast probability of bin k, and is the observed frequency of the outcomes predicted to occur in bin k. Both $\mathbf{f}_k$ and $\bar{\mathbf{o}}_k$ are vector quantities.

As with the Accuracy criterion, in those cases where there is no ground truth to establish a correct answer (including a simulated ground truth), the reasoning process can still be evaluated in terms of how its answers correspond to a gold standard (e.g. SMEs, documentation, etc.) [34].

In addition to providing the users with correct and consistent outputs, users benefit from understanding how and why the fusion system generated those outputs [74]. The data handling criteria of Interpretation and Traceability qualitatively assess this capability. Interpretation is the ability of a fusion system to support a coherent explanation of its conclusions. This is a summary explanation of the key evidence and reasoning process that supports the output. It is a justification for using the output in decision making. Interpretation can be assessed in at least two ways:

- Operationally via a user/stakeholder assessment that a representative range of output interpretations satisfy their information needs.
- Developmentally via fusion system experts' assessment that the interpretation captures the essential information input into or created by the system

Traceability is a diagnostic capability allowing users to follow the system's processes. It assesses the ability of a fusion system to provide an accurate and unbroken historical record of its inputs and the chain of operations that led to its conclusions. It is useful when the user wants an in-depth understanding of how the system came to its conclusions, or when the user suspects something is wrong or out of the ordinary in the output and its interpretation and wants to investigate further. Few fusion systems log intermediate results. But if the system records all inputs, including user requests, and the initial system states, and allows access to intermediate products during execution, system traceability can be conducted off-line. Traceability also applies to knowing exactly what evidence was used. Some sources occasionally find they need to retract evidence that turns out to be in error. Tracing what evidence items exist in one's data base supports this retraction process.

4.4.3. Effects of Fusion System Processing Uncertainty Assessment Criteria

In assessing the uncertainty representation within a fusion system, one must consider the overall ability of the system in reducing the total uncertainty on the reported outputs, the errors introduced by the fusion process, and the cost and fusion limitations imposed by the selected uncertainty representation approach. This is an area where significant work is required to fully understand where and how uncertainty is generated and propagated through the various fusion processes.

4.4.3.1. Uncertainty reduction and introduction of errors and uncertainty reduction

A fusion system is designed to reduce uncertainty by integrating the evidence, using one piece of evidence to reduce uncertainty in another. This requires (at least) conditional independence between the evidences. That is, the only dependencies between the evidence are mediated by the output whose uncertainty one wants to reduce. There are no other causes of correlation between the evidence items. With increasing efforts to increase the degree of L1/L2 fusion at source systems, such as the US Air Force's Distributed Common Ground System [75], It is important for fusion system designers to understand the possibility of multiple counting of common source evidence.

The fusion process can also introduce errors, which can increase the uncertainty in the output. Common errors are in the data extraction/alignment/registration process through incorrect classification/assignments, rounding, and misalignment [56] [20]. Information development processes can introduce errors through misassociation, misclassification, or unwarranted elimination of embedded uncertainty in the source evidence. The uncertainty representational scheme used plays a significant role in establishing the kinds of uncertainties that can be assessed in the information development process. For example, if the incoming data is heavily ambiguous, but the process has no mechanism for representing that ambiguity, the evidence output may be specified as being more definitive than the data warrants. Fusion reasoning elements need to account for possible accuracy, precision and veracity errors in extracted information [38]. For the reasoning processes, expressiveness of the chosen representations is an important consideration. These are:

- Assessment: Establishes what kinds of uncertainties can be addressed in the fusion system.
- Outcomes: Determines whether the outputs can incorporate the residuals of the types of uncertainties in the input data and created by the fusion process.
- Configurality: Determines the range over which a particular uncertainty representation needs to operate.
- Dependency: Determines whether the world segment model and source models incorporate all the dependencies necessary for the fusion model to correctly represent the uncertainties in the world segment and the sources.
- Higher order uncertainty: Determines if the uncertainty representation can include uncertainty about one's uncertainties. This is especially the case for uncertainty about probabilities that are used in reasoning models.

4.4.3.2. Effects on Fusion System Resources

In addition to assessing the range of needed uncertainty representation capabilities, there are a set of criteria to evaluate the effects of the uncertainty representation capabilities on the resource costs and range of capabilities for the fusion system. The first set of criteria identify the effects of different uncertainty representation approaches on the design of the fusion system. They are:

- Computational costs. Different representation schemes have varying demands on the fusion system's computational resources. Truth-functional approaches of possibilistic representations or probabilistic approaches that use canonical models [19] generally have the lowest cost, while random set approaches [50] have the highest. The computational cost will also depend on whether exact or approximate techniques are used, which have their own effects on output uncertainty.
- Performance (sub criteria—throughput and timeliness): Assesses the upper limit on system volume and velocity, determining if the selected uncertainty representation schemes significantly affect the ability of the system to meet the users' needs. These two subcriteria are intertwined with the computational cost criteria.
- Scalability: Effect of the representation to scale the model used. This is of especial interest when the world segment model allows for a significantly varying number of entities with different relationships between them.

The second set look at the constraints the uncertainty representation models place on the use of the system:

- Adaptability: Degree of change allowed to the configuration of the uncertainty representation, allowing it to model variations in the world segment or source models.
- Compatibility: how well the representation allows the use of common data standards within the domain within which the fusion system works (e.g. STANAGS for NATO systems, NIST IT standards for US systems, etc.).
- Knowledge handling: The effect of a particular uncertainty representation on the fusion system's information management capabilities.
- Simplicity: the degree of complexity of the user interface, especially with regards to the system's output explanation capabilities.

The assessment results on the effects on system resources should be incorporated into a larger system performance analysis. This enables a proper trade-off analysis between resource demand and uncertainty handling representation with the context of the overall system requirements.

5. CONCLUSION

This paper provided a broad examination of how the URREF uncertainty handling criteria can be applied to typical HLIF applications. We ground the discussion with a Fusion Process Environment Model to identify where the criteria should be applied. The application of the Framework's criteria to the evaluation of the uncertainties and their representations in a fusion system is shown from different perspectives. As noted, as an uncertainty evaluation framework, URREF must be seen in its current state as a first output of an effort to better understand the representation and effects of uncertainties in a HLIF system. As HLIF technologies advance, understanding and correctly addressing uncertainties will play an important part. Based on the points raised in this paper, we forecast two major directions for this effort in the future. First, comprehensive quantitative and qualitative comparisons among different representation approaches are important to better understand the appropriate applicability of each approach and guide HLIF developers in their design decisions. As probabilistic, possibilistic, and evidential approaches evolve, they gain new capabilities and provide new insights that can be shared across approaches. Second, a deeper understanding of real-world fusion processes is required to select and apply the most appropriate fusion models and systems for each specific situation.

REFERENCES

[1] R. G. Almond
Graphical Belief Modeling
CRC Press, 1995.

[2] M. Bedworth and J. O'Brien
The Omnibus model: a new model of data fusion?
IEEE Aerospace and Electronic Systems Magazine, 15 (2000), 30–36.

[3] P. L. Bernstein
Against the Gods: The Remarkable Story of Risk.
New York: John Wiley & Sons, 1985.

[4] M. Black
Vagueness: An Exercise in Logical Analysis.
Philosophy of Science, 4 (1937), 427–455.

[5] E. Blasch, I. Kadar, J. Salerno, M. M. Kokar, S. Das, G. M. Powell, D. D. Corkill, and E. H. Ruspini
Issues and challenges of knowledge representation and reasoning methods in situation assessment (Level 2 Fusion).
in Proceedings: Signal Processing, Sensor Fusion, and Target Recognition XV, 6235 (2006).

[6] E. Blasch, K. B. Laskey, A. L. Jousselme, V. Dragos, P. C. Costa, and J. Dezert
URREF reliability versus credibility in information fusion (STANAG 2511).
in Proceedings of the 16th International Conference On Information Fusion, 2013, 1600–1607.

[7] E. P. Blasch and S. Plano
Level 5: user refinement to aid the fusion process.
in Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications. International Society for Optics and Photonics, 2003, 288–298.

[8] H. Boström, S. F. Andler, M. Brohede, R. Johansson, A. Karlsson, J. van Laere, L. Niklasson, M. Nilsson, A, Persson, and T. Ziemke
On the Definition of Information Fusion as a Field of Research.
Report No. HS-IKI-TR-07-006, University of Skovde, School of Humanities and Informatics, Skovde, Sweden, 2007.

[9] R. J. Brachman, and H. J. Levesque
Knowledge Representation and Reasoning.
Amsterdam: Morgan Kaufmann, 2004.

[10] G. Brier
Verification of forecasts expressed in terms of probability.
Monthly Weather Review, 28 (1950), 1–3.

[11] H. Carvalho, W. Heinzelman, A. Murphy, and C. Coelho
A general data fusion architecture.
in Proceedings of the 6th International Conference on Information Fusion, 2003, 1465–1472.

[12] F. Castanedo
A review of data fusion techniques.
The Scientific World Journal, 2013.

[13] 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 15th International Conference on Information Fusion, 2012, 2301–8.

[14] P. C. G. Costa and D. M. Buede
Dynamic decision making: a comparison of approaches.
Journal of Multi-Criteria Decision Analysis, 9 (2000), 243–262.

[15] S. K. Das
High-Level Data Fusion.
Artech House, 2008.

[16] B. V. Dasarathy
Sensor fusion potential exploitation-innovative architectures and illustrative applications.
Proceedings of the IEEE, 85 (1997), 24–38.

[17] J. P. de Villiers, K. Laskey, A. L. Jousselme, E. Blasch, A. de Waal, G. Pavlin, and P. Costa
Uncertainty representation, quantification and evaluation for data and information fusion.
in Proceedings of the 18th International Conference on Information Fusion, 2015, 50–57.

[18] J. Dezert and F. Smarandache
Presentation of DSmT.
in J. Dezert and F. Smarandache (Eds.). Advances and Applications of DSmT for Information Fusion. American Research Press, 2004, 3–35.

[19] F. J. Díez and M. J. Druzdzel
Canonical probabilistic models for knowledge engineering.
Report No. CISIAD-06-01, Madrid, Spain: Universidad Nacional de Educación a Distancia, 2006.

[20] X. L. Dong, E. Gabrilovich, G. Heitz, W. Horn, K. Murphy, S. Sun, and W. Zhang
From data fusion to knowledge fusion.
Proceedings of the Very Large Database Endowment, 7 (2014), 881–892.

[21] E. Elles and B. Fitelson
Measuring confirmation and evidence.
The Journal of Philosophy, 97 (2000), 663–672.

[22] M. R. Endsley
Toward a Theory of Situation Awareness in Dynamic Systems.
Human Factors, 37 (1995), 32–64.

[23] M. R. Endsley
Theoretical underpinnings of situation awareness: A critical review.
in M. R. Endsley and D. J. Garland (Eds.) Situation Awareness Analysis and Measurement. Mahwah, NJ: Lawrence Earlbaum Associates, 2000, pp. 3–32.

[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, 14 (2005), 273–281.

[25] Evaluation of Technologies for Uncertainty Reasoning Working Group (ETURWG)
Homepage: http://eturwg.c4i.gmu.edu.

[26] P. H. Foo and G. W. Ng
High-level information fusion: An overview.
Journel of Advances in Information Fusion, 8 (2013), 33–72.

[27] J. García, L. Snidaro, and J. Llinas
Architectural aspects for context exploitation in information fusion.
in Context-Enhanced Information Fusion. Springer, 2016, 185–203.

[28] L. Getoor, N. Friedman, D. Koller, A. Pfeffer, and B. Taskar
Probabilistic relational models.
in L. Getoor and B. Taskar (Eds.) Introduction to Relational Statistical Models. The MIT Press, 2007.

[29] T. Gneiting and A. E. Raftery
Strictly proper scoring rules, prediction, and estimation.
Journal of the American Statistical Association, 102 (2007), 359–378.

[30] D. L. Hall and J. Llinas
An introduction to multisensor data fusion.
Proceedings of the IEEE, 85 (1997), 6–23.

[31] D. L. Hall and J. Llinas
Multi-sensor data fusion.
in M. E. Liggins, D. L. Hall and J. Llinas (Eds.) Handbook of Multisensor Data Fusion: Theory and Practice. Boca Raton, FL: CRC Press, 2009, 1–14.

[32] E. K. Jones, N. Denis, and D. Hunter
Hypothesis management for information fusion.
in Proceedings of the Fifth International Conference On Information Fusion, 2002, 894–899.

[33] A. Jøsang
Subjective Logic.
Springer, 2016.

[34] A. Jøsang, P. C. Costa, and E. Blasch
Determining model correctness for situations of belief fusion.
in Proceedings of the 16th International Conference On Information Fusion, 2013, 1886–1893.

[35] A. Jousselme, P. Maupin, E. Bosse, and D. Lambert
Knowledge, belief and uncertainty.
in E. Bossé, J. Roy, and S. Wark (Eds.) Concepts, Models, and Tools for Information Fusion. Artech House, 2007, 141–150.

[36] A. L. Jousselme
Semantic criteria for the assessment of uncertainty handling fusion models.
in Proceedings of the 19th International Conference On Information Fusion, 2016, 488–495.

[37] A. L. Jousselme, A. C. Boury-Brisset, B. Debaque, and D. Prevost
Characterization of hard and soft sources of information: a practical illustration.
in Proceedings of the 17th International Conference On Information Fusion, 2014, 1–8.

[38] A. L. Jousselme and G. Pallotta
Dissecting uncertainty-based fusion techniques for maritime anomaly detection.
in Proceedings of the 18th International Conference On Information Fusion, 2015, 34–41.

[39] J. A. Karakowski
Towards Visual Data Fusion.
in Proceedings of the National Symposium of Sensor and Data Fusion, 1998, Lexington, MA, USA.

[40] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi
Multisensor data fusion: A review of the state-of-the-art.
Information Fusion, 14 (2013), 28–44.

[41] F. H. Knight
Risk, Uncertainty, and Profit.
Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Co., 1921.

[42] D. Koller, A. Levy, and A. Pfeffer
P-CLASSIC: a tractable probabilistic description logic.
Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence, 1997, 390–397.

[43] P. Krause and D. Clark
Representing Uncertain Knowledge: An Artificial Intelligence Approach.
Kluwer Academic Publishers, Dordrecht, Boston, 1993.

[44] D. A. Lambert
Grand challenges of information fusion.
in Proceedings of the Sixth International Conference on Information Fusion, 2003, 213–220.

[45] D. A. Lambert
The state transition data fusion model.
in E. Blasch, E. Bossé, and D. A. Lambert. (Eds.) High-Level Information Fusion Management and Systems Design. Artech House, 2012, 33–79.

[46] K. B. Laskey, P. C. Costa, and T. Janssen
Probabilistic ontologies for knowledge fusion.
in Proceedings of the 11th International Conference On Information Fusion, 2008, 1–8.

[47] K. J. Laskey, K. B. Laskey, P. C. Costa, M. M. Kokar, T. Martin, and T. Lukasiewicz
Uncertainty Reasoning for the World Wide Web.
W3C Incubator Group Report 31, 2008.

[48] D. Lee (Ed.)
Law of Armed Conflict Deskbook, 5th ed.
The Judge Advocate General's Legal Center and School, U.S. Army, Charlottesville, Virginia, USA, 2015.

[49] T. C. Ly, S. Greenhill, S. Venkatesh, and A. Pearce
Multiple hypotheses situation assessment.
in Proceedings of the Sixth International Conference of Information Fusion, 2003, 972–978.

[50] R. P. Mahler
Statistical Multisource-Multitarget Information Fusion.
Artech House, Inc, 2007.

[51] J. E. Matheson and R. L. Winkler
Scoring rules for continuous probability distributions.
Management Science 22 (1976), 1087–1096.

[52] B. Milch, B. Marthi, S. Russell, D. Sontag, D. L. Ong, and A. Kolobov
BLOG: Probabilistic Models with Unknown Objects.
in L. Getoor and B. Taskar (Eds.) Statistical Relational Learning. The MIT Press, 2007, 373–398.

[53] Opinion
Oxford English Dictionary (Online version).
https://en.oxforddictionaries.com/definition/opinion.

[54] H. Pasula, B. Marthi, B. Milch, S. Russell, and I. Shpitser
Identity uncertainty and citation matching.
in Advances in Neural Information Processing Systems, 2002, 1401–1408.

[55] L. F. Pau
Sensor data fusion.
Journal of Intelligent and Robotic Systems, 1 (1988), 103–116.

[56] J. Piskorski, R. Yangarber
Information extraction: past, present and future.
in Multi-Source, Multilingual Information Extraction and Summarization. Springer, 2013, 23–49.

[57] J. Roy, D. Lambert, F. White, J. Salerno, J. Karakowski, and M. Hinman
Data- and Information-fusion models.
in E. Bosse, J. Roy and S. Wark, (Eds.) Concepts, Models, and Tools for Information Fusion. Artech House, Boston, 2007, 69–118.

[58] J. Salerno
Information fusion: a high-level architecture overview.
in Proceedings of the Fifth International Conference On Information Fusion, 2002, 680–686.

[59] D. A. Schum
The Evidential Foundations of Probabilistic Reasoning.
Northwestern University Press, 1994.

[60] D. A. Schum
Evidence marshaling for imaginative fact investigation.
Artificial Intelligence and Law, 9 (2001), 165–188.

[61] D. A. Schum
A science of evidence: contributions from law and probability.
Law, Probability & Risk, 8 (2009), 197–231.

[62] G. Shafer
A Mathematical Theory of Evidence.
Princeton University Press, 1976.

[63] E. Shahbazian, D. E. Blodgett, and P. Labbé
The extended OODA model for data fusion systems.
Proceedings of the Fourth International Conference On Information Fusion, 2001.

[64] P. P. Shenoy
Valuation-based systems for Bayesian decision analysis.
Operations Research, 40 (1992), 463–484.

[65] F. Smarandache and J. Dezert
Advances and Applications of DSmT for Information Fusion, Vol. IV: Collected Works.
Infinite Study, 2015.

[66] P. Smets
The combination of evidence in the transferable belief model.
IEEE Transactions on Pattern Analysis and Machine Intelligence, 12 (1990), 447–458.

[67] A. Smirnov, T. Levashova, and N. Shilov
Patterns for context-based knowledge fusion in decision support systems.
Information Fusion, 21 (2015), 114–129.

[68] M. Smithson
Ignorance and Uncertainty: Emerging Paradigms.
New York: Cognitive Science. Springer-Verlag, 1989.

[69] L. Snidaro, L. Vati, J. García, E. D. Martí, A. L. Jousselme, K. Bryan, D. D. Bloisi, and D. Nardi
A framework for dynamic context exploitation.
in Proceedings of the 18th International Conference On Information Fusion, 2015, 1160–1167.

[70] J. F. Sowa
Knowledge Representation: Logical, Philosophical, and Computational Foundations.
Pacific Grove: Brooks/Cole, 2000.

[71] A. N. Steinberg
Foundations of situation and threat assessment.
in M. E. Liggins, D. L. Hall, and J. Llinas (Eds.) Handbook of Multisensor Data Fusion: Theory and Practice. Boca Raton, FL: CRC Press, 2009, 691–699.

[72] A. N. Steinberg, C. L. Bowman, and F. E. White
Revisions to the JDL data fusion model.
in Sensor Fusion: Architectures, Algorithms, and Applications III. International Society for Optics and Photonics, 1999, 430–442.

[73] J. Steinhauer, A. Karlsson, and S. F. Andler
Traceable uncertainty.
in Proceedings of the 16th International Conference On Information Fusion, 2013, 1582–1589.

[74] S. C. Thomopoulos
Sensor integration and data fusion.
Journal of Robotic Systems, 7 (1990), 337–372.

[75] U.S. Air Force
Air Force Distributed Common Ground System.
Website: http://www.af.mil/About-Us/Fact-Sheets/Display/Article/104525/air-force-distributed-common-ground-system.

[76] U.S. Congress
Rule 401.
Federal Rules of Evidence, 1975.

[77] D. L. Weed
Weight of evidence: a review of concept and methods.
Risk Analysis, 25 (2005), 1545–1557.

[78] P. D. Whitney, C. Posse, and X. C. Lei
Toward a Unified Approach to Information Integration: a Review Paper on Data/Information Fusion.
Report No. PNNL-15428, Pacific Northwest National Laboratory, Richland, WA, USA, 2005.

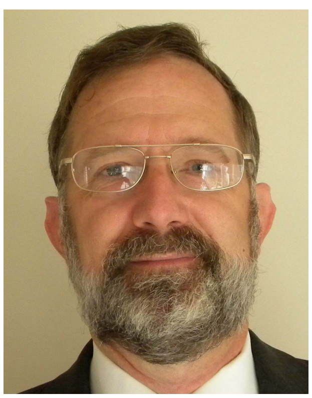

Mark Locher is currently a PhD candidate in System Engineering at George Mason University, Fairfax Virginia. He received a BS in Operations Research and Aviation Science in 1979 from the US Air Force Academy, an MBA from St Mary's University, San Antonio TX, in 1982, and an MS in System Engineering from George Mason University in 2012. He completed a 20-year career with the US Air Force and then 15 years as an industry consultant to the US Department of Defense. His primary focus was developing and integrating Intelligence Surveillance Reconnaissance systems and platforms in multiple domains. His research interests include uncertain reasoning techniques, knowledge representation and logical reasoning, and decision problem aiding. He is a member of the International Council of Systems Engineering.

Paulo Costa is an Associate Professor of Systems Engineering and Operations Research at George Mason University. His teaching and research interests comprise the areas of probabilistic ontologies, multi-sensor information fusion, Bayesian reasoning, predictive analysis, cyber-security and Decision Theory. He is a former fighter pilot with extensive experience in tactical and operational planning, and an expert in requirements engineering for complex systems, such as intelligent transportation and health-care support systems. Dr. Costa is an IEEE senior member, Member of the Board of Directors of the International Society for Information fusion, and member of the International Council of Systems Engineering.

Uncertainty representation and evaluation for modelling and decision-making in information fusion

J. P. DE VILLIERS
G. PAVLIN
A.-L. JOUSSELME
S. MASKELL
A. DE WAAL
K. LASKEY
E. BLASCH
P. COSTA

In this paper, the uncertainties that enter through the life-cycle of an information fusion system are exhaustively and explicitly considered and defined. Addressing the factors that influence a fusion system is an essential step required before uncertainty representation and reasoning processes within a fusion system can be evaluated according to the Uncertainty Representation and Reasoning Evaluation Framework (URREF) ontology.

The life cycle of a fusion system consists primarily of two stages, namely *inception and design*, as well as *routine operation and assessment*. During the inception and design stage, the primary flow is that of abstraction, through modelling and representation of real-world phenomena. This stage is mainly characterised by epistemic uncertainty.

During the routine operation and assessment stage, aleatory uncertainty combines with epistemic uncertainty from the design phase as well as uncertainty about the effect of actions on the mission in a feedback loop (another form of epistemic uncertainty). Explicit and accurate internal modelling of these uncertainties, and the evaluation of how these uncertainties are represented and reasoned about in the fusion system using the URREF ontology, are the main contributions of this paper for the information fusion community. This paper is an extension of previous works by the authors, where all uncertainties pertaining to the complete fusion life cycle are now jointly and comprehensively considered. Also, uncertainties pertaining to the decision process are further detailed.

Manuscript received January 22, 2019; revised March 20, 2019; released for publication April 17, 2019.

Refereeing of this contribution was handled by Sten Andler.

Authors' addresses: J. P. de Villiers, University of Pretoria, Pretoria, South Africa, Council for Scientific and Industrial Research, Pretoria, South Africa (E-mail: pieter.devilliers@up.ac.za). G. Pavlin, D-CIS Lab, Thales Research and Technology, Delft, The Netherlands (E-mail: gregor.pavlin@d-cis.nl). A.-L. Jousselme, NATO STO Centre for Maritime Research and Experimentation, La Spezia, IT (E-mail: anne-laure.jousselme@cmre.nato.int). S. Maskell, University of Liverpool, Liverpool, UK (E-mail: s.maskell@liverpool.ac.uk). A. de Wall, University of Pretoria, Pretoria, South Africa, Center for Artificial Intelligence Research (CAIR), Cape Town, South Africa (E-mail: alta.dewaal@up.ac.za). K. Laskey, P. Costa, George Mason University, Fairfax, VA, USA (E-mail: {pcosta,klaskey}@gmu.edu). E. Blasch, Air Force Research Lab, Arlington, VA, USA (E-mail: erik.blasch.1@us.af.mil).

1557-6418/18/$17.00 © 2018 JAIF

I. INTRODUCTION

The characterisation of uncertainty is required for pragmatic decision making when sensor data and other forms of information from several sources are fused in decision support systems. Uncertainty characterisation requires implicit and explicit forms of abstraction to model the problem, represent entities and concepts within the world, associate entities to uncertainties, and to reason about decision consequences. Uncertainties propagate through the life cycle of an information fusion system (hereafter referred to as a fusion system), from the problem statement and modelling phases to design and implementation. Ideally a fusion system life cycle should include:

a) the exhaustive characterisation of uncertainties throughout the life cycle of a fusion system;
b) the explicit (i.e., direct, solvable) representation of these uncertainties within the fusion system; and,
c) the implicit (i.e., indirect, iterative) evaluation of these uncertainties.

Two life cycle stages which have been previously considered are the *modelling phase* [1] (representing uncertainty) and the *operation phase* (performing the decision loop) [2]. This paper will consolidate the uncertainty evaluation of these phases, as well as include the *inception and design phase*, presented in [3]. Although subsets of uncertainties are considered during the design and use of all fusion systems, in this paper, and for the first time, all uncertainties that enter throughout the complete fusion life cycle are jointly and comprehensively considered.

This paper provides concepts that, in combination with the evaluation criteria defined in the Uncertainty Representation and Reasoning Evaluation Framework (URREF) [4], facilitate the development of verifiable operational fusion systems. Entity abstraction provides a clear mapping between the physical phenomena of interest and the abstract models used in the fusion system. The development process (or flow of abstraction) is partitioned into activities that focus on isolation abstraction, process abstraction, data generation abstraction, datum abstraction and agent abstraction. The flow of information, on the other hand, introduces a taxonomy of operational elements, which facilitate the development of a system that satisfies the functional and performance requirements. The concepts introduced by abstraction and information flows support both, the analysis in the inception phase (where the problem statement is defined) and the development of concrete solutions in the design phase of a URREF driven development life cycle [3] shown in Fig. 1. Fig. 1 defines the system partitions that enable logical allocations of various URREF evaluation criteria.

Although preliminary works [1], [2] classify several types of uncertainty, there are two types of uncertainty prevalent in the literature. The two types are epistemic

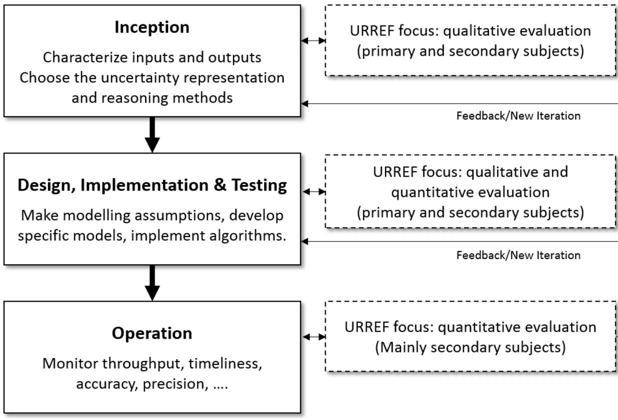

Fig. 1. URREF roles in a development life cycle [3] depicting the inception phase, the design implementation and testing phase, and the operation phase.

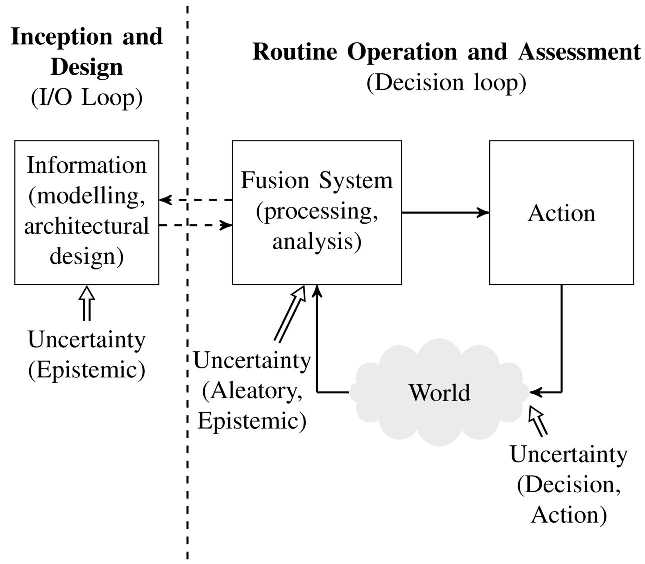

Fig. 2. The two main phases of a fusion system, namely the inception and design phase (input/output loop), and the routine operation phase (decision loop) are depicted. The double arrows depict where uncertainty enters the two phases, and the dashed arrows depict implementation and design refinement. Apart from aleatory and epistemic uncertainty, decision uncertainty captures the uncertainty of the effect of an action on the world.

and aleatory uncertainty [5], [6]. *Epistemic uncertainty* is derived from the Greek word "episteme" and relates to uncertainty owing to a lack of knowledge or ignorance about the modelled process or entity. Therefore this uncertainty lies *outside* of the entity or process being modelled. *Aleatory uncertainty* is derived from the Latin word "alea" which refers to the casting of dice. Aleatory uncertainty refers to random events *within* the entity or process being modelled. As such, both epistemic and aleatory uncertainties are encountered throughout the life cycle of an information fusion system. The focus of this paper will be to unify uncertainties that enter during abstraction, design, and modelling [1], [3] with those during explanation, operation, and decision making [2].[1]

There exists a significant body of knowledge on the quantification of uncertainty inherent in models of physical processes [5]–[9]. In these works, *uncertainty classification* is organized as being *forward* or *inverse* [9]. On the one hand, forward uncertainty quantification considers how uncertainty propagates through a model from the input to the output of the model. On the other hand, inverse uncertainty quantification involves not only the characterisation of the discrepancy between the experimental results and the predictions of the mathematical model, but also the estimation of parameter values [10].

The ISIF Evaluation Techniques for Uncertainty Representation Working Group (ETURWG) investigates challenges associated with uncertainty reasoning, analysis, and usability in information fusion processes. An ongoing effort of the working group is the design of the URREF ontology, which captures primary and secondary concepts that relate to uncertainty representation and reasoning in information fusion systems, as well as the links between the concepts [4]. The evolution of the concepts, links and definitions of the URREF ontology has reached a stable form and is utilised to evaluate uncertainty related aspects in a variety of fusion problems e.g., [11]–[17].

Over the years, a comprehensive "joint uncertainty" formulation (or a globally complete consideration of uncertainty) has been identified as a need by several International Society of Infomation Fusion (ISIF) panels [18]. The purpose of this paper is to define, within the context of the URREF ontology, all the stages at which there is potential for uncertainty to enter the full life cycle of an information fusion system as well as to classify these uncertainties. These uncertainties are referred to as the *subjects of evaluation* of the URREF ontology, as discussed in [19]. Siloed approaches to uncertainty representation and reasoning (traditional approaches) could fail in many applications. Table I (column 3) provides some examples of processes of abstraction (modelling) that could fail if the joint uncertainty is not considered. For example, in [20] the author focused on the scheduling based on the time available. Time available is a good choice, but uncertainty is also needed to get to a "value" function. If one radar's performance starts decreasing (meaning possibly more uncertainty), then scheduling needs to adapt. Furthermore, different types of uncertainty (described semantically) can affect the end utility/policy.

The rest of the paper is ordered as follows. Section II presents the information fusion life cycle. Section III articulates details of an information fusion system design. Section IV complements Section III with the information fusion operation. Section V contains a discussion on use cases and Section VI a discussion of evaluation

[1]Note the duality between: abstraction, design, modelling; and explanation, operation, decision-making.

using the URREF within the context of atomic decision processes. Section VII concludes the paper.

II. INFORMATION FUSION SYSTEM LIFE CYCLE

According to the taxonomy presented in this paper, there are two phases where uncertainty can enter into a fusion system. These are the inception/design and operation/assessment phases. These phases are presented in the subsections below, and Fig. 2 provides further clarification.

A. Inception and design—Abstraction flow

The first phase of an information fusion system is the Inception and Design (IAD) during which the architecture is specified and the mathematical models are assembled. The IAD process is concerned with the *flow of abstraction*, i.e., where real world entities and processes (RWEPs) are modelled, and epistemic and aleatory uncertainties are represented in a mathematical formalism. The abstraction flow takes place on a relatively large time scale (e.g., months), while feedback spiral processes in the systems engineering requirements specification and design can result in incremental improvements in the system in shorter time scales (e.g., days).

B. Routine operation and assessment—Information flow

The second phase of an information fusion systems is the Routine Operation and Assessment (ROA) during which the system functions as a decision process, akin to the Observe, Orient, Decide and Act (OODA) loop of Boyd [21]. The ROA phase is mainly concerned with the *flow of information*, where the information is collected from transducers (sensors) that convert real-world observable phenomena into categorical quantities, associated uncertainties, and representation processes (such as probability, fuzzy logic, belief functions, etc.). The objective of the information fusion system is to reduce uncertainty and improve inference for informed decision making.

III. FUSION SYSTEM INCEPTION AND DESIGN

The modelling of fusion systems involve abstracting RWEPs and the mechanisms whereby they generate observable phenomena, to result in mathematical and uncertainty models of RWEPs of interest. These observable phenomena are, for example in a multisensor radar tracking system, the electromagnetic characteristics of the skin of moving aircraft and how it interacts with radar pulses to form a series of detections, whereby the first objective is to determine the state vector of all the aircraft in some area of regard. The second objective is to make informed decisions, using the inferred state vectors, such as in the case of air traffic control.

Fig. 3 is a symbolic depiction of the process of modelling with the objective of performing information fusion. Fig. 3 has been extended when compared to Fig. 1 in [1] in that the uncertainties that enter during the abstraction and modelling of the decision process resulting in the "Decision Model" have been appended. The objective of presenting such a detailed view, is to provide the fusion system designer with an explicit and exhaustive view of where uncertainties enter the design and modeling process through the adoption of several assumptions.

There is a clear flow of abstraction from left to right. The real world is depicted by the shaded cloud as a series of RWEPs that generate observable phenomena. To be explicit, the nth RWEP denoted by RWEPn generates a real world datum $D_{n,k}$ at time instant k. A datum is defined as an observable real-world effect, such as a radio frequency transmission, a visible light reflection off a target, etc. The nth real world process has physical properties that are represented by the symbol Ω_n. The way in which observable effects are generated by the RWEP, is represented by the transformation $\{D_n \mid \Omega_n\}$, and can be read as D_n given Ω_n, analogous to as if it would have been conditioned on Ω_n in the statistical sense. Furthermore, these real world entities can interact with each other, forming the *situation* and *impact* levels of the Joint Director of the Laboratory/Data Fusion Information Group (JDL/DFIG) fusion models [22]–[25]. The different types of uncertainties that enter through the abstraction process are represented by different variables, which are summarised in the first column of Table I.

A. Isolation Abstraction

If the objective of a specific fusion system is considered, then there are typically only a few RWEPs that are of interest for a specific decision making problem. For example, in the air traffic control application, the controller is only interested in air targets within a certain area of regard, and also not surface targets, unless these are at an airport. This is the first element of abstraction that takes place, and is referred to as *isolation abstraction*. Uncertainties enter during this type of abstraction whereby assumptions are made that outside influences are ignored or simplified, and boundary conditions are specified. These uncertainties are labeled isolation uncertainties and are denoted by γ. Since all models and processes downstream from this decision are influenced by γ, and to simplify notation, dependence on γ will not be explicitly shown, although it should be kept in mind. Isolation abstraction uncertainty γ is epsitemic in nature (indicated by † in Fig. 3).

B. Process Abstraction

Typically, RWEPs contain some properties that are hidden or latent, but which are needed for decision making purposes. It is for this reason that models are needed to describe as accurately as possible how these processes and entities behave and evolve over time. The procedure for assembling such models is labeled as *process abstraction*, and result in a *process or plant model* (PM) for the nth RWEP. Such models are time dependent, and describe the stochastic evolution of cur-

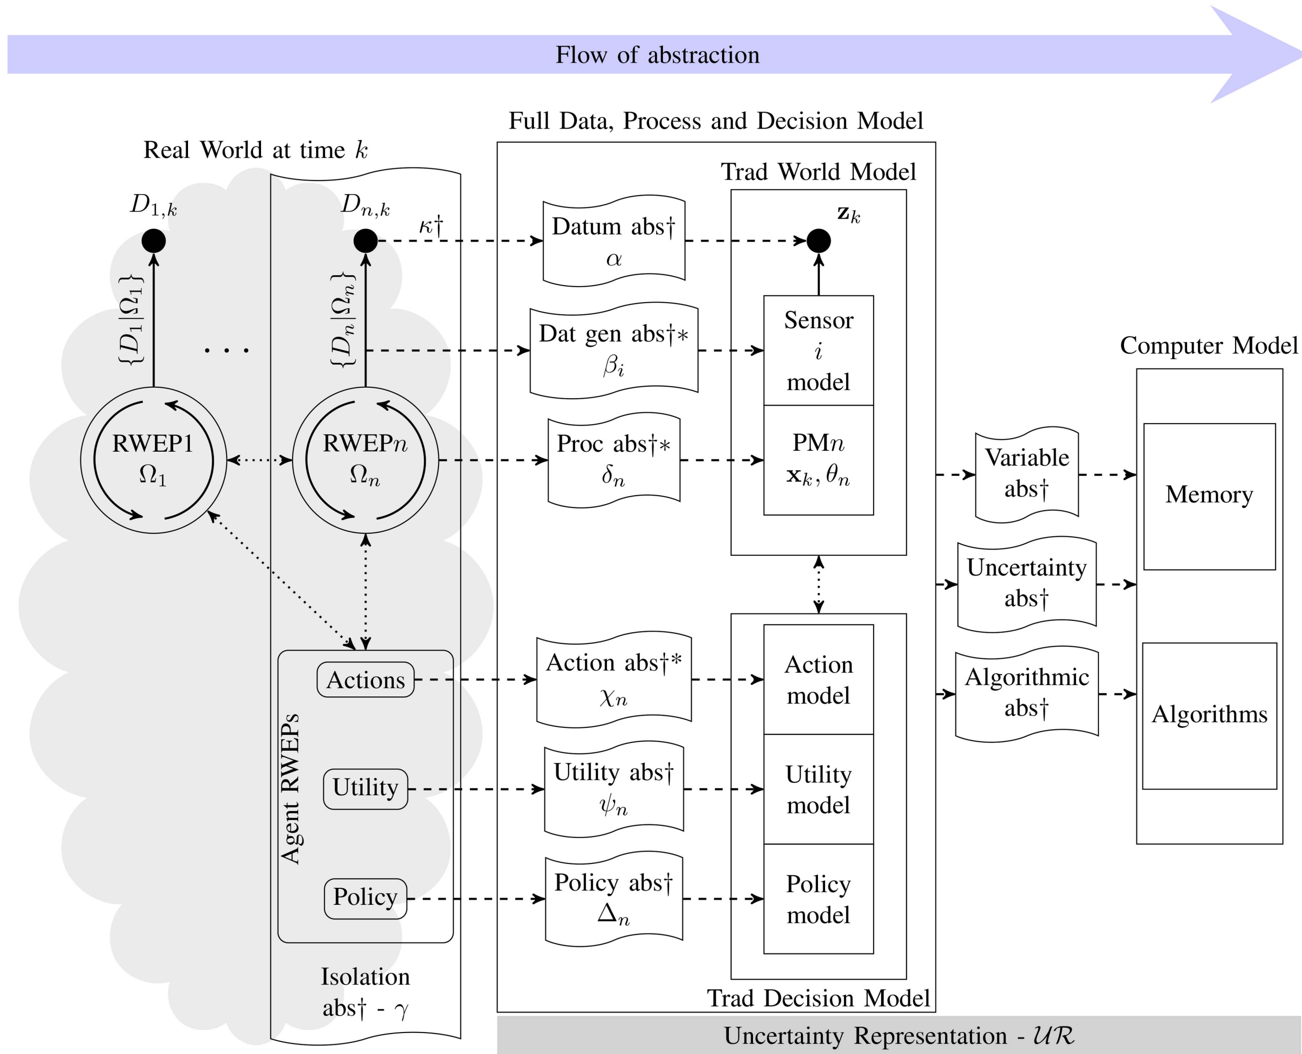

Fig. 3. The modelling (abstraction) of a fusion system making a measurement at time k is depicted. The principal components depicted are a) real world entities and processes (RWEPs), b) agents acting in the world (a specific type of RWEP), and c) models/abstractions of these RWEPs. Solid arrows indicate how data is generated. Dotted arrows indicate that real world or model processes influence each other. Dashed arrows indicate the flow of abstraction during the modelling process. Ribbons indicate processes of abstraction (i.e. representing RWEPs as mathematical objects). The symbol † indicates epistemic uncertainty, whereas the symbol $*$ indicates aleatory uncertainty. The shaded bar in the lower right of the figure shows that the uncertainty representation cross-cuts the modelling and implementation of a fusion system. The index i denotes the sensor index and n is the nth real-world entity/process being modelled.

rent (and future) states $\mathbf{x}_{k:k+N}$ based on past states $\mathbf{x}_{0:k-1}$ and model parameters θ, which are time invariant. These states and parameters are typically abstractions of the real world physical attributes contained in Ω_n. In traditional Bayesian tracking, the evolution of the uncertainty relation in the PM is represented by $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \theta_n)$. The modelling of how RWEPs generate data, and as such, how observed phenomena relate to hidden (unobserved) processes, are encapsulated by the sensor/data model.[2] Hidden uncertainty processes are discussed in the next section.

A process model relates parameters and states to each other over time. Epistemic uncertainty enters into the PM through incomplete knowledge about the corresponding RWEP. Aleatory uncertainty enters into the model through random perturbations in the time evolution of the model. Consider, for example, a discrete time varying equation $\mathbf{x}_k = f(\mathbf{x}_{k-1}) + \epsilon$, where $\mathbf{x}_k$ is the system state at discrete time step k and ϵ some random quantity. In many cases both epistemic and aleatory uncertainties are (possibly incorrectly) lumped together in a single random quantity ϵ. The framework presented here provides for their explicit separation via an additional variable δ_n to capture epistemic uncertainty.

C. Data Generation Abstraction

Data generation abstraction involves the modelling of how observable effects relate to unobservable (hidden or latent) processes with states $\mathbf{x}_k$ and parameters θ_n. The output of data generation abstraction is both a model of how a specific measurement is related to an unobserved parameter or state, and also a sensor/data model, which

[2] This is also known as a measurement or observation model.

TABLE I
Different types of abstraction in the modelling process, their descriptions and examples

Abstraction Type/Related Uncertainty Variable	Abstraction Process	Description	Example
Isolation γ	Choosing system boundaries, making assumptions	Isolating the RWEP or multiple RWEPs by choosing the domain, processes and entities of interest in the real world	The features, dynamics and sensing of multiple targets that are observable or can be inferred indirectly from measurements within the coverage area of multiple radars. This isolation could explicitly be represented by an ontology.
Datum α	Define mathematical variable type and uncertainty representation	Choosing a mathematical or numeric representation of a measurement $\mathbf{z}_k$ and associated uncertainty to represent a real world datum $D_{n,k}$ or data	Integer, natural number, real number, vector, matrix, complex number, tensor, norm, first order logic expression, etc.
Data generation β_i	Define data/sensor model	Choosing a mapping between RWEPs, and data and an uncertainty representation for representing uncertainty in the data generation process *as well as* characterising the real world data generation process	Choosing a probabilistic uncertainty representation and specifying a Gaussian model of data generation with mean and covariance parameters to model the generation of range and Doppler measurements by a radar.
Process δ_n	Define process model	Choosing states, parameters, a mapping between parameters and states* and an uncertainty representation for states, parameters and mappings	Choosing a hidden Markov model to represent the time evolution of a target state, where the plant noise captures both uncertainties in knowledge of the motion model and real world randomness such as air pockets, and imprecise control inputs by the pilot of an aircraft.
Action χ_n	Define model of actions	Define the actions available to an agent. Define a mapping between available actions, and the evolution of world (and agent) states.	Defining the available scan patterns and tracking tasks in an Active Electronically Scanned Array (AESA) radar, and how these tasks influence future tasks of the radar.
Utility ψ_n	Define a utility/reward model	Choosing a mapping between agent/world states and their desirability as perceived by the agent/system user	Define a reward function which balances the effort spent by the AESA radar tracking existing targets as opposed to scanning for possibly undetected targets.
Policy Δ_n	Define a policy representation	Choose a mapping between the world state as perceived by the agent and the most appropriate action for being in that perceived state	Choose a pre-defined rule for time spent on tracking vs scanning, which maximises the expected sum of future discounted rewards.

*An example of a mapping between parameters and states is how a probability distribution over target mass maps to a probability distribution over accelerations.

specifies how data are generated and transduced by the ith sensor. These are two sides of the same coin. In the case of traditional probabilistic modelling, these relations are characterised by the quantity $p_i(\mathbf{z}_k \mid \mathbf{x}_k, \theta_n)$. If the measurement $\mathbf{z}_k$ is known and $\mathbf{x}_k, \theta_n$ are variable, the function $p_i(\mathbf{z}_k \mid \mathbf{x}_k, \theta_n)$ represents the likelihood $L_z(\mathbf{x}_k, \theta_n)$ and is a function, not a probability distribution. However if $\mathbf{x}, \theta_n$ are known and $\mathbf{z}_k$ is the variable, then $p_i(\mathbf{z}_k \mid \mathbf{x}_k, \theta_n)$ represents the probabilistic model of data generation, and it is a proper probability distribution. Note that $p(\mathbf{z}_k \mid \mathbf{x}_k, \theta_n)$ typically includes the sensor model or the model of perception, as the sensor forms part of the RWEPs and also generates data. Therefore, $p_i(\mathbf{z}_k \mid \mathbf{x}_k, \theta_n)$ could serve as both a model for estimation/inference (for example maximum likelihood) which is related to inverse uncertainty quantification or a model for data generation (a generative model) which is related to forward uncertainty quantification.

The uncertainty in data generation abstraction for sensor i is denoted by the symbol β_i. The procedure of data generation abstraction causes epistemic uncertainty, since there may be lack of knowledge about the nature of the transformation from a RWEP to a datum. In addition to epistemic uncertainty, aleatory uncertainty (denoted by a $*$ in Fig. 3) is expressed through the random nature by which data are generated and sensed. Hence the measurement process is depicted in Fig. 3 to contain both epistemic and aleatory uncertainties.

D. Datum Abstraction

The datum D_n is a real world effect that is observed. It cannot be used in any kind of reasoning, since a process of abstraction is needed to convert it into a mathematical quantity such as a integer, real number, complex vector, a first order logic statement, etc. This process is labelled *datum abstraction*. In some cases, a datum may already be abstracted, such as output of another fusion process (such as the output of a filter), and as such, dependencies exist between data points. In a subset of these cases, datum abstraction may not be needed, unless some form of conversion takes place. A datum should also not be confused with a *measurement* (in this taxonomy denoted by $\mathbf{z}_k$) which has already been transduced by a sensor into an instantiation of a mathematical quantity.

Uncertainties that enter with the process of datum abstraction (i.e., the numerical, ordinal or logical representation of observable physical phenomena), are denoted by the symbol α and is epistemic in nature (indicated by † in Fig. 3.). An example would be for α to represent the fact that a continuous variable is discretised, and as such may not sufficiently capture the important or relevant properties of the datum, resulting in significant quantisation noise. Epistemic uncertainties associated with representing the *uncertainty relations/functions* (probability densities, belief functions) of a datum D_n are also contained within α, and a loss may occur if, for example, an imprecise language statement is represented by a discrete probability distribution. This is an example of second order uncertainty (uncertainty about uncertainty).

E. Agent abstraction

The decision process, fusion resource management, and mission actions need to be modelled if a fusion system needs to be automatically steered to produce desired states of the world. In Fig. 3, a model is depicted as an agent. Although an agent is simply another type of RWEP, whose actions and influences can be observed as data by sensors, they merit explicit mention, as being an integral part of the decision loop. An agent in the real world is motivated by some utility or reward, which captures the desirability of a world state at a time instance. If all time is considered, a (discounted) accumulation of utilities (sum of rewards) over all time is of relevance. The agent would then act according to a general set of rules (or policy) which would ideally maximise the discounted accumulation of utilities/rewards over a possibly infinite time horizon. Agent actions are the general premise of the fields of linear Gaussian quadratic (LGQ) control [26], [27], reinforcement learning [28], Markov decision processes (MDPs) [28], [29], partially observed Markov decision processes (POMDPs) [28], [30], and model predictive control [31]. Being central to the decision making process, this setting needs to be modelled—first mathematically and then be instantiated algorithmically, for automated decision making. These processes of abstraction are depicted in Fig. 3, which capture the main components of the agent. The processes include: *action abstraction* μ_n, which models the effect of actions on the evolution of world states, *utility abstraction* ψ_n, which models the desirability of world states, and *policy abstraction* Δ_n, which models the rule set by which to act given a world state. Action abstraction may introduce aleatory and epistemic uncertainty—"aleatory" owing to how actions may influence the world state in a "noisy" sense, and "epistemic" owing to lack of knowledge how actions are represented and how they influence the world state. The utility and policy abstraction processes typically exhibit epistemic uncertainty, since the uncertainty pertains to how the desirability of states, and the mapping of perceived states (otherwise known as belief states) to actions are modelled (represented by some function). Owing to the vastness of policies for most belief state spaces, several methods exist to compress these policies, leading to epistemic uncertainty owing to representation approximations. These include belief compression [32], certainty equivalence [28], and symbolic policy approximation [33] to name a few. Current and recent research has, for example, looked to extend the scalability [34] of these approaches and apply them in pertinent contexts such as automotive applications [35].

F. Association Uncertainty

The association problem in information fusion is concerned with knowing which entity or process generated which observable datum $D_{n,k}$ at some time k. This ambiguity is depicted as the diagonal dotted lines between different RWEPs and D's. The association uncertainty will also be assigned a symbol, and will be denoted by κ. Association uncertainty κ is epistemic in nature, because it is due to a lack of knowledge.

G. The Computer Model

The final layers of abstraction, when proceeding from the mathematical model to a computer model is very briefly discussed here, and quotes the discussion in [1]. "In the case of digital computers, the use of established scientific libraries and vector-matrix mathematical programming environments make *variable abstraction* fairly well characterised. Uncertainties may enter through *algorithmic abstraction* in the form of possible incorrect implementation, numerical instabilities or strange behaviour in untested states. However, most cases of numerical instabilities in digital computer code are well characterised [36], and examples include the inversion of an ill-conditioned matrix, or numerical instabilities owing to Euler numerical integration. In this case incorrect implementation would be owing to oversight by the programmer. *Uncertainty abstraction* is characterised by pseudo number generators and Taylor series expansions to represent continuous probability distributions. Uncertainties for this type of abstraction are also well characterised in the literature. If on the other hand, analogue computers were used, this abstraction would have needed particular care in characterising uncertainties, as the results would be noisy."

H. Towards a full data, process and decision model

Epistemic modelling uncertainties (i.e., those that occur when going through the different processes of abstraction) are sometimes not sufficiently accounted for or explicitly modelled in traditional models. Traditional models are depicted as "Trad World Model" and "Trad Decision Model" in Fig. 3. Explicit consideration of modelling uncertainties are thus accounted for as in Ch 3 of [37]). A full data, process and decision model is therefore proposed, extended from [1]. Although it might be that the fusion system designer may choose to discount some of the uncertainties in Fig. 3, it is better that it is a conscious decision with consideration for the implications thereof, rather than an act of omission.

In traditional statistical modelling, $\mathbf{z_k}$ is considered to be the "datum" and $p(\mathbf{z} \mid \mathbf{x}, \theta_n)$ is considered to be the complete uncertainty model of $\mathbf{z}$. However, $\mathbf{z_k}$ is itself an abstraction of $D_{n,k}$, and similarly $p(\mathbf{z} \mid \mathbf{x}, \theta_n)$ is an abstraction of $\{D_n \mid \Omega_n\}$. As such, any uncertainties associated with these abstraction processes are ignored in traditional models. This steers the discussion towards higher order uncertainty (uncertainty about uncertainty). Higher-order uncertainty is modelled by imprecise probability models, belief functions or credal sets. For instance: rather than a single probability distribution, a set of probability distributions is considered, and the probability of an event is defined by upper and lower bounds.

A complete model of data generation must have the form $p(\Gamma \mid \mathbf{x}_k, \theta_n, \alpha)$, where $\Gamma = \{\mathbf{z}_k, \alpha\}$ is a mathematical model for $\mathbf{z}_k$ *as well as* the uncertainties associated with constructing $\mathbf{z}_k$, denoted by α. Furthermore, the uncertainty representation denoted by $p(\cdot \mid \mathbf{x}_k, \theta_n, \beta)$ must be a mathematical model of both the data generation process, as well as the uncertainties β associated with its construction. Such an uncertainty representation analogous to the *generalised likelihood* in [37].

The complete process model $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \theta_n, \delta)$ (which describes the time evolution of the world state) should encapsulate the aleatory uncertainty in the evolution of states as well as the epistemic uncertainties δ associated its construction. This is opposed to the traditional process model $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \theta_n)$ which is not conditioned on δ.

A similar approach should be followed for the decision model, where epistemic and aleatory uncertainties should be explicitly considered and incorporated into models where appropriate.

IV. FUSION SYSTEM OPERATION

In contrast with the inception, design and implementation of a fusion system in Fig. 3, the system operation at runtime is depicted in Fig. 4. Fig. 4 depicts the operation of the fusion system within the context of a decision loop. There are two principal flows that are identified in Fig. 4. The first is the flow of information, from RWEPs which generate observable phenomena, observed by sensors (or sources in general), combined in the fusion system, resulting in inference of world states and parameters. The second flow, the flow of decisions/actions involves the interpretation of inferences of the fusion system through a system which balances uncertainties with risks, rewards and utilities (such as Bayes' risk). The result of this process is a *decision* which is fed to a resource management algorithm, which in turn generates *actions* or *controls* that instruct sensors and mission actors to execute instructions. The principal taxonomies of such a decision process are addressed in [38], [11] and [19] as elementary constructs of conceptually indivisible *atomic decision processes* or ADPs.

The following sections will make the uncertainties that propagate through the fusion system explicit, so that each of them can be addressed if necessary. These sections are organised in the same order as the OODA loop, and Fig. 4 depicts the fusion decision loop. This loop contains the fusion system, which in turn comprises the conceptual *fusion elements* (FEs). These elements are conceptual, since in certain fusion methods they may all be present but not necessarily separable—for example a certain uncertainty representation cannot be separated from its inference method. Furthermore, it shows where different types of uncertainties enter the fusion system and propagate through the system. Fig. 4 is adapted from [2], where the elements of the fusion system, denoted by FE-1 to FE-4 have replaced ADP-1 to ADP-4 that were presented in [2]. The fusion elements include information source (FE-1), the instantiated model (FE-2), the inference and prediction (FE-3) as well as the decision method and resource management (FE-4).

A. Observe

Clues to the state of the world can be obtained by observations. Such observations can be obtained using sensors in the form of electronic transducers or human observers. Observations are required under the premise that "all decisions are based on observations of the evolving situation tempered with implicit filtering of the problem being addressed" [21]. In the subsections below a distinction is made between a) physical effects that *could* be observed by humans or sensors (observable real world data), and b) source reports by either humans or transducers (sensor data) that *have* observed the aforementioned physical effects.

1) Observable real world data:

Referring to Fig. 4, as in Fig. 3 observations originate from observable phenomena generated by RWEPs that interact with each other. A part of the world is isolated for which decisions are to be made (as in the case of modelling phase). Sensors make measurements of phenomena in the isolated area of interest. Reports from these sensors could assist in making inferences that may inform decisions. In the taxonomy of the decision loop in Fig. 4, not only the nth RWEP generates a datum $D_{i,k}$ which is sensed by sensor i, but $D_{i,k}$ may also be influenced by other RWEPs. An example is the use case of a

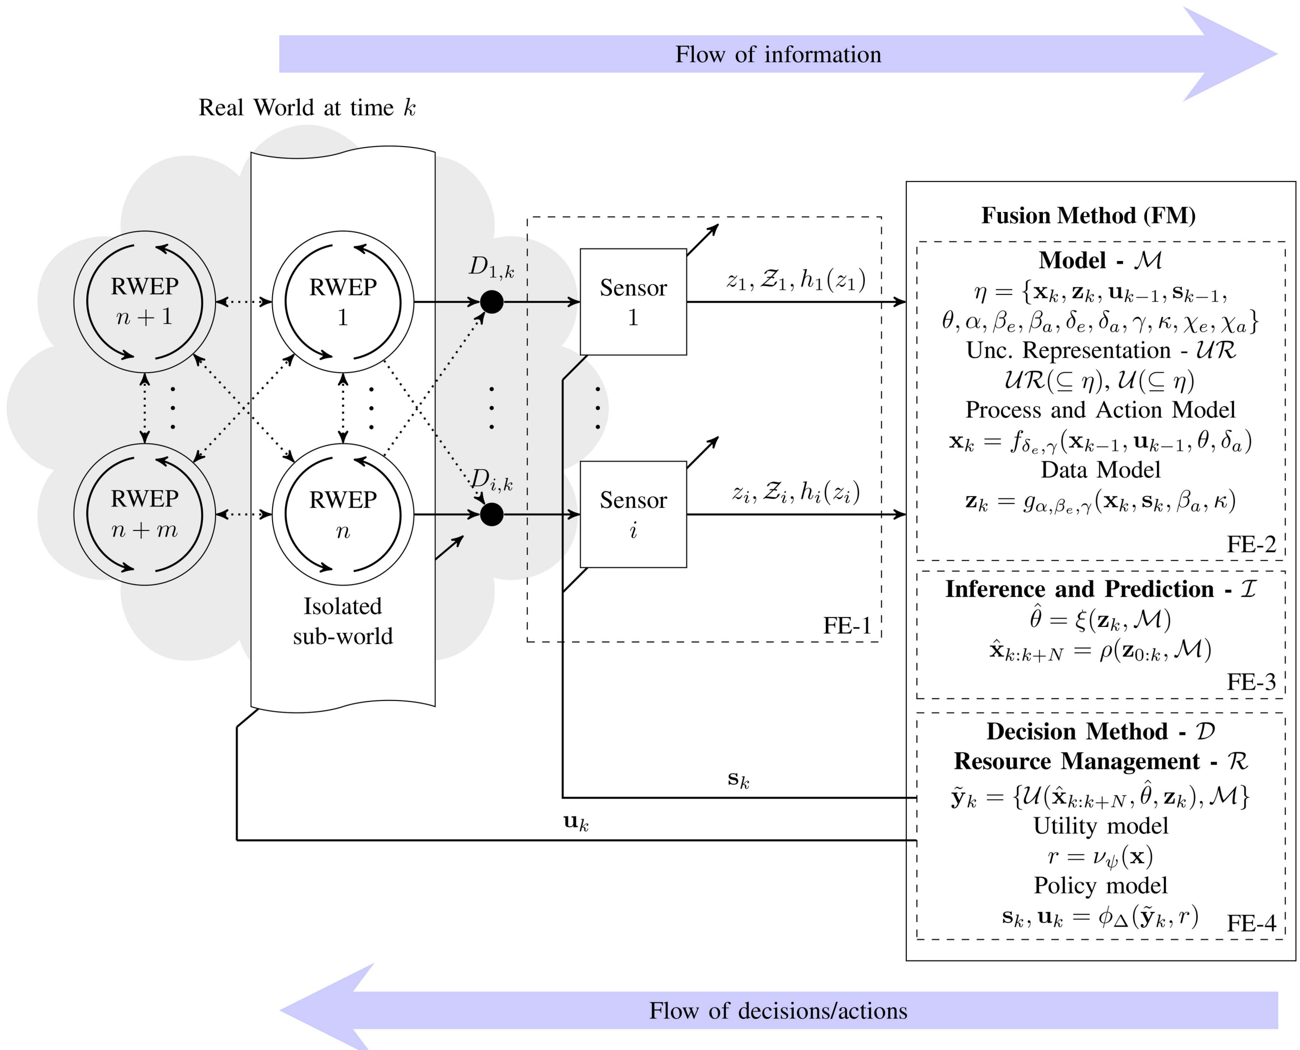

Fig. 4. The fusion decision (e.g., OODA) loop depicting the flow of information through sensors (observe) and the Fusion Method (FM) (orient), and the flow of decisions (decide) and actions (act) out of the decision method and resource management block. These actions in turn influence the real world. Although this figure looks similar to Fig. 3, it has some distinct and important differences. It describes uncertainties that enter the FM during *runtime* (routine operation phase), as opposed to Fig. 3, which describes uncertainties that enter during *modelling* (inception and design phase). The flow of abstraction in Fig. 3 takes place on a large time scale (months/years), whereas the flow information/decisions/actions takes place on a relatively short time scale (seconds or less).

single radar sensor i sensing multiple targets (RWEPs) in an area of regard. Thus, the datum $D_{i,k}$ might be composite and represents the set of observable effects by all RWEPs visible to sensor i. As assumed in [2], this is a generalisation of what is presented in Section III and [1]. Specifically, we let the datum $D_{i,k}$ be conditioned upon $\omega_i \subseteq \{\Omega_{1,k}, \ldots, \Omega_{n,k}\}$ since the observable datum depends on the properties of the physical entities which sensor i can observe. Consequently the datum conditioned upon its physical properties, ω, is written as $\{D_{i,k} \mid \omega_i\}$ or $D_{i,k}$ *given* ω_i.

2) Sensor data:

Consider real word data $\{D_{1,k}, \ldots, D_{n,k}\}$. Measurements are made of $\{D_{1,k}, \ldots, D_{n,k}\}$ by sensors 1 to i and converted into mathematical representations, which not only represent the quantities themselves (z_1 to z_i), but also supplement them with an uncertainty representation $\mathcal{Z}_1$ to $\mathcal{Z}_i$, and associated uncertainty relations $h_1(\cdot)$ to $h_i(\cdot)$. Examples for quantities z_1 to z_i include integers, real numbers, vectors, complex numbers, tensors, norms, logic expressions, etc. Examples for uncertainty representations $\mathcal{Z}_1$ to $\mathcal{Z}_i$ include probabilistic, evidential or fuzzy based representations. Examples of uncertainty relations $h_1(\cdot)$ to $h_i(\cdot)$ include probability density functions, belief functions or fuzzy membership functions An *uncertainty representation* could be defined as a set containing an uncertainty nature (aleatory or epistemic), uncertainty theory (e.g., Bayesian probability theory, evidence theory, fuzzy set theory), an uncertainty model (e.g., Markov model, Bayesian network, Kalman filter), a semantic interpretation (e.g., causality, frequentist), uncertain variables (e.g., random variables, fuzzy variables) and joint uncertainty relations over these variables as described above (e.g,. probability distribution functions, belief functions, fuzzy membership functions).

It is noted that sensors have a broad definition and may include transducers, humans that enter language statements into a computer, and also information from other fusion systems, along the lines of the distributed fusion architectures of [39], [40]. Note that a distinction should be made between uncertainty representations from the sensors, $\mathcal{Z}_1$ to $\mathcal{Z}_i$, which may differ from each other (in the case of heterogeneous sensors) and the uncertainty representation internal to the FM, which is typically common to all variables in the engine.

Relation to the ADPs: The "observe" part of the decision loop may influence the *universe of discourse* elementary construct of the ADP [38], [19] within the modelling phase, as the definition of the universe of discourse for uncertain variable of interest may be guided not only by some design concern fixing the granularity of the problem (i.e., to ensure fast computation) but also by the limitation of the sensors.

B. Orient

According to [21], the orient part of the loop serves "as the repository of our genetic heritage, cultural tradition, and previous experiences." In a semi-autonomous or autonomous fusion system, the orient phase would be the internal model of the fusion system (our understanding of the functioning of the world), which contains representations of RWEPs (process models, agents, rewards and policies), representations of data generation (data/sensor models), representations of quantities in the real world (variables), and a representation of uncertainties, both of the model (epistemic) and of the RWEPs and sensors (aleatory). In addition, the "orient" part of the decision loop also involves making inferences from the sensor data. The orient part of the decision loop corresponds to the FM in Fig. 4. To summarise, the FM contains mathematical models and algorithms for the purpose of data association, data and information fusion, and inference.

In the subsections below, the overarching system model $\mathcal{M}$ is described followed by a discussion on the distinction between physical models and uncertainty models. Uncertain variables and the relations between them are then discussed, followed by the concepts of a composite uncertainty model and second order uncertainty. The process and data models are then considered. The "orient" phase of the decision loop is concluded by a subsection discussing inference and prediction in a fusion system.

1) Fusion System Model:

Considering the FM in more detail, we define the model $\mathcal{M}$ as the overarching fusion system model, which contains several sub models for RWEPs (object models), models for their observation (sub-object models), models for groups of RWEPs (situation models), models for the current and future impact of situations (impact models) and models for agents (process refinement models). Sub-object models correspond to level 0 of the JDL/DFIG taxonomy [22]–[24], [41], object models of level 1, situation models of level 2, impact models of level 3 and process refinement models of level 4.

Inside the FM, the combined sensor measurement vector of all sensors at time k are collected together in a composite variable $\mathbf{z}_k$, which may be an array, vector, set, etc. and their uncertainty relations in the composite variable $\mathbf{h}$. It is important to note that $\mathbf{z}_k$ and $\mathbf{h}$ are distinct from z_1 to z_i and $h_1(\cdot)$ to $h_i(\cdot)$ respectively, since heterogeneous sensor reports may have different uncertainty representations, whereas $\mathbf{z}_k$ and $\mathbf{h}$ would have typically been converted to a single uncertainty representation $\mathcal{UR}$ such that an specific uncertainty calculus can be applied within the FM. The uncertainty of such a conversion is a component of the variable α introduced earlier. This removes the necessity of the uncertain variable ρ in [2], since by definition in Section III-D, it is contained in the uncertain variable α.

2) Physical and uncertainty models:

In the taxonomy of Fig. 4, a distinction is made between *physical models*, which explain RWEPs and the data, and *uncertainty models*, which represent uncertainties that enter into the FM, either during design or during routine operation (runtime). The physical models consist of a process model $f(\cdot)$ and a sensor/data model $g(\cdot)$, which are characterised by uncertainties during modelling, and encompasses several processes of abstraction as explained in [1]. A discussion on the uncertainty representation $\mathcal{UR}$ follows, after which the effect of these uncertainties upon the physical models $f(\cdot)$ and $g(\cdot)$ are discussed. FE-2 refers to the collection of the physical and uncertainty models, i.e., the overarching model $\mathcal{M}$.

3) Uncertainty representation and relations:

Following the definition of [2], consider an explicit set η of all known uncertain variables (see Table II) that represent different types of uncertainty (e.g., in a probabilistic representation, these may be random variables). The uncertainty representation $\mathcal{UR}$ is the internal characterisation of all uncertainty elements of the fusion system (uncertainty natures, theories, relations, semantic interpretations), for a subset of η, i.e., $\mathcal{UR}(\subset \eta)$ since not all sources of uncertainty may be explicitly represented within the fusion system model $\mathcal{M}$. Similarly, uncertainty relations $\mathcal{U}(\cdot)$ (e.g., probability density functions, or belief functions) may be defined for a subset of η, i.e., $\mathcal{U}(\subset \eta)$. For example, in a fusion system implementing Bayesian reasoning, a joint distribution might not be available for *all* random variables, since in a traditional model, many sources of uncertainty are typically omitted.

The notation $\mathcal{U}(\eta)$ indicates as the most general case a *joint uncertainty relation* over all uncertain variables in the FM. An example is a joint probability distribution if the uncertainty representation is probabilistic. At the very least, most traditional Bayesian based fusion systems will represent an uncertainty relation over inputs $\mathbf{x}_k, \mathbf{z}_k, \beta_a$ and δ_a and outputs $\hat{\mathbf{x}}_k, \hat{\theta}$, i.e. $\mathcal{U}(\mathbf{x}_k, \mathbf{z}_k, \beta_a, \delta_a, \hat{\mathbf{x}}_k, \hat{\theta}, \kappa)$.

4) Composite (joint) uncertainty variable—η:

The first, second and third components of η represent the uncertain hidden state $\mathbf{x}_k$ of the process model to be inferred, the uncertain measurement $\mathbf{z}_k$, and the composite process model parameter variable θ respectively. The variables $\{\alpha, \beta, \delta, \gamma\}$ are the abstraction uncertainty variables as defined before, and the subscripts *e* and *a* in Fig. 4 make a distinction between epistemic and aleatory components of the underlying variable. Note that there may be distinct β_e and β_a variables for every sensor, unless the sensors and processes generating the data are identical. Similarly there may be distinct δ_e and δ_a variables for every RWEP of interest and χ_e and χ_a for the actions of agent RWEPs of interest, unless the entities and processes in the real world can be explained using a single model. The variable χ_e represents epistemic uncertainty about how sensor controls $\mathbf{s}_k$ and the mission controls $\mathbf{u}_k$ influence the fusion system and the real world respectively. The variable χ_a represents aleatory uncertainty about how world states evolve because of $\mathbf{s}_k$ and $\mathbf{u}_k$ owing to random effect inherent to the world. The following subsections explain the components of η that follow from the modelling (abstraction) processes. Finally, γ represents association uncertainty. i.e. uncertainty about which RWEP generated which datum $D_{n,k}$.

5) Second order uncertainty:

Although second order uncertainty is not represented explicitly in Fig. 4, this concept warrants a brief discussion. There will be uncertainty about whether the uncertainty representation $\mathcal{UR}$ and its corresponding relation $\mathcal{U}$ adequately represent all the uncertainties listed in Table II. This is a second order uncertainty (uncertainty about uncertainty) and cannot be represented within the model $\mathcal{M}$, since it involves a shortcoming of the uncertainty representation $\mathcal{UR}$.

6) Process model:

Consider the equation for the process model in the FM of Fig. 4. The state evolution of RWEPs is governed by the function $f(\cdot)$. In Fig. 4 the evolution is first order (i.e., the current state $\mathbf{x}_k$ is a function of only the previous state $\mathbf{x}_{k-1}$ and the previous control input $\mathbf{u}_{k-1}$). The Markovian state evolution may be generalised to higher orders if required. The current state is also a function of the uncertain static parameters θ of the sub-world and the aleatory uncertainties associated with the state evolution (e.g., the process noise). Since $f(\cdot)$ is influenced by epistemic uncertainties associated with the model $\mathcal{M}$, the subscripts δ_e and γ in $f_{\delta_e,\gamma}$ indicate that the model is influenced by uncertainties in the abstraction of how RWEPs operate (δ_e), and the abstraction of isolating part of the world (γ). In Fig. 4, $f_{\delta_e,\gamma}$ is not shown to explicitly consider them (i.e., they are not explicitly a function of these epistemic uncertainties), since most typical systems do not; however in a complete model of Fig. 3, they should be considered. Most models typically take aleatory uncertainty δ_a (randomness or noise) in the state evolution equation $f(\cdot)$, and hence $f_{\delta_e,\gamma}$ is a function of δ_a.

7) Data/sensor model:

In Fig. 4, the data/sensor model is given by a function $g(\cdot)$ under the heading "Data Model" in the FM. The measurement or observation vector $\mathbf{z}_k$ at discrete time k, is a function of the hidden state $\mathbf{x}_k$, the sensor control vector[3] $\mathbf{s}_k$, aleatory measurement uncertainty (e.g., sensor noise) β_a and association uncertainty κ. The influences of epistemic uncertainties such as the datum uncertainty α, data/sensor model uncertainty β_e and isolation abstraction uncertainty γ are again not typically considered in most models, unless a full model is used. As such $g_{\alpha,\beta_e,\gamma}(\cdot)$ is not shown to explicitly consider these uncertainties (i.e. it is not shown as a function of them). As with the process model, aleatory measurement uncertainty β_a (for example measurement noise) typically does form part of $g(\cdot)$, and as such, $g_{\alpha,\beta_e,\gamma}(\cdot)$ is a function of β_a.

8) Inference and Prediction:

Models are mathematical representations of reality and uncertainties owing to inherent randomness in reality or incomplete knowledge of humans. These models are used to infer hidden states and parameters that are needed for informed decision making. In Fig 4, inferred or estimated states and parameters of some model $\mathcal{M}$ by an inference engine $\mathcal{I}$ are denoted by $\hat{\mathbf{x}}_k$ and $\hat{\theta}$ respectively, and are obtained by inference procedures such as Bayesian filtering in time varying systems [42]–[44] (i.e., Kalman, particle, or Poisson point process). The parameter inference procedure is denoted by $\xi(\mathbf{z}_k, \mathcal{M})$ and the state inference procedure by $\rho(\mathbf{z}_{0:k}, \mathcal{M})$, where the subscript $0:k$ indicates that all measurements up to time k are used. In the probabilistic case, the outputs of the fusion system are probability distributions, meaning that $\mathcal{U}$ takes the form of a joint probability distribution over system inputs $\mathbf{z}_k$ and outputs $\hat{\mathbf{x}}_k$, $\hat{\theta}$, i.e., $\mathcal{U}(\hat{\mathbf{x}}_k, \hat{\theta}, \mathbf{z}_k)$. This corresponds to the joint uncertainty relations between different inputs, different outputs, and also between inputs and outputs as in [15]. Often state and parameter inference is performed jointly, and as such the functions $\xi(\cdot)$ and $\rho(\cdot)$ are conflated. The inference part of the fusion system, corresponds to FE-3, which are

[3]The sensor control vector $\mathbf{s}_k$ is a set of sensor controls that can change the measurement function $g(\cdot)$. In a networked radar system, $\mathbf{s}_k$ could be a vector of several azimuth and elevation values to steer the beams of multiple radars.

the inference or *reasoning* parts of the atomic decision process.

C. Decide

Inferences may take the form of multiple competing hypotheses of world states and parameters, but in the end a single final decision needs to be made, which balances the costs/rewards/utilities of making a decision with probabilities of certain outcomes. The "Decision Method/Resource Management" block in Fig. 4 represents the balancing of competing decisions, actions and outcomes. The outputs of the inference engine at time k are the inferences about RWEPs, situations and impacts and their uncertainties, and are represented by $\tilde{\mathbf{y}}_k$. This quantity is fed into the decision method $\mathcal{D}$. In a system where the uncertainty representation is frequentist (non-Bayesian) statistics, the decision involves the thresholding of some uncertainty relation to end up with a non-probabilistic estimate of the world states and/or parameters (i.e., a single hypothesis of states and parameters). In the case of using Bayes risk for decisions, the decision method and resource management blocks combine, since $\mathbf{s}_k$ and $\mathbf{u}_k$ are optimised directly such that a utility function is optimised. The decision method is then concerned with balancing the reward/cost of events with the probability of them occurring, for example by maximising the expected reward (or minimising the expected cost). The decision method and its output correspond to FE-4 in the atomic decision process, namely the *decision method and output information*. The model $\mathcal{M}$ will be used to make predictions under different actions $\mathbf{s}_k$ and $\mathbf{u}_k$ with the inferred $\hat{\mathbf{x}}_k$, $\hat{\theta}$ in order to optimise the decision and maximise some utility/reward r, or alternatively minimise some cost/loss function. The utility/reward function $\nu_\psi(\mathbf{x}_k$ is a function which maps a state $\mathbf{x}_k$ to a reward r, and is characterised by the epistemic utility uncertainty ψ. In a reinforcement learning or model predictive control setting, a policy $\phi_\Delta(\cdot)$ would be defined/learned which would maximise the discounted sum of rewards over a (possibly infinite) time horizon. The uncertainty associated with a particular policy representation would be characterised by the epistemic policy uncertainty variable Δ.

The fusion system user might sensibly consider a different abstraction in making a decision to that which was used to provide inferences of the current situation awareness picture. For example, "belief compression" [32] a technique for summarising probability distribution functions (lowering their dimensionality) in the rollout over a sliding window into the future. More generally, there are different requirements placed on the models used here than in the "Orient" part of the decision loop. Therefore, $\mathcal{D}$ and the associated decision mapping $\nu(\cdot)$ may be rooted in a different formalism than the FM. As such another form of uncertainty may be introduced through, for example, dimensionality reductions, which may be easily overlooked.

TABLE II
Table of variables representing currently known forms of uncertainty that enter or exist within the Fusion Method (the elements of η and uncertainties pertaining to the utility and policy models)

Uncertain variable	Description
$\mathbf{x}_k$	State at time k
$\mathbf{z}_k$	Measurement at time k
θ	Time invariant parameters of process model
$\mathbf{s}_k$	Sensor controls at time k with uncertain effect
$\mathbf{u}_k$	Mission controls at time k with uncertain effect
α	Datum abstraction variable (pertaining to quantities, associated uncertainty representations and relations)
β_e	Epistemic data/sensor model variable, representing that the process of generating data is poorly understood (one for each sensor type)
β_a	Aleatory data model variable, for noisiness of the data source (sensor or uncertainty in the way the RWEP generates $D_{n,k}$). Typically one exists for each sensor type and/or mechanism which generates data in the real world)
δ_e	Epistemic process model variable (one per process representation, unless different models for different processes are used)
δ_a	Aleatory process model variable (one per process, unless different models for different processes are used)
γ	Isolation abstraction variable
κ	Association uncertainty variable, capturing uncertainty about which RWEP generated which datum $D_{k,n}$
χ_e	Epistemic action uncertainty variable, capturing uncertainty about how state evolution is modelled because of actions $\mathbf{s}_k$ and $\mathbf{u}_k$
χ_a	Aleatory action uncertainty variable, capturing uncertainty about state evolution because of some inherent random effects of actions $\mathbf{s}_k$ and $\mathbf{u}_k$
ψ	Epistemic utility uncertainty variable, capturing uncertainty about the proper representation of the agent's mapping from a perceived state to a utility or reward
Δ	Epistemic policy uncertainty variable, capturing uncertainty about the proper representation of the agent's mapping from a perceived state to appropriate actions $\mathbf{s}_k$ and $\mathbf{u}_k$ that maximises, for example, the discounted sum of future rewards

D. Act

Once a decision is made, it is converted to action by some resource management function in Fig. 4. It affects controls $\mathbf{s}_k$ over sensors and controls $\mathbf{u}_k$ over missions. As discussed, there will be uncertainty in how decisions and actions will influence RWEPs in the real world. These are represented by χ_e and χ_a and are considered in the PM, which models world state evolution. It should be noted that although the information fusion system (including the sensors) is explicitly indicated in Fig. 4 as being separate from the real world, this is not actually the case. In a real setting, the fusion system is part of the real world. However, in the presented formulation, it is assumed that the fusion system affects the real world only through the quantities $\mathbf{s}_k$ and $\mathbf{u}_k$, and that all other effects are deemed to be negligible. Whether this is the case depends on the accuracy of the understanding of the effect of $\mathbf{s}_k$ and $\mathbf{u}_k$ on the real world in the model $\mathcal{M}$, the decision method $\mathcal{D}$ and resource management function $\mathcal{R}$, through an understanding of δ_e and χ_e.

V. EXAMPLE USE CASES

For the sake of brevity, a single example use case is presented (the same as in [2]), which demonstrates the fusion uncertainty evaluation taxonomy presented here. RWEPs represent aircaft that can be sensed by a network of radars (for example as in [20] and [45]). The radars are intelligent sensors, in that they already provide processed information to the fusion system in the form of target tracks and associated filtering covariances. Consequently, the FM combines the tracks from several radars to result in one fused track for each target, all contained within the joint inferred state vector $\hat{\mathbf{x}}_k$. This vector and its associated uncertainty support is used in the decision method and resource management functional blocks to a) to search an area and detect targets, b) balance the search requirement with the requirement to direct the radars through $\mathbf{s}_k$ to minimise (for example) the sum of covariances of all existing tracks and c) to decide and communicate through $\mathbf{u}_k$ whether to scramble fighters to intercept targets deemed to be serious threats based on some cost/benefit analysis. The reader can consult [2] for an additional anti-rhino poaching use case example. The example (captured in Table III) should hopefully be self explanatory, but for a brief description, the reader can consult [2].

VI. EVALUATION USING THE URREF ONTOLOGY

The Uncertainty Representation and Reasoning Evaluation Framework (URREF) includes an ontology, the URREF ontology, that captures primary and secondary concepts related to uncertainty representation and reasoning in information fusion systems, the criteria for their evaluation, as well as the links between the concepts.[4] One of the main objectives of the URREF ontology is to define and articulate the criteria which enable the systematic *reasoning about* and *evaluation of* uncertainty *representation* (instantiated or theoretical, for example a specific probability distribution or the underlying uncertainty formalism e.g., probability, belief based representations, fuzzy representations) and *reasoning* (inference in general e.g., Bayes' rule, Dempster's combination rule) in information fusion systems. These are the *primary subjects* of evaluation [19]. The

[4]The latest version of the ontology can be viewed at the webpage with the following URL: http://eturwg.c4i.gmu.edu/?q=URREFv3. The OWL file of the URREF ontology can be opened using the free, open-source ontology software "protégé."

TABLE III
Table of symbols together with examples from multi-sensor multi-target tracking with track fusion use case.

Symbol	Example
RWEP	An aircraft that can be sensed by radars
Isolated sub-world	Area that is within range of radar network
$D_{i,k}$	All EM returns at time k from targets sensed by radar i
Sensor i	The ith radar in a network of air surveillance radars
$\Omega_{n,k}$	Dynamical characteristics (mass, powerplant, airfoil etc.) of the nth aircraft
ω_i	Dynamical characteristics (mass, powerplant, airfoil etc.) of all aircraft, as well as dynamical characteristics owing to interactions between aircraft, all observed by sensor i
z_i	All radar tracks at time k from radar i
$\mathcal{Z}_i$	Bayesian probability (sensors), Fuzzy natural language (human report)
$h_i(\cdot)$	Probability density function of filtering densities parameterised by means and covariances
$\mathbf{x}_k$	Combined state of all targets after track fusion
$\mathbf{z}_k$	Combined state vectors of all tracks before fusion
$f_{\delta_e,\gamma}(\cdot)$	Almost constant velocity dynamical model
$g_{\alpha,\beta_e,\gamma}(\cdot)$	Gaussian filtering probability densities for radar tracks
ρ	N/A, since $\mathcal{Z}$ and $\mathcal{UR}$ are both probabilistic
$\mathbf{u}_k$	Message to fighter to intercept target
$\mathbf{s}_k$	Message to increase scan rate of a radar
θ	New track density
α	Uncertainty associated with quantisation error in radar digital to analog converter
β_e	Uncertainy owing to Gaussian approximation of measurement noise in rectangular coordinates
β_a	Measurement noise
δ_e	Uncertainy owing to Gaussian approximation of plant noise to represent target manoeuvres
δ_a	Plant noise
γ	Uncertainty owing to ignoring targets out of range of the radar network
$\mathcal{UR}$	Bayesian probabilistic representation
$\mathcal{U}$	Probability distribution
$\hat{\theta}$	Inferred new track density
$\hat{\mathbf{x}}_{k:N}$	Inferred distribution of the states of all targets after fusion at time k, and state distribution predictions from time $k+1$ up to a future horizon of $k+N$
$\nu_\psi(\cdot)$	A mapping from a perceived state to a utility/reward. In a target tracking system, this could be the reciprocal of the sum of track covariances.
$\phi_\Delta(\cdot)$	A mapping from a perceived state and predicted future states to actions. In the case of a target tracking example, this could be a function which defines the amount of time spent by radars on scanning as opposed to tracking, given the sum of track covariances and the recency of scan coverage of an area. This would be to balance current and future track accuracy as opposed to detecting possibly undetected targets at a time and into the future.

primary subjects cannot stand on their own, and as such, the evaluation of *secondary subjects* is also catered for in the URREF ontology. The secondary subjects are defined as the source of information (sensors), the piece of information (sensor output), the fusion method (implemented by the fusion algorithm) and the mathematical model (the process and sensor/data model, both represented by $\mathcal{M}$).

A. FE-1 (sources of information)

Sources (sensors) that produce information, whether they are humans or transducers should be evaluated according to source criteria. These are secondary subjects of evaluation, and fall under *DataCriterion* the current view of the ontology, with the relevant subclasses being *Quality* (specifically relating to source quality distinct from information quality) and *Credibility*. Note that since in this paper FE-1 to FE-4 replace ADP-1 to ADP-4 that was presented [2], the criteria specified are different.

B. FE-2 (input information and model)

Here the information criteria are relevant for the input information, and representation criteria are relevant for the model $\mathcal{M}$ and uncertainty representation $\mathcal{UR}$ and associated uncertainty relations $\mathcal{U}$. In the URREF ontology the information criteria are under the classes *DataCriterion* and *DataHandlingCriterion*. The associated subclasses can be used to evaluate the input information. *Quality* can also be used, but here relate to information quality as opposed to source quality used in FE-1. The model $\mathcal{M}$ and uncertainty representation $\mathcal{UR}$ and associated uncertainty relations $\mathcal{U}$ can be evaluated using the class *RepresentationCriterion* and all the associated subclasses.

C. FE-3 (reasoning and combined information)

This element of the FM (the inference engine $\mathcal{I}$ in Fig. 4), is evaluated according to *ReasoningCriteria*, which consist of *ComputationalCost*, *Scalability*, *Performance* and *Consistency*. The output of the reasoning component (or inference engine) can again be evaluated according to the *DataCriteria*, as with the input of the FM. The output of a FM may form the input of another FM in the case of distributed fusion.

D. FE-4 (decision method and output information)

The uncertainty about the effect of actions $\mathbf{s}_k$ and $\mathbf{u}_k$ on the real world in the model $\mathcal{M}$ is a form of epistemic process abstraction uncertainty, represented by χ_e. It reduces the optimality of the decision process. This is epistemic uncertainty may be evaluated according to *RepresentationCriteria*, and is the uncertainty owing to imperfect modelling contained in the model $\mathcal{M}$. Furthermore, the decision process is a form of reasoning (through optimisation), and can be therefore be evaluated according to *ReasoningCriteria*. Maximising the expected utility combines uncertainty with utility, and the utility part carries an element of subjectivity related to a desired outcome. In many cases, a desired outcome is the combination of conflicting and competing objectives with relative weightings. Therefore, some *DataCriteria* such as *Objectivity*, *RelevanceToProblem* and *WeightOfEvidence* may be used.

VII. CONCLUSIONS

In this paper, the flow of abstraction in fusion system inception, design and implementation is contrasted to the flow of information and the flow of decisions/actions during the routine operation of a fusion system. Without a complete list of uncertainties that enter during these two phases of the fusion system life cycle, the fusion system practitioner might not consider the implications of certain design choices relating to chosen variables of interest, uncertainty representations, reasoning formalisms, and simplifying assumptions. As mentioned in [3], engineers and system designers are biased towards a default uncertainty representation or reasoning methods, namely the methods they know and are comfortable with. As such, the cost for them to learn new formalisms that could possibly be better suited to a particular application should also be evaluated. Consulting a list of explicit uncertainty types that are a result of fusion system development and routine operation, would minimise errors of omission and oversight, and simplifying assumptions and design choices can be properly characterised.

REFERENCES

[1] J. P. de Villiers, K. Laskey, A.-L. Jousselme, E. Blasch, A. D. Waal, G. Pavlin, and P. Costa
"Uncertainty representation, quantification and evaluation for data and information fusion,"
in *Information Fusion (FUSION), 18th International Conference on*, July 2015, pp. 50–57.

[2] J. P. de Villiers, A.-L. Jousselme, A. de Waal, G. Pavlin, K. Laskey, E. Blasch, and P. Costa
"Uncertainty evaluation of data and information fusion within the context of the decision loop,"
in *Information Fusion (FUSION), 19th International Conference on*, July 2016, pp. 766–773.

[3] G. Pavlin, A.-L. Jousselme, J. de Villiers, P. Costa, and P. de Oude
"Towards the rational development and evaluation of complex fusion systems: a URREF-driven approach,"
in *Information Fusion (FUSION), 21st International Conference on*, July 2018.

[4] P. C. G. Costa, K. B. Laskey, E. Blasch, and A.-L. Jousselme
"Towards unbiased evaluation of uncertainty reasoning: The URREF ontology,"
in *Information Fusion (FUSION), 15th International Conference on*, July 2012, pp. 2301–2308.

[5] A. D. Kiureghian and O. Ditlevsen
"Aleatory or epistemic? does it matter?"
Structural Safety, vol. 31, no. 2, pp. 105–112, 2009, risk Acceptance and Risk Communication Risk Acceptance and Risk Communication. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167473008000556.

[6] H. G. Matthies
"Quantifying uncertainty: Modern computational representation of probability and applications,"
in *Extreme Man-Made and Natural Hazards in Dynamics of Structures*, ser. NATO Security through Science Series, A. Ibrahimbegovic and I. Kozar, Eds. Springer Netherlands, 2007, pp. 105–135. [Online]. Available: http://dx.doi.org/10.1007/978-1-4020-5656-7_4.

[7] G. Lin, D. Engel, and P. Eslinger
"Survey and evaluate uncertainty quantification methods,"
US Department of Energy, Internal report, February 2012.

[8] M. C. Kennedy and A. O'Hagan
"Bayesian calibration of computer models,"
Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 63, no. 3, pp. 425–464, 2001. [Online]. Available: http://dx.doi.org/10.1111/1467-9868.00294.

[9] A. Abusam, K. Keesman, and G. van Straten
"Forward and backward uncertainty propagation: an oxidation ditch modelling example,"
Water Research, vol. 37, no. 2, pp. 429–435, 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0043135402002889.

[10] P. D. Arendt, D. W. Apley, and W. Chen
"Quantification of model uncertainty: Calibration, model discrepancy, and identifiability,"
Journal of Mechanical Design, vol. 134, no. 10, p. 100908, 2012.

[11] J. de Villiers, G. Pavlin, P. Costa, K. Laskey, and A.-L. Jousselme
"A URREF interpretation of bayesian network information fusion,"
in *Information Fusion (FUSION), 17th International Conference on*, July 2014, pp. 1–8.

[12] E. Blasch, A. Josang, J. Dezert, P. Costa, and A.-L. Jousselme
"URREF self-confidence in information fusion trust,"
in *Information Fusion (FUSION), 17th International Conference on*, July 2014, pp. 1–8.

[13] J. Ziegler and F. Detje
"Application of empirical methodology to evaluate information fusion approaches,"
in *Information Fusion (FUSION), 16th International Conference on*, July 2013, pp. 1878–1885.

[14] E. Blasch, K. Laskey, A.-L. Jousselme, V. Dragos, P. Costa, and J. Dezert
"URREF reliability versus credibility in information fusion (STANAG 2511),"
in *Information Fusion (FUSION), 16th International Conference on*, July 2013, pp. 1600–1607.

[15] A.-L. Jousselme and G. Pallotta
"Dissecting uncertainty-based fusion techniques for maritime anomaly detection,"
in *Information Fusion (FUSION), 8th International Conference on*, 2015, pp. 34–41. [Online]. Available: http://ieeexplore.ieee.org/document/7266541/.

[16] A.-L. Jousselme
"Semantic criteria for the assessment of uncertainty handling fusion models,"
in *Information Fusion (FUSION), 19th International Conference on*, July 2016, pp. 488–495.

[17] A. de Waal, H. Koen, P. de Villiers, H. Roodt, N. Moorosi, and G. Pavlin
"Construction and evaluation of Bayesian networks with expert-defined latent variables,"
in *Information Fusion (FUSION), 19th International Conference on*. IEEE, 2016, pp. 774–781.

[18] E. P. Blasch, D. A. Lambert, P. Valin, M. M. Kokar, J. Llinas, S. Das, C. Chong, and E. Shahbazian
"High level information fusion (HLIF): Survey of models, issues, and grand challenges,"
IEEE Aerospace and Electronic Systems Magazine, vol. 27, no. 9, pp. 4–20, Sep. 2012.

[19] A.-L. Jousselme, J. de Villiers, G. Pavlin, P. C. G. Costa, K. Laskey, V. Dragos, and E. Blasch
"Subjects under evaluation with the URREF ontology,"
in *Information Fusion (FUSION), 20th International Conference on*, July 2017, pp. 1–8.

[20] R. W. Focke
"Investigating the use of interval algebra to schedule mechanically steered multistatic radars,"
Ph.D. dissertation, University of Cape Town, 2015.

[21] F. P. Osinga
Science, strategy and war: The strategic theory of John Boyd.
Routledge, 2007.

[22] A. N. Steinberg, C. L. Bowman, and F. E. White
"Revisions to the JDL data fusion model,"
B. V. Dasarathy, Ed., vol. 3719, no. 1. SPIE, 1999, pp. 430–441.

[23] J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White
"Revisiting the JDL Data Fusion Model II,"
in *Information Fusion (FUSION), 7th International Conference on*, Jun. 2004.

[24] E. Blasch, A. Steinberg, S. Das, J. Llinas, C. Chong, O. Kessler, E. Waltz, and F. White
"Revisiting the JDL model for information exploitation,"
in *Information Fusion (FUSION), 16th International Conference on*. IEEE, 2013, pp. 129–136.

[25] E. P. Blasch and S. A. Israel
"Situation/threat context assessment,"
in *Information Fusion (FUSION), 18th International Conference on*. IEEE, 2015, pp. 1168–1175.

[26] H. Kwakernaak and R. Sivan
Linear optimal control systems.
Wiley-Interscience New York, 1972, vol. 1.

[27] K. J. Åström
Introduction to stochastic control theory.
Courier Corporation, 2012.

[28] R. S. Sutton and A. G. Barto
Reinforcement learning: An introduction.
MIT press, 2018.

[29] R. Bellman
"A markovian decision process,"
Journal of Mathematics and Mechanics, pp. 679–684, 1957.

[30] K. J. Astrom
"Optimal control of markov decision processes with incomplete state estimation,"
Journal of mathematical analysis and applications, vol. 10, pp. 174–205, 1965.

[31] E. F. Camacho and C. B. Alba
Model predictive control.
Springer Science & Business Media, 2013.

[32] N. Roy, G. Gordon, and S. Thrun
"Finding approximate POMDP solutions through belief compression,"
Journal of artificial intelligence research, vol. 23, pp. 1–40, 2005.

[33] J. Kubalík, E. Alibekov, and R. Babuška
"Optimal control via reinforcement learning with symbolic policy approximation,"
IFAC-PapersOnLine, vol. 50, no. 1, pp. 4162–4167, 2017.

[34] C. Amato and F. A. Oliehoek
"Scalable planning and learning for multiagent POMDPs."
in *AAAI*, 2015, pp. 1995–2002.

[35] S. Brechtel, T. Gindele, and R. Dillmann
"Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs,"
in *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*. IEEE, 2014, pp. 392–399.

[36] N. J. Higham
Accuracy and Stability of Numerical Algorithms, 2nd ed.
Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2002.

[37] R. P. S. Mahler
Statistical Multisource-Multitarget Information Fusion.
Norwood, MA, USA: Artech House, Inc., 2007.

[38] A.-L. Jousselme and P. Maupin
"A brief survey of comparative elements for uncertainty calculi and decision procedures assessment,"
in *Information Fusion (FUSION), 15th International Conference on*, 2012, panel Uncertainty Evaluation: Current Status and Major Challenges.

[39] M. E. Liggins, C.-Y. Chong, I. Kadar, M. G. Alford, V. Vannicola, S. Thomopoulos et al.
"Distributed fusion architectures and algorithms for target tracking,"
Proceedings of the IEEE, vol. 85, no. 1, pp. 95–107, 1997.

[40] J. Llinas
"Assessing the performance of multisensor fusion processes,"
in *Handbook of Multisensor Data Fusion: theory and practice*, M. Liggins II, D. Hall, and J. Llinas, Eds. CRC press, 2008, pp. 655–675.

[41] E. Blasch
"One decade of the data fusion information group (DFIG) model,"
pp. 94 990L–94 990L–10, 2015. [Online]. Available: http://dx.doi.org/10.1117/12.2176934.

[42] R. E. Kalman
"A new approach to linear filtering and prediction problems,"
Transactions of the ASME—Journal of Basic Engineering, vol. 82, no. Series D, pp. 35–45, 1960.

[43] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp
"A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,"
IEEE Transactions on Signal Processing, vol. 50, no. 2, pp. 174–188, Feb 2002.

[44] R. Streit, C. Degen, and W. Koch
"The pointillist family of multitarget tracking filters,"
arXiv preprint arXiv:1505.08000, 2015.

[45] R. W. Focke, J. P. de Villiers, and M. R. Inggs
"Interval algebra—an effective means of scheduling surveillance radar networks,"
Information Fusion, vol. 23, pp. 81–98, 2015.

Pieter de Villiers is an associate professor at the University of Pretoria, South Africa and was a principal researcher at the Council for Scientific and Industrial Research (CSIR) until October 2017. He obtained his Bachelors and Masters degrees at the University of Pretoria, South Africa, and a PhD in 2008 at the University of Cambridge, UK, in statistical signal processing (particle filtering) under the supervision of Simon Godsill. From 2010 until 2018 he was performing research into data fusion at the Radar and Electronic Warfare competency at the CSIR. His research interests include data fusion, target tracking, Bayesian inference, nonlinear filtering, pattern recognition, graphical models and machine learning. Pieter has been regularly attending the International Conference of Information Fusion since 2010 and his ISIF activities include membership of the technical program committees, a tutorial selection committee and acting as session chairs over the years. He is the general co-chair for the 23rd International Conference on Information Fusion to be held in 2020 in South Africa. Pieter is a member of the official ISIF Evaluation Techniques for Uncertainty Representation and Reasoning Working Group (ETURWG). He is also a guest editor for a special issue at the Journal of Advances in Information Fusion (JAIF).

Gregor Pavlin received the M. Sc. degree in theoretical engineering and the Ph.D. degree in computer science from Graz University of Technology, Austria in 1995 and 2001, respectively. He has extensive industrial experience in safety critical software systems as well as complex AI-driven solutions. His current research interests are (i) robust algorithms and architectures supporting distributed probabilistic AI, (ii) machine learning and (iii) interoperability in complex service oriented processing systems. Since 2006 he has been a senior researcher and project manager at a corporate research lab of the Thales Group in Delft, the Netherlands. Between 2006 and 2015, he was also a part-time visiting researcher at the Intelligent Autonomous Systems lab, University of Amsterdam. He also has an extensive experience with the coordination of European and national collaborative projects. He served in the organizing committee of the Fifth International Symposium on Intelligent Distributed Computing in Delft (IDC 2011) and is also a member of the organizing committee of the 23rd International Conference on Information Fusion to be held in South Africa in 2020. He is also a member of the official ISIF Evaluation Techniques for Uncertainty Representation and Reasoning Working Group (ETURWG).

Anne-Laure Jousselme received her PhD degree from the Electrical Engineering Department of Laval University in Quebec City (Canada) and the Institut National Polytechnique de Grenoble (France) in 1997. Formely with Defense Research and Development Canada (DRDC), she is now with the NATO STO Centre for Maritime Research and Experimentation (CMRE) in La Spezia (Italy), where she conducts research activities on reasoning under uncertainty, high-level and hard & soft information fusion, information quality assessment and serious gaming applied to maritime situational awareness and anomaly detection. She is area editor of the International Journal of Approximate Reasoning and associate editor of the Perspectives on Information Fusion magazine. She is a member of the Boards of Directors of the International Society of Information Fusion (ISIF) where she serves as VP membership and of the Belief Functions and Applications Society (BFAS) where she serves as Secretary. She serves on program committees of the International Conference of Information Fusion (FUSION) and the International Conference on Belief Functions (BELIEF). She was Tutorial Chair of FUSION 2007 in Quebec City (CA), International Co-chair of FUSION 2015 in Washington and Technical Co-chair of FUSION 2019 in Ottawa (CA). She was general Chair of the Canadian Tracking and Fusion Conference (CTFG) in 2014 in Ottawa (CA) and Local Organizer of the International Conference of Scalable Uncertainty Management (SUM) in 2015 in Quebec City (CA).

Simon Maskell is a Professor of Autonomous Systems at the University of Liverpool. His research interests centre on using Bayesian inference to address real-world challenges across a range of different applications. Often, doing so ends up demanding the development of advanced numerical Bayesian methods, e.g., novel extensions of particle filters. Many of the projects Simon works on relate to conventional applications of data fusion (e.g., increasing the ability of track-before-detect to improve the detection performance of a radar and highlighting anomalous long-term behaviours in wide area motion imagery). However, Simon is also currently engaged in projects that, for example, involve extracted data using machine learning applied to textual records and thereby identifying side effects that result from interactions between drugs.

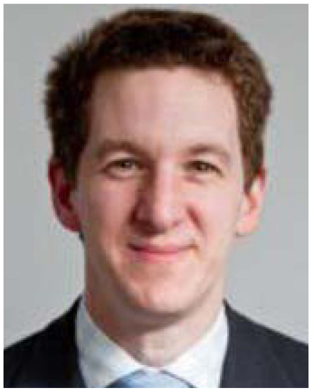

Prior to starting at the University of Liverpool in 2013, Simon worked for 13 years for the defence and security company, QinetiQ. Simon has also been a non-executive director for an insurance company and is currently an honorary employee of a mental health trust. For many years, Simon has also worked closely with UK government and, in particular, Dstl (the UK Ministry of Defence's research lab): Simon understands industrial, academic and governmental perspectives on Fusion.

Simon was the general chair for Fusion 2010 (in Edinburgh) and was, more recently, co-chair for Fusion 2018 (in Cambridge). Simon has also been on the organisational committee of numerous Fusion conferences, served as a member of ISIF's Board of Directors (BoD) from 2009 to 2012 and is now serving on the ISIF BoD from 2019–2022.

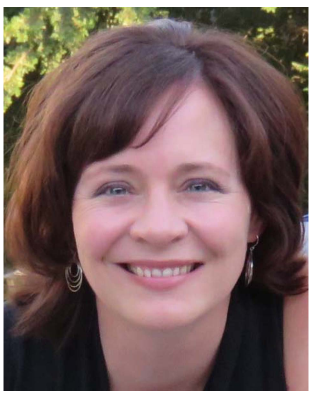

Alta de Waal currently holds a senior lecturer position in the Department of Statistics, University of Pretoria. Previously she was employed as a senior researcher at the Council for Scientific and Industrial Research (CSIR). She received her MSc (Mathematical Statistics) cum laude from the University of the Free State, and her PhD (Engineering Science) from the North West University in South Africa. Alta is a principal researcher at the Center for Artificial Intelligence Research (CAIR), South Africa.

Her research focus is in the field of expert knowledge and Bayesian network modelling as well as statistical pattern recognition with special focus on unsupervised modelling of natural language.

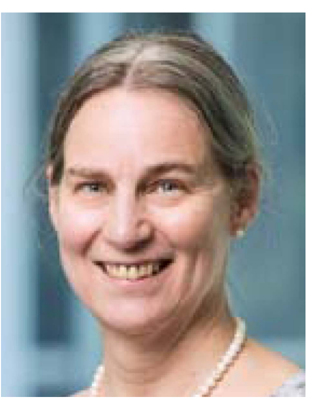

Kathryn Blackmond Laskey, Ph.D., is Professor of Systems Engineering and Operations Research at George Mason University and Associate Director of the Center of Excellence in Command, Control, Communications, Computing and Intelligence (C4I Center). She teaches and performs research on multisource information fusion, decision theoretic knowledge representation and reasoning methodology, data analytics, and decision support. A major focus of her research has been knowledge representation and reasoning for higher level multi-source fusion to support situation awareness and decision support. She has performed research in diverse application areas, including modeling the emplacement of improvised explosive devices, detecting insider threats, predicting aircraft delays, managing terrorist risk at public facilities, and planning military engagements. Dr. Laskey developed multi-entity Bayesian networks (MEBN), a language and logic that extends classical first-order logic to support probability. She was a key contributor to the development of the PR-OWL language for representing uncertainty in OWL ontologies. She serves on the ISIF Board of Directors and has is co-founder and active participant in the ISIF Evaluation of Techniques for Uncertainty Management Working Group (ETURWG). She serves on the Board of Directors of the Washington Metropolitan Area chapter of INCOSE and is past board chair of the Association for Uncertainty in Artificial Intelligence. Dr. Laskey served on several boards and committees of the United States National Academy of Sciences.

Erik Blasch is a program officer at the Air Force Research Laboratory (AFRL)—Air Force Office of Scientific Research (AFOSR) in Arlington, VA. Previously he was he was a principal scientist at the AFRL Information Directorate in Rome, NY, USA (2012–2017), exchange scientist to the Defence Research and Development Canada (DRDC) in Valcartier, Quebec (2010–2012), and Information Fusion Evaluation Tech Lead for the AFRL Sensors Directorate—COMprehensive Performance Assessment of Sensor Exploitation (COMPASE) in Dayton, OH (2000–2009). Additional assignments include Col from 28-year USAF Reserve Officer. He was an adjunct associate professor in Electrical and Biomedical Engineering (2000–2010) at Wright State University and the Air Force Institute of Technology (AFIT) teaching classes in signal processing, electronics, and information fusion as well as research adjunct appointments at the Univ. of Dayton (2001–2014), Binghamton University (2012–2017), and Rochester Institute of Technology (2015–2017).

Dr. Blasch was a founding member of the International Society of Information Fusion (ISIF), (www.isif.org), 2007 President, and Board of Governors (2000–2010). He served on the IEEE Aerospace and Electronics Systems Society (AESS) Board of Governors (2011–2016), distinguished lecturer (2012–2018), co-chair of 5 conferences, and associate editor of 3 academic journals. He has focused on information fusion, target tracking, robotics, and pattern recognition research compiling 800+ scientific papers and book chapters. He holds 25 patents, received 33 team-robotics awards, presented 60+ tutorials, and provided 9 plenary talks. His coauthored books include High-Level Information Fusion Management and Systems Design (Artech House, 2012), Context-enhanced Information Fusion (Springer, 2016), Multispectral Image Fusion and Colorization (SPIE, 2018), and Handbook of Dynamic Data Driven Applications Systems (Springer 2018).

Dr. Blasch received his B.S. in Mechanical Engineering from the Massachusetts Institute of Technology ('92) and Masters' Degrees in Mechanical ('94), Health Science ('95) and Industrial Engineering (Human Factors) ('95) from Georgia Tech and attended the University of Wisconsin for a MD/PhD Neuroscience/Mechanical Engineering until being call to military service in 1996 to the United States Air Force. He completed an MBA ('98), MS Econ ('99), and PhD ('99) in Electrical Engineering from Wright State University and is a graduate of Air War College ('08). He is the recipient of the IEEE Bioengineering Award (Russ-2008), IEEE AESS magazine best paper Award (Mimno-2012), Military Sensing Symposium leadership in Data Fusion Award (Mignogna-2014), Fulbright scholar selection (2017), and 15 research/technical and team awards from AFRL. He is an American Institute of Aeronautics and Astronautics (AIAA) Associate Fellow, Society of Photonics and Industrial Engineers (SPIE) Fellow, and Institute of Electrical and Electronics Engineers (IEEE) Fellow.

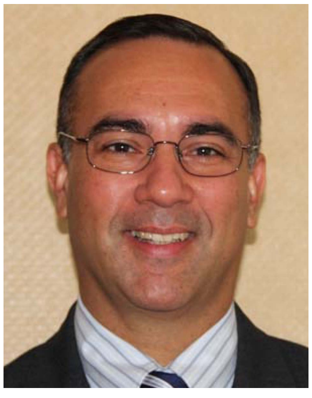

Paulo C. G. Costa is an Associate Professor of Systems Engineering and Operations Research at George Mason University, and Associate Director of the C4I & Cyber Center's Radio and Radar Engineering Laboratory.

His teaching and research interests comprise the areas of probabilistic ontologies, multi-sensor information fusion, Bayesian reasoning, predictive analysis, cybersecurity and Decision Theory.

He is a former fighter pilot with extensive experience in tactical and operational planning, and an expert in requirements engineering for complex systems, such as intelligent transportation and health-care support systems. Dr. Costa is an IEEE senior member, a member of the International Council of Systems Engineering, and currently serves as President of the International Society for Information Fusion (ISIF—tenures 2019 and 2020).

Tracking uncertainty propagation from model to formalization: Illustration on trust assessment

VALENTINA DRAGOS
JEAN DEZERT
KELLYN REIN

This paper investigates the use of the URREF ontology to characterize and track uncertainties arising within the modeling and formalization phases. Estimation of trust in reported information, a real-world problem of interest to practitioners in the field of security, was adopted for illustration purposes. A functional model of trust was developed to describe the analysis of reported information, and it was implemented with belief functions. When assessing trust in reported information, the uncertainty arises not only from the quality of sources or information content, but also due to the inability of models to capture the complex chain of interactions leading to the final outcome and to constraints imposed by the representation formalism. A primary goal of this work is to separate known approximations, imperfections and inaccuracies from potential errors, while explicitly tracking the uncertainty from the modeling to the formalization phases. A secondary goal is to illustrate how criteria of the URREF ontology can offer a basis for analyzing performances of fusion systems at early stages, ahead of implementation. Ideally, since uncertainty analysis runs dynamically, it can use the existence or absence of observed states and processes inducing uncertainty to adjust the tradeoff between precision and performance of systems on-the-fly.

Manuscript received January 17, 2018; revised October 2, 2018; released for publication March 20, 2019.

Refereeing of this contribution was handled by Anne-Laure Jousselme.

Authors' addresses: V. Dragos, ONERA—The French Aerospace Lab, Chemin de la Hunière, Palaiseau, France (E-mail: valentina.dragos@onera.fr). J. Dezert, ONERA—The French Aerospace Lab, Chemin de la Hunière, Palaiseau, France (E-mail: jean-dezert@onera.fr). K. Rein, Fraunhofer FKIE, Wachtberg, Germany (E-mail: kellyn.rein@fraunhofer.fkie.de).

I. INTRODUCTION

A key element when designing information fusion systems is the way the system designer isolates and analyzes real world phenomena. A model is abstracted into a simpler representation, in which components, modules, interactions, relationships and data flows are easier to express. Uncertainty tracking highlights approximations induced by model construction and its formalization, as well as providing a checklist to ensure that all uncertainty factors have been identified and considered ahead of system implementation.

This paper illustrates the use of the uncertainty representation and reasoning framework (URREF) ontology [13] to identify and assess uncertainties arising during the modeling and formalization phases of an information fusion system intended to estimate trust in reported information.

Trust assessment is a real-world problem grounded in many applications relying on reported items, with different persons observing and then reporting on objects, individuals, actions or events. For such contexts, using inaccurate, incomplete or distorted items can result in unfortunate consequences and analysts need to ensure the consistency of reported information by collecting multiple items from several sources.

From the perspective of an information analyst, trust can be analyzed along two dimensions: the subjective evaluation of items reported by the source itself, called self-confidence, and the evaluation of source by the analyst, called reliability. While self-confidence encompasses features of subjectivity, the reliability of a source is related to the quality of previously reported items, the competence of the source for specific topics, and the source's capacity for misleading intentions. Trust estimation aims at capturing, in an aggregated value, the combined effects of self-confidence and reliability on the perceived quality of information. The model is represented with belief functions, a formalism which offers a sound mathematical basis to implement fusion operators which estimate trust by combining self-confidence and reliability.

The model developed for trust assessment focuses on the global characterization of information and provides a better understanding of how trust is to be estimated from various dimensions. The overall process has humans as a central element in both the production and the analysis of information.

Trust in reported information offers a good illustration for tracking uncertainty: the phenomenon is complex, so any model adopted is generally a simplification of the real world interactions. Uncertainties can be made explicit not only for static elements of the model, such as sources or items, but also for the dynamic processes of combining items with one another. Moreover, adopting belief functions as representation formalism will have an impact on the way an information system could be implemented and on the accuracy of its results.

The contribution of this paper is twofold: first, it presents a trust estimation model which combines the reliability of sources and self-confidence of reported items, and, second, the paper analyzes types of uncertainty occurring during modeling and formalization by relating elements of the model to uncertainty criteria defined by the URREF ontology.

The remainder of this paper is divided into 8 sections: section II discusses related approaches for trust modeling and uncertainty assessment. The problem tackled in this paper in presented in section III. Section IV describes the model developed for trust estimation, while its implementation with belief functions is presented in section V. The analysis of uncertainty is discussed in VI, while examples and scenarios for trust assessment are presented in section VII. Strengths and limitations of belief-based formalization are discussed in section VIII and section IX concludes this paper.

II. RELATED APPROACHES

The work presented in this paper is related to approaches for trust modeling and assessment as well as solutions for uncertainty analysis for information fusion systems. Trust modeling is not a new research topic; it spans diverse areas such as agent systems [30] and logical modeling and argumentation [50]. The Internet and social media offer new application contexts for trust assessment; this topic is addressed in relation to service provision on the Internet [36], social networks analysis [57], and crowdsourcing applications [64]. Trust analysis is also of interest in the military field where techniques have been developed in order to identify clues of veracity in interview statements [63].

The concept of trust in these communities varies in how it is represented, computed and used. Although having an obvious social dimension, trust is not only understood with regard to other humans, but also towards information pieces [64], information sources [44], Internet sites [21], algorithms for data and knowledge fusion [20], intelligent agents [30], and services for the Internet of things [31].

While definitions of trust vary from one domain to another, there are some common elements. The first commonality for all research areas cited above is to consider trust as a user-centric notion that needs to be addressed in integrated human-machine environments which rely heavily on information collected by humans, even if further processing can be executed automatically. Moreover, all definitions associate some degree of uncertainty with trust, which is then captured by concepts such as subjective certainty [27] and subjective probability [10].

Trust goes hand in hand with the concepts veracity [4] and deception. [45] addresses veracity along the dimensions of truthfulness/deception, objectivity/subjectivity and credibility/implausibility. The authors developed a veracity index ranging from true/objective/credible to untrustworthy/subjective/implausible to characterize texts in the context of big data analysis. Deception is defined as a message knowingly transmitted with the intent to foster false beliefs or conclusions. The topic is addressed in studies from areas such as interpersonal psychology and communication [9], [33] and it is also considered in the field of natural language processing, as part of a larger research direction tackling subjectivity analysis and the identification of private states (emotions, speculations, sentiments, beliefs). These solutions stem from the idea that humans express various degrees of subjectivity [55] that are marked linguistically and can be identified with automatic procedures [54].

Contributions on trust estimation keep the distinction between analyzing the source of information, the item reported and reasoning about trust. Approaches developed for trust in information sources consider that trust is not a general attribute of the source but rather related to certain properties: competence [29], sincerity and willingness to cooperate [50]. On this basis, it becomes possible to consider the competence of a source not in general but with respect to specific topics [28]. Trust can be also analyzed in relation to roles, categories or classes [34].

Research efforts on reasoning about trust analyze information sources from past behaviors rather than directly from their properties [46], or they infer trust from estimations already computed for a set of properties [1]. These approaches generally focus on building trust by using argumentation [62] or beliefs functions [26], or investigating the joint integration of those techniques [52]. Taking this work a step further, [51] identified several patterns for reasoning about trust and its provenance while the notion of conflict in handling trust is discussed in [65].

As shown by approaches above, trust is a multifaceted concept and, in practice, this complex notion can be decomposed into two components: communication or interaction trust, and data trust [48]. The model developed deals with data trust and keeps the distinction between sources and items provided by those sources, although several approaches consider these elements as a whole [26], estimating the trust of information sources [1], [65] rather than information items. The model does not require statistical data to infer the behavior of the source [46] and introduces reliability to characterize the source. More specifically, reliability encompasses not only competence [34], [29] and reputation [28]—two attributes already considered by previous approaches—but also intentions which constitute an original aspect of the model. Intention is of important significance in the context of human-centered systems, including open-sources, and supports the analysis of emerging phenomena such as on-line propaganda or disinformation. Another original aspect of the model is consideration of the characterization of items by the source itself, thus overcoming a main limitation of the solution presented in [12]. Our approach can be considered as partially overlapping solutions investigating trust propagation in direct and indirect reporting [51], [62], and the model enables a particular kind of trust estimation, based both on more or less complete characterizations of the source by the analyst, and more or less accurate characterizations of the items by the source. The model also addresses disagreement and the fusion of diverging opinions, not in a panel of experts as described in [52], but rather between items showing high levels of confidence according to the source and sources having low reliability according to the analyst. By ascribing characterizations to both information sources and reported items, the model allows analysts to make use of both prior experience and their own beliefs in order to assess various degrees of trust.

From a different perspective, the evaluation of uncertainty regarding the inputs, reasoning and outputs of the information fusion is the goal of Evaluation Techniques for Uncertainty Representation Working Group[1] (ETURWG). The group developed an ontology for this purpose [13]. The URREF ontology defines the main subjects under evaluation [18], such as uncertainty representation and reasoning components of fusion systems. Furthermore, the frame also introduces criteria for secondary evaluation subjects: sources and pieces of information, fusion methods and mathematical formalisms. URREF criteria have generic definitions and therefore can be instantiated for applications with coarser or finer granularity levels. This means evaluation metrics can be defined for data analysis [17], increased particularity for data specific types [22] or attributes, reliability and credibility [7], self-confidence [8] or veracity [5].

In addition to allowing a continuous analysis of uncertainty representation, quantification and evaluation, as described in [15], URREF criteria are detailed enough to capture model-embedded uncertainties [37], imperfection of knowledge representations [25], and their propagation in the context of the decision loop [16]. The frame also offers a basis to compare different fusion approaches [24]. URREF criteria were used for uncertainty tracking and investigation in several applications: vessel identification for maritime surveillance [38], activity detection for rhino poaching [43] and imagery analysis for large area protection [6].

Beyond developing a model for trust estimation, this paper also fills a gap within the ETURWG community by illustrating how uncertainty analysis tracks imperfections occurring from problem definition to model abstraction and formalization.

III. HUMAN SOURCES AND REPORTED INFORMATION

Many applications rely on human sources which are used to continuously supply observations, hypotheses, subjective beliefs and opinions about what they sense or learn. In such applications reports are often wrong, due to environment dynamics, simple error, malicious act or intentions, [58]. From the analyst standpoint, decisions have to be made based on indirect reporting and trust relies upon the in-depth investigation of items and sources, thus the analysis of reported items is a critical step. This analysis is a multilevel process, relying on the ability of analysts to understand the content of messages and assess their quality from additional clues. The use cases described below highlight levels of indirection occurring when collecting information and their with impact on trust estimation.

Fig. 1. Assertions and opinions in human messages.

A. Assertions, opinions and reported information

For illustration, let's consider X, the analyst receiving information provided by a human source Y.

Case 1: direct reporting X is an analyst collecting evidence in order to decide whether or not an individual is involved in terrorist activities. In particular, he takes into account reports submitted by Y, a human source. Those reports usually consist on a mixed set of assertions (e.g., descriptions of events or states observed by Y) and opinions (i.e., judgments, assessments, or beliefs) expressed by Y about assertion which give the analyst an insight into how strongly the source commits to the assertion, see Fig. 1.

In the statement contained in Fig. 1, the source Y lets us know that she does not commit her full belief to the assertion that *John is a terrorist*, otherwise the reporter would have used phrasing such as *I am completely convinced* or *it is without doubt* or simply reported *John is a terrorist* as an unadorned statement.

The information item is the sentence, which contains the assertion *John is a terrorist* and the uncertainty degree to be assigned because the analyst knows that Y is not completely certain about her own statements. The analyst must make a judgment about the veracity of John being a terrorist based upon factors such as previous experience with Y's assessments in the past, or, perhaps, on the fact that other sources are relating the same information.

Case 2: indirect reporting Again, let X be an analyst collecting evidence in order to decide whether or not an individual is involved in terrorist activities. In this case, he takes into account reports submitted by Y, a human source who is herself relating information obtained from a secondary source named Mary, see Fig. 2.

The source Y does not report on her direct observations or her deductions or beliefs, but conveys information received from a second source, in this case Mary, in the statement in Fig. 2.

[1] http://eturwg.c4i.gmu.edu/

Fig. 2. Hearsay, assertions and opinions in human messages.

In this report the information item is again the sentence containing the assertive part *John is a terrorist* but this use case introduces more levels of complexity in uncertainty to deal with. The information that the assertion comes from Mary, who has added her own opinion, is a distancing mechanism on the part of the source Y as (unlike in Fig. 1), she is neither claiming the opinion nor the assertion.

This case introduces yet more layers of uncertainty. How sure can we be that the reporter Y has accurately repeated what Mary said? For example, did Mary really say *it is likely* or did the reporter insert this (intentionally or unintentionally) based upon the reporter's assessment of the reliability of Mary as a source of information? Or perhaps, subtly, Y is expressing her own uncertainty by putting words in Mary's mouth. Furthermore, it is possible Mary made this statement under circumstances which would strengthen or weaken this statement, but those conditions have not been passed on by the reporter.

The goal of the analyst is to take this assertion into account, but also to encode his own belief about the quality of the source further in the analysis. All these different attitudes have to be evaluated by the analyst, who may have additional background information or prior evaluation of the source that have to be considered.

In both cases discussed above, the outcome of the analyst is the assertive part of the information item, augmented with a coefficient that helps to measure and track the different levels of trust for their future exploitation. For the purpose of this work, this quality is called *trust in reported information*.

B. Concepts and notions for trust assessment

This section introduces several notions that are relevant for trust analysis.

Trustworthiness of information sources is considered, for the purpose of this work, as confidence in the ability and intention of an information source to deliver correct information, see [3]. Trustworthiness is an attribute of information sources who have the competences to report information, and who can be relied upon to share sincerely and clearly their beliefs on the uncertainty level of reported information. An item provided by such a source is then trusted by analysts.

Self-confidence [8] captures the explicit uncertainty assigned to reported assertions by the source. Statements may include the source's judgments when lacking complete certainty; these judgments are generally identified through the use of various lexical clues such as *possibly, probably, might be, it is unlikely, undoubtedly*, etc., all of which signal the source's confidence (or lack thereof) in the veracity of the information being conveyed. It should be noted that self-confidence, in our usage understood as the linguistic dimension of the certainty degree that the source assigns to reported items, is an aspect exhibited by the source, but it will be considered from the analyst's standpoint during trust analysis.

Reliability of sources indicates how strongly the analyst is willing to accept items from a given source at their face-value. As an overall characterization, reliability is used in this work to rate how much a source can be trusted with respect to their reputation, competence and supposed intentions.

Reputation of sources [11] captures a commonly accepted opinion about how the source performs when reporting information, and is generally understood as the degree to which prior historical reports have been consistent with fact. For human sources, reputation is considered by the analyst for each source based on previous interactions with the source and on the source's history of success and failure in delivering accurate information. Reputation relies, to a large extent, upon negative and positive experiences provided to the analyst by the source in the past.

Competence of sources [29] is related to a source's possession of the skills and knowledge in reporting on various topics: This aspect defines to what extent a human source can understand the events they report on, whether the source has the ability to accurately describe those events, and how capable the source is of following the logic of processes producing the information.

Intentions correspond to specific attitudes toward the effect of one's actions or conduct. Reporting information can become *more a means to manipulate others than a means to inform them* [14] and thus can be carried out with the express purpose of inducing changes in another person's beliefs and understanding. Intentions are specific to human sources as only humans have the capacity to deliberately provide false or misleading information. Sensors may provide erroneous data due to a number of factors such as device failure or environmental conditions, but never due to intention.

In addition to the above facets, *credibility of information* and *reliability of sources* are two notions introduced by the STANAG 2511 [49], which standardizes the terminology used in analysis of intelligence reports used by NATO Forces with distinct focus on sources and information provided. STANAG reliability is understood with respect to the quality of information that has been delivered by sources in the past. STANAG credibility relies on the intuition that a joint analysis of items in combination with each other will likely reveal inconsistencies, contradictions or redundancies. Reliability and credibility are independent criteria for evaluation. Definitions for both reliability and credibility are in natural language.

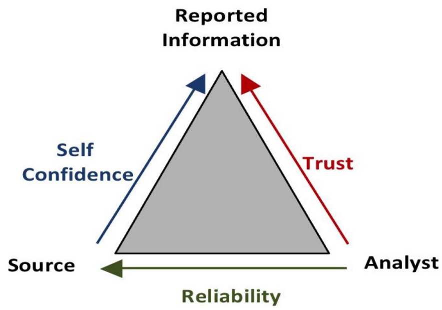

Fig. 3. Model for trust analysis.

Attributes of sources and information items adopted for the model of trust are related to the notions introduced by the STANAG 2511 but are addressed differently: reliability of sources is understood here in terms of source competence, reputation and intentions, while credibility is restricted to features of self-confidence as described above.

IV. A FUNCTIONAL MODEL OF TRUST

This section introduces the model developed to estimate trust in reported information by taking into account the reliability of the source and the source's own characterization of reported items. The advantage of this distinction is to better dissociate the impact of both beliefs of sources and opinions of analysts on the source on the information provided.

Even if the primary function of a source is to provide information, we keep the distinction between the source and the information by considering separate dimensions for each element. The rationale behind this is the observation that even reliable sources can sometimes provide inaccurate or imprecise information from one report to another, which is even more plausible in the case of human sources.

The model, illustrated in Fig. 3., is composed of a source which provides an information item augmented with a degree of uncertainty captured by self-confidence to an analyst. Based upon his direct assessment of the reliability of the source, the analyst constructs his own estimation of trust in the item reported.

In the following section, the model is discussed using a granularity that is detailed enough to describe its elements, but still rough enough to avoid the adoption of a representation formalism.

A. Elements of the trust model

The model is composed of two elements: an information source and reported items from that source. The analyst is considered to be outside the model, although she has multiple interactions with its elements.

Definition of information source: an information source is an agent who provides an information item along with a characterization of its level of uncertainty. "Source" is a relative notion, depending on the perspective of analysis. In general, information is propagated within a chain relating real world information to some decision maker, and agents along the path can be both trained observers, whose job is to provide such reports, as well as witnesses or lay observers who may add items, in spite of not being primarily considered as information sources, but rather as opportunistic ones.

The notion of source is central in many information fusion applications and numerous research efforts aimed at modeling the properties of those applications. A general analysis of sources is undertaken by [32], who identify three main classes: S-Space, composed of physical sensors, H-Space for human observers and I-Space for open and archived data on the Internet. In [39], a unified characterization of hard and soft sources is described, along with a detailed description of their qualities and processing capabilities.

Processing hard sensor information is widely covered [42] in the research community, and can be considered quite mature, while the integration of human sources brings many new challenges. Our model addresses human sources, and reported items can refer to actions, events, persons or locations of interest.

Information reported by humans is unstructured, vague, ambiguous and subjective, and thus is often contrasted with information coming from physical sensors, described as structured, quantitative and objective. While humans can deliberately change the information or even lie, sensors are also prone to errors and therefore hard information items are not always accurate.

For human agents, the source is part of the real world, (a community, a scene, an event) and can be either directly involved in the events reported, or just serving as a witness.

Definition of reported information: Reported information is a couple $(I, \chi(I))$, where I is an item of information and $\chi(I)$ the confidence level as assigned by the source. Items are information pieces that can be extracted from natural language sentences, although the extraction and separation from subjective content are out of the scope for the model developed. Each item I has assertive i_a and subjective i_s components conveying factual and subjective contents respectively.

The analysis of reported information continues to be an open topic as the fusion of information from soft sources receives increasing attention in recent years. Although some authors have developed logic-based approaches for modelling distortions of items exchanged between agents who have both the intention and the ability to deceive [12], there are still more challenges arising when the information is analyzed in its textual form.

Features of uncertainty, as expressed in natural language statements, are analyzed in [2] while [23] provides a broader discussion of pitfalls and challenges related to soft data integration for information fusion.

B. Functions of the trust model

The model introduces several functions estimating features of reliability, self-confidence and trust, as described hereafter.

Definition of a reliability function: a reliability function is a mapping which assigns a real value to an information source.

This real value is a quantitative characterization of the source, inferred with respect to the source's previous failures, its reputation and the relevance of its skills for specific domains. For this model, the reliability of human sources combines three features: competence, reputation and intention. Competence captures the intuition that the quality of information reported by a source depends on the level of training and expertise, which may be designated as satisfactory or not, depending upon the task. Reputation is the overall quality of a source, estimated by examination of the history of its previous failures. Intentions refer to attitudes or purposes, often defined with respect to a hidden purpose or plan to achieve.

Reliability is a complex concept and, from a practical standpoint, it is difficult to have complete information about the global reliability of a source. Thus, this model describes reliability along the three attributes (competence of a source, its reputation and its intentions) described above. In practical applications, this solution allows for compensation for insufficient information on one or several aspects of reliability and to conduct, if necessary, the analysis of reliability based on just one attribute.

Evaluation of reliability Assessing reliability is of real interest when opportunistic sources are considered because the analyst has neither an indication of how the source might behave nor the ability to monitor or control either the human providing the information or the environment in which the source operates. Various methods can be developed to estimate competence, reputation and intentions of the source. For example, competence is closely related to the level of training of an observer or can be defined by domain knowledge. Values can be expressed either in a linguistic form (*bad, good, fair, unknown*) or by a number. Reputation is an attribute which can be constructed not just by examining previous failures of the source but also by considering its level of conflict with other sources; this too can be expressed by numeric or symbolic values.

While reputation and competence can be, at least in some cases, estimated from prior knowledge, characterizing the intentions of a source is subject to human perception and analysis. Judgment of human experts is needed not just because there usually is no *a priori* characterization of the source with respect to its intentions but also because it is important to assess those aspects from the subjective point of view of an expert in the form of binary values only.

From a practical standpoint, it is suitable to provide an expert with a description of source competence, reputation and intentions as assessed independently. This way, experts can have the opportunity to develop different strategies of using reliability: they can decide to assign different importance to those attributes under different contexts or can use their own hierarchy of attributes. For instance, an expert may consider as irrelevant the information provided by a source whose competences is lower than a specific threshold or if he suspects the source of having malicious intentions.

Definition of a self-confidence function: a self-confidence function is a mapping linking a real value and an information item. The real value is a measure of the information credibility as evaluated by the sensor itself and is of particular interest for human sources, as often such sources provide their own assessments of the information conveyed. Identifying features of self-confidence requires methods related to a research task of natural language processing: the identification of assertions and opinions in texts. In this field, the commonly adopted separation of those notions considers assertions as statements that can be proven true or false, while opinions are hypotheses, assumptions and theories based on someone's thoughts and feelings and cannot be proven.

Evaluation of self-confidence: Estimation of self-confidence aims at assigning a numerical value which captures how strongly the author stands behind assertions in the statement, on the basis of lexical clues he has included in the utterance. More generally, markers of an author's commitment are in the form of hedges, modal verbs and forms of passive/active language. A hedge is a mitigating word that modifies the commitment to the truth of propositions, i.e., certainly, possibly. Its impact can be magnified by a booster (highly likely) or weakened by a downtoner (rather certain).

Modal verbs indicate if something is plausible, possible, or certain (*John could be a terrorist, you might be wrong*). Moreover, in some domains sentences making use of the passive voice are considered as an indicator of uncertainty, in the sense that author seeks to distance himself from the assertions in the items reported through use of passive voice. Quantifying self-confidence is a topic of particular interest for intelligence analysis, and it was early addressed by Kent in 1962, [40] who created a standardized list of words of estimative probability which were widely used by intelligence analysts. This list has continued to be a common basis to be used by analysts to produce uncertainty assessments.

Kesselman describes in [41] a study conducted to analyze the way the list was used by analysts over the past, and identifies new trends to convey estimations and proposes a new list having the verb as a central element. Given the variety of linguistic markers for uncertainty, the estimation of a numerical value based on every possible combination seems unrealistic, as the same sentence oftencontains not just one but multiple expressions of uncertainty. Additionally, assigning numerical values to lexical expressions is not an intuitive task, and Rein shows that there are no universal values to be associated in a unique manner to hedges or other uncertainty markers, see [53]. As the author argues further, it is, however, possible to order those expressions and use this relative ordering as a more robust way to compare combinations of uncertainty expressions, and thus highlight different levels of uncertainty in natural language statements.

Using the model for trust analysis: The model proposed in this work proposed in this work combines various attributes of the source (discussed previously under "reliability") with "self-confidence" in order to capture trust of information as conveyed by the human. The model is source-centric predominantly focused on the source's ability to correct, alter or qualify the information report Although the rules for ranking, prioritizing and combining the attributes introduced by the model can be drafted empirically, the estimation of a trust value requires a formal representation of the model.

A possible solution for estimating a unified value for trust is to consider reliability and self-confidence within the framework of an uncertainty theory and to rely on the set of combination rules the theory defines—for example, those developed in probability theory, in possibility theory, or in belief functions theory. All these theories provide various operators to combine reliability and self-confidence in order to estimate trust.

In the following the model is represented by using belief functions and several scenarios are used to illustrate trust estimation.

V. TRUST FORMALIZATION WITH BELIEF FUNCTIONS

The aim of trust formalization is to provide a formal representation of the model, combining the capability to exploit the structure and relationship of elements of the model with the ability to express degrees of uncertainty about those elements. Of particular interest to this paper is the observation that the developed model introduces a cognitive view of trust as a complex structure of beliefs that are influenced by the individual's opinions about certain features and elements, including their own stances. Such a structure of beliefs determines various degrees of trust, which are based on personal choices made by analyst, on the one hand, and the source, on the other hand. Therefore, the formalization requires a formalism that is more general than probability measures or fuzzy category representation, which are more suitable for applications considering trust in the context of interactions between agents. Moreover, the limitations of using subjective probabilities to formalize trust from this cognitive standpoint are clearly stated in [10]. As a result, the model was represented with belief functions, a formalism that is consistent with the cognitive perspective of trust adopted by the model. This belief-based representation provides the most direct correspondence with elements of the model and their underlying uncertainty, while being able to quantify subjective judgments.

After introducing main concepts of belief functions, this section shows how the formalism is used to represent the trust model.

A. Basic Belief Assignment

Belief Functions (BF) have been introduced by Shafer in his his mathematical theory of evidence [56], also referred to Dempster-Shafer Theory (DST), to model epistemic uncertainty. The frame of discernment (FoD) of the decision problem under consideration, denoted Θ, is a finite set of exhaustive and mutually exclusive elements. The powerset of Θ denoted 2^Θ is the set of all subsets of Θ, empty set included. A body of evidence is a source of information characterized by a Basic Belief Assignment (BBA), or a mass function,which is the mapping $m(.) : 2^\Theta \rightarrow [0,1]$ that satisfies $m(\emptyset) = 0$, and the normalization condition $\sum_{A\in 2^\Theta} m(A) = 1$. The belief (a.k.a credibility) Bel(.) and plausibility Pl(.) functions usually interpreted as lower and upper bounds of unknown (subjective) probability measure $P(.)$, are defined from $m(.)$ respectively by

$$\mathrm{Bel}(A) = \sum_{B\subseteq A|B\in 2^\Theta} m(B) \tag{1}$$

$$\mathrm{Pl}(A) = \sum_{B\cap A\neq\emptyset|B\in 2^\Theta} m(B) \tag{2}$$

An element $A\in 2^\Theta$ is called a focal element of the BBA $m(.)$, if and only if $m(A) > 0$. The set of all focal elements of $m(.)$ is called the core of $m(.)$ and is denoted $\mathcal{K}(m)$. This formalism allows for modeling a completely ignorant source by taking $m(\Theta) = 1$. The Belief Interval (BI) of any element A of 2^Θ is defined by

$$BI(A) \triangleq [Bel(A), Pl(A)] \tag{3}$$

The width of belief interval of A, denoted $U(A) = Pl(A) - Bel(A)$ characterizes the degree of imprecision of the unknown probability $P(A)$, often called the uncertainty of A. We define the uncertainty (or imprecision) index by

$$U(m) \triangleq \sum_{A\in\Theta} U(A) \tag{4}$$

to characterize the overall imprecision of the subjective (unknown) probabilities committed to elements of the FoD bounded by the belief intervals computed with the BBA $m(.)$.

Shafer proposed using Dempster's rule of combination for combining multiple independent sources of evidence [56] which is the normalized conjunctive fusion rule. This rule has been strongly disputed in the

BF community after Zadeh's first criticism in 1979, and since the 1990s many rules have been proposed to combine (more or less efficiently) BBAs; the reader is advised to see discussions in [59], in particular the proportional conflict redistribution rule number 6 (PCR6). To combine the BBAs we use the proportional conflict redistribution (PCR) rule number 6 (denoted PCR6) proposed by Martin and Osswald in [59] because it provides better fusion results than Dempster's rule in situations characterized by both high and low conflict as explained in detail in [19], [35].

The PCR6 rule is based on the PCR principle which transfers the conflicting mass only to the elements involved in the conflict and proportionally to their individual masses, so that the specificity of the information is entirely preserved. The steps in applying the PCR6 rule are:

1) apply the conjunctive rule;
2) calculate the total or partial conflicting masses; and
3) redistribute the (total or partial) conflicting mass proportionally on non-empty sets.

The general PCR6 formula for the combination of $n > 2$ BBAS is very complicated (see [59] Vol. 2, Chap. 2). For convenience's sake, we give here just the PCR6 formula for the combination of only two BBAs. When we consider two BBAs $m_1(.)$ and $m_2(.)$ defined on the same FoD Θ, the PCR6 fusion of these two BBAs is expressed as $m_{PCR6}(\emptyset) = 0$ and for all $X \neq \emptyset$ in 2^Θ

$$m_{PCR6}(X) = \sum_{\substack{X_1,X_2\in 2^\Theta \\ X_1\cap X_2 = X}} m_1(X_1)m_2(X_2) + \sum_{\substack{Y\in 2^\Theta\setminus\{X\} \\ X\cap Y=\emptyset}} \left[\frac{m_1(X)^2 m_2(Y)}{m_1(X)+m_2(Y)} + \frac{m_2(X)^2 m_1(Y)}{m_2(X)+m_1(Y)}\right] \tag{5}$$

where all denominators in (5) are different from zero. If a denominator is zero, that fraction is discarded. A very basic (not optimized) Matlab code implementing the PCR6 rule can be found in [59] and [61], and also in the toolboxes repository on the web.[2]

Instead of working with quantitative (numerical) BBA, it is also possible to work with qualitative BBA expressed by labels using the linear algebra of refined labels proposed in Dezert-Smarandache Theory (DSmT), [59] (Vol. 2 & 3).

B. Trust formalization model

Because beliefs are well defined mathematical concepts in the theory of belief functions, we prefer to use self-confidence terminology to represent the confidence declared by a source Y on its own assertion A. Let's denote by A the assertion given by the source, for instance $A =$ *John is a terrorist*. With respect to elements of the model, A (the assertion) corresponds to i_a, the assertive part of the item I and $v(A)$ is a numeric estimation of the subjective i_s component of I.

The valuation $v(A)$ made by the source Y about the assertion A can be done either quantitatively (by a probability or a BBA) or qualitatively (by a label associated to a linguistic form). This paper considers quantitative representation of $v(A)$ for simplicity.[3]

The basic information items provided by a source Y consists of A (the assertion), and $v(A)$ (its valuation). To be as general as possible, we suppose that $v(A)$ is a basic belief mass assignment defined with respect to the very basic frame of discernment $\Theta_A \triangleq \{A, \bar{A}\}$ where $\bar{A}$ denotes the complement of A in Θ_A, that is $v(A) = (m(A), m(\bar{A}), m(A \cup \bar{A}))$. Note that only two values of the triplet are really necessary to define $v(A)$ because the third one is automatically derived from the normalization condition $m(A) + m(\bar{A}) + m(A \cup \bar{A}) = 1$. So one could also have chosen equivalently $v(A) = [Bel(A), Pl(A)]$ instead of the BBA. In a probabilistic context, one will take $m(A \cup \bar{A}) = 0$ and so $v(A) = P(A)$ because $Bel(A) = Pl(A) = P(A)$ in such a case.

The self-confidence of the source Y is an extra factor $\alpha_Y \in [0, 1]$ which characterizes the self-estimation of the quality of the piece of information $(A, v(A))$ provided by the source itself. $\alpha_Y = 1$ means that the source Y is 100% confident in his valuation $v(A)$ about assertion A, and $\alpha_Y = 0$ means that the source Y is not at all confident in his valuation $v(A)$. In the theory of belief functions, this factor is often referred as the discounting factor of the source because this factor is usually used to discount the original piece of information $(A, v(A))$ into a discounted one $(A, v'(A))$ as follows [56]:

$$m'(A) = \alpha_Y \cdot m(A) \tag{6}$$

$$m'(\bar{A}) = \alpha_Y \cdot m(\bar{A}) \tag{7}$$

$$m'(A \cup \bar{A}) = \alpha_Y \cdot m(A \cup \bar{A}) + (1 - \alpha_Y) \tag{8}$$

The idea of Shafer's discounting technique is to diminish the belief mass of all focal elements with the factor α_Y and redistribute the missing discounted mass $(1 - \alpha_Y)$ to the whole ignorance $A \cup \bar{A}$. Note that the valuation of the discounted piece of information is always degraded because its uncertainty index is always greater than the original one, that is, $U(m') > U(m)$, which is normal.

The reliability factor r estimated by the analyst X on the piece of information $(A, v(A))$ provided by the source Y must take into account both the competence C_Y, the reputation R_Y and the intention I_Y of the source Y. A simple model to establish the reliability factor

[2] http://bfaswiki.iut-lannion.fr/wiki/index.php/Main_Page

[3] Without loss of generality one can always map a qualitative representation to a quantitative one by a proper choice of scaling and normalization (if necessary).

r is to consider that C_Y, R_Y and I_Y factors are represented by numbers $[0,1]$ associated to select subjective probabilities, that is $C_Y = P(Y \text{ is competent})$, $R_Y = P(Y \text{ has a good reputation})$ and $R_Y = P(Y \text{ has a good intention (i.e. is fair)})$. If each of these factors has equal weight, then one could use $r = C_Y \times R_Y \times I_Y$ as a simple product of probabilities. However, in practice, such simple modeling does not fit well with what the analyst really needs to take into account epistemic uncertainties in Competence, Reputation and Intention. In fact, each of these factors can be viewed as a specific criterion influencing the level of the global reliability factor r. This is a multi-criteria valuation problem. Here we propose a method to solve the problem.

We consider the three criteria C_Y, R_Y and I_Y with their associated importance weights w_C, w_R, w_I in $[0,1]$ with $w_C + w_R + w_I = 1$. We consider the frame of discernment $\Theta_r = \{r, \bar{r}\}$ about the reliability of the source Y, where r means that the source Y is reliable, and $\bar{r}$ means that the source Y is definitely not reliable. Each criteria provides a valuation on r expressed by a corresponding BBA. Hence, for the competence criteria C_Y, one has $(m_C(r), m_C(\bar{r}), m_C(r\cup\bar{r}))$, while for the reputation criteria R_Y, one has $(m_R(r), m_R(\bar{r}), m_R(r\cup\bar{r}))$ and for the intention criteria I_Y, one has $(m_I(r), m_I(\bar{r}), m_I(r\cup\bar{r}))$.

To get the final valuation of the reliability r of the source Y, one needs to efficiently fuse the three BBAs $m_C(.)$, $m_R(.)$ and $m_I(.)$, taking into account their importance weights w_C, w_R, and w_I. This fusion problem can be solved by applying the importance discounting approach combined with PCR6 fusion rule of DSmT [60] to get the resultant valuation $v(r) = (m_{PCR6}(r), m_{PCR6}(\bar{r}), m_{PCR6}(r\cup\bar{r}))$ from which the decision (r, or $\bar{r}$) can be drawn (using BI distance, for instance). If a firm decision is not required, an approximate probability $P(r)$ can also be inferred with some lossy transformations of BBA to probability measure [59]. Note that Dempster's rule of combination cannot be used here because it does not respond to the importance discounting, as explained in [60].

The trust model consists of the piece of information $(A, v(A))$ and the self-confidence factor α_Y provided by the source Y, as well as the reliability valuation $v(r)$ expressed by the BBA $(m(r), m(\bar{r}), m(r\cup\bar{r}))$ to infer the trust valuation about the assertion A. For this, we propose using the mass $m(r)$ of reliability hypothesis r of the source Y as a new discounting factor for the BBA $m'(.)$ reported by the source Y, taking into account its self-confidence α_Y. Hence, the trust valuation $v_t(A) = (m_t(A), m_t(\bar{A}), m_t(A\cup\bar{A}))$ of assertion A for the analyst X is defined by

$$m_t(A) = m(r)\cdot m'(A) \tag{9}$$

$$m_t(\bar{A}) = m(r)\cdot m'(\bar{A}) \tag{10}$$

$$m_t(A\cup\bar{A}) = m(r)\cdot m'(A\cup\bar{A}) + (1 - m(r)) \tag{11}$$

or equivalently by

$$m_t(A) = m(r)\alpha_Y \cdot m(A) \tag{12}$$

$$m_t(\bar{A}) = m(r)\alpha_Y \cdot m(\bar{A}) \tag{13}$$

$$m_t(A\cup\bar{A}) = m(r)\alpha_Y \cdot m(A\cup\bar{A}) + (1 - m(r)\alpha_Y) \tag{14}$$

The DSmT framework using the PCR6 fusion rule and the importance discounting technique provides an interesting solution for the fusion of attributes having different degrees of importance while making a clear distinction between those attributes.

The discounting method proposed in this work is directly inspired by Shafer's classical discounting approach [56]. In our application, the classical discounting factor that we propose integrates both the mass of reliability hypothesis $m(r)$ and the self-confidence factor α_Y. It is worth noting that more sophisticated (contextual) belief discounting techniques [47] exist and they could also have been used, in theory, to refine the discounting but these techniques are much more complicated and they require additional computations. The evaluation of contextual belief discounting techniques for such types of application is left for further investigations and research works.

VI. UNCERTAINTY ANALYSIS UNDER URREF CRITERIA

Tracking uncertainties from problem description to model construction and formalization is done under criteria of the uncertainty representation and reasoning evaluation framework.

The goal of URREF is to place the focus on the evaluation of uncertainty representation and reasoning procedures. The URREF ontology defines four main classes of evaluation criteria: Data Handling, Representation, Reasoning and Data Quality. These criteria make distinctions between the evaluation of the fusion system, the evaluation of its inputs and outputs, and the evaluation of the uncertainty representation and reasoning aspects.

Listing all criteria is an extensive task and in this paper the authors will provide one piece of the puzzle by considering criteria that relate to the evaluation of uncertainty induced by the proposed model. In the model developed in this paper, uncertainty is due to imperfections of information gathering and reporting as well as constraints of the representation formalism.

Uncertainty analysis is carried out by assigning uncertainty criteria to elements and functions of the trust model in order to make explicit the uncertainty arising when the problem is abstracted by the model and the model is then simplified in order to fulfill constraints of specific formalism, Fig. 6.

The URREF criteria selected are subclasses of two main concepts: *Credibility*, a subconcept under *DataCriteria*, and *EvidenceHandling*, a subconcept of *RepresentationCriteria*.

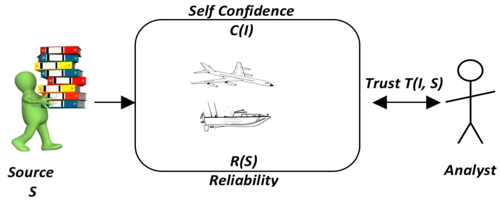

Fig. 4. Trust estimation from source to analyst

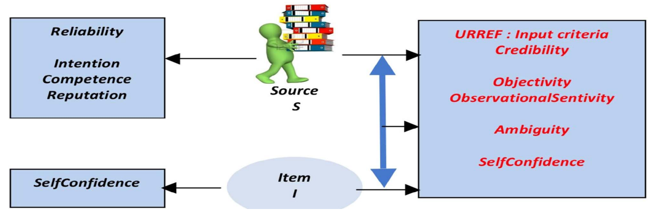

Fig. 5. Mapping of model attributes to URREF criteria

To summarize, uncertainties of the model will be captured by the following URREF criteria:

- **Objectivity**, subconcept of **Credibility**: indicates a source providing unbiased information;
- **ObservationalSensitivity**, subconcept of **Credibility**: characterizes the skills and competences of sources;
- **SelfConfidence**, subconcept of **Credibility**: measures the certainty degree about the piece of information reported, according to the source;
- **Ambiguity**, subconcept of **EvidenceHandling**: captures if the sources provide data supporting different conclusions;
- **Dissonance**, subconcept of **EvidenceHandling**: captures the ability of formalism to represent inconsistent evidence;
- **Completeness**, subconcept of **EvidenceHandling**: is a measure of how much is known given the amount of evidence; and
- **Conclusiveness**, subconcept of **EvidenceHandling**: indicates how strong the evidence supports a conclusion;

Besides selecting uncertainty criteria relevant for trust estimation, the analysis also discusses the mapping of URREF criteria to attributes of the model and sheds a light on imperfect matchings. This mapping offers a basis for identifying the limitations of the URREF ontology, by emphasizing those elements whose characterizations in terms of uncertainty are out of the ontology's reach or beyond the ontology's intended scope.

A. Uncertainties from problem definition to model abstraction

Let M be the model for trust estimation, with elements introduced in paragraph IV: the source Y, the reported item I with its assertive i_a and subjective i_s parts, and $\chi(I)$ the confidence level assigned by the source Y to I.

From an information fusion standpoint, inputs of the model are the source and the information items, along with their uncertainty, captured with the following URREF criteria: *Objectivity, ObservationalSensitivity and SelfConfidence*. These criteria are subclasses of the concept *InputCriteria*.

Objectivity is an attribute of the source, related to its ability to provide factual, unbiased items, without adding their own points of view or opinions. For a source Y providing information item i, having i_s and i_a as the subjective and factual parts respectively, objectivity can be expressed as:

$$Objectivity(Y, I) = \psi_o(i_s, i_a) \tag{15}$$

where $\psi_o(i_s, i_a)$ represents the mathematically quantified expression of the subjective over the factual content of i.

ObservationalSensitivity is an attribute of the source which represents the source's ability to provide accurate reports. In the proposed model, this criterion is an aggregation of competence C and reputation R, two attributes of the model.

$$ObservationalSensitivity(Y, i) = \psi_{os}(C, R) \tag{16}$$

where $\psi_{os}(C, R)$ is a function aggregating values of competence and reputation.

Information items entering the system are described by *SelfConfidence*. Again, considering i_s and i_a as the subjective and factual items conveyed by I, *SelfConfidence* can be expressed as:

$$SelfConfidence(I) = \psi_{sc}(i_s) \tag{17}$$

with $\psi_{sc}(i_s)$ a function quantifying the subjective content of item I.

Fig. 5 shows the mapping between the elements of the model and the set of relevant URREF uncertainty criteria. The mapping shows a perfect match between *SelfConfidence* as introduced by the model and the eponymous URREF criterion as well as several imperfect matches described later in this paper.

At source level, URREF criteria are not able to capture in a distinct manner the features of competence, reputation and intentions, the main attributes of the sources added by the model under Reliability. To some extent, competence and reputation can be related to *ObservationalSensitivity*, but intentions clearly remains out of reach for URREF criteria.

B. Uncertainties from model to formal representation

Let F be the DST formalization of the trust estimation model, with parameters introduced in paragraph V. The formalism induces two types of uncertainty related to its capacity to handle incomplete, ambiguous or contradictory evidence. The uncertainty of evidence handling is captured by *Ambiguity*, *Dissonance*, *Conclusiveness* and *Completeness*. Those criteria are subclasses of the concept *EvidenceHandling*.

Ambiguity measures the extent to which the formalism can handle data sets which support different conclusions.

$$Ambiguity(F) = \phi_a(\alpha_Y, R_Y) \tag{18}$$

where the function $\phi_a(\alpha_Y, R_Y)$ considers the self-confidence factor α_Y provided by the source Y and the reliability of Y provided by the analyst R_Y to estimate the degree of ambiguity. The measure is of particular interest in the case where items having high values of self-confidence are provided by unreliable sources.

Dissonance captures the ability of the formalism to represent inconsistent evidence. For BBA representations, dissonance can be related to the capacity of the formalism to assign belief mass to an element and its negation, and can therefore be assessed for every BBA representation build for the model. For example, the dissonance for a source's competence can be in the form:

$$Dissonance(F) = \phi_d(m_C(r), m_C(\bar{r})) \tag{19}$$

where $\phi_d(m_C(r), m_C(\bar{r}))$ is a function combining the belief mass assigned to whether the source is considered to be competent or incompetent, respectively.

Dissonance is useful for highlighting situations in which there are significant differences in belief masses assigned at the attribute level, such as when a source is considered to be incompetent (low $m_C(r)$, high $m_C(\bar{r})$) but has a good reputation (high $m_R(r)$, low $m_R(\bar{r})$).

Conclusiveness is a measure expressing how strongly the evidence supports a specific conclusion or unique hypothesis:

$$Conc.(F) = \phi_{cc}(m_t(A), m_t(\bar{A}), m_t(A \cup \bar{A})) \tag{20}$$

where $\phi_{cc}(m_t(A), m_t(\bar{A}), m_t(A \cup \bar{A}))$ is a function combining the belief masses estimated for truthful, untruthful and unknown qualifications of assertion A respectively. This measure indicates to which extent the result of inferences can support a conclusion, in this case whether the hypothesis that the assertion under analysis is trustworthy or not. It can be used during the inference process to show how taking into account additional elements such as the competence of the source, its reputation or intentions impact the partial estimations of trust.

Completeness is a measures of the range of the available evidence, and captures the ability of formalism to take into account how much is unknown. The measures is somewhat similar to *Dissonance*, as is can be assessed for every BBA representation build for the model. Thus, completeness of source's reliability is described as:

$$Completeness(F) = \phi_{cp}(m_{(}r \cup \bar{r})) \tag{21}$$

where $\phi_{cp}(m_{(}r \cup \bar{r}))$ is a function depending on the belief mass assigned to unknown.

The measure is used for estimation and analysis before entering the fusion process, in order to have a picture of how complete the evidence describing the various elements of the model is, and to avoid performing

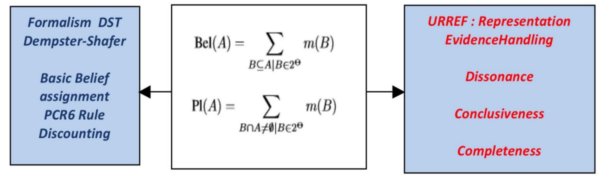

Fig. 6. Mapping of formalism uncertainties to URREF criteria

fusion on highly incomplete data sets. Both *EvidenceHandling* and *KnowledgeHandling* are subclasses of *RepresentationCriteria*.

This section has analyzed the nature of uncertainties arising when going from problem to model definition and then on to formalization with belief functions. The next section shows how uncertainties can be highlighted for particular scenarios of trust estimation.

VII. UNCERTAINTY ANALYSIS FOR TRUST ESTIMATION

A. Running example and method for uncertainty tracking

As a running example, let's consider an assertion A and its valuation $v(A)$ provided by the source Y as follows: $m(A) = 0.7$, $m(\bar{A}) = 0.1$ and $m(A \cup \bar{A}) = 0.2$. Its self-confidence factor is $\alpha_Y = 0.75$. Hence, the discounted BBA $m'(.)$ is given by

$$m'(A) = 0.75 \cdot 0.7 = 0.525$$

$$m'(\bar{A}) = 0.75 \cdot 0.1 = 0.075$$

$$m'(A \cup \bar{A}) = 1 - m'(A) - m'(\bar{A}) = 0.4$$

Let's assume that the BBAs about the reliability of the source based on Competence, Reputation and Intention criteria are given as follows:

$$m_C(r) = 0.8, m_C(\bar{r}) = 0.1, m_C(r \cup \bar{r}) = 0.1$$

$$m_R(r) = 0.7, m_R(\bar{r}) = 0.1, m_R(r \cup \bar{r}) = 0.2$$

$$m_I(r) = 0.6, m_I(\bar{r}) = 0.3, m_I(r \cup \bar{r}) = 0.1$$

with importance weights $w_I = 0.6$, $w_R = 0.2$ and $w_C = 0.2$.

After applying the importance discounting technique presented in [60] which consists of discounting the BBAs with the importance factor and redistributing the missing mass onto the empty set, then combining the discounted BBAs with PCR6 fusion rule, we finally get, after normalization, the following BBA

$$m(r) = 0.9335$$

$$m(\bar{r}) = 0.0415$$

$$m(r \cup \bar{r}) = 1 - m(r) - m(\bar{r}) = 0.025$$

The final trust valuation of assertion A reported by the source Y taking into account its self-confidence

$\alpha_Y = 0.75$ and the reliability factor $m(r) =$ is therefore given by Eqs. (12)–(14) and obtaining

$$m_t(A) = 0.4901$$

$$m_t(\bar{A}) = 0.0700$$

$$m_t(A \cup \bar{A}) = 1 - m_t(A) - m_t(\bar{A}) = 0.4399$$

Note that if $m_C(r) = m_R(r) = m_I(r) = 1$, then we will always get $m(r) = 1$ regardless of the choice of weightings factors, which is normal. If there is a total conflict between valuations of reliability based on Competence, Reputation and Intention criteria, then Dempster's rule cannot be applied to get the global reliability factor $m(r)$ because of 0/0 indeterminacy in the formula of Dempster's rule. For instance, if one has $m_C(r) = m_R(r) = 1$ and $m_I(\bar{r}) = 1$, then $m(r)$ is indeterminate with Dempster's rule of combination, whereas it corresponds to the average value $m(r) = 2/3$ using PCR6 fusion rule (assuming equal importance weights $w_C = w_R = w_I = 1/3$), which makes more sense.

The following subsections explore several scenarios for trust assessment, corresponding to different situations of BBAs distributions, and track the uncertainty according to URREF criteria. Each scenario illustrates specific instances of the model developed for trust estimation.

The method adopted to track uncertainty defines the following measures to estimate URREF criteria:

$$SelfConfidence = \alpha_Y$$

$$Ambiguity = |\alpha_Y - m(r)|$$

$$Objectivity = m_I(r)$$

$$ObservationalSensitivity = min(m_C(r), m_R(r))$$

As shown in previous formulas, URREF criteria are estimated based on features of the BBA formalization and are assigned to the static elements of the model, i.e., the source and the information item. While *Objectivity* and *ObservationalSensitivity* captures imperfections of observations, *SelfConfidence* and *Ambiguity* reflect inaccuracies in reporting information to analysts. These criteria are assessed before entering the fusion phase, and describe the initial uncertainty present in the system before inferences.

In addition, *Dissonance*, *Conclusiveness* and *Completeness* will be estimated at the scenario level by adopting the following formulas:

$$Dissonance = 1 - |m_t(A) - m_t(\bar{A})|$$

$$Conclusiveness = |m_t(A) - m_t(\bar{A})|$$

$$Completeness = 1 - m(A \cup \bar{A})$$

Criteria above will be assessed for elements impacted by the fusion process: the reliability of the source, the updated BBAs of the initial assertion and estimated trust. In the following subsection we illustrate several scenarios for trust estimation and the uncertainty analysis underlying each scenario.

TABLE I.
Consensus: input uncertainty

Uncertainity of inputs		
Observation	Objectivity ObservationalSensitivity	1 1
Reporting	SelfConfidence Ambiguity	1 0

TABLE II.
Consensus: fusion uncertainty

Fusion uncertainty	Dissonance	Conclusiv.	Complet.
Updated BBAs	0	1	1
Reliability	0	1	1
Trust	0	1	1

B. Scenarios for trust assessment and uncertainty analysis

Scenarios introduced below provide examples of trust construction using various operators and highlight the uncertainty assigned to elements of the model and its propagation during the fusion process.

Scenario 1—Consensus: Suppose that Y provides the assertion A, while stating that A certainly holds and that X considers Y to be a reliable source.

In this case, the trust will be constructed on the basis of two consensual opinions: the analyst X that considers Y as a reliable source, and the source's conviction that the information provided is certain. In this case, $m(A) = 1$, $\alpha_Y = 1$ and $m(r) = 1$, so that $m'(A) = 1$ and $m_t(A) = m(r) \cdot m'(A) = 1$. The result will be in the form $(A, v(A))$ initially provided by the source.

This scenario illustrates an ideal situation for trust assessment, where the source is trustworthy and well known to the analyst, and observations are reported in perfect conditions. As shown in table I, there is no uncertainty induced by the source, and once fusion is performed the items impacted show high values for conclusiveness and completeness, while dissonance is 0 for the updates BBAs for values, source's reliability and estimated trust, as shown in table II.

Scenario 2—Uncertain utterances: Y is considered by X to be a reliable source and reports the assertion A, while showing a low level of certainty $v(A)$ about the veracity of A. This example is relevant for situations where a reliable source provides (possibly) inaccurate descriptions of events due to, say, bad conditions for observation. This scenario corresponds by example to

TABLE III.
Uncertain uttering: Input uncertainty

Uncertainty of inputs		
Observation	Objectivity ObservationalSensitivity	0.3 0.9
Reporting	SelfConfidence Ambiguity	0.6 0.38

TABLE IV.
Uncertain utterance: fusion uncertainty

Fusion uncertainty	Dissonance	Conclusiv.	Complet.
Updates BBAs	0.3	0.7	0.9
Reliability	0.02	0.98	0.98
Trust	0.59	0.41	0.54

the following case for inputs: $\alpha_Y = 0.6$

$$m(A) = 0.8, m(\bar{A}) = 0.1, m(A \cup \bar{A}) = 0.1$$

$$m_C(r) = 0.9, m_C(\bar{r}) = 0, m_C(r \cup \bar{r}) = 0.1$$

$$m_R(r) = 0.9, m_R(\bar{r}) = 0, m_R(r \cup \bar{r}) = 0.1$$

$$m_I(r) = 0.3, m_I(\bar{r}) = 0.3, m_I(r \cup \bar{r}) = 0.6$$

and $w_C = 0.5$, $w_R = 0.5$ and $w_I = 0$. This results in

$$m'(A) = 0.48, m'(\bar{A}) = 0.06, m'(A \cup \bar{A}) = 0.46$$

and

$$m(r) = 0.9846, m(\bar{r}) = 0, m(r \cup \bar{r}) = 0.0154$$

Therefore, one finally obtains the trust valuation

$$m_t(A) = 0.47, m_t(\bar{A}) = 0.05, m_t(A \cup \bar{A}) = 0.46$$

This case shows that self-confidence has an important impact on the values of discounted BBA, as $m'(A)$ is decreased from 0.8 to 0.48, and thus the remaining mass is redistributed on $m'(A \cup \bar{A})$.

The combination of competence, reliability and intention are in line with the assumption of the scenario, which states that Y is a reliable source. After normalization, values for trust assessment clearly highlight the impact of uncertain utterances, as the BBA shows a mass transfer from $m_t(A)$ to $m_t(A \cup \bar{A})$. Still, values of trust are close to BBA integrating the self-confidence, which confirms the intuition that when the analyst X considers Y to be a reliable source, the assertion A is accepted with an overall trust level almost equal to the certainty level stated by the source.

This scenario illustrates uncertainty induced by observations failures, as *Objectivity*, and *SelfConfidence* are low, see table III.

While the quality of the source is highlighted by high values of *Conclusiveness* and *Completeness*, showing the analyst's confidence in the reports analyzed, the impact of imperfect observation is shown in the overall estimation of trust, through a combination of *Dissonance*, *Conclusiveness* and *Completeness* which have values close to 0.5, see table IV.

Scenario 3—Reputation: Suppose that Y provides A and $v(A)$ and X has no global description of Y in terms of reliability. As the reliability of Y is not available, Y's reputation will be used instead, as derived from historical data and previous failures. This scenario corresponds by example to the following case for inputs: $\alpha_Y = 1$

$$m(A) = 0.8, m(\bar{A}) = 0.1, m(A \cup \bar{A}) = 0.1$$

$$m_C(r) = 0.1, m_C(\bar{r}) = 0.1, m_C(r \cup \bar{r}) = 0.8$$

$$m_R(r) = 0.9, m_R(\bar{r}) = 0.1, m_R(r \cup \bar{r}) = 0$$

$$m_I(r) = 0.1, m_I(\bar{r}) = 0.1, m_I(r \cup \bar{r}) = 0.8$$

and $w_C = 0.1$, $w_R = 0.8$ and $w_I = 0.1$. Hence, one gets

$$m'(A) = 0.8, m'(\bar{A}) = 0.1, m'(A \cup \bar{A}) = 0.1$$

and

$$m(r) = 0.94, m(\bar{r}) = 0.01, m(r \cup \bar{r}) = 0.03$$

Therefore, one finally obtains the trust valuation

$$m_t(A) = 0.75, m_t(\bar{A}) = 0.09, m_t(A \cup \bar{A}) = 0.14$$

For this scenario, the source is confident about their own assertions, and therefore

$$m(A) = 0.8, m(\bar{A}) = 0.1, m(A \cup \bar{A}) = 0.1$$

and

$$m'(A) = 0.8, m'(\bar{A}) = 0.1, m'(A \cup \bar{A}) = 0.1$$

have identical BBA distributions. The reliability of the source is built namely on its reputation, as there are clues about the competence and intentions of the source. Hence, the overall BBA

$$m(r) = 0.9449, m(\bar{r}) = 0.0196, m(r \cup \bar{r}) = 0.0355$$

is close to the initial reputation distribution

$$m_R(r) = 0.9, m_R(\bar{r}) = 0.1, m_R(r \cup \bar{r}) = 0$$

Values of trust show the impact of using not completely reliable sources, which decreased the certainty level of the initial BBA

$$m'(A) = 0.8, m'(\bar{A}) = 0.1, m'(A \cup \bar{A}) = 0.1$$

to

$$m_t(A) = 0.75, m_t(\bar{A}) = 0.09, m_t(A \cup \bar{A}) = 0.14$$

TABLE V.
Reputation: input uncertainty

Uncertainty of inputs		
Observation	Objectivity ObservationalSensitivity	0.10 0.10
Reporting	SelfConfidence Ambiguity	1 0.60

TABLE VI.
Reputation: fusion uncertainty

Fusion uncertainty	Dissonance	Conclusiv.	Complet.
Updated BBAs	0.30	0.70	0.90
Reliability	0.07	0.93	0.95
Trust	0.34	0.66	0.84

They also support the intution that the trust assigned by the analyst to A will have an upper limit equal to the reputation of the source.

This scenario is similar the previous one as, in both cases, there are incomplete descriptions of the source. For this particular case, a historical recording of source's failures offers a basis to overcome the missing pieces and, in spite of low values for *Objectivity* and *ObservationalSensitivity* (see table V), the final trust evaluation is improved with respect to the previous scenario and shows a better combination of *Dissonance*, *Conclusiveness* and *Completeness*, as shown in table VI.

Scenario 4—Misleading report: In this case, Y provides the assertion A, while stating that it certainly holds and X considers Y to be a completely unreliable source. For this case, the analyst knows that the report is somehow inaccurate, for example, it cannot be corroborated or it contradicts, at least in part. information from other (more reliable) sources. The analyst suspects the source of having misleading intentions, and can therefore assign a maximal uncertainty level to the information reported. This scenario corresponds by example to the following case for inputs: $\alpha_Y = 1$

$$m(A) = 1, m(\bar{A}) = 0, m(A \cup \bar{A}) = 0$$

$$m_C(r) = 0.1, m_C(\bar{r}) = 0.1, m_C(r \cup \bar{r}) = 0.8$$

$$m_R(r) = 0.1, m_R(\bar{r}) = 0.1, m_R(r \cup \bar{r}) = 0.8$$

$$m_I(r) = 0.1, m_I(\bar{r}) = 0.8, m_I(r \cup \bar{r}) = 0.1$$

and $w_C = 0.1$, $w_R = 0.1$ and $w_I = 0.8$. Hence, one gets

$$m'(A) = 1, m'(\bar{A}) = 0, m'(A \cup \bar{A}) = 0$$

and

$$m(r) = 0.02, m(\bar{r}) = 0.91, m(r \cup \bar{r}) = 0.06$$

Therefore, one finally obtains as trust valuation

$$m_t(A) = 0.023, m_t(\bar{A}) = 0, m_t(A \cup \bar{A}) = 0.976$$

TABLE VII.
Misleading report: input uncertainty

Uncertainty of inputs		
Observation	Objectivity ObservationalSensitivity	0.10 0.10
Reporting	SelfConfidence Ambiguity	1.00 0.97

TABLE VIII.
Misleading: fusion uncertainty

Fusion uncertainty	Dissonance	Conclusiv.	Complet.
Assertion	0	1	1
Source	0.11	0.89	0.93
Trust	0.76	0.23	0.03

The values for this scenario reflect the high self-confidence of the source and high accuracy of the assertion provided; therefore, the initial BBA is unchanged after fusion with self-confidence. Nevertheless, the impact of having misleading intention is visible first on the mass distribution assigned to reliability and then on the overall values of trust. With respect to the initial values

$$m(A) = 1, m(\bar{A}) = 0, m(A \cup \bar{A}) = 0$$

and the partially fused ones

$$m'(A) = 1, m'(\bar{A}) = 0, m'(A \cup \bar{A}) = 0$$

the integration of a misleading source transfers the mass assignation almost exclusively to $m_t(A \cup \bar{A})$. Intuitively, the assertion A will be ignored, as the reliability of the source is dramatically decreased by a high mass assignment on misleading intentions.

This scenario illustrates the impact of misleading sources on trust estimation. Hence, the use case has very good values for reporting induced uncertainty, with high *SelfConfidence* and low *Ambiguity* (see table VII)), but the overall trust characterization shows strong *Dissonance*, corroborated with low *Conclusiveness* and near zero *Completeness*, as shown in table VIII.

Scenario 5—Ambiguous report: The source Y provides A and $v(A)$, the uncertainty level. Suppose that $v(A)$ has a low value, as the source is not very sure about the events reported, and that X considers Y to be unreliable. This scenario corresponds by example to the following case for inputs: $\alpha_Y = 0.3$

$$m(A) = 0.6, m(\bar{A}) = 0.2, m(A \cup \bar{A}) = 0.2$$

$$m_C(r) = 0.1, m_C(\bar{r}) = 0.8, m_C(r \cup \bar{r}) = 0.1$$

$$m_R(r) = 0.1, m_R(\bar{r}) = 0.8, m_R(r \cup \bar{r}) = 0.1$$

$$m_I(r) = 0.1, m_I(\bar{r}) = 0.1, m_I(r \cup \bar{r}) = 0.8$$

and $w_C = 0.2$, $w_R = 0.4$ and $w_I = 0.4$. Hence, one gets

$$m'(A) = 0.18, m'(\bar{A}) = 0.06, m'(A \cup \bar{A}) = 0.76$$

TABLE IX.
Ambiguous report: input uncertainty

Uncertainty of inputs		
Observation	Objectivity	0.10
	ObservationalSensitivity	0.10
Reporting	SelfConfidence	0.30
	Ambiguity	0.27

TABLE X.
Ambiguous report: fusion uncertainty

Fusion uncertainty	Dissonance	Conclusiv.	Complet.
Assertion	0.6	0.4	0.8
Source	0.583	0.417	0.47
Trust	0.973	0.027	0.006

and

$$m(r) = 0.02, m(\bar{r}) = 0.43, m(r \cup \bar{r}) = 0.53$$

Therefore, one finally obtains the trust valuation

$$m_t(A) = 0.0040, m_t(\bar{A}) = 0.0013$$

and

$$m_t(A \cup \bar{A}) = 0.9946$$

This scenario is an illustration for the worst practical case and is relevant when the analyst receives a report provided by a source that lacks the skills or competence to provide accurate descriptions of events. In this case, the reports are incomplete, ambiguous, or even irrelevant. In addition to low competence and reliability, the source himself is also unsure about the statement.

The first modification of BBA shows the strong impact of self-confidence, which changes drastically the BBA of the initial assertions, from

$$m(A) = 0.6, m(\bar{A}) = 0.2, m(A \cup \bar{A}) = 0.2$$

to

$$m'(A) = 0.18, m'(\bar{A}) = 0.06, m'(A \cup \bar{A}) = 0.76$$

Unsurprisingly, the overall reliability is low:

$$m(r) = 0.0223, m(\bar{r}) = 0.4398, m(r \cup \bar{r}) = 0.5379$$

and the results of the final combination show an important mass assigned to $m_t(A \cup \bar{A}) = 0.9946$. Intuitively, the information provided is useless, and considered as highly uncertain.

This scenario shows the combined effects of uncertain reporting and incomplete source description for trust estimation. First, the outcome is affected by high values of uncertainty induced during observation and reporting passes, table IX. Then, fusion leads to a trust estimation having high values of *Dissonance*, and very low values of *Conclusiveness* and *Completeness*.

The same criteria estimated for reliability show the main difference with respect to the previous case, which was also based on unreliable sources. While in scenario 4 the source still has important *Completeness*, this measure is drastically decreased for this scenario, as shown in table X.

VIII. STRENGTHS AND LIMITATIONS OF BELIEF-BASED FORMALIZATION FOR TRUST ASSESSMENT

This section discusses the strengths and limitations of the belief-based perspective in trust modeling in the light of results shown by previous scenarios. The main advantage of using belief functions is that the formalism is consistent with the cognitive perspective of trust adopted by the model, thanks to the notion of belief. It also captures uncertainties both of the analyst with respect to the source and of the source with respect to their own statements with different mechanisms. First, self-confidence is implemented thanks to a discounting coefficient, as, in practice, the values of self-confidence may rely upon linguistic clues of certainty/uncertainty that can be translated into numerical values. Second, the formalization introduces weighting factors in order to offer a flexible solution, which allow for situations in which the analyst has more or less complete knowledge about distinct attributes of the source, or wishes to emphasize one particular attribute. Moreover, the formalization is able to handle ignorance on various aspects, including missing data. The overall fusion mechanism performs trust estimation in several steps, which allows for a better traceability of the outcome and the mapping at different processing stages using URREF criteria. The results of these scenarios are in line with their specific hypotheses, reflecting the intuition that the fusion technique is appropriate for estimating trust.

As with any user-centric approach, the main limitation of the solution discussed in this paper is the lack of guidance for choosing the set of numerical values with which to instantiate the model. For example, two different analysts may choose differing mass distribution and weight coefficients with respect to the same source, and they may also use slightly different approaches to infer a numerical value from linguistic clues when handling self-confidence. Thus, the outcome depends crucially on the interventions of users and their ability to build a model able to capture the situation under analysis. Also, the solution requires preexisting knowledge about the source's reputation, competence, and intention, indeed, in practice, it is difficult to have access to information on those aspects. Provided that there is no other metadata or domain knowledge available for use, the model is likely to fail to produce an accurate trust evaluation in some contexts due to the shortage of knowledge on critical aspects.

As such, the belief-based formalization has limited capabilities to explain the outcome. To overcome this limitation, a mapping to URREF uncertainty criteria is used. The mapping highlights when uncertainties are added into the system and which partial results and affected. It facilitates the interpretation of results by adding additional information as to why the item is to be trusted or no; for example, whereas the fusion process outputs low values of trust for a given item, the mapping to URREF criteria allows to underline problems related to evidence collection or reporting, dissonance or incompleteness during the fusion stages.

As shown in previous scenarios, using a belief-oriented formalism and URREF criteria mapping offers a pragmatic approach to develop a more comprehensive and easy to interpret solution for trust estimation.

IX. CONCLUSION

This paper presents a computational model by which an analyst is able to assess trust in reported information based on several possible unknown attributes of the source as well as additional characterization of the informational content by the source itself. The paper also illustrates the use of URREF criteria to track uncertainty affecting the results, from model construction to its formalization with belief functions. First, a model for trust estimation has been developed that combines several attributes of sources and their own assessment of the items reported. The model is implemented using belief functions, and takes advantage of its mathematical background to define fusion operators for trust assessment. Several scenarios are presented to illustrate uncertainty analysis, illustrating when uncertainty occurs and how it affects partial results for different applications.

Tracking uncertainty is suitable for fusion systems in which various human sources send observations of questionable quality and there is a need to continuously update the trust associated with reports to be analyzed. The set of URREF criteria offers a unified basis to analyze inaccuracies affecting trust estimation during different phases: observation, reporting, and fusion. Select use cases clearly illustrated the benefits of managing uncertainties arising during the modeling and formalization phases, with the twofold analysis offering additional details on results and improving their interpretation.

The general approach taken in this paper could be adapted to investigate the general mechanisms by which fusion processes integrate information from multiple sources. The solution is especially useful for comparing different fusion approaches with respect to their implications for uncertainty management.

REFERENCES

[1] L. Amgoud and R. Demolombe
An argumentation-based approach for reasoning about trust in information sources.
Argument & Computation, 5(2–3):191–215, 2014.

[2] A. Auger and J. Roy
Expression of uncertainty in linguistic data.
In *Information Fusion, 2008 11th International Conference on*, pages 1–8. IEEE, 2008.

[3] K. S. Barber and J. Kim
Belief revision process based on trust: Simulation experiments.
In *Proceedings of Autonomous Agents '01 Workshop on Deception, Fraud, and Trust in Agent Societies*. Citeseer, 2001.

[4] L. Berti-Équille and J. Borge-Holthoefer
Veracity of data: From truth discovery computation algorithms to models of misinformation dynamics.
Synthesis Lectures on Data Management, 7(3):1–155, 2015.

[5] E. Blasch and A. Aved
Urref for veracity assessment in query-based information fusion systems.
In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 58–65. IEEE, 2015.

[6] E. Blasch, P. C. Costa, K. B. Laskey, H. Ling, and G. Chen
The urref ontology for semantic wide area motion imagery exploitation.
In *Aerospace and Electronics Conference (NAECON), 2012 IEEE National*, pages 228–235. IEEE, 2012.

[7] E. Blasch, K. B. Laskey, A.-L. Jousselme, V. Dragos, P. C. Costa, and J. Dezert
Urref reliability versus credibility in information fusion (stanag 2511).
In *Information Fusion (FUSION), 2013 16th International Conference on*, pages 1600–1607. IEEE, 2013.

[8] E. Blasch, A. Jøsang, J. Dezert, P. C. Costa, and A.-L. Jousselme
Urref self-confidence in information fusion trust.
In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[9] D. B. Buller and J. K. Burgoon
Interpersonal deception theory.
Communication theory, 6(3):203–242, 1996.

[10] C. Castelfranchi and R. Falcone
Trust is much more than subjective probability: Mental components and sources of trust.
In *System Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on*, pages 10–pp. IEEE, 2000.

[11] C. Castelfranchi, R. Falcone, and G. Pezzulo
Trust in information sources as a source for trust: a fuzzy approach.
In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 89–96. ACM, 2003.

[12] L. Cholvy
How strong can an agent believe reported information?
In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 386–397. Springer, 2011.

[13] P. C. Costa, K. B. Laskey, E. Blasch, and A.-L. Jousselme
Towards unbiased evaluation of uncertainty reasoning: The urref ontology.
In *Information Fusion (FUSION), 2012 15th International Conference on*, pages 2301–2308. IEEE, 2012.

[14] L. Cronk
Communication as manipulation: Implications for biosociological research.
In *Proceedings of the American Sociological Association Annual Meetings*, 1991.

[15] J. P. de Villiers, K. Laskey, A.-L. Jousselme, E. Blasch, A. de Waal, G. Pavlin, and P. Costa
Uncertainty representation, quantification and evaluation for data and information fusion.
In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 50–57. IEEE, 2015.

[16] J. P. de Villiers, A.-L. Jousselme, A. de Waal, G. Pavlin, K. Laskey, E. Blasch, and P. Costa
Uncertainty evaluation of data and information fusion within the context of the decision loop.
In *Information Fusion (FUSION), 2016 19th International Conference on*, pages 766–773. IEEE, 2016.

[17] J. P. de Villiers, R. Focke, G. Pavlin, A.-L. Jousselme, V. Dragos, K. Laskey, P. Costa, and E. Blasch
Evaluation metrics for the practical application of urref ontology: An illustration on data criteria.
In *Information Fusion (Fusion), 2017 20th International Conference on*, pages 1–8. IEEE, 2017.

[18] J. P. de Villiers, G. Pavlin, P. Costa, A.-L. Jousselme, K. Laskey, V. Dragos, and E. Blasch
Subjects under evaluation with the urref ontology.
In *Information Fusion (Fusion), 2017 20th International Conference on*, pages 1–8. IEEE, 2017.

[19] J. Dezert, W. Pei, and A. Tchamova
On the validity of dempster-shafer theory.
In *Information Fusion (Fusion), 2012 15th International Conference on*. IEEE, 2012.

[20] X. L. Dong, E. Gabrilovich, G. Heitz, W. Horn, K. Murphy, S. Sun, and W. Zhang
From data fusion to knowledge fusion.
Proceedings of the VLDB Endowment, 7(10):881–892, 2014.

[21] X. L. Dong, E. Gabrilovich, K. Murphy, V. Dang, W. Horn, C. Lugaresi, S. Sun, and W. Zhang
Knowledge-based trust: Estimating the trustworthiness of web sources.
Proceedings of the VLDB Endowment, 8(9):938–949, 2015.

[22] V. Dragos
An ontological analysis of uncertainty in soft data.
In *Information Fusion (FUSION), 2013 16th International Conference on*, pages 1566–1573. IEEE, 2013.

[23] V. Dragos and K. Rein
Integration of soft data for information fusion: Pitfalls, challenges and trends.
In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[24] V. Dragos, X. Lerouvreur, and S. Gatepaille
A critical assessment of two methods for heterogeneous information fusion.
In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 42–49. IEEE, 2015.

[25] V. Dragos, Z. Juergen, and J. P. de Villiers
Applicaton of urref criteria to assess knowledge representation in cyber threats models.
In *Information Fusion (FUSION), 2018 21 International Conference on*. IEEE, 2018.

[26] P. Everaere, S. Konieczny, and P. Marquis
Belief merging versus judgment aggregation.
In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 999–1007. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[27] R. Falcone and C. Castelfranchi
Social trust: A cognitive approach.
In *Trust and deception in virtual societies*. Springer, 2001.

[28] R. Falcone, M. Piunti, M. Venanzi, and C. Castelfranchi
From manifesta to krypta: The relevance of categories for trusting others.
ACM Transactions on Intelligent Systems and Technology (TIST), 4(2):27, 2013.

[29] R. Falcone, A. Sapienza, and C. Castelfranchi
The relevance of categories for trusting information sources.
ACM Transactions on Internet Technology (TOIT), 15(4):13, 2015.

[30] J. Granatyr, V. Botelho, O. R. Lessing, E. E. Scalabrin, J.-P. Barthès, and F. Enembreck
Trust and reputation models for multiagent systems.
ACM Computing Surveys (CSUR), 48(2):27, 2015.

[31] J. Guo and R. Chen
A classification of trust computation models for service-oriented internet of things systems.
In *Services Computing (SCC), 2015 IEEE International Conference on*, pages 324–331. IEEE, 2015.

[32] D. L. Hall and J. M. Jordan
Human-centered information fusion.
Artech House, 2010.

[33] J. T. Hancock, L. E. Curry, S. Goorha, and M. Woodworth
On lying and being lied to: A linguistic analysis of deception in computer-mediated communication.
Discourse Processes, 45(1):1–23, 2007.

[34] R. Hermoso, H. Billhardt, and S. Ossowski
Trust-based role coordination in task-oriented multiagent systems.
Knowledge-Based Systems, 52:78–90, 2013.

[35] D. Jean, W. Pei, and A. Tchamova
On the validity of dempster's fusion rule and its interpretation as a generalization of bayesian fusion rule.
International Journal of Intelligent Systems, Special Issue: Advances in Intelligent Systems, 29(3):223–252, 2014.

[36] A. Josang, R. Ismail, and C. Boyd
A survey of trust and reputation systems for online service provision.
Decision Support Systems, 43(2):618–644, 2007.

[37] A.-L. Jousselme
Semantic criteria for the assessment of uncertainty handling fusion models.
In *Information Fusion (FUSION), 2016 19th International Conference on*, pages 488–495. IEEE, 2016.

[38] A.-L. Jousselme and G. Pallotta
Dissecting uncertainty-based fusion techniques for maritime anomaly detection.
In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 34–41. IEEE, 2015.

[39] A.-L. Jousselme, A.-C. Boury-Brisset, B. Debaque, and D. Prevost
Characterization of hard and soft sources of information: A practical illustration.
In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[40] S. Kent
Words of estimative probability.
Studies in Intelligence, 8(4):49–65, 1964.

[41] R. F. Kesselman
Verbal probability expressions in national intelligence estimates: a comprehensive analysis of trends from the fifties through post 9/11.
PhD thesis, MERCYHURST COLLEGE, 2008.

[42] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi
Multisensor data fusion: A review of the state-of-the-art.
Information Fusion, 14(1):28–44, 2013.

[43] H. Koen, J. P. de Villiers, G. Pavlin, A. de Waal, P. de Oude, and F. Mignet
A framework for inferring predictive distributions of rhino poaching events through causal modelling.
In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–7. IEEE, 2014.

[44] A. Koster, A. L. Bazzan, and M. de Souza
Liar liar, pants on fire; or how to use subjective logic and argumentation to evaluate information from untrustworthy sources.
Artificial Intelligence Review, 48(2):219–235, 2017.

[45] T. Lukoianova and V. L. Rubin
Veracity roadmap: Is big data objective, truthful and credible?
Advances in Classification Research Online, 24(1):4–15, 2014.

[46] P.-A. Matt, M. Morge, and F. Toni
Combining statistics and arguments to compute trust.
In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1–Volume 1*, pages 209–216. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[47] D. Mercier, B. Quost, and T. DenÅ\uxa
Refined modeling of sensor reliability in the belief function framework using contextual discounting.
Information Fusion, 9(2):246–258, 2008.

[48] M. Momani, S. Challa, and R. Alhmouz
Bnwsn: Bayesian network trust model for wireless sensor networks.
In *Communications, Computers and Applications, 2008. MIC-CCA 2008. Mosharaka International Conference on*, pages 110–115. IEEE, 2008.

[49] NATO
STANAG 2011.
Technical report, Intelligence reports, 2003.

[50] F. Paglieri, C. Castelfranchi, C. da Costa Pereira, R. Falcone, A. Tettamanzi, and S. Villata
Trusting the messenger because of the message: feedback dynamics from information quality to source evaluation.
Computational and Mathematical Organization Theory, 20(2):176–194, 2014.

[51] S. Parsons, K. Atkinson, Z. Li, P. McBurney, E. Sklar, M. Singh, K. Haigh, K. Levitt, and J. Rowe
Argument schemes for reasoning about trust.
Argument & Computation, 5(2–3):160–190, 2014.

[52] G. Pigozzi
Belief merging and judgment aggregation.
Stanford Encyclopedia of Philosophy, 2015.

[53] K. Rein
I think it is likely that it might be so—Exploiting Lexical Clues for the Automatic Generation of evidentiality Weights for Information Extracted from English Text.
Dissertation, University of Bonn, 2016.

[54] V. L. Rubin
Stating with certainty or stating with doubt: Intercoder reliability results for manual annotation of epistemically modalized statements.
In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 141–144. Association for Computational Linguistics, 2007.

[55] R. Saurí and J. Pustejovsky
Are you sure that this happened? assessing the factuality degree of events in text.
Computational Linguistics, 38(2):261–299, 2012.

[56] G. Shafer
A mathematical theory of evidence,
volume 1. Princeton university press Princeton, 1976.

[57] W. Sherchan, S. Nepal, and C. Paris
A survey of trust in social networks.
ACM Computing Surveys (CSUR), 45(4):47, 2013.

[58] M. B. Sinai, N. Partush, S. Yadid, and E. Yahav
Exploiting social navigation.
arXiv preprint arXiv:1410.0151, 2014.

[59] F. Smarandache and J. Dezert
Advances and applications of DSmT for information fusion,
volume 1–4. American Research Press (ARP), 2015.

[60] F. Smarandache, J. Dezert, and J.-M. Tacnet
Fusion of sources of evidence with different importances and reliabilities.
In *Information Fusion (FUSION), 2010 13th Conference on*, pages 1–8. IEEE, 2010.

[61] F. Smarandache, J. Dezert, and J.-M. Tacnet
Fusion of sources of evidence with different importances and reliabilities.
In *Information Fusion (Fusion), 2012 15th International Conference on*. IEEE, 2010.

[62] Y. Tang, K. Cai, P. McBurney, E. Sklar, and S. Parsons
Using argumentation to reason about trust and belief.
Journal of Logic and Computation, 22(5):979–1018, 2011.

[63] D. P. Twitchell, D. P. Biros, M. Adkins, N. Forsgren, J. K. Burgoon, and J. F. Nunamaker
Automated determination of the veracity of interview statements from people of interest to an operational security force.
In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS'06)*, volume 1, pages 17a–17a. IEEE, 2006.

[64] M. Venanzi, A. Rogers, and N. R. Jennings
Trust-based fusion of untrustworthy information in crowdsourcing applications.
In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems*, pages 829–836. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[65] S. Villata, G. Boella, D. M. Gabbay, and L. Van Der Torre
Arguing about the trustworthiness of the information sources.
In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 74–85. Springer, 2011.

Dr. Valentina Dragos is a research scientist, member of the Department of Information Modeling and Systems at ONERA, The French Aerospace Lab in Palaiseau, France. Valentina received Master and PhD degrees in Computer Science from Paris V University and her research interests include artificial intelligence, with emphasis on natural language processing, semantics technologies and automated reasoning. Since joining ONERA in 2010, Valentina contributed to several academic and industrial security-oriented projects, addressing topics such as: semantic interoperability for command and control systems, heterogeneous information fusion, and exploitation of semantic data (HUMINT, OSINT) for situation assessment.

Jean Dezert was born in l'Hay les Roses, France in 1962. He received D.E.A. degree from University Paris VII (Jussieu), Paris, in 1986, and the Ph.D.degree from University Paris XI, Orsay, France, in 1990, all in automatic control and signal processing. He is a Senior Research Scientist with the Information and Fusion Systems Research Team, in Information Processing and Systems (DTIS), ONERA, Palaiseau, France. He gave several invited plenary talks on information fusion in Europe, America, Australia, and China, during latest years. His current research interests include autonomous navigation, estimation, stochastic systems and their applications to multisensormultitarget tracking, information fusion, and plausible reasoning. He is cofounder with Prof. Smarandache of Dezert-Smarandache Theory (DSmT) for information fusion and has published around 180 papers in his research fields in scientific journals and in proceedings of international conferences, and five books. Dr. Dezert has served as a Local Arrangements Organizer for the 2000 3rd International Conference on Information Fusion (Fusion) in Paris, and a Secretary, an Executive Vice-President, and the President for the International Society of Information Fusion (ISIF) in 2001, 2004, and 2016, respectively. He has been involved in the Technical Program Committee of Fusion International Conferences.

Kellyn Rein is a Research Associate at the Research Institute for Communication, Information Processing Ergonomics (FKIE) at Fraunhofer Gesellschaft, Europe's largest applied research organization. She received her PhD from the University of Bonn in Linguistics, after a BA at Michigan State University, and two masters, one in Management and one in Computer Information Systems, both with joint conferral from Boston University (1984) and Vrije Universiteit Brussel in Belgium. Her area of research interest is in the analysis of uncertainty in natural language information, information fusion and information quality. She has been active in EU projects focused variously on crisis management, cross-border organized crime, smart cities, and CBRN terrorism. She has also been involved in numerous NATO Research Task Groups, including one on analysis and communication of uncertainty in intelligence, mission assurance and cyber risk assessment for unmanned autonomous system, and currently chairing a task group on multi-level, multi-source information fusion.

Categories of belief fusion

AUDUN JØSANG

Belief Fusion consists of merging beliefs about a domain of interest from multiple separate sources. No single belief fusion method is adequate for all categories of situations, hence the challenge is to determine which belief fusion method is the most appropriate for a given situation. The conclusion to be drawn from this discussion is that the analyst must first understand the dynamics of the situation at hand in order to find the best fusion method for analysing it. The aim of this article is first to demonstrate that there are appropriate situations to use belief fusion, and that different mathematical fusion operators are required for the different situations. Secondly we propose criteria than can be applied to identify the various categories of fusion situations, and describe specific belief fusion operators that are suitable for modeling the fusion situations in each category.

Manuscript received February 20, 2018; revised October 19, 2018 and October 27, 2018; released for publication January 25, 2019.

Refereeing of this contribution was handled by Johan de Villiers.

Author's address: University of Oslo, Norway (E-mail: josang@mn.uio.no).

1557-6418/18/$17.00

1. INTRODUCTION

When analyzing hypotheses about specific domains of interest there is often a need to combine evidence from multiple sources. This principle belongs to information fusion in general, and is called belief fusion when the evidence is represented as belief. It is important to realise that there is no single fusion method that is suitable for analyzing all situations of belief fusion. It is also quite challenging to determine the best belief fusion method for a specific situation, and there has been considerable confusion around this issue in the literature. It is therefore crucial to have a consistent method for categorising different situations of belief fusion, and to apply this method for selecting the most suitable belief-fusion operator for each category of situations.

Beliefs are represented as subjective opinions throughout this article. A subjective opinion generalises the traditional representations of belief functions by including a base rate distribution over the values of the domain variable. A domain of interest contains the possible hypotheses or states that the analyst is interested in, e.g. for identifying the hypothesis which correspond best with reality. A subjective opinion is denoted ω_X^C, where C represents the source of the opinion and X represents the variable of the opinion's object/target domain.

In general, the source of an opinion can be a human, or it can be a sensor which produces data which in turn can form the basis an opinion. Multiple separate sources, e.g. denoted $C_1, C_2, \ldots C_N$, can produce different and possibly conflicting opinions $\omega_X^{C_1}, \omega_X^{C_2}, \ldots \omega_X^{C_N}$ about the same variable X. In this situation, source fusion consists of merging the different sources into a single source that can be denoted $\diamond(C_1, C_2, \ldots C_N)$, and mathematically fusing their opinions into a single opinion denoted $\omega_X^{\diamond(C_1,C_2,\ldots C_N)}$ which then represents the opinion of the merged sources. The source merger function is here denoted by the symbol '$\diamond$', and the general belief-fusion principle is illustrated in Figure 1.

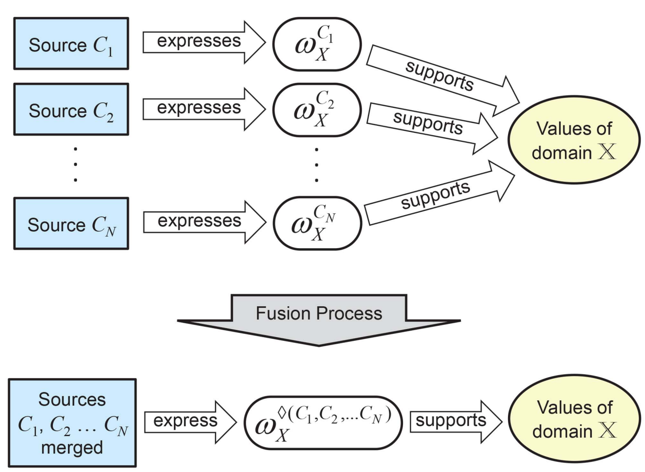

Fig. 1. Belief-fusion principle

Different belief fusion situations can vary significantly and semantically depending on the purpose and nature of the fusion process, and hence require different fusion operators. However, it can be challenging to identify the correct or most suitable fusion operator for a specific situation. In general, a given fusion operator is unsuitable when it produces wrong results in some instances of a situation, even if it produces correct results in most instances of the situation. A fusion operator should produce sound and intuitive results in all realistic instances of the situation to be analysed.

In order to see the importance of using the correct belief fusion method in a given fusion situation it is instructive to consider other situation types where the effect of applying the correct or incorrect formal model and method is more obvious. First, consider the situation of predicting the physical strength of a steel chain, where the classical and correct model is that of the weakest link. Then, consider the situation of determining the competitive strength of a relay swimming team, for which an adequate model is the average strength of each swimmer on the team, in terms of how fast each swimmer can swim.

Applying the weakest-link model (i.e. the slowest swimmer) to predict the overall speed of the relay swimming team is an approximation which might give a relatively good prediction in most instances of high-level swimming championships. However, it is obviously an incorrect model and would produce rather unreliable predictions if there are large variations in speed between the swimmers in a relay swimming team.

Similarly, applying the average strength model for assessing the physical strength of the steel chain represents an approximation which would produce relatively good strength predictions in most instances of high-quality steel chains where the link strength is highly uniform. However, it is obviously a very poor model which would be unreliable in general, and which could have fatal consequences if life depended on it.

These examples illustrate the inadequacy of anecdotal examples for determining whether the weakest-link model is suitable for predicting the strength of relay swimming teams. Similarly it is insufficient to simply use a few anecdotal examples to test whether the averaging principle is adequate for modelling the strength of steel chains. Without a clear understanding of the situation to be modelled, the analyst does not have a basis for selecting the correct and appropriate model. The selection of appropriate models might be obvious for the simple examples above, but it can be challenging to judge whether a fusion operator is suitable for a specific situation of belief fusion [1].

The conclusion to be drawn from this discussion is that the analyst must first understand the dynamics of the situation at hand in order to find the best model for analysing it.

The aim of this article is first to demonstrate that there can be many different categories of situations of belief fusion, and that different mathematical fusion operators are required for the computation of belief fusion in the different categories of belief-fusion situations. Secondly we propose criteria for identifying the various categories of fusion situations, and describe specific belief fusion operators that are suitable for belief fusion in each category.

This work forms part of the effort to define "Evaluation of Techniques for Uncertainty Representation" under the ETUR Working Group [2] where the URREF ontology is one of the reference documents [3], [4]. Previous work on defining categories for belief fusion is described in [1], [5]. The contribution of the current work is to generalise and define new operators for belief fusion, and to clarify the understanding of fusion categories. Belief fusion belongs to the domain of high-level fusion [6] in contrast to other types of low-level data fusion.

Section 2 describes a set of belief-fusion categories The criteria defined in Section 3 then describe how a given fusion situation can be understood and categorised. Section 6 describes corresponding fusion operators for the respective categories. Section 7 provides numerical examples to compare the different fusion operators, and Section 8 discusses the implications of the categories of belief fusion presented in this article.

2. CATEGORIES OF FUSION SITUATIONS

Situations of belief fusion take belief arguments from multiple sources through a fusion process to produce a single belief argument. More specifically, a fusion situation is characterised by a domain of two or more state values, and the various sources' different belief arguments about these values. The domain of state values can be interpreted as a set of competing hypotheses, where it is assumed that only one value/hypothesis is TRUE at any one time. Each belief argument can assign belief mass to one or several state values, which thereby represents support for those values in terms of which values are believed to be TRUE. The purpose of belief fusion is to produce a new belief argument that reflect the sources' collective set of belief arguments in the most fair or correct way. It is then assumed that the fused belief argument supports the most correct, acceptable or most preferred value, when seen from the perspective of the collective set of sources.

It is often challenging to determine the correct or the most appropriate fusion operator for a specific situation. Our approach of addressing this challenge is to define categories of similar situations according to their typical characteristics, which then allows to determine a suitable belief fusion operator for each category. Four

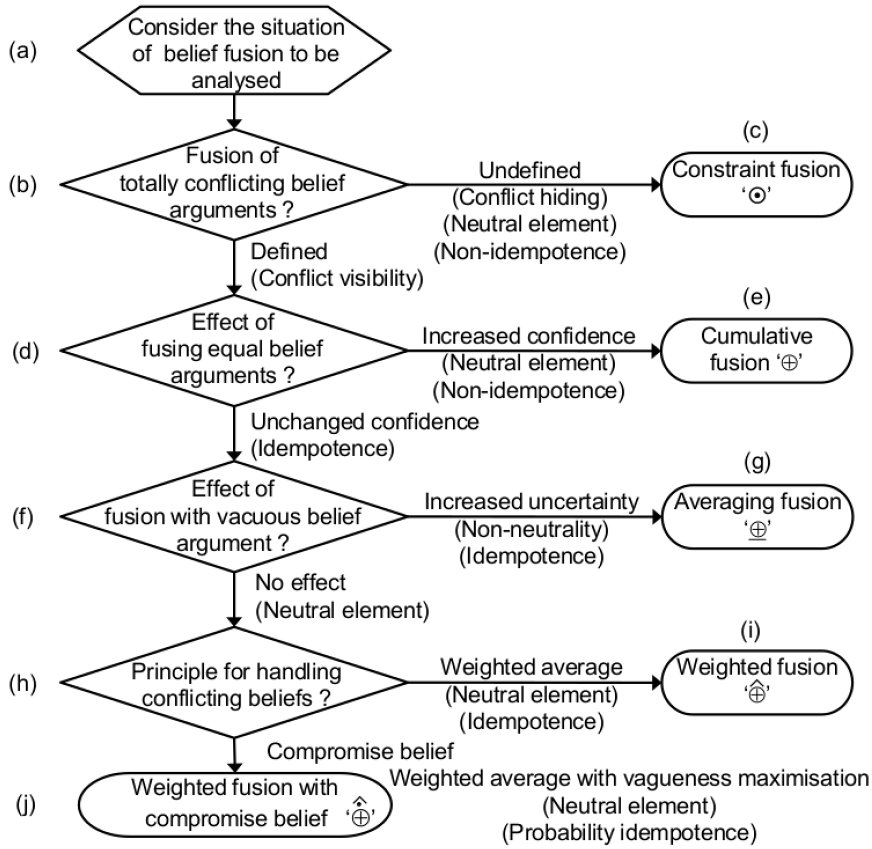

Fig. 2. Procedure for selecting a suitable belief-fusion operator for each category

distinct categories as well as one hybrid category of fusion situations are described below.

- **Belief Constraint Fusion** (BCF) is suitable when assuming that: 1) belief arguments must not be wrong (sources are totally reliable), and 2) there is no compromise in case of totally conflicting arguments, hence the fusion result is not defined in that case. In some situations these properties are desirable. An example is when two persons try to agree on seeing a movie at the cinema. If their preferences share common movies they can decide to watch one of them. Yet, if their preferences have no movies in common then there is no solution, with the rational consequence they will not watch any movie together. BCF is described in Section 6.1.
- **Cumulative Belief Fusion** (CBF) is suitable when assuming that the amount of independent evidence increases by including more and more sources. For example, when different independent biometric sensors (e.g. fingerprint, voice, face) are being used to authenticate a person, the results from each sensor can be fused with CBF, which produces an opinion with decreasing uncertainty (increased assurance) about the identity of the person. CBF has the vacuous opinion as neutral element, but is not idempotent. CBF is described in Section 6.2. A modification of CBF is when it is assumed or desired that the fusion process produces uncertainty maximised opinions. It is then possible to apply uncertainty maximisation after CBF, which is called CBF-UM for short. This could e.g. be when witnesses express their opinions about whether Oswald shot Kennedy, which when fused with CBF-UM produces an epistemic opinion about who shot him. CBF-UM is described in Section 6.5.
- **Averaging Belief Fusion** (ABF) is suitable when dependence between sources is assumed, so that including more sources does not necessarily add more evidence behind the fused belief, it just changes the average distribution of evidence. In case of equal belief arguments, the fused result should be the same, which means that idempotence is assumed. An example of this type of situation is when a jury tries to reach a verdict after having observed the court proceedings. It is also assumed that a vacuous belief argument does have an influence on the fused result, which means that ABF does not have a neutral element. This is interpreted in the sense that the source of the vacuous belief argument says: *"I do not see any evidence and therefore do not have any belief about this, and I want my vacuous argument belief to be reflected in the fused output belief"*. ABF is described in Section 6.3.
- **Weighted Belief Fusion** (WBF) is also suitable when dependence between sources is assumed, so that adding more sources does not necessarily add more evidence in total. Equal belief arguments should produce equal fused belief, meaning that idempotence

is assumed. However, it is assumed that a vacuous belief argument has no influence on the fused result, meaning that WBF does have a neutral element in the form of vacuous belief. This is interpreted in the sense that the source of a vacuous belief argument says: *"I do not see any evidence and therefore do not have any belief about this, and I will let the sources that do have evidence and belief about this determine the fused belief without me"*. An example of this type of situation is when experts (e.g. medical doctors) express multinomial opinions about a set of hypothesis (e.g. diagnoses). WBF is described in Section 6.4. In case of hyper-opinions WBF does not identify shared (vague) belief on overlapping (composite) values in the domain, and simply computes the weighted average.

- **Weighted Belief Fusion with Vagueness Maximisation** (WBF-VM) can be used when the analyst naturally wants to preserve shared beliefs from different sources, and to transform conflicting beliefs into vague belief. In this way shared belief is preserved when it exists, and compromise vague belief is formed when necessary. In the case of totally conflicting beliefs, then the resulting fused belief becomes vague. WBF-VM is probability-idempotent, commutative and has the vacuous belief argument as neutral element. Probability-idempotence means that the projected probability distribution is preserved when fusing equal opinions, but the fused opinion will in general have different vague belief. A situations where WBF-VM is suitable is when experts (e.g. medical doctors) express hyper-opinions about a set of hypothesis (e.g. diagnoses). WBF-VM takes into account shared (vague) belief on overlapping (composite) values, and is therefore suitable for preserving shared beliefs when fusing hyper-opinions. WBF-VM is described in Section 6.6.

The subtle differences between the fusion situations above illustrate the challenge of modelling them correctly. For instance, consider the task of determining the location of a mobile phone subscriber at a specific point in time by collecting location evidence from a base station, in which case it seems natural to use belief constraint fusion. If two adjacent base stations detect the subscriber, then the belief constraint operator can be used to locate the subscriber within the overlapping region of the respective radio cells. However, if two base stations far apart detect the subscriber at the same time, then the result of belief constraint fusion is not defined so there is no conclusion. With additional assumptions, it would still be reasonable to think that the subscriber is probably located in one of the two cells, but not which one in particular, and that the case needs further investigation because the inconsistent signals might be caused by an error in the system. Some method of trust revision [7] can be applied in this situation.

3. CRITERIA FOR IDENTIFYING FUSION CATEGORIES

While having multiple fusion categories can help in scoping the solution space, there is still the issue of determining which category a specific situation belongs to. In order to select the correct or most adequate fusion method the analyst must consider a set of assumptions about the fusion situation to be analysed and for each assumption judge whether it is applicable. The most adequate fusion method is then identified as a function of the set of assumptions that applies to the situation to be analysed. This procedure for identifying and selecting the most appropriate fusion operator is illustrated in Figure 2. The steps in the selection procedure are further described below.

(a) The analyst first needs a good understanding of the situation to be modelled in order to select the most suitable fusion operator. This includes being able to make the binary choices of (b), (d), (f) and (h) below.
(b) Shall it be possible to fuse totally conflicting beliefs?
(c) In case it is assumed that two totally conflicting belief arguments should leave no room for compromise, then BCF (Belief Constraint Fusion) is probably the most suitable operator. BCF is not defined in case of totally conflicting belief or preference arguments, which reflects the assumption that there is no compromise solution in case of total conflict.
(d) Is idempotence assumed, i.e. should two equal belief arguments produce the same output belief?
(e) In case idempotence is not assumed, then CBF (Cumulative Belief Fusion) is probably the most suitable operator. CBF is suitable when non-idempotent is assumed, meaning that equal belief arguments represent independent support for specific values of the variable, which thereby contribute to reducing the uncertainty in the output belief. In addition to being non-idempotent, CBF can handle totally conflicting opinions, as required for this category.
(f) Should a vacuous belief argument have any influence on the output fusion result?
(g) In case it is assumed that a vacuous belief arguments shall influence the output, then no neutral element exists, which indicates that ABF (Averaging Belief Fusion) is a suitable operator. ABF can be meaningful e.g. for making a survey of opinions where vacuity (lack of belief) in a belief argument shall be reflected as less confidence in the output fused belief.
(h) How should conflicting belief be handled?
(i) The simplest belief conflict management principle is to compute the weighted average of conflicting

belief mass. WBF (Weighted Belief Fusion) is suitable for fusing multinomial opinions, but less so for fusing hyper-opinions because the operator is blind to common belief between two vague belief arguments which assign belief mass to partially overlapping composite values.

(j) In case it is assumed that conflicting belief mass should be transformed into compromise (vague) belief then WBF-VM is suitable, i.e. it would be adequate to apply vagueness maximisation (VM) after the weighted belief fusion (WBF). In contrast to simple WBF, the post-processing with vagueness maximisation takes into account and reflects common belief aspects between different opinion arguments, which often better reflects human intuition.

It can be difficult to tell which category a specific situation belongs to. In addition, the choice of fusion operator can also be influenced by the type of fusion result the analyst wants to obtain, which e.g. could be to have an uncertainty-maximised or vagueness-maximised fused opinion.

The various belief fusion operators corresponding to each category in Figure 2 are described in Section 6 below. Before delving into the the formalism of belief fusion operators it is necessary to first describe the representation of subjective opinions and the corresponding Dirichlet PDF (Probability Density Function).

4. SUBJECTIVE OPINIONS

This section describes subjective opinions which represent beliefs over random variables in subjective logic.

In the formalism of subjective logic, a *domain* is a state space of values which can represent e.g. observable or hidden states, events, hypotheses or propositions [5]. A variable X associated with a domain $\mathbb{X}$ can take values $x \in \mathbb{X}$. A variable with an associated probability distribution over its domain is called a *random variable*.

The different values of the domain are assumed to be mutually exclusive and exhaustive, which means that the variable can take only one value at any time, and that all possible values of interest are included in the domain.

Available evidence may indicate that the variable takes a value in a given subset of values, but it is unclear which specific value in particular. For this reason it is meaningful to consider subsets as composite values, where the *hyperdomain* contains all the singletons as well as composites values. It is then possible to have a belief mass distribution over all these values, instead of only having a probability distributions over singleton values.

A subjective opinion distributes a *belief mass* over the values of the hyperdomain. The sum of the belief masses is less than or equal to 1, and is complemented with an *uncertainty mass* which reflects the opinion's confidence level. Subjective opinions also contain a *base rate* probability distribution expressing prior knowledge about the specific class of random variables, so that in case of significant uncertainty about a specific variable, the base rates provide a basis for default likelihoods. We give formal definitions of these concepts in what follows.

Let X be a variable over a domain $\mathbb{X} = \{x_1, x_2, \ldots, x_k\}$ of cardinality k, where x_i $(1 \leq i \leq k)$ represents a specific value from the domain. Let $\mathcal{P}(\mathbb{X})$ be the powerset of $\mathbb{X}$. The *hyperdomain* is the reduced powerset of $\mathbb{X}$, denoted by $\mathcal{R}(\mathbb{X})$, and defined as:

$$\mathcal{R}(\mathbb{X}) = \mathcal{P}(\mathbb{X}) \setminus \{\mathbb{X}, \emptyset\}. \tag{1}$$

All proper subsets of $\mathbb{X}$ are values of $\mathcal{R}(\mathbb{X})$, but $\mathbb{X}$ and $\emptyset$ are not, because they are not considered as possible observations to which belief mass can be assigned. Since $\mathbb{X}$ and $\emptyset$ are excluded the hyperdomain has cardinality $2^k - 2$. We use the same notation for the values of a domain and its hyperdomain, and say that X is a *hypervariable* when it takes values from the hyperdomain.

Let A denote a source which can be a human, a sensor, etc. A *subjective opinion* ω_X^A of the source A on the variable X is a tuple

$$\omega_X^A = (\mathbf{b}_X^A, u_X^A, \mathbf{a}_X^A), \tag{2}$$

where $\mathbf{b}_X^A : \mathcal{R}(\mathbb{X}) \to [0,1]$ is a belief mass distribution, the parameter $u_X^A \in [0,1]$ is an uncertainty mass, and $\mathbf{a}_X^A : \mathbb{X} \to [0,1]$ is a base rate probability distribution satisfying the following additivity constrains:

$$u_X^A + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X^A(x) = 1, \tag{3}$$

$$\sum_{x \in \mathbb{X}} \mathbf{a}_X^A(x) = 1. \tag{4}$$

In the notation of the subjective opinion ω_X^A, the superscript is the source A, while the subscript is the object target variable X. An explicit source notation makes is possible to express the fact that different sources produce different opinions on the same variable. The source can be omitted in the opinion notation whenever the source is implicit or irrelevant, for example when there is only one source in the modelled situation.

The belief mass distribution $\mathbf{b}_X^A$ has $2^k - 2$ parameters, whereas the base rate distribution $\mathbf{a}_X^A$ only has k parameters. The uncertainty parameter u_X^A is a simple scalar. A general opinion thus contains $2^k + k - 1$ parameters. However, given that Eq. (3) and Eq. (4) remove one degree of freedom each, an opinion over a domain of cardinality k only has $2^k + k - 3$ degrees of freedom. Note that it is possible to express base rates over composite values as expressed by Eq. (5) below.

$$\mathbf{a}_X(x_i) = \sum_{\substack{x_j \in \mathbb{X} \\ x_j \subseteq x_i}} \mathbf{a}_X(x_j), \quad \forall x_i \in \mathcal{R}(\mathbb{X}). \tag{5}$$

A subjective opinion in which $u_X = 0$, i.e. an opinion without uncertainty, is called a *dogmatic opinion*. A

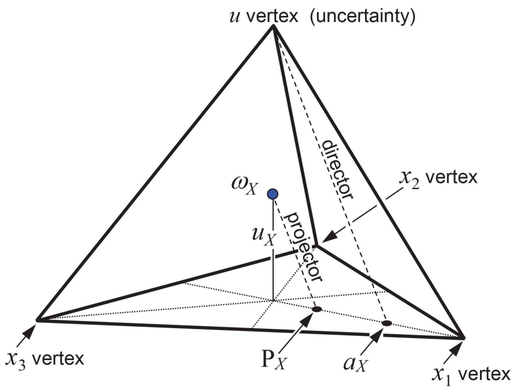

Fig. 3. Example trinomial opinion

dogmatic opinion for which $b_X(x) = 1$, for some x, is called an *absolute opinion*. In contrast, an opinion for which $u_X = 1$, and consequently, $b_X(x) = 0$, for every $x \in \mathcal{R}(\mathbb{X})$, i.e. an opinion with total uncertainty, is called a *vacuous opinion*.

Every subjective opinion 'projects' to a probability distribution $\mathbf{P}_X$ over $\mathbb{X}$ defined through the following function:

$$\mathbf{P}_X(x_i) = \sum_{x_j \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x_i \mid x_j)\mathbf{b}_X(x_j) + \mathbf{a}_X(x_i)u_X, \quad (6)$$

where $a_X(x_i \mid x_j)$ is the *relative base rate* of $x_i \in \mathbb{X}$ with respect to $x_j \in \mathcal{R}(\mathbb{X})$ defined as follows:

$$\mathbf{a}_X(x_i \mid x_j) = \frac{\mathbf{a}_X(x_i \cap x_j)}{\mathbf{a}_X(x_j)}, \quad (7)$$

where a_X is extended on $\mathcal{R}(\mathbb{X})$ additively. For the relative base rate to be always defined, it is enough to assume $a_X^A(x_i) > 0$, for every $x_i \in \mathbb{X}$. This means that everything we include in the domain has a non-zero probability of occurrence in general.

Binomial opinions apply to binary random variables where the belief mass is distributed over the two values in a binary domain. Multinomial opinions apply to random variables in n-ary domains, and where the belief mass is distributed over the values of the domain. Figure 3 visualises a ternary multinomial opinion as a point inside a tetrahedron.

General opinions, also called *hyper-opinions*, apply to hypervariables where belief mass is distributed over values in a hyperdomain which is the reduced powerset of an n-ary domain. Given a hyper-opinion, it is possible to project it onto a multinomial opinion. Assume a hyper opinion ω_X and let $\underset{\cdot}{\mathbf{b}}_X$ be the belief mass distribution defined by the sum in Eq. (6), i.e.

$$\underset{\cdot}{\mathbf{b}}_X(x) = \sum_{x' \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x \mid x')\mathbf{b}_X(x'), \quad (8)$$

then it is easy to check that $\underset{\cdot}{\mathbf{b}}_X : \mathbb{X} \to [0,1]$, and that $\underset{\cdot}{\mathbf{b}}_X$ together with u_X satisfies the additivity property in Eq. (3). The multinomial opinion denoted $\underset{\cdot}{\omega}_X = (\underset{\cdot}{\mathbf{b}}_X, u_X, \mathbf{a}_X)$ is the projected opinion from the hyper-opinion of ω_X. By defining the unary operator $\underset{\cdot}{\downarrow}$ to represent hyper-to-multinomial projection we can write:

$$\text{Hyper-to-Multinomial Projection:} \quad \underset{\cdot}{\omega}_X = \underset{\cdot}{\downarrow}(\omega_X). \quad (9)$$

From Eq. (6) and Eq. (8) we obtain $\mathbf{P}(\omega_X) = \mathbf{P}(\underset{\cdot}{\omega}_X)$. This means that every hyper-opinion can be approximated with its projected multinomial opinion which by definition has the same projected probability distribution as the initial hyper-opinion.

A binomial opinion is equivalent to a Beta probability density function, a multinomial opinion is equivalent to a Dirichlet probability density function, and a hyper-opinion is equivalent to a Dirichlet hyper-probability density function [8]. Binomial opinions thus represent the simplest opinion type, which can be generalised to multinomial opinions, which in turn can be generalised to hyper-opinions. Simple visualisations for binomial and trinomial opinions are based on barycentric coordinate systems as illustrated in Figures 3 and 4.

Consider a domain $\mathbb{X}$ with its hyperdomain $\mathcal{R}(\mathbb{X})$ and powerset $\mathcal{P}(\mathbb{X})$. Recall that $\{\mathbb{X}\} \in \mathcal{P}(\mathbb{X})$. Let x denote a specific value of $\mathcal{R}(\mathbb{X})$ or of $\mathcal{P}(\mathbb{X})$.

In DST (Dempster-Shafer Theory) [9], the belief mass on value x is denoted $\mathbf{m}(x)$, and the belief mass distribution is called a *basic belief assignment* (bba). It is possible to define a direct bijective mapping between the bba of DST and the belief mass distribution and uncertainty mass of subjective opinions, as expressed by Eq. (10):

$$\begin{array}{c}\text{Mapping between the}\\ \text{bba of DST and the}\\ \text{belief/uncertainty masses}\\ \text{of subjective opinions:}\end{array} \begin{cases} \mathbf{m}(x) = \mathbf{b}_X(x), & \forall x \in \mathcal{R}(\mathbb{X}), \\ \mathbf{m}(\mathbb{X}) = u_X. \end{cases} \quad (10)$$

Technically, the bba of DST and the belief/uncertainty representation of subjective opinions are thus equivalent. Their interpretations however are different. Subjective opinions can not assign belief mass to the domain $\mathbb{X}$ itself. This interpretation corresponds to the (hyper-) Dirichlet model, where only observations of values of $\mathbb{X}$ (or $\mathcal{R}(\mathbb{X})$) are counted as evidence. The domain $\mathbb{X}$ itself can not be an observation in the (hyper-) Dirichlet model, and hence can not be counted as evidence. The difference between the belief representation in DST and the opinion representation in SL is that the DST belief representation does not take base rates into account. As a result the projected (called *'pignistic'*) probability in DST [9] can only be computed with default base rates equal to the relative cardinalities of (hyper) values in the domain, whereas the projected probability of subjective opinions can be computed with any base rate distribution.

5. DIRICHLET REPRESENTATION OF BELIEFS

A hyper-opinion is equivalent to a Dirichlet HPDF (hyper probability density function) over a hyperdomain

$\mathcal{R}(\mathbb{X})$, according to the bijective mapping described in Section 5.2. For self-containment, we briefly outline the Dirichlet hypernomial model below, and refer to [10] for details about the Dirichlet model, and to [5] for details about the Dirichlet HPDF. The Dirichlet HPDF can be projected to a Hyper-Dirichlet PDF [11] which is useful for visualisation, but the Hyper-Dirichlet PDF is out of the scope of this presentation.

5.1. The Dirichlet Hypernomial Model

Multinomial probability density over a domain $\mathbb{X}$ of cardinality k is expressed by the k-dimensional Dirichlet PDF, where the special case of a probability density over a binary domain (where $k = 2$) is expressed by the Beta PDF. As a generalisation, hypernomial probability over the hyperdomain $\mathcal{R}(\mathbb{X})$ of cardinality $\kappa = 2^k - 2$ is expressed by the κ-dimensional Dirichlet HPDF [11].

The set of input arguments to the Dirichlet HPDF over $\mathcal{R}(\mathbb{X})$ then becomes a sequence of strength parameters of the κ possible (composite) values $x \in \mathcal{R}(\mathbb{X})$ represented as κ positive real numbers $\alpha_X(x_i)$, $i = 1 \dots \kappa$, each corresponding to one of the possible values $x \in \mathcal{R}(\mathbb{X})$. Because this is a Dirichlet PDF over a hypervariable, it is called a Dirichlet Hyper-PDF, or Dirichlet HPDF for short.

Definition 1 (**Dirichlet HPDF**). Let $\mathbb{X}$ be a domain consisting of k mutually disjoint values, where the corresponding hyperdomain $\mathcal{R}(\mathbb{X})$ has cardinality $\kappa = (2^k - 2)$. Let α_X represent the strength vector over the κ values $x \in \mathcal{R}(\mathbb{X})$. The hyper-probability distribution $\mathbf{p}_X^{\mathrm{H}}$ and the strength vector α_X are both κ-dimensional. The Dirichlet hyper-probability density function over $\mathbf{p}_X^{\mathrm{H}}$, called Dirichlet HPDF for short, is denoted $\mathrm{Dir}_X^{\mathrm{H}}(\mathbf{p}_X^{\mathrm{H}}; \alpha_X)$, and is expressed as

$$\mathrm{Dir}_X^{\mathrm{H}}(\mathbf{p}_X^{\mathrm{H}}; \alpha_X) = \frac{\Gamma\left(\sum_{x \in \mathcal{R}(\mathbb{X})} \alpha_X(x)\right)}{\prod_{x \in \mathcal{R}(\mathbb{X})} \Gamma\left(\alpha_X(x)\right)} \prod_{x \in \mathcal{R}(\mathbb{X})} \mathbf{p}_X^{\mathrm{H}}(x)^{(\alpha_X(x)-1)}, \tag{11}$$

where $\alpha_X(x) \geq 0$, with the restrictions that $p_X^{\mathrm{H}}(x) \neq 0$ if $\alpha_X(x) < 1$.

The strength vector α_X represents the prior as well as the observation evidence, now assumed applicable to values $x \in \mathcal{R}(\mathbb{X})$.

Since the values of $\mathcal{R}(\mathbb{X})$ can contain multiple singletons from $\mathbb{X}$, a value of $\mathcal{R}(\mathbb{X})$ has a base rate equal to the sum of the base rates of the singletons it contains, as expressed by Eq. (5). The strength $\alpha_X(x)$ for each value $x \in \mathcal{R}(\mathbb{X})$ can then be expressed as

$$\begin{array}{c} \forall x \in \mathcal{R}(\mathbb{X}), \\ \alpha_X(x) = \mathbf{r}_X(x) + \mathbf{a}_X(x)W, \end{array} \quad \text{where} \quad \begin{cases} \mathbf{r}_X(x) \geq 0, \\ \mathbf{a}_X(x) = \sum_{\substack{x_j \subseteq x \\ x_j \in X}} \mathbf{a}(x_j), \\ W = 2. \end{cases} \tag{12}$$

The Dirichlet HPDF over a set of κ possible states $x_i \in \mathcal{R}(\mathbb{X})$ can thus be expressed as a function of the observation evidence $\mathbf{r}_X$ and the base rate distribution $\mathbf{a}_X(x)$, where $x \in \mathcal{R}(\mathbb{X})$. The constant W represents the non-informative prior weight which as a convention is set to $W = 2$ [5] (p.33). The superscript 'eH' in the notation $\mathrm{Dir}_X^{\mathrm{eH}}$ indicates that it is expressed as a function of the evidence parameter vector $\mathbf{r}_X$ (not the strength parameter vector α_X), and that it is a Dirichlet HPDF (not a traditional Dirichlet PDF). The evidence-based Dirichlet HPDF is expressed as

$$\begin{aligned} \mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X, \mathbf{a}_X) = &\frac{\Gamma\left(\sum_{x \in \mathcal{R}(\mathbb{X})} (\mathbf{r}_X(x) + \mathbf{a}_X(x)W)\right)}{\prod_{x \in \mathcal{R}(\mathbb{X})} \Gamma\left(\mathbf{r}_X(x) + \mathbf{a}_X(x)W\right)} \\ &\times \prod_{x \in \mathcal{R}(\mathbb{X})} \mathbf{p}_X^{\mathrm{H}}(x)^{(\mathbf{r}_X(x) + \mathbf{a}_X(x)W - 1)}, \end{aligned} \tag{13}$$

where $(\mathbf{r}_X(x) + \mathbf{a}_X(x)W) \geq 0$, with the restriction that $\mathbf{p}_X^{\mathrm{H}}(x) \neq 0$ if $(\mathbf{r}_X(x) + \mathbf{a}_X(x)W) < 1$.

$\mathrm{Dir}_X^{\mathrm{eH}}$ in Eq. (13) is the expression for probability density over hyper-probability distributions p_X^{H}, where each value $x \in \mathcal{R}(\mathbb{X})$ has a base rate according to Eq. (7).

Because a value $x_j \in \mathcal{R}(\mathbb{X})$ can be composite, the expected probability of any value $x \in \mathbb{X}$ is not only a function of the direct probability density on x, but also of the probability density of all other values $x_j \in \mathcal{R}(\mathbb{X})$ that contain x. More formally, the expected probability of $x \in \mathbb{X}$ results from the probability density of each $x_j \in \mathcal{R}(\mathbb{X})$ where $x \cap x_j \neq \emptyset$.

Given the Dirichlet HPDF of Eq. (13), the expected probability of any of the k values $x \in \mathbb{X}$ can be written as

$$\mathbf{E}_X(x) = \frac{\sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x \mid x_i)\mathbf{r}(x_i) + W\mathbf{a}_X(x)}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}(x_i)} \quad \forall x \in \mathbb{X}. \tag{14}$$

The mapping between the hyper-opinion and the Dirichlet HPDF is based on defining the expected probability distribution of a Dirichlet HPDF expressed by Eq. (14) to be equal to the projected probability of hyper-opinions expressed by Eq. (6), i.e. $\mathbf{E}_X = \mathbf{P}_X$.

5.2. Mapping Between a Hyper-opinion and a Dirichlet HPDF

Figure 4 is a screenshot of the visualisation of the mapping between binomial opinions $\omega_X^{C_1}$ and $\omega_X^{C_2}$ on the left and the corresponding Beta PDFs on the right.

In general, a hyper-opinion is equivalent to a Dirichlet HPDF according to the mapping defined below.

Definition 2 (**Mapping: Hyper-opinion $\leftrightarrow$ Dirichlet HPDF**). Let $\mathbb{X}$ be a domain consisting of k mutually disjoint values, where the corresponding hyperdomain $\mathcal{R}(\mathbb{X})$ has cardinality $\kappa = (2^k - 2)$, and let X be a hypervariable in $\mathcal{R}(\mathbb{X})$. Let ω_X be a hyper-opinion on X, and let $\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X, \mathbf{a}_X)$ be a Dirichlet HPDF over the hyper-probability distribution $\mathbf{p}_X^{\mathrm{H}}$. The hyper-opinion ω_X

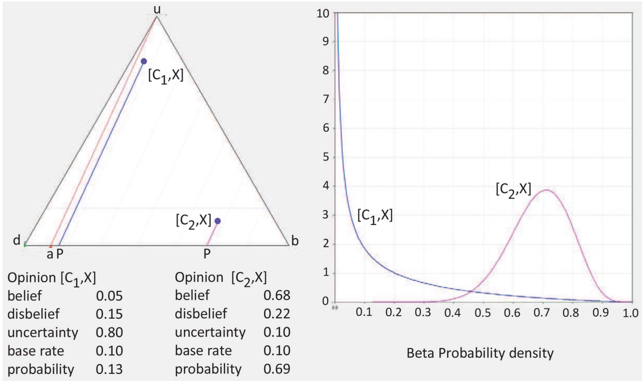

Fig. 4. Mapping opinions $\omega_X^{C_1}$ and $\omega_X^{C_2}$ to Beta PDFs

and the Dirichlet HPDF $\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}};\mathbf{r}_X,\mathbf{a}_X)$ are equivalent through the following mapping:

$$\forall x \in \mathcal{R}(\mathbb{X}) \quad \begin{cases} \mathbf{b}_X(x) = \dfrac{\mathbf{r}_X(x)}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X(x_i)}, \\ u_X = \dfrac{W}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X(x_i)}, \end{cases} \Leftrightarrow$$

$$\left(\begin{array}{cc} \text{For } u_X \neq 0: & \text{For } u_X = 0: \\ \begin{cases} \mathbf{r}_X(x) = \dfrac{W\mathbf{b}_X(x)}{u_X}, \\ 1 = u_X + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X(x_i), \end{cases} & \begin{cases} \mathbf{r}_X(x) = \mathbf{b}_X(x) \cdot \infty, \\ 1 = \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X(x_i). \end{cases} \end{array} \right) \tag{15}$$

The advantage of the Dirichlet HPDF is to provide an interpretation and equivalent representation of hyper-opinions.

This equivalence is very powerful because tools and methods used in Bayesian statistics can be applied to subjective opinions. In addition, the operators of subjective logic, such as conditional deduction, the subjective Bayes' theorem [12] and abduction, can be applied to statistical representations of data based on the Dirichlet model.

6. BELIEF FUSION OPERATORS

There are different categories of belief fusion situations, and each category requires its own operator for the computation of belief fusion [1]. In this article we focus on five different fusion categories, namely *constraint fusion*, *cumulative fusion*, *averaging fusion*, *weighted fusion* and *weighted fusion with vagueness* which are described below.

6.1. Belief Constraint Fusion

A typical application of belief theory in the literature is belief fusion with the classical Dempster's rule [9]. There has been considerable confusion and controversy around the adequacy of belief fusion operators, especially regarding Dempster's rule [13]. The confusion started with Zadeh's example from 1984 [14] where Dempster's rule is applied to a situation for which it is unsuitable and therefore produces erratic results. The controversy followed when authors failed to realise that it is not a question of whether Dempster's rule is correct or wrong, but of recognising the type of situations for which Dempster's rule is suitable.

As an analogy of the controversy around Dempster's rule, imagine a world where the swim vest (analogy of Dempster's rule) has been invented as a safety device (analogy of a belief fusion operator). Then somebody demonstrates with an example that swim vests provide very poor protection in a car crash (analogy of Zadeh's example). Some researchers explain this by saying that swim vests perform poorly only in the case of high speed (analogy of high conflict) car crashes, and suggest to reduce the driving speed to make swim vests perform better. Other researchers propose the seat belt as an alternative safety device because it works well in car crashes, but this proposal is met with criticism by people who claim that seat belts provide poor protection in a sinking boat, in which case swim vests provide good protection. Many other safety devices are invented, and each device is promoted with an anecdotal example where it provides relatively good protection. In this confusing discussion nobody seems to understand that different safety hazards require different safety devices for protection, and that there is no single safety device that can provide adequate protection in all situations.

In an analogous fashion, the fact that different belief fusion situations require different belief fusion operators has often been ignored in the belief theory literature, and has been a significant source of confusion for many years [13]. There is nothing wrong with Dempster's rule *per se*; there are situations where it is perfectly appropriate, and there are situations where it is clearly inappropriate. No single belief fusion operator is suitable in every situation.

Dempster's rule is traditionally presented as a method for (cumulative) fusion of beliefs from different (independent) sources [9] with the purpose of identifying the most 'correct' hypothesis value from the domain. However, many authors have demonstrated that Dempster's rule is not an appropriate operator for this type of fusion [14]. Motivated by the apparent inconsistency of results produced by Dempster's rule numerous authors have proposed alternative belief fusion operators [15], [16], [17], [18], [19], [20], [21], [22], but the authors often fail to specify which type of situations they model.

We argue that Dempster's rule is better suited as a method for *belief constraint fusion* [13], [23], as shown in Figure 2. Situations of this type are e.g. when agents express different preferences with regard to a common decision that the agents must agree on [23] or when the analyst is presented with specific hints that are guaranteed to be valid [24], which is expressed by saying that the sources are 'reliable'.

It is common to see situations where people with different preferences try to agree on a single choice, or situations where evidence is presented as factual hints.

This must not be confused with fusion of belief from different agents to determine the most likely correct hypothesis or actual event, because the beliefs can not be taken as factual. Multi-agent preference combination assumes that each agent has already made up her mind, and then that they together want to determine the most acceptable decision or choice for all. Similarly, the fusion of hints assumes that the truth is known to the sources, but that they only reveal parts of the truth in the form of hints. Preferences and hints over a variable can be expressed in the form of subjective opinions. The constraint fusion operator of subjective logic can be applied as a method for merging preferences and hints from multiple sources into a single conclusion for the group of sources. This operator is expressive and flexible, and produces perfectly intuitive results. Preference can be represented as belief mass, and indifference can be represented as uncertainty mass. Positive and negative beliefs are considered as symmetric concepts, so they can be represented in the same way and combined using the same operator. Vacuous belief has no influence on the conclusion, and thereby represents the neutral element.

6.1.1. Method of Belief Constraint Fusion

The BCF (Belief Constraint Fusion) operator described next is an extension of Dempster's rule. The notation is also generalised to cover multiple sources, not only two sources.

DEFINITION 3 (**The Constraint Fusion Operator**). Assume the domain $\mathbb{X}$ and its hyperdomain $\mathcal{R}(\mathbb{X})$, and assume the hypervariable X which takes its values from $\mathcal{R}(\mathbb{X})$. Let $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$ denote a set of N independent sources. Let $C \in \mathbb{C}$ denote a specific source, and let ω_X^C denote its opinion about the variable X.

The respective opinions can be mathematically merged using the BCF (Belief Constraint Fusion) operator denoted '$\odot$' which can be expressed as

$$\begin{aligned}\omega_X^{\&(\mathbb{C})} &= \underset{C\in\mathbb{C}}{\odot}(\omega_X^C)\\ &= \omega_X^{C_1} \odot \omega_X^{C_2} \odot \cdots \omega_X^{C_N}.\end{aligned} \tag{16}$$

Source combination denoted '&' thus corresponds to belief fusion with '$\odot$'. The multi-source expression for BCF is given by Eq. (17):

$$\forall x \in \mathcal{R}(\mathbb{X}), \quad \omega_X^{\&(\mathbb{C})}:$$

$$\begin{cases}\mathbf{b}_X^{\&(\mathbb{C})}(x) = \dfrac{\mathrm{Har}(x)}{(1-\mathrm{Con})},\\[2ex] u_X^{\&(\mathbb{C})} = \dfrac{\prod_{C\in\mathbb{C}} u_X^C}{(1-\mathrm{Con})},\\[2ex] \mathbf{a}^{\&(\mathbb{C})}(x) = \dfrac{\sum_{C\in\mathbb{C}} \mathbf{a}_X^C(x)(1-u_X^C)}{N-\sum_{C\in\mathbb{C}} u_X^C}, \quad \exists u_X^C < 1,\\[2ex] \mathbf{a}^{\&(\mathbb{C})}(x) = \dfrac{\sum_{C\in\mathbb{C}} \mathbf{a}_X^C(x)}{N}, \quad \forall u_X^C = 1.\end{cases} \tag{17}$$

The term Har(x) represents the relative *harmony* between the constraint opinion ω_X^C (in terms of overlapping belief mass) on x. The term Con represents the relative *conflict* between constraints (in terms of non-overlapping belief mass) between the constraint opinions ω_X^C. DST's notation $\mathbf{m}(x)$ for belief-mass of $x \in \mathcal{P}(\mathbb{X})$ given by Eq. (10) gives the most compact notation for computing 'Har' and 'Con':

$$\mathrm{Har}(x) = \sum_{\substack{\cap x^C = x\\ x^C\in\mathcal{P}(\mathbb{X})}} \prod_{C\in\mathbb{C}} \mathbf{m}_X^C(x^C), \tag{18}$$

$$\mathrm{Con} = \sum_{\substack{\cap x^C = \emptyset\\ x^C\in\mathcal{P}(\mathbb{X})}} \prod_{C\in\mathbb{C}} \mathbf{m}_X^C(x^C). \tag{19}$$

The divisor $(1 - \mathrm{Con})$ in Eq. (17) normalises the belief mass and uncertainty mass; i.e. it ensures their additivity. The application of the BCF operator is mathematically possible only if the constraint opinions ω_X^C are not totally conflicting, i.e., if $\mathrm{Con} \neq 1$.

The BCF operator is commutative and non-idempotent. Associativity is preserved when the base rate is equal for all agents. Associativity in case of different base rates requires that all preference opinions be combined in a single operation which requires that Eq. (17) is applied for all input arguments in a single operation, which then represents semi-associativity.

The base rates of the two arguments are normally equal, but different base rates can be used in case of base rate disagreement between the sources, in which case the fused base rate distribution is the confidence-weighted average base rate.

Associativity in case of different base rates requires that all arguments opinions be combined in a single operation according to Definition 3. A totally indifferent opinion acts as the neutral element for constraint fusion, formally expressed as

$$\begin{aligned}&\text{IF } (\omega_X^A \text{ is indifferent, i.e. } u_X^A = 1)\\ &\text{THEN } (\omega_X^A \odot \omega_X^B = \omega_X^B).\end{aligned} \tag{20}$$

Having a neutral element in the form of the totally indifferent (i.e. vacuous) opinion can be useful when modelling situations of preference combination.

The rich format of subjective opinions makes it simple to express positive and negative preferences within the same framework, as well as indifference/uncertainty. Because preferences can be expressed over arbitrary subsets of the domain, this is in fact a multi-polar model for expressing and combining preferences. Even in the case of totally conflicting dogmatic opinions the belief constraint fusion operator produces meaningful results, namely that the preferences are incompatible. Examples in Sections 6.1.2–6.1.5 demonstrates the usefulness of this property.

TABLE 1
Example preferences and corresponding subjective opinions

Example Type	Domain & Opinion	Expression
"Ingredient x is mandatory"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Hard positive	Binomial opinion	$\omega_x : (1, 0, 0, \frac{1}{2})$
"Ingredient x is totally out of the question"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Hard negative	Binomial opinion	$\omega_x : (0, 1, 0, \frac{1}{2})$
"I prefer x with rating 0.3"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Quantitative	Binomial opinion	$\omega_x : (0.3, 0.7, 0.0, \frac{1}{2})$
"I prefer x or y, but z is also acceptable"	Ternary domain	$\Theta = \{x, y, z\}$
Qualitative	Trinomial opinion	$\omega_\Theta : (b(\{x, y\}) = 0.6,$ $b(z) = 0.3,\ u = 0.1,$ $a(x_1), a(x_2), a(x_3) = \frac{1}{3})$
"I like x, but I like y even more"	Binary domains	$\mathbb{X} = \{x, \bar{x}\}$ and $\mathbb{Y} = \{y, \bar{y}\}$
Positive rank	Binomial opinions	$\omega_x : (0.6, 0.3, 0.1, \frac{1}{2})$, $\omega_y : (0.7, 0.2, 0.1, \frac{1}{2})$
"I don't like x, and I dislike y even more"	Binary domains	$\mathbb{X} = \{x, \bar{x}\}$ and $\mathbb{Y} = \{y, \bar{y}\}$
Negative rank	Binomial opinions	$\omega_x : (0.3, 0.6, 0.1, \frac{1}{2})$, $\omega_y : (0.2, 0.7, 0.1, \frac{1}{2})$
"I'm indifferent about x, y and z"	Ternary domain	$\Theta = \{x, y, z\}$
Neutral	Trinomial opinion	$\omega_\Theta : (u_\Theta = 1.0,$ $a(x_1), a(x_2), a(x_3) = \frac{1}{3})$
"I'm indifferent but most people prefer x"	Ternary domain	$\Theta = \{x, y, z\}$
Neutral with bias	Trinomial opinion	$\omega_\Theta : (u_\Theta = 1.0,\ a(x) = 0.6,$ $a(y), a(z) = 0.2)$

6.1.2. Expressing Preferences with Subjective Opinions

Preferences can be expressed as soft or hard constraints, qualitative or quantitative, ordered or partially ordered, etc. It is possible to specify a mapping between qualitative verbal tags and subjective opinions, which enables easy solicitation of preferences [25]. Table 1 describes examples of how preferences can be expressed.

All the preference types of Table 1 can be interpreted in terms of subjective opinions, and further combined by considering them as constraints expressed by different sources/agents. The examples which comprise two binary domains could equally well have been modelled with a quaternary product domain with a corresponding quatronomial product opinion. In fact, to compute product opinions over product domains is an alternative approach of simultaneously considering preferences over multiple variables.

TABLE 2
Fusion of film preferences

	Belief preferences of:			Fusion results:	
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	*A&B* $\omega_X^{A\&B}$	*A&B&C* $\omega_X^{A\&B\&C}$
$b(x_1)$	0.99	0.00	0.00	0.00	0.00
$b(x_2)$	0.01	0.01	0.00	1.00	1.00
$b(x_3)$	0.00	0.99	0.00	0.00	0.00
$b(\{x_2, x_3\})$	0.00	0.00	1.00	0.00	0.00

Default base rates are specified in all but the last example, which indicates total indifference, but with a bias that expresses the average preference in the population. Base rates are useful in many situations, such as for default reasoning. Base rates influence the computed results only in case of significant indifference or uncertainty.

6.1.3. Example: Going to the Cinema, First Attempt

Assume three friends, Alice, Bob and Clark, who want to see a film together at the cinema one evening, and that the only films showing are *Black Dust* (x_1), *Grey Matter* (x_2) and *White Powder* (x_3), represented as the ternary domain $X = \{x_1, x_2, x_3\}$. Assume that the friends express their preferences in the form of the opinions of Table 2.

Alice and Bob have strong and conflicting preferences. Clark, who strictly does not want to watch *Black Dust* (x_1), and who is indifferent about the two other films, is not sure whether he wants to come along, so Table 2 shows the results of applying the belief/preference constraint fusion operator, first without him, and then when including him in the party.

By applying belief constraint fusion, Alice and Bob conclude that the only film they are both interested in seeing is *Grey Matter* (x_2). Including Clark in the party does not change that result because he is indifferent to *Grey Matter* (x_2) and *White Powder* (x_3) anyway, he just does not want to watch *Black Dust* (x_1).

The belief mass values of Alice and Bob in the above example are in fact equal to those that Zadeh [14] used to demonstrate the unsuitability of Dempster's rule for fusing beliefs by showing how they produce counter-intuitive results. Zadeh's example describes a medical case where two medical doctors express their expert opinions about possible diagnoses, which typically should not have been modelled with Dempster's rule (BCF), but with the weighted belief fusion (WBF) operator [1], and possibly followed by vagueness maximisation (WBF-VM). In order to select the appropriate operator, it is crucial to fully understand the nature of the situation to be modelled. The failure to understand that Dempster's rule does not represent an operator for cumulative or averaging belief fusion, combined with the unavailability of the general cumulative, averaging and weighted fusion operators during that period

TABLE 3
Fusion of film preferences with indifference and non-default base rates

	Belief preferences of:			Fusion results:	
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	$A\&B$ $\omega_X^{A\&B}$	$A\&B\&C$ $\omega_X^{A\&B\&C}$
$\mathbf{b}(X_1)$	0.98	0.00	0.00	0.490	0.000
$\mathbf{b}(x_2)$	0.01	0.01	0.00	0.015	0.029
$\mathbf{b}(x_3)$	0.00	0.98	0.00	0.490	0.961
$\mathbf{b}(\{x_2,x_3\})$	0.00	0.00	1.00	0.000	0.010
u	0.01	0.01	0.00	0.005	0.000
$\mathbf{a}(x_1)$	0.6	0.6	0.6	0.6	0.6
$\mathbf{a}(x_2) = a(x_3)$	0.2	0.2	0.2	0.2	0.2

(1976 [9]–2013 [1]), has often led to inappropriate applications of Dempster's rule to cases of belief fusion [13]. However, when specifying the same numerical values as in [14] in a case of preference constraints such as in the example above, the belief constraint fusion operator (which is a simple extension of Dempster's rule) is the correct fusion operator which produces perfectly intuitive results.

6.1.4. Example: Going to the Cinema, Second Attempt

In this example Alice and Bob soften their preference with some indifference in the form of $u = 0.01$, as specified by Table 3. Clark has the same opinion as in the previous example, and is still not sure whether he wants to come along, so Table 3 shows both the results without him, and with his preference included.

Having some indifference in the preferences would mean that Alice and Bob should pick film *Black Dust* (x_1) or *White Powder* (x_3), because in both cases, one of them actually prefers one of the films, and the other finds it acceptable. Neither Alice nor Bob prefers *Grey Matter* (x_2), they only find it acceptable, so it would be a bad choice for both of them. When taking into consideration the base rates $a(x_1) = 0.6$ for *Black Dust* and $a(x_3) = 0.2$ for *White Powder*, the expected preference levels according to Eq. (6) are such that

$$\mathbf{P}_X^{A\&B}(x_1) > \mathbf{P}_X^{A\&B}(x_3). \tag{21}$$

More precisely, the preference probabilities from Eq. (6) are

$$\mathbf{P}_X^{A\&B}(x_1) = 0.493, \qquad \mathbf{P}_X^{A\&B}(x_3) = 0.491. \tag{22}$$

Because of the higher base rate, *Black Dust* (x_1) also has a higher expected preference than *White Powder* (x_3), so the rational choice would be to watch *Black Dust* (x_1).

However, when including Clark, who does not want to watch *Black Dust* (x_1), the base rates no longer dictate the result. In this case constraint fusion with Eq. (6) produces $\mathbf{P}^{A\&B\&C}(x_3) = 0.966$ so the obvious choice is to watch *White Powder* (x_3).

TABLE 4
Combination of film preferences with hard and conflicting preferences

	Belief preferences of:			Fusion results:	
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	$A\&B$ $\omega_X^{A\&B}$	$A\&B\&C$ $\omega_X^{A\&B\&C}$
$b(x_1)$	1.00	0.00	0.00	Undefined	Undefined
$b(x_2)$	0.00	0.00	0.00	Undefined	Undefined
$b(x_3)$	0.00	1.00	0.00	Undefined	Undefined
$b(\{x_2,x_3\})$	0.00	0.00	1.00	Undefined	Undefined

6.1.5. Example: Not Going to the Cinema

Assume now that Alice and Bob have totally conflicting preferences as specified in Table 4, i.e. Alice has a hard preference for *Black Dust* (x_1) and Bob has a hard preference for *White Powder* (x_3). As before, Clark still does not want to watch *Black Dust* (x_1), and is indifferent about the other two films.

In this case, the belief constraint fusion operator can not be applied because Eq. (17) involves a division by zero. The conclusion is that the friends will not go to the cinema to see a film together that evening. The test for detecting this situation is to observe $Con = 1$ in Eq. (19). It makes no difference to include Clark in the party, because a conflict can not be resolved by including additional preferences. However it would have been possible for Bob and Clark to watch *White Powder* (x_3) together without Alice.

6.2. Cumulative Belief Fusion

Cumulative Belief Fusion (CBF) is when it is assumed that the amount of evidence increases by including additional sources of independent evidence. An example of this type of situation is when different witnesses express their opinions about whether they saw the accused at the crime scene, and where their independent testimonies can be fused to produce an opinion about whether the accused really was there.

Assume a hyperdomain $\mathcal{R}(\mathbb{X})$ and a process where the outcome variable X takes values from $\mathcal{R}(\mathbb{X})$. Assume further that the outcome can be observed by different independent sources which can be expressed as $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$. Let $C \in \mathbb{C}$ denote a specific source, and let ω_X^C denote its opinion about the variable X. Assume that the sources in $\mathbb{C}$ produce independent opinions about the same variable X.

Observations can be vague, meaning that sometimes the sources observe an outcome which might be one of multiple possible singletons in $\mathbb{X}$, but the sources are unable to identify the observed outcome uniquely.

For example, assume that sources C_1 and C_2 observe coloured balls being picked from an urn, where the balls can have one of four colours: black, white, red or green. Assume further that the observer C_2 is colour-blind, which means that in poor light conditions he is unable see the difference between red and green balls, although he can always tell the other colour combinations apart. As a result, his observations can be vague, meaning that sometimes he perceives a specific ball to be either red or green, but is unable to identify the ball's colour precisely. This corresponds to the situation where X is a hypervariable which can take composite values from $\mathcal{R}(\mathbb{X})$.

The symbol '$\diamond$' denotes the fusion of independent sources $C \in \mathbb{C}$ into a single cumulative merged source denoted $\diamond(\mathbb{C})$.

Let $\mathbb{C} = \{C_1, C_2, \dots C_N\}$ be a frame of N sources with the respective opinions $\omega_X^{C_1}, \omega_X^{C_2}, \dots \omega_X^{C_N}$ over the same variable X. Let C denote a specific source $C \in \mathbb{C}$. The cumulative merger of all the sources in the source frame $\mathbb{C}$ is denoted $\diamond(\mathbb{C})$. The opinion $\omega_X^{\diamond(\mathbb{C})} \equiv (\mathbf{b}_X^{\diamond(\mathbb{C})}, u_X^{\diamond(\mathbb{C})}, \mathbf{a}_X^{\diamond(\mathbb{C})})$ is the cumulative fused opinion expressed as:

Case I: $\quad u_X^C \neq 0, \forall C \in \mathbb{C}$:

$$
\begin{aligned}
\mathbf{b}_X^{\diamond(\mathbb{C})}(x) &= \frac{\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right) - (N-1)\prod_{C\in\mathbb{C}} u_X^C},\\
u_X^{\diamond(\mathbb{C})} &= \frac{\prod_{C\in\mathbb{C}} u_X^C}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right) - (N-1)\prod_{C\in\mathbb{C}} u_X^C},\\
\mathbf{a}_X^{\diamond(\mathbb{C})}(x) &= \frac{\sum_{C\in\mathbb{C}}\left(\mathbf{a}_X^C\prod_{C_j\neq C} u_X^{C_j}\right) - \sum_{C\in\mathbb{C}}\mathbf{a}_X^C\cdot\prod_{C\in\mathbb{C}} u_X^C}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right) - N\prod_{C\in\mathbb{C}} u_X^C},\\
\mathbf{a}_X^{\diamond(\mathbb{C})}(x) &= \frac{\sum_{C\in\mathbb{C}}\mathbf{a}_X^C}{N}, \quad \forall u_X^C = 1,
\end{aligned}
\tag{23}
$$

Case II: $\quad \exists u_X^C = 0$, define $\mathbb{C}^{\mathrm{dog}} = \{C \text{ where } u_X^C = 0\}$:

$$
\begin{aligned}
\mathbf{b}_X^{\diamond(\mathbb{C})}(x) &= \sum_{C\in\mathbb{C}^{\mathrm{dog}}} \gamma_X^C \mathbf{b}_X^C(x),\\
u_X^{\diamond(\mathbb{C})} &= 0,\\
\mathbf{a}_X^{\diamond(\mathbb{C})}(x) &= \sum_{C\in\mathbb{C}^{\mathrm{dog}}} \gamma_X^C \mathbf{a}_X^C(x),
\end{aligned}
\tag{24}
$$

where

$$
\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\mathrm{dog}}}\to 0} \frac{u_X^C}{\sum_{C_j\in\mathbb{C}^{\mathrm{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\mathrm{dog}}. \tag{25}
$$

The notation $u_X^{\mathbb{C}^{\mathrm{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\mathrm{dog}}$. The cumulative fused opinion $\omega_X^{\diamond(\mathbb{C})}$ results from fusing the respective opinions ω_X^C of the sources $C \in \mathbb{C}$. The symbol '$\oplus$' denotes the cumulative belief fusion operator, hence we define

$$
\begin{aligned}
\omega_X^{\diamond(\mathbb{C})} &\equiv \bigoplus_{C\in\mathbb{C}} (\omega_X^C) && (26)\\
&\equiv \omega_X^{C_1} \oplus \omega_X^{C_2} \oplus \cdots \omega_X^{C_N}. && (27)
\end{aligned}
$$

It can be verified that the cumulative fusion operator is commutative, associative and non-idempotent. In Case II of Eq. (24), the associativity depends on preserving the relative weights of intermediate results with the additional weight parameter γ. In this case, the cumulative fusion operator is equivalent to the weighted average of probabilities.

The argument base rate distributions are normally equal. When that is not the case the fused base rate distribution over X is specified to be the evidence-weighted average base rate.

In case of N dogmatic arguments ω_X^C where $C \in \mathbb{C}$ it can be assumed that the limits in Eq. (24) are defined as $\gamma_X^C = 1/N$.

6.2.1. Justification for the Cumulative Fusion Operator

The cumulative belief fusion operator of Eq. (23) is derived by mapping the argument belief opinions to evidence parameters through the bijective mapping of Eq. (15). Cumulative fusion of evidence opinions simply consists of summing up the evidence parameters, where the sum is mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the justification of the cumulative fusion operator of Eq. (23). A more detailed explanation is provided below.

Let the sources $C \in \mathbb{C}$ have respective belief opinions expressed as ω_X^C. The corresponding Dirichlet PDFs $\mathrm{Dir}_X^{\mathrm{e}}(\mathbf{p}_X; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence vectors $\mathbf{r}_X^C$.

The cumulative fusion of these evidence vectors consists of vector summation of $\mathbf{r}_X^C$ where $C \in \mathbb{C}$, expressed as

$$
\mathbf{r}_X^{\diamond(\mathbb{C})} = \sum_{C\in\mathbb{C}} \mathbf{r}_X^C. \tag{28}
$$

For each value $x \in \mathcal{R}(\mathbb{X})$ the evidence sum $\mathbf{r}_X^{\diamond(\mathbb{C})}(x)$ is

$$
\begin{aligned}
\mathbf{r}_X^{\diamond(\mathbb{C})}(x) &= \sum_{C\in\mathbb{C}} \mathbf{r}_X^C(x) && (29)\\
&= \sum_{C\in\mathbb{C}} \frac{W\mathbf{b}_X^C(x)}{u_X^C} && (30)\\
&= \frac{W\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})}{\prod_{C\in\mathbb{C}} u_X^C}. && (31)
\end{aligned}
$$

The cumulative fused belief opinion $\omega_X^{\diamond(\mathbb{C})}$ of Eq. (23) results from mapping the fused evidence belief mass of Eq. (28) back to a belief opinion by applying the bijective mapping of Eq. (15).

$$
\mathbf{b}_X^{\diamond(\mathbb{C})}(x) = \frac{\mathbf{r}_X^{\diamond(\mathbb{C})}(x)}{W + \sum_{x\in\mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\diamond(\mathbb{C})}(x)} \tag{32}
$$

$$
= \frac{\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})}{\prod_{C\in\mathbb{C}} u_X^C + \sum_{x\in\mathcal{R}(\mathbb{X})}\left(\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})\right)} \tag{33}
$$

$$= \frac{\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right) - (N-1)\prod_{C\in\mathbb{C}} u_X^C},$$

$$\exists u_X^C \neq 0. \tag{34}$$

The transition from Eq. (32) to Eq. (33) results from inserting Eq. (31) into Eq. (32). The transition from Eq. (33) to Eq. (34) results from applying Eq. (3).

$$u_X^{\diamond(\mathbb{C})} = \frac{W}{W + \sum_{x\in\mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\diamond(\mathbb{C})}(x)} \tag{35}$$

$$= \frac{\prod_{C\in\mathbb{C}} u_X^C}{\prod_{C\in\mathbb{C}} u_X^C + \sum_{x\in\mathcal{R}(\mathbb{X})}\left(\sum_{C\in\mathbb{C}}(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j})\right)} \tag{36}$$

$$= \frac{\prod_{C\in\mathbb{C}} u_X^C}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right) - (N-1)\prod_{C\in\mathbb{C}} u_X^C},$$

$$\text{where } \exists \mathrm{u}_{\mathrm{X}}^{\mathrm{C}} \neq 0. \tag{37}$$

The transition from Eq. (35) to Eq. (36) results from inserting Eq. (31) into Eq. (35). The transition from Eq. (36) to Eq. (37) results from applying Eq. (3).

6.3. Averaging Belief Fusion

Averaging Belief Fusion (ABF) is when dependence between sources is assumed. In other words, including more sources does not mean that more evidence is supporting the conclusion. An example of this type of situations is when a jury tries to reach a verdict after having observed the court proceedings. The assumption is that the correctness of the verdict does not increase as a function of the number of jury members, because the amount of evidence is fixed by what was presented in court.

Let $\mathbb{C}$ denote a group of N separate sources which can be expressed as $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$. Assume that the sources in $\mathbb{C}$ produce separate opinions based on the same evidence about the same variable, so their opinions are necessarily dependent. Still, their perceptions might be different, e.g. because their cognitive capabilities are different. For example, assume that sources C_1 and C_2 together observe the picking of coloured balls from an urn, where the balls can have one of four colours: black, white, red or green. Assume that observer C_2 is colour-blind, which means that sometimes he has trouble distinguishing between red and green balls, although he can always distinguish between the other colour combinations. Observer C_1 has perfect colour vision, and normally can always tell the correct colour when a ball is picked. As a result, when a red ball is picked, observer C_1 almost always identifies it as red, but observer C_2 identifies it as green relatively frequently. This can lead to C_1 and C_2 having different and conflicting opinions about the same variable, although their observations and opinions are totally dependent. The averaging belief fusion operator is perfectly suitable for this fusion situation.

Let $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$ be a frame of N sources with the respective opinions $\omega_X^{C_1}, \omega_X^{C_2}, \ldots \omega_X^{C_N}$ over the same variable X. Let C denote a specific source $C \in \mathbb{C}$. The averaging merger of all the sources in the source frame $\mathbb{C}$ is denoted $\underline{\diamond}(\mathbb{C})$. The opinion $\omega_X^{\underline{\diamond}(\mathbb{C})} \equiv (\mathbf{b}_X^{\underline{\diamond}(\mathbb{C})}, u_X^{\underline{\diamond}(\mathbb{C})}, \mathbf{a}_X^{\underline{\diamond}(\mathbb{C})})$ is the averaging-fused opinion expressed as:

Case I: $u_X^C \neq 0, \forall C \in \mathbb{C}$:

$$\mathbf{b}_X^{\underline{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C\in\mathbb{C}}\left(\mathbf{b}_X^C(x)\prod_{C_j\neq C} u_X^{C_j}\right)}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right)},$$

$$u_X^{\underline{\diamond}(\mathbb{C})} = \frac{N\prod_{C\in\mathbb{C}} u_X^C}{\sum_{C\in\mathbb{C}}\left(\prod_{C_j\neq C} u_X^{C_j}\right)},$$

$$\mathbf{a}_X^{A\underline{\diamond}B}(x) = \frac{\sum_{C\in\mathbb{C}} \mathbf{a}_X^C(x)}{N}, \tag{38}$$

Case II: $\exists u_X^C = 0$, define $\mathbb{C}^{\text{dog}} = \{C \text{ where } u_X^C = 0\}$:

$$\mathbf{b}_X^{\underline{\diamond}(\mathbb{C})}(x) = \sum_{C\in\mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{b}_X^C(x),$$

$$u_X^{\underline{\diamond}(\mathbb{C})} = 0,$$

$$\mathbf{a}_X^{\underline{\diamond}(\mathbb{C})}(x) = \sum_{C\in\mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{a}_X^C(x), \tag{39}$$

where

$$\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\text{dog}}}\to 0} \frac{u_X^C}{\sum_{C_j\in\mathbb{C}^{\text{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\text{dog}}. \tag{40}$$

The notation $u_X^{\mathbb{C}^{\text{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\text{dog}}$. The averaging-fused opinion $\omega_X^{\underline{\diamond}(\mathbb{C})}$ results from averaging fusion of the respective opinions ω_X^C of the sources $C \in \mathbb{C}$. By using the symbol '$\underline{\oplus}$' to designate the averaging belief fusion operator, we define

$$\omega_X^{\underline{\diamond}(\mathbb{C})} \equiv \underset{C\in\mathbb{C}}{\underline{\oplus}} (\omega_X^C). \tag{41}$$

It can be verified that the averaging belief fusion operator is commutative, idempotent, and non-associative. The non-associativity means that

$$(\omega_X^{C_1} \underline{\oplus} \omega_X^{C_2}) \underline{\oplus} \omega_X^{C_3} \neq \omega_X^{C_1} \underline{\oplus} (\omega_X^{C_2} \underline{\oplus} \omega_X^{C_3}). \tag{42}$$

However, semi-associativity exists as expressed by Eq. (41) where the argument order is irrelevant because all the arguments are fused in one single operation. The only way to apply averaging fusion to more than two arguments is thus by fusing all arguments in one operation as described in Eq. (38) and expressed by the notation of Eq. (41). For three argument sources, this is expressed as:

$$\omega_X^{\underline{\diamond}(C_1,C_2,C_3)} \equiv \underline{\oplus}(\omega_X^{C_1}, \omega_X^{C_2}, \omega_X^{C_3}). \tag{43}$$

The argument base rate distributions are normally equal. When that is not the case the fused base rate distribution is specified to be the average base rate distribution. In case the opinions of the N sources in $\mathbb{C}$ are all dogmatic opinions, then the limits in Eq. (39) can be set to $\gamma_X^C = 1/N$.

6.3.1. Justification for the Averaging Fusion Operator

The averaging belief fusion operator of Eq. (38) is derived by mapping the argument belief opinions to evidence opinions through the bijective mapping of Eq. (15). Averaging fusion of evidence opinions simply consists of computing the average of the evidence parameters. The fused evidence opinion is then mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the justification of the averaging fusion operator of Eq. (38). A more detailed explanation is provided below.

Let the sources $C \in \mathbb{C}$ have respective belief opinions expressed as ω_X^C. The corresponding Dirichlet PDFs $\mathrm{Dir}_X^{\mathrm{e}}(\mathbf{p}_X; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence vectors $\mathbf{r}_X^C$.

The averaging fusion of these evidence vectors consists of vector averaging of $\mathbf{r}_X^C$ where $C \in \mathbb{C}$, expressed as

$$\mathbf{r}_X^{\underline{\diamond}(\mathbb{C})} = \frac{\sum_{C\in\mathbb{C}} \mathbf{r}_X^C}{N}. \tag{44}$$

For each value $x \in \mathcal{R}(\mathbb{X})$ the average evidence $\mathbf{r}_X^{\underline{\diamond}(\mathbb{C})}(x)$ is

$$\mathbf{r}_X^{\underline{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C\in\mathbb{C}} \mathbf{r}_X^C(x)}{N} = \frac{\sum_{C\in\mathbb{C}} W \mathbf{b}_X^C(x)/u_X^C}{N} \tag{45}$$

$$= \frac{W \sum_{C\in\mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j}\right)}{N \prod_{C\in\mathbb{C}} u_X^C}. \tag{46}$$

The averaging-fused belief opinion $\omega_X^{\underline{\diamond}(\mathbb{C})}$ of Eq. (38) results from mapping the fused evidence belief mass of Eq. (44) back to a belief opinion by applying the bijective mapping of Eq. (15).

$$\mathbf{b}_X^{\underline{\diamond}(\mathbb{C})}(x) = \frac{\mathbf{r}_X^{\underline{\diamond}(\mathbb{C})}(x)}{W + \sum_{x\in\mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\underline{\diamond}(\mathbb{C})}(x)} \tag{47}$$

$$= \frac{\sum_{C\in\mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j\neq C} u_X^{C_j}\right)}{N \prod_{C\in\mathbb{C}} u_X^C + \sum_{x\in\mathcal{R}(\mathbb{X})} \left(\sum_{C\in\mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j\neq C} u_X^{C_j}\right)\right)} \tag{48}$$

$$= \frac{\sum_{C\in\mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j\neq C} u_X^{C_j}\right)}{\sum_{C\in\mathbb{C}} \left(\prod_{C_j\neq C} u_X^{C_j}\right)},$$

$$\text{where } \exists \mathrm{u}_{\mathrm{X}}^{\mathrm{C}} \neq 0. \tag{49}$$

The transition from Eq. (47) to Eq. (48) results from inserting Eq. (46) into Eq. (47). The transition from Eq. (48) to Eq. (49) results from applying Eq. (3).

$$u_X^{\underline{\diamond}(\mathbb{C})} = \frac{W}{W + \sum_{x\in\mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\underline{\diamond}(\mathbb{C})}(x)} \tag{50}$$

$$= \frac{N \prod_{C\in\mathbb{C}} u_X^C}{N \prod_{C\in\mathbb{C}} u_X^C + \sum_{x\in\mathcal{R}(\mathbb{X})} \left(\sum_{C\in\mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j\neq C} u_X^{C_j}\right)\right)} \tag{51}$$

$$= \frac{N \prod_{C\in\mathbb{C}} u_X^C}{\sum_{C\in\mathbb{C}} \left(\prod_{C_j\neq C} u_X^{C_j}\right)},$$

$$\text{where } \exists \mathrm{u}_{\mathrm{X}}^{\mathrm{C}} \neq 0. \tag{52}$$

The transition from Eq. (50) to Eq. (51) results from inserting Eq. (46) into Eq. (50). The transition from Eq. (51) to Eq. (52) results from applying Eq. (3).

6.4. Weighted Belief Fusion

The weighted belief fusion (WBF) operator produces averaging beliefs weighted by the opinion confidences.

The confidence c_X of an opinion ω_X is computed as:

$$c_X = 1 - u_X. \tag{53}$$

WBF is suitable for fusing source opinions in situations where the confidence should determine the opinion weight in the fusion process, which e.g. means that a vacuous opinion (i.e. an without confidence) has no effect on the fusion result.

When the arguments are conflicting multinomial opinions the fused result will be a dissonant multinomial opinion. This property could be seen as counter-intuitive when fusing opinions from human expert sources, because humans would tend to leverage belief on overlapping values and prefer vagueness over dissonance [26]. WBF is therefore best suited for frequentist situations where dissonance is preferred over vagueness. When vagueness is preferred the WBF-VM operator described in Section 6.6 can be used because it transforms dissonance into vagueness.

The definition of 2-source WBF specified in [5] was extended to multi-source WBF in [27] which is expressed below.

Definition 4 (**The Weighted Belief Fusion Operator**). Assume a hyperdomain $\mathcal{R}(\mathbb{X})$ and a situation where the variable X takes values from the domain $\mathcal{R}(\mathbb{X})$. Assume further that the different sources from a frame of N sources $\mathbb{C} = \{C_1, C_2, \dots C_N\}$ have their respective independent opinions on X. A specific source is denoted by $C \in \mathbb{C}$, and its opinion about the variable X is denoted ω_X^C.

Let $\omega_X^{\hat{\diamond}(\mathbb{C})}$ be the opinion such that

$$\omega_X^{\hat{\diamond}(\mathbb{C})} = (\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}, u_X^{\hat{\diamond}(\mathbb{C})}, \mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}), \quad \text{where} \tag{54}$$

Case I: $(\forall C \in \mathbb{C} : u_X^C \neq 0) \wedge (\exists C \in \mathbb{C} : u_X^C \neq 1)$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{b}_X^C(x)(1-u_X^C) \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i}}{\left(\sum_{C \in \mathbb{C}} \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i}\right) - N \prod_{C \in \mathbb{C}} u_X^C},$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = \frac{\left(N - \sum_{C \in \mathbb{C}} u_X^C\right) \prod_{C \in \mathbb{C}} u_X^C}{\left(\sum_{C \in \mathbb{C}} \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i}\right) - N \prod_{C \in \mathbb{C}} u_X^C},$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)(1-u_X^C)}{N - \sum_{C \in \mathbb{C}} u_X^C}, \tag{55}$$

Case II: $\exists C \in \mathbb{C}$: $u_X^C = 0$. Let $\mathbb{C}^{\mathrm{dog}} = \{C \in \mathbb{C} : u_X^C = 0\}$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\mathrm{dog}}} \gamma_X^C \mathbf{b}_X^C(x),$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = 0,$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\mathrm{dog}}} \gamma_X^C \mathbf{a}_X^C(x), \tag{56}$$

where

$$\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\mathrm{dog}}} \to 0} \frac{u_X^C}{\sum_{C_j \in \mathbb{C}^{\mathrm{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\mathrm{dog}}.$$

Case III: $\forall C \in \mathbb{C} : u_X^C = 1$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = 0,$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = 1,$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)}{N}. \tag{57}$$

The notation $u_X^{\mathbb{C}^{\mathrm{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\mathrm{dog}}$. $\omega_X^{\hat{\diamond}(\mathbb{C})}$ denotes the WBF (Weighted Belief Fusion) opinion resulting from the opinions ω_X^C provided by the sources $C \in \mathbb{C}$. By using the symbol '$\hat{\oplus}$' to denote this belief operator, we define

$$\omega_X^{\hat{\diamond}(\mathbb{C})} \equiv \underset{C \in \mathbb{C}}{\hat{\oplus}} (\omega_X^C). \tag{58}$$

It can be verified that WBF is commutative, idempotent and has the vacuous opinion as neutral element. Semi-associativity requires that three or more arguments must first be combined together in the same operation.

The argument base rate distributions are normally equal among the sources. When that is not the case the fused base rate distribution over X is specified to be the confidence-weighted average base rate distribution. In case of dogmatic arguments assume the limits in Eq. (56) to be $\gamma_X^C = 1/N$ where $N = |\mathbb{C}|$.

The WBF operator is equivalent to updating Dirichlet PDFs as the confidence-weighted average of source agents' evidence to produce posterior Dirichlet PDFs. The derivation of the confidence-weighted fusion operator is based on the bijective mapping between the belief and evidence notations described in Eq. (15).

THEOREM 1 *The weighted belief fusion operator of Definition 4 is equivalent to confidence-weighted averaging of the evidence parameters of the Dirichlet HPDF in Eq. (14).*

PROOF 1. The weighted belief fusion operator of Definition 4 is derived by mapping the argument belief opinions to evidence opinions through the bijective mapping of Eq. (15). Weighted belief fusion of evidence opinions simply consists of computing the confidence-weighted average of the evidence parameters. The fused evidence opinion is then mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the proof of Theorem 1. A more detailed explanation is provided below.

Let the N sources $C \in \mathbb{C}$ have the respective belief opinions ω_X^C. The corresponding evidence opinions $\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence parameters $\mathbf{r}_X^C$.

The weighted fusion of these bodies of evidence simply consists of weighted vector averaging of the parameters in the evidence opinions $\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X^C, \mathbf{a}_X^C)$:

$$\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}, \mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}) = \underset{C \in \mathbb{C}}{\hat{\oplus}} \mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X^C, \mathbf{a}_X^A). \tag{59}$$

More specifically, for each value $x \in \mathcal{R}(\mathbb{X})$ the confidence-weighted fusion evidence $\mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}(x)$ is computed as

$$\mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{r}_X^C(x)(1-u_X^C)}{N - \sum_{C \in \mathbb{C}} u_X^C}. \tag{60}$$

The weighted fusion opinion $\omega_X^{\hat{\diamond}(\mathbb{C})}$ of Definition 4 results from mapping the fused evidence belief mass of Eq. (59) back to a belief opinion as defined in Definition 4 by applying the bijective mapping of Eq. (15).

6.5. Uncertainty Maximisation

Uncertainty maximisation consists of transforming belief mass of an opinion ω_X into uncertainty mass while preserving the projected probability distribution $\mathbf{P}_X$.

Given a specific multinomial opinion ω_X, the corresponding uncertainty-maximised opinion is denoted $\ddot{\omega}_X = (\ddot{\mathbf{b}}_X, \ddot{u}_X, \mathbf{a}_X)$. Obviously, the base rate distribution $\mathbf{a}_X$ is not affected by uncertainty-maximisation.

The theoretical maximum uncertainty mass $\ddot{u}_X$ is determined by converting as much belief mass as possible into uncertainty mass, while preserving consistent projected probabilities. This process is illustrated in Figure 5 which shows an opinion ω_X as well as the corresponding uncertainty-maximised opinion $\ddot{\omega}_X$.

The projector line defined by the equations

$$\mathbf{P}_X(x_i) = \mathbf{b}_X(x_i) + \mathbf{a}_X(x_i) u_X, \quad i = 1, \ldots k, \tag{61}$$

which by definition is parallel to the base rate director line, and which joins $\mathbf{P}_X$ and $\ddot{\omega}_X$ in Figure 5, defines possible opinions ω_X for which the projected probability distribution is constant. As the illustration shows, the opinion $\ddot{\omega}_X$ is the uncertainty-maximised opinion when

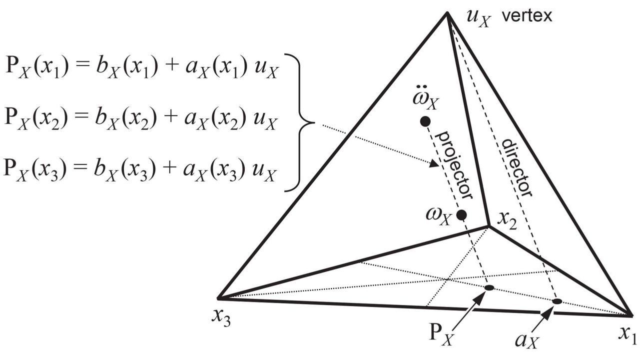

Fig. 5. Uncertainty-maximised opinion $\ddot{\omega}_X$ of multinomial opinion ω_X

Eq. (61) is satisfied and at least one belief mass of $\ddot{\omega}_X$ is zero, since the corresponding point would lie on a side of the simplex. In general, not all belief masses can be zero simultaneously, except for vacuous opinions. The example of Figure 5 shows the case where $\ddot{\mathbf{b}}_X(x_1) = 0$.

The candidate maximum uncertainty mass $\breve{u}_X(x_i)$ at each point where the projector intersects a side plane defined by $\mathbf{b}_X(x_i) = 0$ can be determined by Eq. (62):

$$\breve{u}_X(x_i) = \frac{\mathbf{P}_X(x_i)}{\mathbf{a}_X(x_i)}. \tag{62}$$

All belief masses determined according to Eq. (65) must be non-negative, which is satisfied through the constraint of Eq. (63):

$$\breve{u}_X(x_i) \leq \frac{\mathbf{P}_X(x)}{\mathbf{a}_X(x)}, \quad \forall x \in \mathbb{X}. \tag{63}$$

Under the constraint of Eq. (63) the maximised uncertainty $\ddot{u}_X$ is the minimum candidate uncertainty from Eq. (62):

$$\ddot{u}_X = \min_{x_i \in \mathbb{X}} [\breve{u}_X(x_i)]. \tag{64}$$

The belief masses under uncertainty maximisation emerge from Eq. (65) which is simply a transformation of Eq. (6):

$$\ddot{\mathbf{b}}_X(x) = \mathbf{P}_X(x) - \mathbf{a}_X(x)\ddot{u}_X. \tag{65}$$

The uncertainty-maximised opinion consists of the components denoted $\ddot{\omega}_X = (\ddot{\mathbf{b}}_X, \mathbf{a}_X, \ddot{u}_X)$. By defining $\ddot{\upharpoonleft}$ to be the unary operator for uncertainty maximisation we can write:

$$\text{Uncertainty Maximisation:} \quad \ddot{\omega}_X = \ddot{\upharpoonleft}(\omega_X). \tag{66}$$

A natural application of uncertainty maximisation is to produce epistemic opinions during opinion fusion. For that it is necessary to first generate a fused opinion, and subsequently to apply vagueness maximisation. In the case of e.g. CBF (Cumulative Belief Fusion) the combination with uncertainty maximisation is called CBF-UM (Cumulative Belief Fusion with Uncertainty Maximisation). An situation where it would be natural to apply CBF-UM could be when different witnesses express highly confident and highly conflicting opinions about whether Oswald shot Kennedy in 1968, which when fused with e.g. CBF would produce an opinion with high confidence. Since the combined testimonies in this case would be inconclusive it could be natural to apply uncertainty maximisation to the result of CBF to produce CBF-UM, as shown in the example of Section 7.

6.6. Vagueness Maximisation

In situations where people give different hypotheses it is fair to acknowledge that anyone can be wrong, and that a good consensus might be to agree that one of the hypotheses probably is right. This would typically be the situation in Zadeh's example [14] where two medical doctors give different diagnoses to explain a patient's symptoms, so that it would be natural for the doctors to agree that one of the diagnoses is correct, but that they are unable to identify which diagnosis in particular is correct. In this situation the combination of the two doctors result in a *vague* diagnosis.

Composite values $x \in \mathcal{R}(\mathbb{X})$ are state values containing multiple singleton values which e.g. can be different hypotheses such as medical diagnoses. Vague belief is belief mass assigned to a composite value, meaning that the belief mass applies to multiple singletons simultaneously. Vague belief mass thus reflects that the source believes that one of the singletons in the composite value is TRUE, without being able to identify which singleton in particular is TRUE. Vagueness is relevant for belief fusion, especially for WBF because vagueness can express compromise belief between conflicting sources. Vagueness maximisation consists of transforming belief masses on multiple singleton values into belief mass on a composite value, while preserving the projected probability distribution of Eq. (6).

In case the fused opinion ω_X is hypernomial we need to first apply Eq. (8) to compute the projected multinomial opinion $\underset{\cdot}{\omega}_X$.

Vagueness maximisation consists of transforming belief masses on multiple singleton values into a vague belief mass on the composite value containing the singletons. In case $\underset{\cdot}{\omega}_X$ has belief mass on every singleton $x \in \mathbb{X}$ then a transformation into belief mass on $\mathbb{X}$ would not be meaningful because this is the same as uncertainty mass, and the transformation would break the assumption of preserving the amount of belief mass. We must identify the value(s) $x_i \in \mathbb{X}$ that should not be subject to vague belief mass, which can be done by computing the uncertainty-maximised opinion $\ddot{\omega}_X$ as described in Section 6.5 above.

The method of uncertainty maximisation described above forms the basis for the computation of vagueness-maximised opinions which is described below in the form of 4 consecutive steps. Note that this method of vagueness maximisation applies to multinomial opinions. Hence, if the goal is to apply vagueness maximisation to a hyper-opinion, a necessary preliminary step is to first project it to a multinomial opinion according to Eq. (9)

Step 1:
Compute $\ddot{u}_X$ according to the procedure for uncertainty-maximisation described in Section 6.5. Let $\mathbb{X}^{[1]}_{\text{cut}}$ be the cut-out set of values x_i for which $\check{u}_X(x_i) = \ddot{u}_X$ with reference to Eq. (62) and Eq. (64). Note that $\mathbb{X}^{[1]}_{\text{cut}}$ may contain a single or multiple values.

Case A: $|\mathbb{X}^{[1]}_{\text{cut}}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ of the singleton value $x \in \mathbb{X}^{[1]}_{\text{cut}}$ and proceed to Step 2.

Case B: $1 < |\mathbb{X}^{[1]}_{\text{cut}}| < |\mathbb{X}|$. The composite value $x^{[1]}_{\text{vag}} = \{x \in \mathbb{X}^{[1]}_{\text{cut}}\}$ gets assigned the vague belief mass $\mathbf{b}_X(x^{[1]}_{\text{vag}})$ according to Eq. (67).

$$\mathbf{b}_X(x^{[1]}_{\text{vag}}) = \sum_{x \in \mathbb{X}^{[1]}_{\text{cut}}} \mathbf{b}_X(x). \tag{67}$$

Then proceed to Step 2.

Case C: $|\mathbb{X}^{[1]}_{\text{cut}}| = |\mathbb{X}|$: Split $\mathbb{X}$ into two exclusive sets $\mathbb{X}^{[1]}_{\text{res}}$ and $\mathbb{X}^{[2]}_{\text{res}}$ for which the respective sums of projected probability $\mathbf{P}(\mathbb{X}^{[1]}_{\text{res}})$ and $\mathbf{P}(\mathbb{X}^{[2]}_{\text{res}})$ are (approximately) equal. While this is a form of the knapsack problem we propose to simply sum up the greatest projected probabilities until the sum is greater than 0.5, and assign the corresponding set of values to $\mathbb{X}^{[1]}_{\text{res}}$, and the remaining values to $\mathbb{X}^{[2]}_{\text{res}}$. Define the composite values $x^{[1]}_{\text{vag}} = \{x \in \mathbb{X}^{[1]}_{\text{res}}\}$ and $x^{[2]}_{\text{vag}} = \{x \in \mathbb{X}^{[2]}_{\text{res}}\}$. Assign the vague belief masses $\mathbf{b}_X(x^{[1]}_{\text{vag}}) = \sum_{x \in \mathbb{X}^{[1]}_{\text{res}}} \mathbf{b}_X(x)$ and $\mathbf{b}_X(x^{[2]}_{\text{vag}}) = \sum_{x \in \mathbb{X}^{[2]}_{\text{res}}} \mathbf{b}_X(x)$. Proceed to the Final Step.

Step 2:
We exclude $\mathbb{X}^{[1]}_{\text{cut}}$ to produce the residual set $\mathbb{X}^{[2]}_{\text{res}}$:

$$\mathbb{X}^{[2]}_{\text{res}} = \mathbb{X} \setminus \mathbb{X}^{[1]}_{\text{cut}}. \tag{68}$$

Case A: $|\mathbb{X}^{[2]}_{\text{res}}| = 0$. Proceed to the Final Step.

Case B: $|\mathbb{X}^{[2]}_{\text{res}}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ on the singleton value $x \in \mathbb{X}^{[2]}_{\text{res}}$. Proceed to the Final Step.

Case C: $|\mathbb{X}^{[2]}_{\text{res}}| \geq 2$. Now we focus exclusively on values $x_i \in \mathbb{X}^{[2]}_{\text{res}}$ when applying the constraint of Eq. (63). The next synthetic maximum uncertainty mass is:

$$\ddot{u}^{[2]}_X = \min_{x_i \in \mathbb{X}^{[2]}_{\text{res}}} [\check{u}_X(x_i)]. \tag{69}$$

Eq. (70) gives the corresponding synthetic belief masses:

$$\ddot{\mathbf{b}}^{[2]}_X(x) = \mathbf{P}_X(x) - \mathbf{a}_X(x)\ddot{u}^{[2]}_X, \quad \forall x \in \mathbb{X}^{[2]}_{\text{res}}. \tag{70}$$

We define the composite value $x^{[2]}_{\text{vag}} = \{x \in \mathbb{X}^{[2]}_{\text{res}}\}$. The vague belief mass $\mathbf{b}_X(x^{[2]}_{\text{vag}})$ can then be assigned according to Eq. (71)

$$\mathbf{b}_X(x^{[2]}_{\text{vag}}) = \sum_{x \in \mathbb{X}^{[2]}_{\text{res}}} (\mathbf{b}_X(x) - \ddot{\mathbf{b}}^{[2]}_X(x)). \tag{71}$$

Let the iterative step index be denoted η. Set $\eta = 3$ and proceed to Step η.

Step η:
Let $\mathbb{X}^{[\eta-1]}_{\text{cut}}$ be the set of values x_i for which $\check{u}_X(x_i) = \ddot{u}^{[\eta-1]}_X$ with reference to Eq. (62) and Eq. (64). We exclude $\mathbb{X}^{[\eta-1]}_{\text{cut}}$ from $\mathbb{X}^{[\eta-1]}_{\text{res}}$ to produce the residual set $\mathbb{X}^{[\eta]}_{\text{res}}$:

$$\mathbb{X}^{[\eta]}_{\text{res}} = \mathbb{X}^{[\eta-1]}_{\text{res}} \setminus \mathbb{X}^{[\eta-1]}_{\text{cut}}. \tag{72}$$

Case A: $|\mathbb{X}^{[\eta]}_{\text{res}}| = 0$. Proceed to the Final Step.

Case B: $|\mathbb{X}^{[\eta]}_{\text{res}}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ on the singleton value $x \in \mathbb{X}^{[\eta]}_{\text{res}}$. Proceed to the Final Step.

Case C: $|\mathbb{X}^{[\eta]}_{\text{res}}| \geq 2$. Now we focus exclusively on values $x_i \in \mathbb{X}^{[\eta]}_{\text{res}}$ when applying the constraint of Eq. (63). The next synthetic maximum uncertainty mass is:

$$\ddot{u}^{[\eta]}_X = \min_{x_i \in \mathbb{X}^{[\eta]}_{\text{res}}} [\check{u}_X(x_i)]. \tag{73}$$

The computation of the belief masses emerges from Eq. (74):

$$\ddot{\mathbf{b}}^{[\eta]}_X(x) = \mathbf{P}_X(x) - \mathbf{a}_X(x)\ddot{u}^{[\eta]}_X. \tag{74}$$

We define the composite value $x^{[\eta]}_{\text{vag}} = \{x \in \mathbb{X}^{[\eta]}_{\text{res}}\}$. The vague belief mass $\mathbf{b}_X(x^{[\eta]}_{\text{vag}})$ can then be assigned according to Eq. (75)

$$\mathbf{b}_X(x^{[\eta]}_{\text{vag}}) = \sum_{x \in \mathbb{X}^{[\eta]}_{\text{res}}} (\ddot{\mathbf{b}}^{[\eta-1]}_X(x) - \ddot{\mathbf{b}}^{[\eta]}_X(x)). \tag{75}$$

Increment the step index η as $\eta := \eta + 1$, then repeat Step η.

Final Step:
Finally, the components of the vagueness-maximised opinion $\dot{\omega}_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ can be assembled, consisting of the computed vague belief masses $\mathbf{b}_X(x^{[\eta]}_{\text{vag}})$, and whenever applicable the singular belief masses $\mathbf{b}_X(x_i)$, in addition to the original uncertainty mass u_X and base rate distribution $\mathbf{a}_X$. This ends the process of vagueness maximisation.

By defining the unary operator $\dot{\uparrow}$ to represent vagueness maximisation we can write

$$\text{Vagueness Maximisation:} \quad \dot{\omega}_X = \dot{\uparrow}(\omega_X). \tag{76}$$

A natural application of vagueness maximisation is to produce compromise belief when fusing opinions from multiple (conflicting) sources. To this end it is necessary to first generate a fused opinion with WBF, and subsequently to apply vagueness maximisation. This combination is called WBF-VM (Weighted Belief Fusion with Vagueness Maximisation) and is denoted '$\hat{\dot{\oplus}}$'.

As an alternative to WBF-VM for belief fusion with compromise, the belief fusion operator CCF (Consensus & Compromise Fusion) has been described with a simple two-source version [5] as well as with a multi-source version [27]. The definition of multi-source CCF is rather complex [27], whereas multi-source WBF-VM is rather simple in comparison. In situations where it

TABLE 5
Zadeh's numerical example applied to belief constraint fusion (BCF), cumulative belief fusion (CBF), cumulative belief fusion with uncertainty maximisation (CBF-UM), averaging belief fusion (ABF), weighted belief fusion (WBF) and weighted belief fusion with vagueness maximisation (WBF-VM)

		Source opinions:		Fused opinions resulting from applying:					
		A	B	BCF	CBF	CBF-UM	ABF	WBF	WBF-VM
$b_X(x_1)$	=	0.99	0.00	0.00	0.495	0.485	0.495	0.495	0.000
$b_X(x_2)$	=	0.01	0.01	1.00	0.010	0.000	0.010	0.010	0.010
$b_X(x_3)$	=	0.00	0.99	0.00	0.495	0.485	0.495	0.495	0.000
$b_X(x_1,x_2)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.000
$b_X(x_1,x_3)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.990
$b_X(x_2,x_3)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.000
u_X	=	0.00	0.00	0.00	0.000	0.030	0.000	0.000	0.000

TABLE 6
A variation of Zadeh's example applied to belief constraint fusion (BCF), cumulative belief fusion (CBF), cumulative belief fusion with uncertainty maximisation (CBF-UM), averaging belief fusion (ABF), weighted belief fusion (WBF) and weighted belief fusion with vagueness maximisation (WBF-VM)

		Source opinions:		Fused opinions resulting from applying:					
		A	B	BCF	CBF	CBF-UM	ABF	WBF	WBF-VM
$b_X(x_1)$	=	0.98	0.00	0.889	0.890	0.880	0.882	0.889	0.806
$b_X(x_2)$	=	0.01	0.01	0.011	0.010	0.000	0.010	0.010	0.010
$b_X(x_3)$	=	0.00	0.90	0.091	0.091	0.081	0.090	0.083	0.000
$b_X(x_1,x_2)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.000
$b_X(x_1,x_3)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.166
$b_X(x_2,x_3)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.000
u_X	=	0.01	0.09	0.009	0.009	0.039	0.018	0.018	0.018

is suitable to apply a fusion operator with belief compromise, the most practical choice is therefore to apply WBF-VM which is included in the example of Section 7.

7. COMPARISON OF FUSION OPERATORS

The fusion example in Table 5 takes as input arguments the numerical belief masses from Zadeh's example [14]. In this example, the sources are two medical doctors who each have an opinion about the hypothesis space of three possible diseases, and Dempster's rule (called BCF (Belief Constraint Fusion) in subjective logic) is applied for fusing the two opinions. The counter-intuitive results produced by Dempster's rule (BCF) demonstrate that Dempster's rule is unsuitable for this particular category of situations. A more suitable operator for the situation of the two doctors is WBF-VM (Weighted Belief Fusion with Vagueness Maximisation), because it preserves common belief and produces compromise belief from conflicting belief sources.

Exactly the same pair of argument opinions can of course occur in other fusion situations as well. Table 5 shows the results of fusion with each operator described in the previous sections, where the the fused result opinion produced by a given operator is sound and intuitive according to the corresponding situation category described in Section 2.

On an abstract level, sources A and B provide opinions about the hypothesis space $\mathbb{X} = \{x_1, x_2, x_3\}$ with variable X. The base rate distributions are assumed to be equal and uniform, expressed as $\mathbf{a}_X^A = \mathbf{a}_X^B = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$.

Each operator produces intuitive results given respective relevant situations for which the operators are suitable. For example, in the medical situation of the original Zadeh's example where two medical doctors A and B have conflicting opinions about the diagnosis of a patient, WBF-VM produces vague belief in the form of $\mathbf{b}_X^{A\hat{\diamond}B}(x_1, x_3) = 0.99$ which seems natural until the doctors can agree on a single diagnosis for the patient. The BCF operator produces a sound and intuitive fused opinion with the same argument opinions when e.g. assuming a situation where two friends express preferences for watching a film at the cinema.

Fusion of dogmatic conflicting opinions, i.e. where $u_X = 0$, is defined for all operators except for BCF. If the fusion situation is determined to be in the BCF category the interpretation of fusing dogmatic conflicting opinions is that there is no solution, which is perfectly logic. See Section 6.1.5 for an example of this situation.

Zadeh's example as in Table 5 does not clearly expose the difference between the various belief fusion operator because many fusion operators produce equal results when the sources are dogmatic as in this case. The modified example in Table 6 brings greater differentiation in the fusion results by introducing unbalanced levels of uncertainty in the argument opinions. The difference between the arguments of Table 5 and Table 6 can be interpreted and explained through the assumptions of the various belief-fusion categories with regard to how conflicting belief arguments are handled in the belief fusion process.

8. DISCUSSION AND CONCLUSION

We argue that the main research question in belief fusion is not about finding the single most correct belief fusion operator, because no single operator is suitable for all situations. Instead, the interesting question and the biggest challenge is how to select the most suitable belief fusion operator for a given situation of belief fusion. For this purpose we propose to classify situations of belief fusion into different categories, where a set of belief-fusion assumptions can be used as criteria for selecting the category to which a specific belief fusion situation belongs.

This article illustrates the importance of selecting a belief fusion operator that adequately matches the situation to be modelled and analyzed. It is scientifically misguided to follow the approach of always applying the favourite belief fusion operator with which the analyst or scientist happens to be familiar, without regard to the nature of the situation to be modelled. By using the selection criteria to categorise a given belief-fusion situation and applying the corresponding belief fusion operator the analyst is able to obtain sound and useful results more consistently than by simply making an uninformed choice when selecting a belief fusion operator for a given application.

ACKNOWLEDGMENT

We thank the anonymous reviewers of the Journal of Advances in Information Fusion for their deep insight into the topic, and for their valuable suggestions to improve the manuscript.

REFERENCES

[1] A. Jøsang, P. C. Costa, and E. Blash
"Determining model correctness for situations of belief fusion."
in *Proceedings of the 16th International Conference on Information Fusion (FUSION 2013)*. IEEE, Los Alamitos, 2013, pp. 1225–1232.

[2] "ETURWG (Evaluation of Techniques for Uncertainty Representation Working Group),"
Available from http://eturwg.c4i.gmu.edu/, accessed June 2018.

[3] P. C. G. Costa, K. B. Laskey, E. P. Blasch, and A.-L. Jousselme
"Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.

[4] E. P. Blasch, P. C. G. Costa, K. B. Laskey, D. Stampouli, G. W. Ng, J. Schubert, R. Nagi, and P. Valin
"Issues of Uncertainty Analysis in High-Level Information Fusion: Fusion 2012 Panel Discussion."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.

[5] A. Jøsang
Subjective Logic: A Formalism for Reasoning Under Uncertainty.
Springer, Heidelberg, 2016.

[6] P. H. Foo and G. W. Ng
"High-level information fusion: An overview,"
J. Adv. Inf. Fusion, vol. 8, no. 1, pp. 33–72, 2013.

[7] A. Jøsang, M. Ivanovska, and T. Muller
"Trust Revision for Conflicting Sources."
in *Proceedings of the 18th International Conference on Information Fusion (FUSION 2015)*. IEEE, Los Alamitos, 2015.

[8] A. Jøsang and R. Hankin
"Interpretation and fusion of hyper-opinions in subjective logic."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*. IEEE, Los Alamitos, Singapore, July 2012.

[9] G. Shafer
A Mathematical Theory of Evidence.
Princeton University Press, 1976.

[10] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin
Bayesian Data Analysis, 3rd ed.
Chapman and Hall/CRC, Boca Raton, 2013.

[11] R. K. Hankin
"A generalization of the Dirichlet distribution,"
Journal of Statistical Software, vol. 33, no. 11, pp. 1–18, February 2010.

[12] A. Jøsang
"Generalising Bayes' Theorem in Subjective Logic."
in *International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2016)*. IEEE, Los Alamitos, 2016.

[13] A. Jøsang and S. Pope
"Dempster's rule as seen by little colored balls,"
Computational Intelligence, vol. 28, no. 4, pp. 453–474, November 2012.

[14] L. A. Zadeh
"Review of Shafer's 'A Mathematical Theory of Evidence',"
AI Magazine, vol. 5, pp. 81–83, 1984.

[15] M. Daniel
"Associativity in combination of belief functions,"
in *Proceedings of 5th Workshop on Uncertainty Processing*, J. Vejnarová, Ed. Ediční oddelění VŠE 2000, Praha, 2000, pp. 41–54.

[16] D. Dubois and H. Prade
"Representation and combination of uncertainty with belief functions and possibility measures,"
Comput. Intell., vol. 4, no. 3, pp. 244–264, 1988.

[17] A. Jøsang
"The consensus operator for combining beliefs,"
Artificial Intelligence, vol. 142, no. 1–2, pp. 157–170, October 2002.

[18] A. Jøsang, J. Diaz, and M. Rifqi
"Cumulative and averaging fusion of beliefs,"
Information Fusion, vol. 11, no. 2, pp. 192–200, 2010, doi:10.1016/j.inffus.2009.05.005.

[19] E. Lefevre, O. Colot, and P. Vannoorenberghe
"Belief functions combination and conflict management,"
Information Fusion, vol. 3, no. 2, pp. 149–162, June 2002.

[20] C. K. Murphy
"Combining belief functions when evidence conflicts,"
Decision Support Systems, vol. 29, pp. 1–9, 2000.

[21] P. Smets
"The Combination of Evidence in the Transferable Belief Model,"
IEEE Transansactions on Pattern Analysis and Machine Intelligence, vol. 12, no. 5, pp. 447–458, 1990.

[22] R. Yager
"On the Dempster-Shafer framework and new combination rules,"
Information Sciences, vol. 41, pp. 93–137, 1987.

[23] A. Jøsang
"Multi-agent preference combination using subjective logic,"
in *International Workshop on Preferences and Soft Constraints (Soft'11)*, Perugia, Italy, 2011, pp. 61–75.

[24] J. Kohlas and P. Monney
A Mathematical Theory of Hints. An Approach to Dempster-Shafer Theory of Evidence,
ser. Lecture Notes in Economics and Mathematical Systems. Springer-Verlag, 1995, vol. 425.

[25] S. Pope and A. Jøsang
"Analsysis of Competing Hypotheses using Subjective Logic,"
in *Proceedings of the 10th International Command and Control Research and Technology Symposium (ICCRTS)*. United States Department of Defense Command and Control Research Program (DoDCCRP), 2005.

[26] A. Jøsang, J.-H. Cho, and F. Chen
"Uncertainty Characteristics of Subjective Opinions,"
in *Proceedings of the 21st International Conference on Information Fusion (FUSION 2018)*. Available at: http://folk.uio.no/josang/papers/JCC2018-FUSION.pdf: IEEE, Los Alamitos, 2018.

[27] R. W. van der Heijden, H. Kopp, and F. Kargl
"Multi-Source Fusion Operations in Subjective Logic,"
in *Proceedings of the 21st International Conference on Information Fusion (FUSION 2018)*. Available at: https://arxiv.org/abs/1805.01388: IEEE, Los Alamitos, 2018.

Audun Jøsang is a Professor in information security at the University of Oslo, and is also an Adjunct Professor at QUT in Australia. Prof. Jøsang is the main author behind subjective logic which is widely applied by researchers and practitioners in the areas of uncertainty representation, belief fusion, subjective Bayesian reasoning, reputation systems, and computational trust. He holds a Master's degree in information security from Royal Holloway College, University of London and a PhD from NTNU in Norway.

Uncertainty in avionics analytics ontology for decision-making support

CARLOS C. INSAURRALDE
ERIK BLASCH

With the growing congestion in the airspace, Air Traffic Management (ATM) requires advances in massive data processing, sophisticated avionics techniques, coordination with weather updates, and assessment of multiple types of uncertainty. The complex situation overwhelms pilots and ATM controllers. To provide dependable artificial decision-making support for ATM and Unmanned Aerial System Traffic Management (UTM) systems, ontologies are an attractive knowledge technology. This paper proposes an Avionics Analytics Ontology (AAO) to bring together different types of uncertainties including semantic from operators, sensing from navigation, and situation from weather modeling updates. The approach is aligned with the Uncertainty Representation and Reasoning Evaluation Framework (URREF), that develops an uncertainty ontology. The degree of uncertainty to improve effectiveness in ATM/UTM decision-making processes quantifies information veracity; in addition to accuracy, timeliness, and confidence. Application examples are presented that involves two ATM/UTM operation scenarios where Unmanned Aerial Vehicles (UAVs) fly nearby commercial aircraft and/or airports which requires situation awareness safety response. As compared to a baseline approach without Automatic Dependent Surveillance-Broadcast (ADS-B), results from recorded ADS-B data demonstrate a over 0.75 veracity improvement) from Newark Liberty International Airport.

Manuscript received March 8, 2018; Revised January 14, 2019; released for publication February 4, 2019.

Refereeing of this contribution was handled by Paulo Costa.

Authors' addresses: C. C. Insaurralde, Department of Engineering Design and Mathematics, University of the West of England, Coldharbour Lane, Bristol BS16 1QY, UK (E-mail: carlos.c.insaurralde@gmx.com). E. Blasch, Air Force Office of Scientific Research, Arlington, VA 22203, USA (E-mail: erik.blasch@gmail.com).

The paper is an extension to a DASC 2017 submission.

1. INTRODUCTION

There has been a growth in the use of ontologies for communities such as medical diagnostics, target assessment, and chemical composition. An area that can benefit from an unified ontology is that of avionics; with only limited reporting for groups interested in supporting the the Federal Aviation Administration (FAA) Next Generation Air Transportation System (NextGen) and the Single European Sky ATM Research (SESAR) systems. The use of ontologies would enhance the coordination between physics-based sensing (e.g., positing and navigation), human-derived communications (e.g., call sign and Notice to Airmen—NOTAMS), and situation reporting (e.g., weather map updates on the cockpit displays). The ontologies support a common taxonomy for reporting to help pilots and Air Traffic Controllers (ATC) make difficult decisions in the context of data, feature, and information uncertainty.

Air Traffic Management (ATM) is growing in complexity as avionics systems are getting sophisticated, airspaces are densely occupied, and air transport is flying in more adverse weather conditions. Overwhelmed aviators, air traffic controllers, and air transport businesses have to prioritize dependability (safety, security, reliability, etc.) in aviation procedures while sharing the airspace with other types of aircraft such as unmanned aerial vehicles (UAVs). Due to the emergence of inexpensive UAVs, accessible from a diverse set of users from the scientific, recreational, commercial, civil and military aviation communities, there is need for a common set of rules (or procedures) for Unmanned aerial system Traffic Management (UTM). ATM/UTM aerospace information management systems need be (1) *efficient* with larger amounts of data, (2) *effective* with combining information from different sources such as weather forecasts, flight profiles, airports, and UAVs, and (3) *relevant* through reducing uncertainty in decision support systems (DSS).

An attractive approach to support decision making in advanced ATM/UTM systems is the implementation of *Ontologies for NextGen Avionics Systems* (ONAS). Ontologies are meant to model cognitive processes by representing and reasoning on knowledge. Following this direction, a proof of concept for an ONAS solution was proposed [1], which has a knowledge-based ATM/UTM architecture for avionics analytics. In this *Avionics Analytics Ontology* (AAO), an ontological database captures information (data along with meaning) as to concepts, entities, and relations in order to build knowledge related to weather, flights, and airspace. The ontology enables artificial reasoning to make decisions based on the knowledge stored and the current situation estimates.

The AAO supports Decision-Support System (DSS) for ATM/UTM to dependably minimize human intervention by making decisions simultaneously based on multiple information inputs. A key issue when designing DSSs is the credibility, reliability, and veracity of the

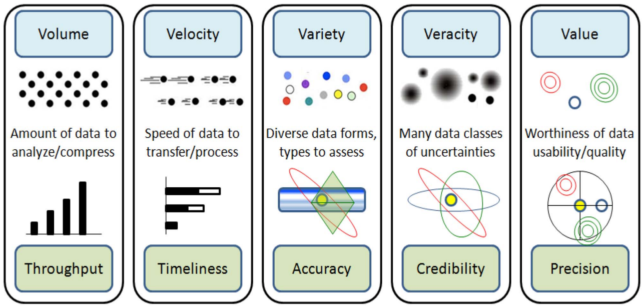

Fig. 1. Big Data Constructs and Uncertainty Metrics

gathered information. Veracity is an element of big data that assesses the truthful of the data and is include in the 4 V's of big data: volume, variety, velocity, and veracity; while other options include value, volatility, and visualization. Veracity can be used to assess the truth of data for such cases are aircraft sensor failures [2]. The alignment of the big data V's and the Uncertainty Representation and Reasoning Evaluation Framework (URREF) are shown in Fig. 1.

This paper proposes to endow the above AAO with semantic uncertainty for input information to improve DSS effectiveness in the decisions taken. The proposed approach is based on the Uncertainty Representation and Reasoning Evaluation Framework (URREF). It deals with the URREF input criterion (i.e., weight of evidence, relevance to problem, and credibility). The ontology development presents *veracity* as part of information *credibility*. The AAO construct captures not only information on concepts, entities, and relations; but also uncertainty of the input information as to its veracity for metadata information. The AAO considers the degree of uncertainty by means of quantitative metrics of throughput, timeliness, confidence, and accuracy. Veracity then includes qualitative metrics such as reliability, credibility, and quality mapped to precision and recall. The URREF assessment enhances avionics DDS analytics when considering semantic and physical data sources. Ultimately, it will enhance Situation AWareness (SAW) as well as Situation Assessment (SA) in information fusion [3].

This paper presents application examples that involve two ATM/UTM operation scenarios where UAVs are flying nearby commercial aircraft and/or airports. The closeness of UAV proximity has an impact on the ATM/UTM decisions taken by the DSS. The DSS provided by the URREF-based AAO takes into account semantics from updates of weather maps, airport maps, and route maps as well as information uncertainty (veracity of the above updates, in particular from flights). The scenarios are meant to represent realistic flight situations since they make use of real-time airspace information provided by a flight tracking service (Flightradar24 [4]).

The rest of the paper is organized as follows. Section 2 recalls existing approaches for ATM. Section 3 reviews supporting and existing technologies and concepts regarding SAW and SA. Section 4 introduces the URREF. Section 5 discusses the AAO foundations for ontological decision-making support in avionics, and the uncertainty scope and considerations for veracity metrics. Section 6 presents applications examples by means of three application examples. The final section presents the conclusions and future research steps.

2. EXISTING APPROACHES FOR AIR TRAFFIC MANAGEMENT

Air Traffic Management evolved with air services and current incorporates three methods: Air Traffic Control (ATC), Air Traffic flow Management (ATFM), and aeronautical information services (AIS).

The approaches to decision support improved with technology, collaboration, visualization, and mandates. For example, in 1982, Pararas developed a modular system using Mixed Integer Linear Programming Language (MILP) modular automation approach for ATM/C that afforded aircraft dynamics, a flexible controller interface, and a real-time terminal area simulation [5]. Many approaches in the 90s sought to use automation for optimization of airspace data to support visualization. In 2000, Ball et al [6] reported on efforts for collaborative decision making using the distribution of the National Airspace System (NAS) status information and the management of en-route traffic flow through optimization with a ground delay program, convective weather forecast, and LAADR (Low Altitude Arrival and Departure Routes) for congestion avoidance. The FAA methods were documented to include decision making, capacity performance, traffic flow, and weather support [7]. Access to the information services in a unified display assists controllers, pilots and dispatchers for a flight management system, as demonstrated by the NASA Multi Aircraft Control System (MACS) [8]. A key element for ATM is the International Civil Aviation Organization (ICAO) air traffic management [9] information that includes traffic flow requirements, separation rules, flight information, coordination routines, message format, phraseology, ADS services, and Controller/pilot Data Link Communications (CPDLC).

ATM decision support systems design sought advances in airspace dynamics developed for monitoring, capacity flow, and scheduling for system wide information management (SWIM), that did not focus on the information services. In 2008, the Sky-Scanner project sought to develop LIDAR sensing for monitoring as an improved decision support system for ATM [9]. The data was utilized with a risk-based approach from the airspace rules to augment capacity flow [10]. Further, the Next Generation Air Traffic Management (NG-ATM) operational concepts were sought for the Single European Sky Air Traffic Management Research (SESAR) and the United States' Next Generation of Air Transportation System (NextGen) programs which included a 4-Dimensional Trajectory Negotiation and Validation System [11]. The system was to support safety, capacity, efficiency, and the environment. An optimization method for spatial-temporal airspace use was developed to assist in scheduling for intent negotiation. Efforts continue to provide techniques for ATM including: performance based operations, capacity and flow control, efficiency and environmental impact, departure and arrival management, Terminal area (TMA) and surface operation interactions, complexity management, and planning quality.

An analysis of text messages was conducted using a Conflict Probe which predicts potential airspace impending separation violations and a Trajectory Predictor suggesting a more accurate aircraft position [12] The Common Message Set (CMS) relays flight plan, altitude, radar tracking and other data. The message data includes: Flight Plan Information (FH), Flight Plan Amendment (AH), Cancellation Information (CL), Interim Altitude Information (LH), Departure Information (DH), and Converted Route Information (HX). A Java En Route Development Initiative (JEDI) software was used to translate the message types for separation error prediction [13]. However, researchers have yet to focus on the semantic analysis of the meaning of the messages as an information service. The need for an ontology was highlighted by Koelle and Strijland [14]. NASA sought to development an otology as evidenced in the slides [15] and the current version is released as the NASA Air Traffic Management Ontology (atmonto) [16]. To the best of our knowledge, no reports can be found of a literature publication using the NASA ATM ontology.

3. DECISION-MAKING SUPPORT IN AVIONICS ANALYTICS

This section reviews supporting and existing technologies and concepts regarding SAW and SA in support of the analysis towards the URREF.

A. Situation Awareness

The decision-making process is based on the four-stage loop called Observe-Orient-Decision-Act (OODA) [17]. The OODA loop is essential for situation aawareness assessment in information fusion [18]. Fig. 2 shows a SAW model.

SAW allows systems to understand dynamic and complex environments, and operate with them. Cognitive SAW can be divided into three separate levels: perception of the elements in the environment, comprehension of the current situation, and projection of future status [18].

The concepts of the OODA loop enable a processing of information. The Observation stage is the SAW perception level. The Orientation stage takes into account the information acquired from the Observation stage and the knowledge represented by the ontology, to understand the situation (SAW comprehension level).

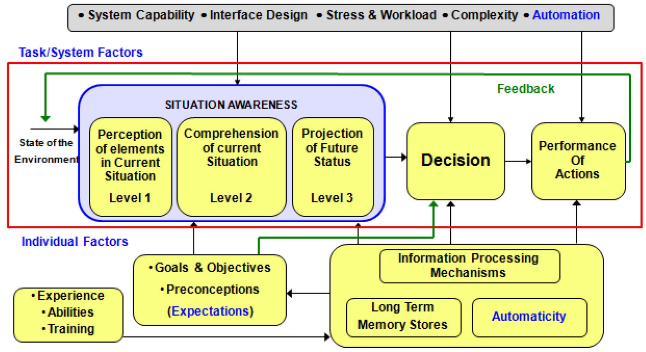

Fig. 2. Situation Awareness (SAW) Model

Level 0 – Data Assessment
Level 1 – Object Assessment
Level 2 – Situation Assessment
Level 3 – Impact Assessment
Level 4 – Process Refinement
Level 5 – User Refinement
Level 6 – Mission Management

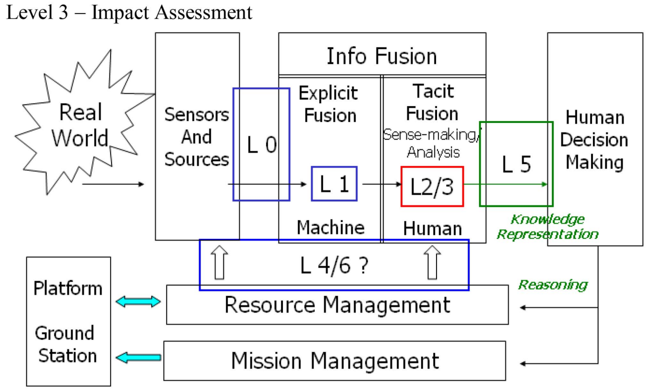

Fig. 3. Data Fusion Information Group (DFIG) model

The Decision stage is carried out at the SAW projection level. The Action stage closes the OODA loop by carrying out actions according to the adaption made in the previous stage.

SAW involves the events, states, condition, and activities of the environment dynamics as to time and space from which some situations arise (in particular those changes that occurred in the environment over some time interval). A *situation* is defined by a specific state after a sequence of events (with intermediate states, and activities with pre and post conditions). The situation is concerned with the comprehension of the environment features, and with the evolvement of these features over time.

SAW decision making mechanisms are critical for problem-solving processes that are preformed every time step for a situation from which data is collected at level 0 information fusion according to the Data Fusion Information Group Model [19], [20].

B. Situation Assessment

Situation assessment takes place at level 2 (SAW comprehension) in data fusion models. The Data Fusion Information Group Model levels include (Fig. 3):

In the DFIG model, the goal was to separate the information fusion (IF) (L0–L3) and resource management (RM) functions (L4–L6) [21], [22].

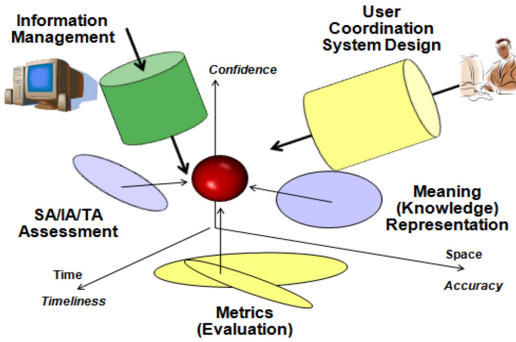

Fig. 4. Situation Uncertainty

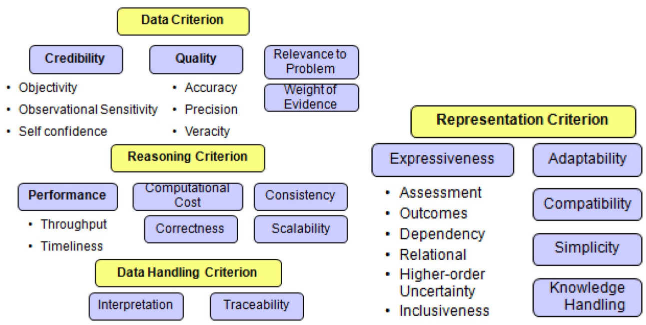

Fig. 5. URREF Categories

For UTM systems, there is both the resource management across sensors, users, and the mission (SUM) to coordinate with the objects, situations, and threats. The elements of the airspace need to be provided to air traffic controllers for enhanced SAW. Two integral concepts for Level 5 "User Refinement" information fusion are displays to support usability [23] and information management systems that are trustworthy [24].

Uncertainty of a situation is based on information, assessment, and knowledge (shown in Fig. 4).

A binding element between the levels of fusion to reduce uncertainty is an ontology [25], [26]. The URREF model provides an ontology that supports the interaction between low-level information fusion (LLIF) and high-level information fusion [27].

4. UNCERTAINTY REPRESENTATION AND REASOING EVALUATION FRAMEWORK

The URREF was developed and used for analysis over imagery [28], detection [29], and text data [30]. The URREF supports uncertainty analysis [31] such as for trust [32] applications. The URREF can advance methods for image quality [33], object recognition [34], and object tracking [35]. Inherently, it is the ontology of metrics of uncertainty that can support DSS.

Recent efforts include applications for rhino poaching assessment [36], maritime anomaly detection [37], and cyber analytics [38]. The URREF developments are meant to support decision making [39] and context [39]. The ontology can resolve the decades old problem of

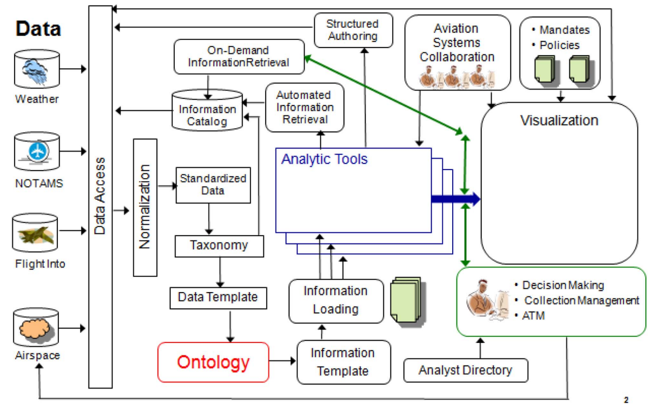

Fig. 6. Use of Ontologies for avionics analytics

relevance metrics in Simultaneous Tracking and Identification (STID) methods [40, 42].

A. URREF Ontology

The key elements from the URREF include data quality issues of accuracy, precision, and veracity (as shown in the current categories of the URREF in Fig. 5) within an OODA architecture. While accuracy and prediction have been explored, *veracity* requires further inspection. More details on the URREF are available at the Evaluation of Technologies for Uncertainty Representation (ETUR) working group (http://eturwg.c4i.gmu.edu/).

B. Ontologies for Air Traffic Management

There is an emergence of interest of the use of ontologies for ATM and aerospace technologies [43–48]. Examples include the Federal Aviation Administration (FAA) Next Generation Air Transportation System (NextGen) [49] and the Single European Sky ATM Research (SESAR) [50] systems. In order to frame the discussion, Fig. 6 highlights an example of how ontologies are included in an avionics system analysis. Using the incoming data from weather, flight profiles, and airports; that data needs to be accessed and normalized. Structuring the data is enabled with templates and ontologies. The structured ontology organizes the information (including syntactic and semantic metadata) for analytic tools. The resulting analytics supports visualization for aviators and Air Traffic Controllers (ATCs). Examples include mandates, current reports, and airspace information. Hence, ontologies afford a common method to organize, process, and share data.

For air traffic management, System Wide Information Management (SWIM) including the ATM Information Reference Model (AIRM) [51], the Information Service Reference Model (ISRM), and the SWIM Technical Infrastructure (SWIM-TI) are being developed [52]. The concept of SWIM is an emerging concept to manage information for aviation systems for various ATM networks [53]. The SWIM approach defines concepts for ATM as well as specifies what kind

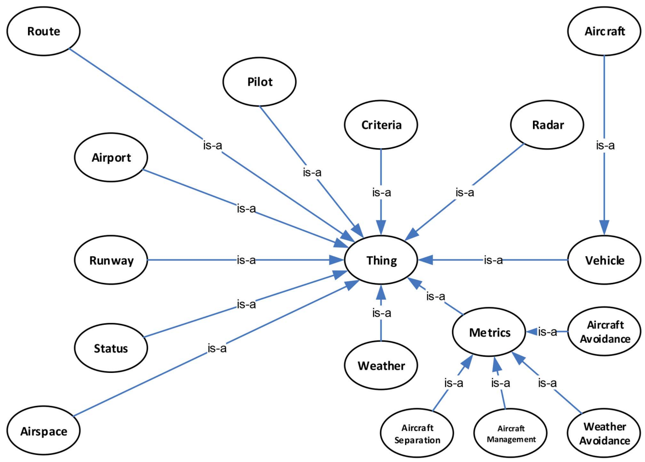

Fig. 7. Main OWL components of the AAO

of information has to be shared, and what stakeholders have to share such information [54]. An AIRM example requiring ontologies is semantic filtering of notices to airmen [55].

Key developments for SESAR and NextGen include the potential for ontological capabilities.

5. ONTOLOGY AND UNCERTAINTY

This section presents foundations of the AAO and the integration of URREF in the AAO to include uncertainty.

A. Avionics Analytics Ontology

The syntax (symbols and rules) of the Avionics Analytics Ontology (AAO) is based on the Description Logic (DL) syntax structure. However, the implementation language for the ontology ultimately defines the syntax to semantically specify and describe ontology elements. The OWL and the Protégé tool [56] are selected to realize the ontology for the approach proposed.

The main OWL components to be created are the concepts (classes), properties for individuals, and instances of classes (individuals) are shown in Fig. 7. These components are set for AAO as follows:

- *Classes (concepts)*. They are conceptually defined as classes (special datatype) in object-oriented programming languages. Thus, they can be atomic classes (stand-alone ones) or associate classes (subclasses) along with “is-a” links. The main AAO classes are: vehicle (aircraft), radar, criteria, pilot, route, airport, runway, status, airspace, weather, and metrics. Fig. 5 shows the above is-a relations between classes.
- *Properties (roles)*. They are basically relationships between classes (or eventually individuals). The OWL allows for properties on objects (based on classes) or data (specific values). The first version of AAO only includes properties for objects as follows: hasRadar, hasPilot, hasRoute, hasTakeoff, hasLanding, hasAirspace, hasRunnay, hasStatus, hasVeracity, and hasWeather.
- *Individuals (instances)*. They are instances of classes (objects), e.g. a Boeing 747-800 is an individual (instance of the class “aircraft”).

The main AAO classes are:

- *Aircraft* (as a subclass of *Vehicle*): any type of aircraft falls into this category, including manned and unmanned fixed-wing or rotatory-wing air vehicles.
- *Route*: all the air corridors (as a collection of waypoints) for different airspace regions for aircraft falls into this category. They are defined by departure point to arrival point. However, no specification of waypoints is required for this first version of AAO.
- *Airport* (as a subclass of Aerodrome): all the aerodromes mostly for commercial air transport fall into this category. They are distinct from aviation airfields and military airbases.
- *Runway*: any runway from aerodromes falls into this category, runways have an identification code.
- *Status*: the class “Status” in the AAO is only defined to define the condition of runways.
- *Airspace*: any aerial region above a territory (portion of the atmosphere) controlled by a country.
- *Weather*: all weather conditions falls into this category.
- *Criteria*: the criteria defined in section 4.B is notated in this category. This class is from the URREF and it is actually the link between the AAO and the URREF. A subclass of *Criteria* is *Veracity*.
- *Radar*: all types of radar used in aeronautics falls into this category.
- *Metrics*: the metric assessment as defined in section 4.B falls into this category.

Table I presents AAO classes, examples of their instances, and some properties associated to them. Appendix A shows details of the hierarchical structure of the AAO.

DL operators are considered as different types of property restrictions in ontologies: quantifier restrictions such as existential and universal restrictions, hasValue restrictions (counting operators such as “less than or equal to” and “more than or equal to”), as well as cardinality restrictions such minimum and maximum cardinality restrictions. Also, complex classes can be created by means of simpler classes described based on logical operators like “*or*” and “*and*”.

Property restrictions along with classes and individuals are the building block to define axioms. Terminological axioms (usually based on operators such as inclusion, equivalence, etc.) are in the TBox, e.g., “Aircraft_A *subclass of* AircraftcannotLand *and* AircraftcanTakeoff”, and “ClearSky *subclass of* GoodWeather *and* VeryGoodWeather”. A set of assertional axioms (facts or assertions) are in ABox, e.g., “AircraftcanLand *equivalent to* Aircraft *and* (hasRoute *only* Landing)”, and “VeryBadWeather *equivalent to* Weather *and* (Tornado *or* microburst)”.

TABLE I.
Examples of the AAO Classes, Instances, and Properties

Class	Subclass	Instance	Property: *Class*	Property: *Instance*
Vehicle	Aircraft (Aircraft_x is a subclass of Aircraft)	e.g. B787 is an instance of Aircraft_x	hasRoute hasPilot hasPeople hasRadar hasSystem	hasWingspanValue
Route	Route_A	e.g. LAX-DWF is an instance of Route_A	hasAirspace hasTakeoff hasLanding hasNearbyAirport hasAirport	
Airport	Airport_II	e.g. LAX is an instance of Airport_II	hasRunway	
Runway	Runway_IA	e.g. 18L/36R is an instance of Runway_IA	hasStatus hasLanding hasTakeoff hasAirspace	
Status	Available Unavailble			
Airspace	Airspace_IV	e.g. USAAirspace airspace is an instance of Airpace_IV	hasWeather	hasSeparation
Weather	BadWeather	e.g. NewarkWeather		
Criteria	Veracity	e.g. Very low veracity	hasVeracity	
Radar	LWRS SWRS AWRS T-KFJK F-KNEL F-KEWR	e.g. JFKRadar	hasStatus	hasSensitivity
Metrics	WeatherAvoidance AircraftAvoidance AircraftManagment AircraftSeparation		hasWingspanValue	

The ABox and the TBox form the AAO knowledge base and are shown in Fig. 8. Details of the TBox and ABox axioms are shown in Appendix B.

Reasoners are the engine for the knowledge-based queries. They not only apply inference rules but also check semantic consistency on ontologies. These reasoning engines are able to deduce logical questions from axioms defined in ontologies. Fig. 9 shows the asserted classes of the AAO, including the added concepts (Criteria and Radar) and their relations.

Aircraft have radars (that detect them) which in turn have veracity for the information provided by them. Aircraft also have aairports and routes.

Fig. 10 shows the inferred classes of the AAO as result of executing the reasoner. This figure shows some example of AAO inferences as follows (from top to bottom). Airport I, II, and III are take-off and landing airports (aircraft can take off and land). Airspace I, II, and IV are flying airspace. Route C and D have landing. However, Route D has no take-off. Aircraft C and D can land in their corresponding airports. Bad weather includes storms and thunderstorms.

B. Semantic Uncertainty

Semantic uncertainty in the context of this paper is achieved by means of endowing the AAO with the URREF. Thus, the AAO is combined with the URREF ontology. The focus is on the input information coming from the ATM sensing systems (in particular, the land radars) which is taken into account through the URREF *InputCriteria* concept. The approach particularly targets *Veracity* (in sensed data) as one of the key concept from the URREF to establish the *Credibility* concept (URREF class).

The DSS provides ATC operators with ontological decision-making support based on the sensed data and

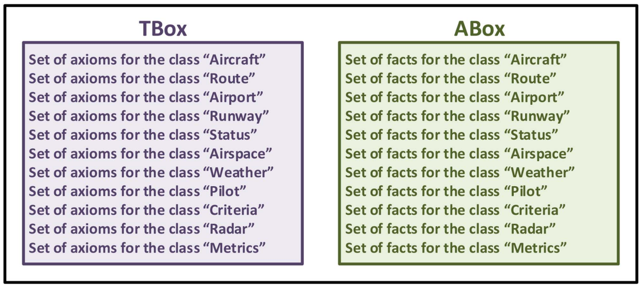

Fig. 8. AAO knowledge base: TBox and Abox (only new axioms and facts)

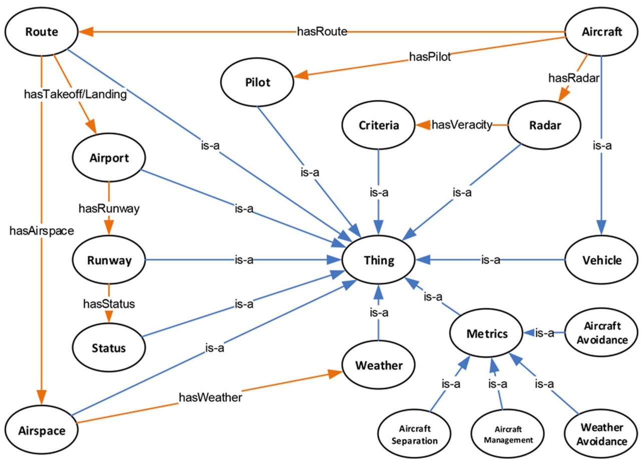

Fig. 9. Asserted AAO classes

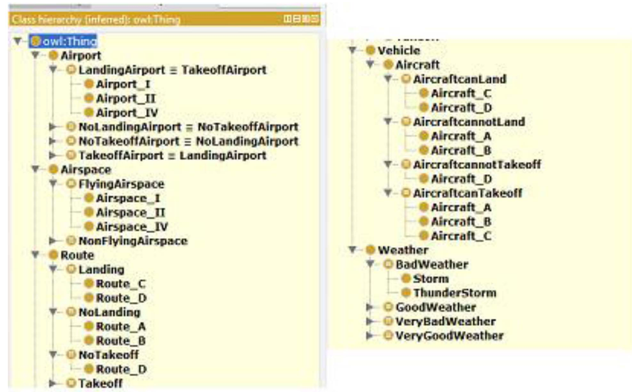

Fig. 10. Inferred AAO classes

processed information. The veracity of the input has an impact on decision outcomes, and it is the main driver for right decisions to be taken. Hence, veracity is important to be researched for valid analysis in the AAO. Thus, the URREF-endowed AAO is expected to improve DSS accuracy, and ultimately DSS effectiveness when decisions are taken by the AAO to support ATC operators.

Veracity metrics are based on the confusion matrix. This matrix is a "true" table that allows for definition and specification of true positives, false positives, true negatives, and false negatives when classifying possible outcomes from a process. Confusion matrices are useful for assessment of sensing systems, in particular for detection of objects/targets, e.g. radars detecting aircraft.

The above statistical classification approach is well known and used in other domains such as machine learning to analyse system accuracy by deferring and identifying elements. Hence, the confusion matrix approach of veracity assessment is attractive for predictive analytics and its statistical measures utilizing well-known attributes of: Sensitivity, Accuracy, Precision, Credibility, and Timeliness.

Typically, these statistical metrics include correlation and normalization for a probabilistic measure. Using probability theory affords Bayesian estimation, and filtering techniques.

The metric approach considered in this paper only focus on the source sensitivity (i.e., the sensor's which comes along with proximity range to the target) to estimate veracity. Thus, the veracity of the sensed data is estimated based on the sensor's sensitivity (*ObservationalSensitivity* subclass of the URREF *Credibility* class), and the range (distance between the sensor and the weather condition). A radar is the sensor in question in this paper. The spectrum defined for ObservationalSensitivity is as follows:

- 0–5%, Very low sensitivity.
- 5–25%, Low sensitivity.
- 25–70%, Regular sensitivity.
- 70–95%, High sensitivity.
- 95–100%, Very high sensitivity.

The radar's range is as follows:

- <50 Km, Very close range.
- 50–150 Km, Close range.
- 150–250 Km, Medium range.
- 250–400 Km, Far range.
- >400 Km, Very far range.

The above radar's sensitivities are combined with the radar's ranges as radars are located at different distances from what is sensed. This combination allows for the estimation of the veracity of the gathered information. This actually has an impact on the veracity metrics. The veracity metric is calculated as follows:

$$V_R = S_R \times R_R \tag{1}$$

Where V_R is the veracity, S_R is the sensitivity, and R_R is the range of the radar. This veracity is notated in the AAO in the URREF class *Veracity* by means of the following subclasses: *VeryLowVeracity*, *LowVeracity*, *RegularVeracity*, *HighVeracity*, and *VeryHighVeracity*. Likewise, the individual veracity for each radar is assigned to the property (object property) *hasVeracity*.

6. APPLICATION EXAMPLES

This section presents application examples of the approach proposed in this paper. They are based on realistic scenarios.

Fig. 11. Area of interest in the US airspace

TABLE II.
FLIGHT BA185 DETAILS

Dataset Record						
Time (UTC)	Latitude	Longitude	Altitude	Heading	Speed	Radar
17:01:59	40.7683	−74.5569	5675	177	299	F-KNEL1
17:03:56	40.6111	−74.543	5100	168	269	F-KNEL1
17:07:57	40.4859	−74.348	3075	62	183	F-KNEL1
17:10:09	40.5561	−74.251	2900	25	173	F-KNEL1
17:10:24	40.5667	−74.2442	2650	25	170	F-KNEL1
17:11:00	40.5925	−74.228	2075	26	161	F-KNEL1
17:11:27	40.6088	−74.2176	1700	25	135	F-KNEL1
17:11:54	40.6248	−74.2076	1375	25	138	F-KNEL1
17:12:40	40.651	−74.191	800	26	133	T-KJFK16
17:16:01	40.6991	−74.1669	0	275	30	F-KEWR1

A. Operation Context

The case study is meant to be as realistic as possible. It involves a dataset from a flight tracking service (Flightradar24 [4]). The dataset records all flights of aircraft with ADS-B transponders. It has 390,607 records generated between 17:00 and 18:00 UTC on 1st April 2017 (approx. 109 records streamed per second). Nevertheless, there is about a revisit rate of 30 second on every aircraft.

The airspace area of interest (Fig. 11) is that from the US airspace, entailing arrivals/departures from the north of Newark Liberty International Airport (code EWR). Three radar systems are considered for the case study: the F-KNEL1 radar from Lakehurst Maxfield Field Airport (code NEL), the T-KJFK16 radar from John F. Kennedy International airport (code JFK), and the F-KEWR1 radar from EWR.

Each of the above radars can cover the above area. However, they usually track aircraft depending on how far aircraft are from the radars and what is the destination of the flights. F-KNEL1 belongs to a military airfield in New Jersey and usually tracks aircraft approaching from or departing to the US east coast. T-KJFK16 tracks landed or arriving/departing flights in JFK. F-KEWR1 tracks landed or arriving/departing flights in EWR. Additionally, EWR has weather updates (from weather forecast and radars) to assist aircraft when lading or departing.

The case study considers Flight BA185, a British Airways flight from London Heathrow (code LHR) to EWR, that is planned to land in EWR. The FAA defines airplane design groups according to aircraft wingspans. The BA185 airplane is a Boeing 777-200, which belongs to group V (52–65 m of wingspan). Table II presents the details for Flight BA185 obtained from the Automatic Dependent Surveillance-Broadcast (ADS-B) dataset.

Two airspace situations are considered when Flight BA185 is approaching EWR: Scenario 1 entails weather conditions ahead of Flight BA185, and Scenario 2 entails potential collision of Flight BA185 with UAVs.

B. Scenario 1: Aircraft in Weather

The first scenario considers Flight BA185 approaching EWR for landing (17:01:59–17:10:09 UTC in Table I). The airplane has descended (altitude 5675 feet) down to 2900 feet in such a period. Flight BA185 took off from LHR and is scheduled to land in EWR. Weather conditions are assumed to be deteriorated in the north of the US east coast. However, the weather is good for landing in EWR.

The information provided by the weather forecast from Satellite Weather Radar Systems (SWRSs) and Land Weather Radar Systems (LWRSs) are considered accurate and true. They have a high sensitivity although the later are considered to have slightly lower sensitivity than the former. Additionally, commercial airplanes are equipped with an Airborne Weather Radar systems (AWRSs) (located in the aircraft nose) which allows for detection of the intensity of convective weather conditions such as massive hails, powerful lighting, and excessive precipitation (strong downdraft), e.g. microbursts. This alternative weather radar source is considered to have the highest sensitivity (of the three weather radars) when the weather in question comes from the area ahead the airplane. Thus, it is used as a very credible reference for the calculation of veracity metrics and the sensitivities for weather forecast are assumed as follows:

- AWRS sensitivity ($S_{AWRS} = 0.99$)
- SWRS sensitivity ($S_{SWRS} = 0.75$)
- LWRS sensitivity ($S_{LWRS} = 0.55$)

The above weather radar sensitivities are combined with the radar ranges as radars are located at different distances from the weather condition. Thus, veracity metric is calculated as follows:

$$V_{xWRS} = S_{xWRS} \times R_{Xwrs} \quad (2)$$

Where V_{xWRS} is the veracity, S_{Xwrs} is the sensitivity, and R_{xWRS} is the range of the type of radar x (A: airborne, S: satellite, and L: land).

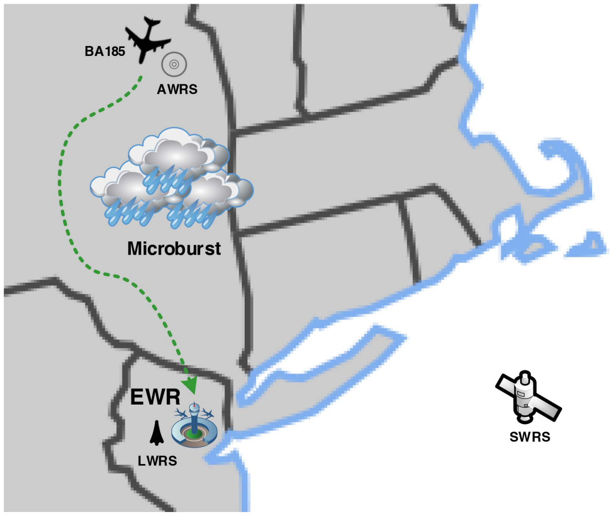

Fig. 12. Scenario 1 in the airspace area of interest

Absolute weights (100) for radar ranges are considered in scenario 1 depending on their distance to the weather condition. The following weights are assumed for AWRS range (the closest to the weather condition).

R_{AWRS} = {Very close = 100, Close = 0, Medium = 0, Far = 0, Very Far = 0}

Then,

$$V_{AWRS} = S_{AWRS} \times R_{AWRS} = 0.99 \times 100 = 99\%$$

Therefore, for AWRS 99% (True) and 1% (false).

The following weights are assumed for SWRS range (the closest to the weather condition).

R_{SWRS} = {Very close = 0, Close = 100, Medium = 0, Far = 0, Very Far = 0}

Then,

$$V_{SWRS} = S_{SWRS} \times R_{SWRS} = 0.75 \times 100 = 75\%$$

Therefore, for SWRS 75% (True) and 25% (false).

The following weights are assumed for LWRS range (the closest to the weather condition).

R_{LWRS} = {Very close = 0, Close = 0, Medium = 100, Far = 0, Very Far = 0}

Then,

$$V_{LWRS} = S_{LWRS} \times R_{LWRS} = 0.55 \times 100 = 55\%$$

Therefore, for LWRS 55% (True) and 45% (false).

Scenario 1 also supposes Flight BA185 and the ATC in EWR are concerned about the weather condition (microburst) when approaching the EWR airport from the northeast. Fig. 12 shows the above scenario 1.

The information provided by the AAO can be visualized by ATCs to support their decisions on the above situation (also, aviators and pilots of remotely-piloted aircraft could make use of this information). They can run AAO queries as to the weather condition in proximity (ahead) of Flight BA185 airway. This also provides suggestions about what to do with Flight BA185 to avoid any potential risk that jeopardize the flight safety. The query is regarding possibilities for an airplane (Flight BA185 in this case) to encounter adverse weather conditions that make aircraft change their route. The rerouting possibilities are:

- Very low chances of re-routing (0–19%),
- Low chances of re-routing (20–39%),
- Medium chances of re-routing (40–59%),
- High chances of re-routing (60–79%), and
- Very high chances of re-routing (80–100%).

The above rerouting possibilities are directly related with the radar veracities as calculated for V_{xWRS}. Thus, V_{AWRS} means a very high chance of re-routing, V_{SWRS} means a high chance of re-routing, and V_{LWRS} means a medium chance of re-routing if a microburst is detected by the above radars.

Fig. 13 shows AAO query results including veracity metrics (top of the figure) for scenario 1 along with AAO queries for each of the radars that detects the weather condition (bottom of the figure). The weather information is provided by three weather radars: AWRS (onboard the Boeing 777-200, i.e. Flight BA185), SWRS (weather forecast), and LWRS (from EWR). AWRS is the most veracious radar (S_{AWRS} = 0.99 and R_{AWRS} = 100; very close)) for this weather condition (NewarkWather_1) since such a radar is closely placed near the weather situation. SWRS is less sensitive and is further (from the weather condition) than AWRS (S_{SWRS} = 0.75 and R_{SWRS} = 100; close), and LWRS is the least sensitive and the furthest one (from the weather condition) of the three weather radars (S_{LWRS} = 0.55 and R_{LWRS} = 100; Medium). Therefore, the veracity of the query is 100% when the weather information is from AWRS, the veracity of the query is 75% when the weather information is from SWRS, and the veracity of the query is 55% when the weather information is from LWRS.

The query inference results (from Fig. 14) suggest that (from left to right):

1. Flight BA185 must slightly change route (to avoid weather condition; NewarkWather_1) on its way to EWR for landing (very high chance of rerouting). The veracity of this query is based on a veracity of 99%, i.e., when AWRS detects the microburst weather condition ahead of Flight BA185. This suggestion is the most veracious out of the three suggestions.
2. Flight BA185 should slightly change route on its way to EWR for landing (high chance of rerouting). The veracity of this query is based on a veracity of 75% when SWRS detects the microburst weather condition ahead of Flight BA185. This suggestion is less veracious than suggestion 1.
3. Flight BA185 could slightly change route on its way to EWR for landing (high chance of rerouting). The veracity of this query is based on a veracity of 55% when LWRS detects the microburst weather

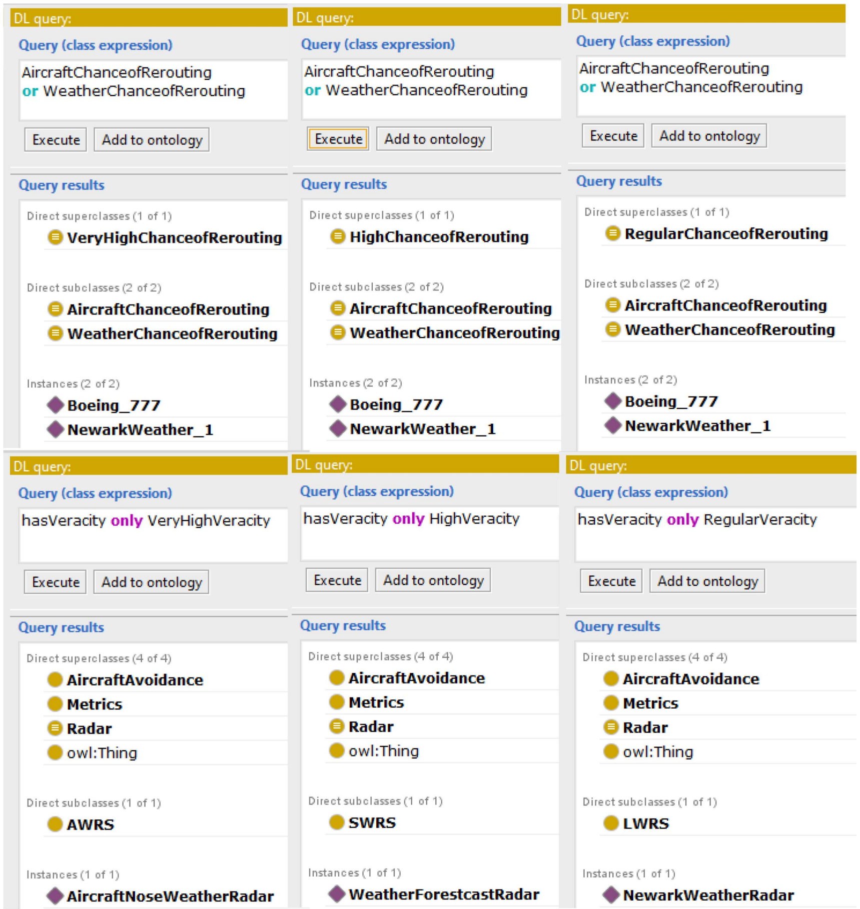

Fig. 13. Querying results to assess situation in scenario 1

condition ahead of Flight BA185. This suggestion is most less veracious of the three suggestions.

Fig. 14 shows the reasoning path (in green colour) followed by the reasoner to determine the AAO query results (“aircraft chance of Rerouting” and (“weather chance of Rerouting”).

C. Scenario 2: UAVs in AirSpace

The second scenario considers the approach of Flight BA185 to EWR for landing as in scenario 1, although a different time segment is considered (17:11:27–17:16:01 UTC in Table I). The airplane has descended (altitude 1700 feet) down to 0 feet in such a period. The weather condition (microburst) from scenario 1 has been left behind. However, Flight BA185 is supposed to face a new challenge (before landing in EWR) which is airspace invasion due to three Unmanned Air Vehicles (UAVs) flying nearby EWR.

The three drones are: a small UAV (sUAV) that is a small-unmanned quadcopter which wheelbase is 0.5 m, a medium UAV (mUAV) that is a medium-unmanned airplane with 1.5 m of wingspan, and a huge UAV (hUAV) that is a large-unmanned airplane which wingspan is 20 m. The UAVs are flying at different altitudes and locations around the EWR airport during the landing of Flight BA185. These UAVs fly high enough to dangerously come close to Flight BA185 while descending from 1700 down to 0 feet in about four and a half minutes. The sUAV has a non-contactable remote pilot, and it is less than 300 m away from Flight BA185. The mUAV is more than 1100 m away from Flight BA185. The hUAV is less than 900 m away from Flight BA185. The mUAV and the hUAV have contactable remote pilots.

The information provided by the dataset for F-KNEL1, T-KJFK16, and F-KEWR1 radars (as specified in the ADS-B) dataset) is considered fully accurate and true since they come from real measurements. These radars are used as a very credible reference for the calculation of veracity metrics. Hence, the sensitive of the above radars is 0.99 when they manage to track the aircraft of interest. The remaining radars (that do not track the aircraft in the dataset) are considered to have smaller sensitivities. This makes sense since they do not track the above aircraft. Such a sensitivity difference along with the range of the radar has an impact on the veracity metrics.

The sensitivities for aircraft detection are assumed as follows (from 17:01:59 to 17:11:54 where F-KNEL1 tracks Flight BA185):

- F-KNEL1 sensitivity ($S_{\text{F-KNEL1}} = 0.99$)
- T-KJFK16 sensitivity ($S_{\text{T-KJFK16}} = 0.80$)

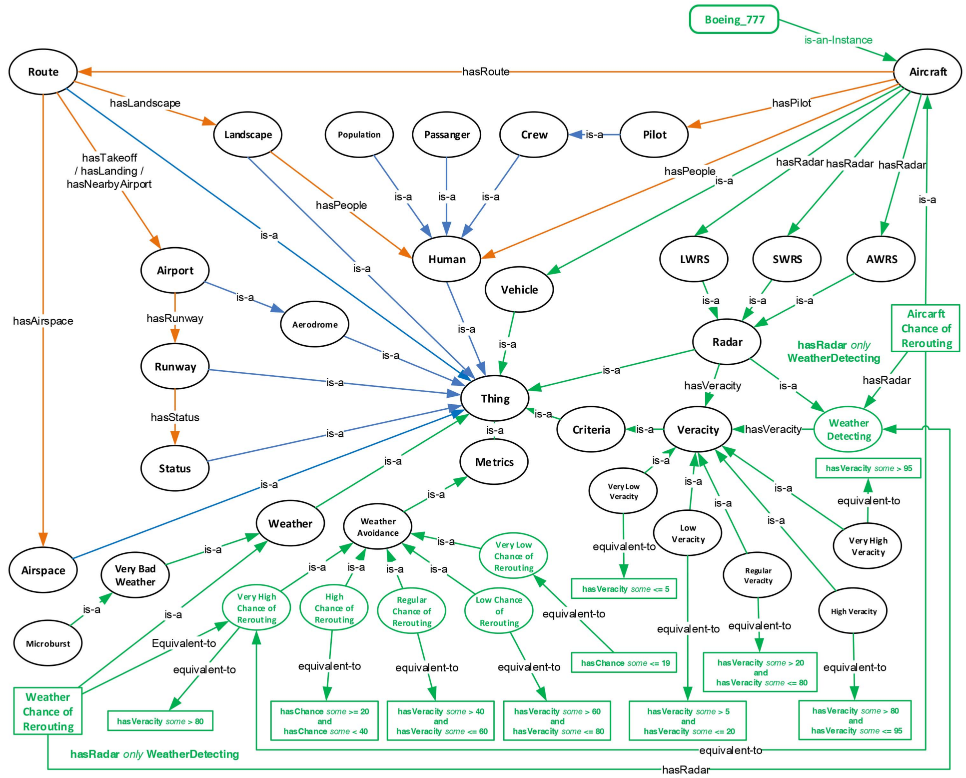

Fig. 14. Reasoning behind the queries for airspace situation in scenario 1

- F-KEWR1 sensitivity ($S_{\text{F-KEWR1}} = 0.60$)

The following weights are assumed for radar ranges:
$R_{\text{F-KNEL1}}$ = {Very close = 100, Close = 0, Medium = 0, Far = 0, Very Far = 0}
$R_{\text{T-KJFK16}}$ = {Very close = 0, Close = 100, Medium = 0, Far = 0, Very Far = 0}
$R_{\text{F-KNWR1}}$ = {Very close = 0, Close = 0, Medium = 100, Far = 0, Very Far = 0}

The calculation of the veracity metric is based on equation (1), similar to the calculation in scenario 1.

The above three radars track Flight BA185 in the period considered by the case study (Table I). They have $S_{\text{BA185}} = 0.99$ (when they track Flight BA185) so they are a fully-truthful source for both radars. However, the tracking of the UAVs (i.e., sUAV, mUAV, and hUAV) is assumed to be done by any of the above radars that have difference veracities (V_{sUAV}, V_{mUAV}, and V_{hUAV}). The combination of two or more veracities given by the multiplication of the veracities. Table III shows examples of the impact of having different veracities when detecting aircraft based on the radar used for detection.

Fig. 15 shows the above scenario 2.

The information provided by the AAO can be visualized by ATCs to support their decisions on the above situation (also, aviators and pilots of remotely-piloted aircraft could make use of this information). They can run AAO queries as to the impact of the proximity of the UAVs on Flight BA185. This also provides suggestions about what to do with Flight BA185 or the

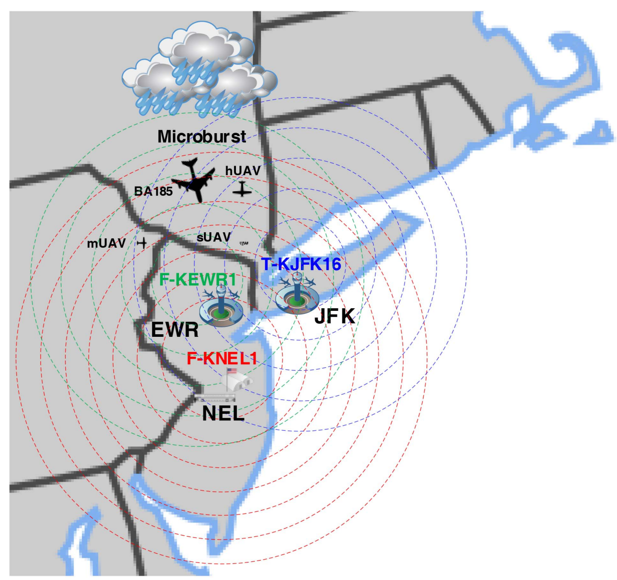

Fig. 15. Scenario 1 in the airspace area of interest

UAVs to avoid any potential air collision. The query is regarding chances for air collision: very low risk of collision (0–19%), low risk of collision (20–39%), medium risk of collision (40–59%), high risk of collision (60–79%), and very high risk of collision (80–100%). These collision possibilities are directly related with the radar veracities as calculated for V_{BA185} and V_{xUAV}.

Scenario 2 considers six different airspace situations (veracities are taken from Table III):

1. F-KNEL1 detects the three UAVs, and Flight BA185.

TABLE III.
VERACITY METRICS (TRUES IN %)

S_{BA185}	S_{xUAV} (solo)			S_{sUAV}, S_{mUAV} & S_{hUAV} (all)					
	KNEL1	KJFK16	KEWR1	KNEL1 & KNEL1 & KNEL1	KJFK16 & KJFK16 & KJFK16	KEWR1 & KEWR1 & KEWR1	KNEL1 & KJFK16 & KJFK16	KNEL1 & KEWR1 & KJFK16	KNEL1 & KEWR1 & KEWR1
F-KNEL1	98	79.2	59.4	96.06	51.2	21.38	62.73	47.05	35.28
T-KJFK16	79.2	64	48	80	40.96	17.28	50.69	38.02	28.52
F-KEWR1	59.4	48	36	60	30.72	12.96	38.02	28.52	21.38

The veracity of this query is 98% ($V_{xUAV} = 99 * V_{BA185} = 99$). The most veracious radar from the ADSB dataset.

2. F-KJFK16 detects sUAV, mUAV and hUAV, and KNEL1 detects Flight BA185. The veracity of this query is 51.2% ($V_{xUAV} = 51.72 * V_{BA185} = 99$).
3. F-KEWR1 detects sUAV, mUAV and hUAV, and KNEL1 detects Flight BA185. The veracity of this query is 21.38% ($V_{xUAV} = 21.6 * V_{BA185} = 99$).
4. F-KNEL1 detects the sUAV and Flight BA185, F-KJFK16 detects the mUAV and the hUAV. The veracity of this query is 62.73% (S_{sUAV} and $S_{BA185} = 98 * S_{mUAV}$ and $S_{hUAV} = 64$).
5. F-KNEL1 detects the sUAV and Flight BA185, F-KEWR1 detects the mUAV, and F-KJFK16 detects the hUAV. The veracity of this query is 47.05% (S_{sUAV} and $S_{BA185} = 98 * S_{mUAV} = 60 * S_{hUAV} = 80$).
6. F-KNEL1 detects the sUAV and Flight BA185, F-KEWR1 detects the mUAV and the hUAV. The veracity of this query is 35.28% ($S_{sUAV} = S_{BA185} = 98 * S_{mUAV}$ and $S_{hUAV} = 36$).

Fig. 16 shows the inferred *AircraftChanceofCollision* class (top) and AAO query results (bottom) for each of the radars that detects the UAVs for airspace situation 1, 2, and 3, including veracity metrics for scenario 2. Fig. 17 shows AAO query results for airspace situation 4, 5, and 6.

The query inference results suggest that (from left to right):

1. sUAV and hUAV have clear chances of collision (very high risk of collision) and the veracity of this query is based on a sensitivity of 99% (and proximity of the radar to the aircraft) when F-KNEL1 detects the UAVs and Flight BA185 (top-left query in Fig. 16). Detection makes by means of NELRadar (F-KNEL1), bottom-left query in Fig. 16. These inference and query suggestion are the most veracious out of the six suggestions. Actually, the real one.
2. sUAV and hUAV have some chances of collision (medium risk of collision) and the veracity of this query is based on a sensitivity of 51.2% (and proximity of the radar to the aircraft). This veracity is not high enough to make a trusted decision when F-KJFK16 detects the UAVs and F-KNEL1 detects Flight BA185 (top-center query in Fig. 13). Detection makes by means of JFKRadar (T-KFJK16), bottom-center query in Fig. 16.
3. Bottom-right query in Fig. 16: sUAV and hUAV have low chances of collision (low risk of collision) and the veracity of this query is based on a sensitivity of 21.6% (and proximity of the radar to the aircraft). This veracity is very low to make a trusted decision when F-KEWR1 detects the UAVs and F-KNEL1 detects Flight BA185 (top-right query in Fig. 16). Detection makes by means of EWRRadar (F-KEWR1), bottom-right query in Fig. 16.
4. sUAV and hUAV have some chances of collision (medium risk of collision) and the veracity of this query is based on a sensitivity of 64% (and proximity of the radar to the aircraft). This veracity is not high enough to make a trusted decision when F-KNEL1 detects the sUAV and Flight BA185, F-KEWR1 detects the mUAV, and F-KJFK16 detects the hUAV. Query on the left of Fig. 17.
5. sUAV and hUAV have some chances of collision (medium risk of collision) and the veracity of this query is based on a sensitivity of 48% (and proximity of the radar to the aircraft). This veracity is not high enough to make a trusted decision when F-KNEL1 detects the sUAV and Flight BA185, F-KEWR1 detects the mUAV, and F-KJFK16 detects the hUAV. Query on the left of Fig. 17.
6. sUAV and hUAV have some chances of collision (medium risk of collision) and the veracity of this query is based on a sensitivity of 51.2% (and proximity of the radar to the aircraft). This veracity is not high enough to make a trusted decision when F-NEL1 detects the sUAV and Flight BA185, F-KEWR1 detects the mUAV and the hUAV. Query on the left of Fig. 17.

Query on the right of Fig. 17 suggests the mUAV has no risk of collision. The inference and query results make sense since the chance of collision is diminished as the veracity of the radars is decreased. However, the real chance of collision is very high (the one suggested in 1.), and it is actually the most veracious.

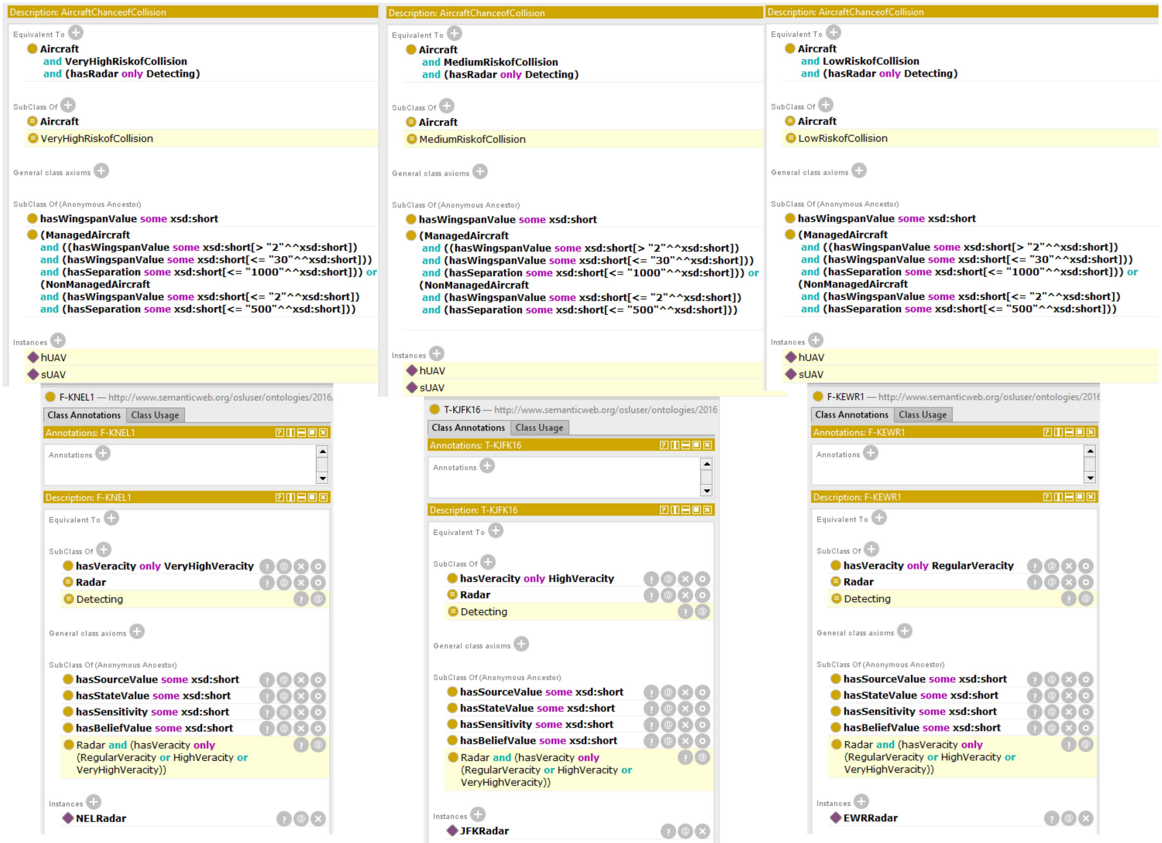

Fig. 16. Querying results and inferred classes to assess situation 1, 2, and 3.

D. Bayes' Risk Assessment

The assessment as to determine whether to alert the pilot is based on the information fusion analysis of Bayes' risk. The Bayesian estimate a posterior is the measurement given a possible collision.

$$P(\theta_j \mid x) = \frac{P(x \mid \theta_j)P(\theta_j)}{P(x)} \tag{1}$$

where $P(\theta_j)$ is the prior sensitivity of the radar configurations for each case $j = 1, \ldots 6$ (as those shown from left to right in Table III for multiple detection of all the UAVs, i.e., S_{sUAV}, S_{mUAV} & S_{hUAV}), and the conditional likelihood $P(x \mid \theta_j)$ is for a collision or no-collision given the radar measurements. To determine whether to send a semantic alert a pilot is based on the measurement, the potential range (distance), and type of the UAV. To determine the Bayes' risk, a loss function L was developed if no action (e.g., send an alert) was taken.

$$R(\alpha_j \mid x) = \sum_{j=1}^{6} L(\alpha_j \mid \theta_j)P(\theta_j \mid x) \tag{2}$$

where $L(\alpha_j \mid \theta_j)$ represents the three cases for loss if the range is j = {close, near, far}. The results were normalized. Given scenario 2, if there is a chance of collision, the best action is to alert the pilot. If there is lower chance of collision, the results still suggest sending a warning to the pilot of a potential collision.

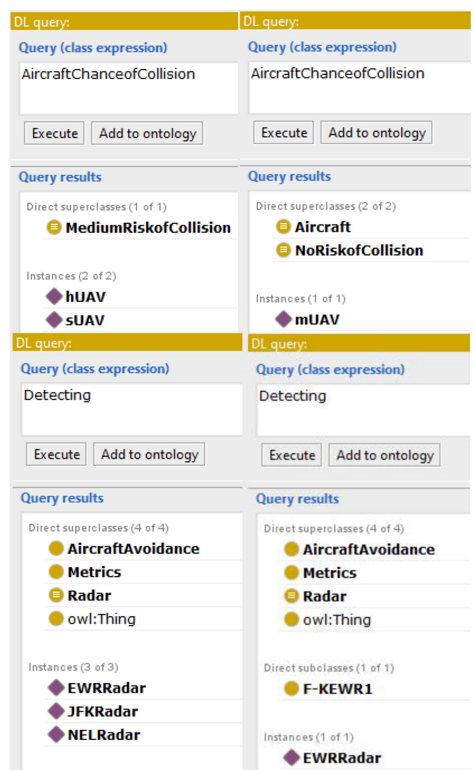

Fig. 17. Querying results to assess situation 4, 5, and 6 in scenario 2.

TABLE IV.
LOSS FUNCTION

	Case 1 High	Case 2 Medium	Case 3 Low	Case 4 Medium	Case 5 Medium	Case 6 Medium
Close	0	1	8	0	1	1
Near	0	2	1	3	3	3
Far	10	7	1	7	6	6

TABLE V.
BAYES' RISK

	Collision	No collision
Close	0.494548297	1.579050007
Near	1.416107104	2.413458212
Far	8.089344598	5.825970614
Sum	10	9.818479

From the first row of Table III {0.9606, 0.512, 0.2138, 0.273, 0.4705, 0.3528}, they total up 3.1364. Then, the prior probabilities (sensitivity/veracity) for the six cases are (by dividing each of them by the total): $P(\theta_j) = \{0.306083, 0.163244, 0.068167, 0.200006, 0.150013, 0.112486\}$.

The likelihood $P(x \mid \theta_j)$ for collision are {0.9606, 0.512, 0.2138, 0.273, 0.4705, 0.3528}, and for no collision is {0.0394, 0.488, 0.7862, 0.3727, 0.5295, 0.6472}.

The prior probabilities (Bayes denominator) $P(x) = P(\theta_j)$. $P(x \mid \theta_j)$ for collisions are {0.2938, 0.08358, 0.01457, 0.1255, 0.0706, 0.03968} which total up 0.627725, and for no collision is {0.0122, 0.07966, 0.0536, 0.0745, 0.0794, 0.0728} which total up 0.372275.

The posterior probabilities $P(\theta_j \mid x) = P(x)/0.627725$ for collisions are {0.4681, 0.1331, 0.0232, 0.1999, 0.1124, 0.0632}, and $P(\theta_j \mid x) = P(x)/0.372275$ for no collisions are {0.0329, 0.2140, 0.1440, 0.2002, 0.2134, 0.1955}.

The loss function $L(\alpha_j \mid \theta_j)$ is defined in Table IV.

The Bayes' risk $R(\alpha_j \mid x)$ (which formula is (2)) for the three cases for loss are shown in Table V.

Presenting the information in a semantically meaningful way by normalizing them based on the sum 10 and 9.818479 for collision and no collision, the Bayes' risk was inverted so as to represent the results as shown in Fig. 18.

For case collision assessments, the results are used as (1 is high, 2, 4, 5, 6 is medium, and 3 is low). The Bayes risk assessment is consistent with the ontology from which $> 0.95\%$ would be a collision confirmed; 0.95–0.85 for collision likely, and < 0.85 for collision possible. For values < 0.5, it is unlikely there would be a collision. From Fig. 18, when a collision is detected within a close range, the best action (reduce risk) is to confirm an alert. Likewise, when the UAV is near, a collision is likely, so a warning should be sent. If a collision range is detected far away, the normalized action is that there is enough time for future measurements to

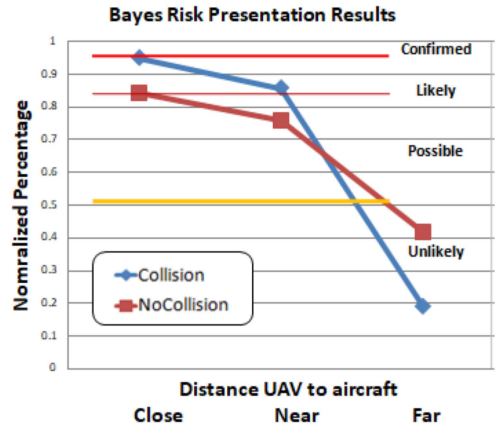

Fig. 18. Bayes' risk assessment results.

determine if a collision would result. On the other hand, the case of a no-collision also presents a semantically interesting result, as if the radars are sensitive and detect a UAV in close proximity to the pilot, a likely warning would result.

7. CONCLUSION AND FUTURE WORK

The paper proposed an Avionics Analytics Ontology (AAO) based on the Uncertainty Representation and Reasoning Evaluation Framework (URREF). The AAO is developed to provide situation awareness updates for aviators, air traffic controllers, and airport security personnel in support of ATM/UTM decision-making processes. The congestion of the airspace with UAVs was presented as use cases to demonstrate the workload reduction through an information fusion ontology methodology. Veracity was the measured degree of uncertainty to support credible reporting and airspace collisions. Examples involving two ATM/UTM operation scenarios where F-KNEL1, T-KJFK16, and F-KEWR1 radars (as specified in the ADS-B) determine the commercial aircraft (Flight BA185) collision analysis from a set of UAVs. The AAO results present a useful approach towards providing an integration method among uncertainties including semantic from operators, sensing from navigation, and situation from weather modeling updates.

Future research work will involve methods to improve veracity metrics. One of the relevant approach as an interesting veracity metric to be considered for further investigation is the big data veracity index [57]. It is based on three main dimensions to define veracity: objectivity (subjectivity), truthful (deception), and credibility (implausibility). The index approach deserves attention, but some research is required to deal with artificial autonomy (DSS) since the potential tools to support such a metric index are too human-oriented. The challenge is to develop supporting tools that allow for machine veracity metrics, e.g. radars. On the other hand, some future refinement on the integration of veracity (uncertainty) into the AAO will enhance usefulness. One of the inspiring methodologies (to deal with

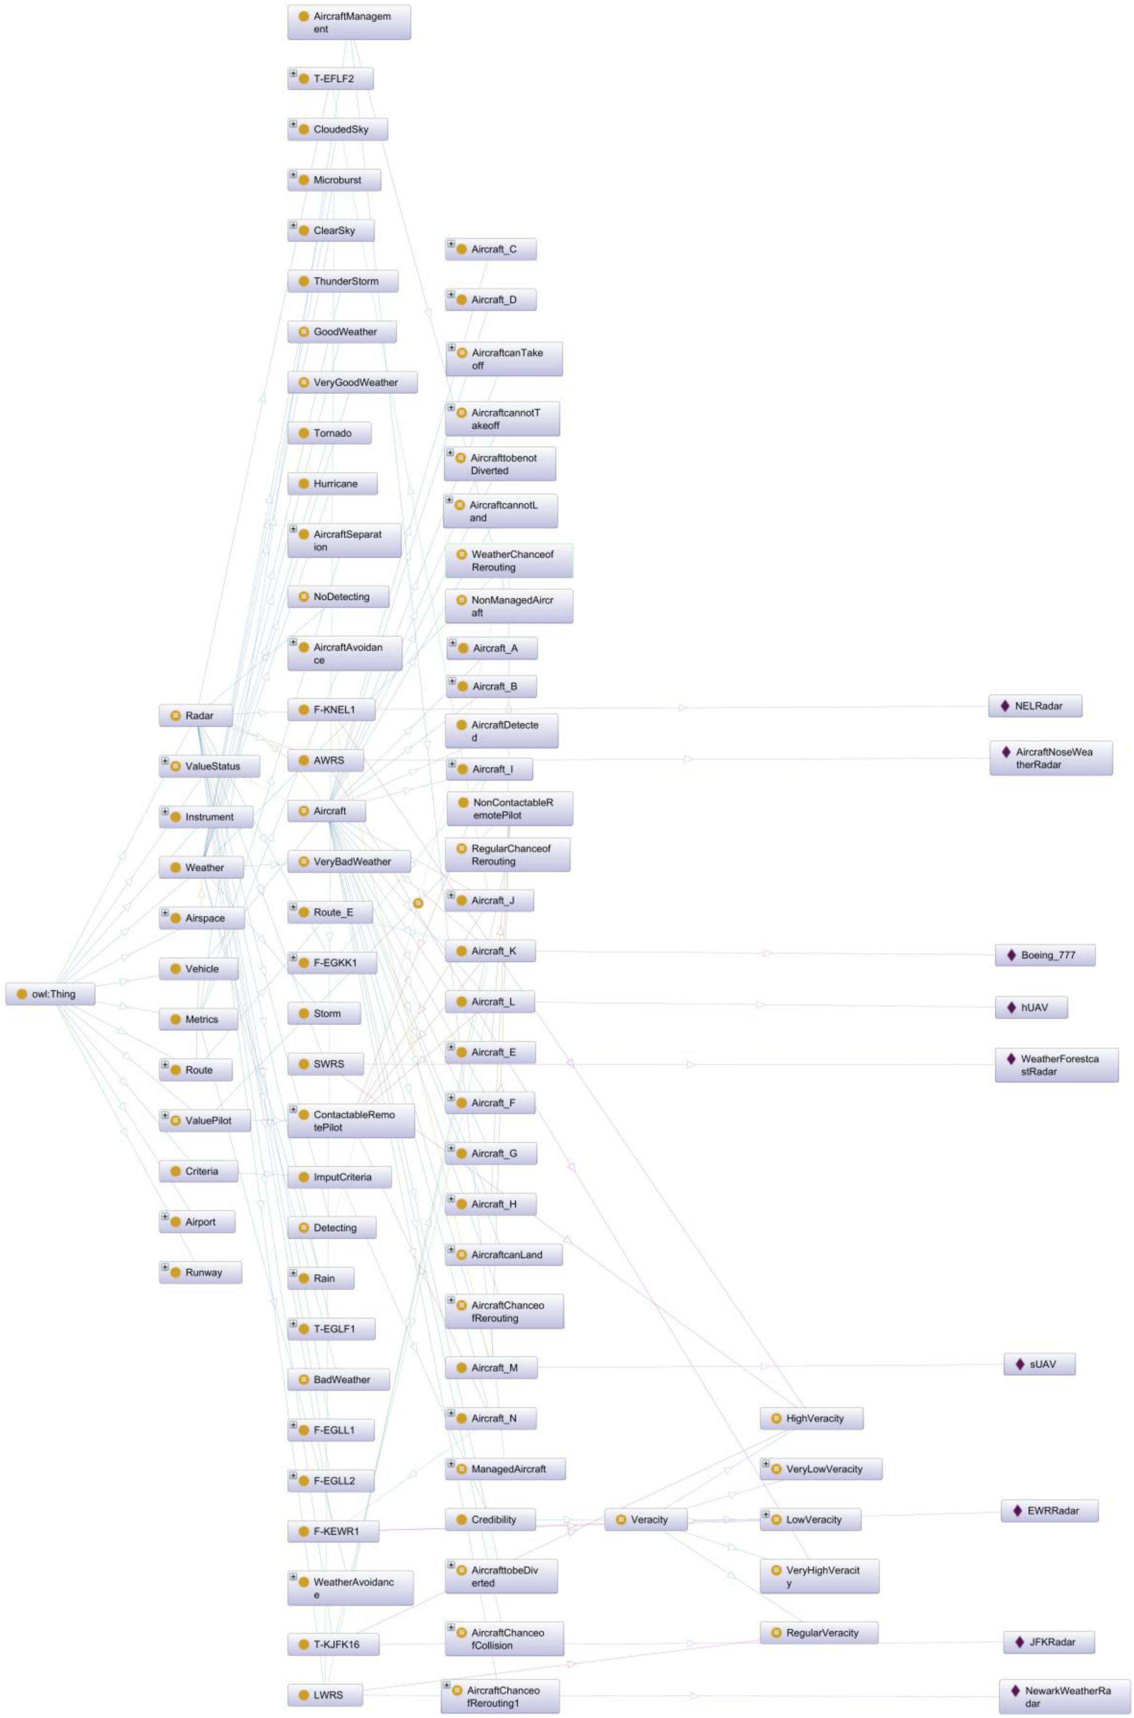

Fig. 19. Structure of hierarchy of the AAO

probabilistic uncertainty when making decision) is the Bayesian networks, e.g. BasesOWL [58] which is suitable for ontologies.

Future methods would also include physics-based and human-derived (PHIF) graphical information fusion methods where graph ontologies can be matched, associated, and extended for narratives [59, 60]. The application scenarios discussed in this paper are meant to easily demonstrate the benefits of the AAO-based DSS proposed. They are simple but realistic. However, further development of the AAO will consider demonstrations involving and targeting ATM operational performance indexes as those discussed by Civil Air Navigation Services Organisation (CANSO) [61] and SESAR Key Performance AREA [62]. For example, capacity and efficiency are listed are listed as operational metrics; while less defined metrics of societal metrics include safety, security, and environmental sustainability [63].

APPENDICES

APPENDIX A: AAO HIERARCHICAL STRUCTURE

Fig. 19 shows the structural hierarchy of the classes in the AAO.

APPENDIX B: TBOX AND ABOX

The axioms of the AAO TBox are shown below.

Aircraft (subclass of Vehicle)
Aircraft_K subclass of AircraftcannotLand and AircraftcanTakeoff
Aircraft_L subclass of AircraftcannotLand and AircraftcanTakeoff
Aircraft_M subclass of AircraftcanLand and AircraftcanTakeoff
Aircraft_N subclass of AircraftcannotLand and AircraftcannotTakeoff
AircraftcanLand subclass of Aircraft
AircraftChanceofCollision subclass of Aircraft
AircraftChanceofRerouting subclass of Aircraft

Route
Route_A subclass of Landing and Takeoff
Route_B subclass of NoLanding and Takeoff
Route_C subclass of Landing and Takeoff
Route_D subclass of Landing and NoTakeoff

Airport
Airport_I subclass of LandingAirport and TakeoffAirport
Airport_II subclass of LandingAirport and TakeoffAirport
Airport_III subclass of NoLandingAirport and NoTakeoffAirport
Airport_IV subclass of LandingAirport and TakeoffAirport

Airspace
Airspace_I subclass of FlyingAirspace
Airspace_II subclass of FlyingAirspace
Airspace_III subclass of NoFlyingAirspace
Airspace_IV subclass of FlyingAirspace

Weather
ClearSky subclass of GoodWeather and VeryGoodWeather
CloudedSky subclass of VeryBadWeather
Hurricane subclass of VeryBadWeather
Rain subclass of GoodWeather
Storm subclass of BadWeather
Thunderstorm subclass of BadWeather
Tornado subclass of VeryBadWeather
Microburst subclass of VeryBadWeather

Metrics
AircraftAvoidance subclass of Metrics
AircraftManagment of Metrics
AircraftSeparation subclass of Metrics
WeatherAvoidance subclass of Metrics

Criteria
ImputCriteria subclass of Criteria
Credibility subclass of ImputCriteria
Veracity subclass of Credibility
VeryLowVeracity subclass of Veracity
LowVeracity subclass of Veracity
RegularVeracity subclass of Veracity
HighVeracity subclass of Veracity
VeryHighVeracity subclass of Veracity

Radar
AWRS subclass of hasVeracity only VeryHighVeracity and Radar
SWRS subclass of hasVeracity only HighVeracity and Radar
LWRS subclass of hasVeracity only RegularVeracity and Radar
F-KNEL1 subclass of hasVeracity only VeryHighVeracity and Radar
F-KJFK16 subclass of hasVeracity only HighVeracity and Radar
F-KNEW1 subclass of hasVeracity only RegularVeracity and Radar

The facts of the AAO ABox are shown below.

Aircraft
AircraftChanceofCollision equivalent to Aircraft and (RiskofCollision and (hasRadar only Detecting)
AircraftChanceofRerouting equivalent to Aircraft and (ChanceofRerouting and (hasRadar only Detecting)
AircraftcanLand equivalent to Aircraft and (hasRoute only Landing)
AircraftcannotLand equivalent to Aircraft and (hasRoute only NoLanding)
AircraftcanTakeoff equivalent to Aircraft and (hasRoute only Takeoff)
AircraftcannotTakeoff equivalent to Aircraft and (hasRoute only NoTakeoff)

Route
Landing equivalent to Route and (hasLanding only LandingAirport)
NoLanding equivalent to Route and (hasLanding only NoLandingAirport)
Takeoff equivalent to Route and (hasTakeoff only TakeoffAirport)
NoTakeoff equivalent to Route and (hasTakeoff only NoTakeoffAirport)

Airport
LandingAirport equivalent to Airport and (has Airspace only FlyingAirspace)
NonLandingAirport equivalent to Airport and (has Airspace only NonFlyingAirspace)
TakingoffAirport equivalent to Airport and (has Airspace only FlyingAirspace)
NonTakingoffAirport equivalent to Airport and (has Airspace only NonFlyingAirspace)

Airspace
FlyingAirspace equivalent to Airspace and (not (NonFlyingAirspace))
NonFlyingAirspace equivalent to Weather and (hasWeather only VeryBadWeather)

Weather
VeryGoodWeather equivalent to Weather and (ClearSky or CloudedSky)
GoodWeather equivalent to Weather and (CloudedSky or Rain)
BadWeather equivalent to Weather and (Storm or ThuderStorm)
VeryBadWeather equivalent to Weather and (Hurrican or Tornado)

Metrics
RiskofCollision equivalent to (ManagedAircraft and ((hasWingspanValue some xsd:short[> ″2″^^xsd:short]) and (hasWingspanValue some xsd:short[<= ″30″^^xsd:short])) and (hasSeparation some xsd:short[<= ″1000″^^xsd:short])) or (NonManagedAircraft and (hasSeparation some xsd:short[<= ″500″^^xsd:short]) and (hasWingspanValue some xsd:short[<= ″2″^^xsd:short]))

Criteria
VeryLowVeracity equivalent to hasVeracity some xsd:short[>= ″5″^^xsd:short]
LowVeracity equivalent to (hasVeracity some xsd:short[> ″5″^^xsd:short]) and (hasVeracity some xsd:short[< ″25″^^xsd:short])
RegularVeracity equivalent to (hasVeracity some xsd:short[> ″25″^^xsd:short]) and (hasVeracity some xsd:short[< ″70″^^xsd:short]))
HighVeracity equivalent to (hasVeracity some xsd:short[> ″70″^^xsd:short]) and (hasVeracity some xsd:short[< ″95″^^xsd:short]))
VeryHighVeracity equivanlent to hasVeracity some xsd:short[> ″95″^^xsd:short]

Radar
Detecting equivalent to Radar and (hasVeracity only VeryHighVeracity)
NoDetecting equivalent to not(Detecting)

DECLARATION

The Avionics Analytics Ontology (AAO) used in this paper has been developed for specific airspace situations. It is based on intuitive knowledge gathered from an investigation done on trusted sources such FAA regulations. The AAO is at its early development stage (prototype) and it is a living approach as it is been continuously updated. It has not been validated yet. However, there is a plan to integrate the NASA ontology into the AAO, which will require validation for further development.

ACKNOWLEDGMENT

Erik Blasch was supported under an AFOSR grant in Dynamic Data-Driven Applications Systems. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the United States Air Force.

REFERENCES

[1] C. C. Insaurralde, and E. Blasch
"Ontological Knowledge Representation for Avionics Decision-Making Support,"
35th Digital Avionics Systems Conference (DASC), Sacramento, CA, USA, Sep 2016.

[2] S. Imai, E. Blasch, A. Galli, W. Zhu, F. Lee, and C. A. Varela
"Airplane Flight Safety Using Error-Tolerant Data Stream Processing,"
IEEE Aerospace and Electronics Systems Magazine, Issue 32, No. 4, 2017.

[3] E. P. Blasch, E. Bosse, and D. A. Lambert
High-Level Information Fusion Management and Systems Design,
Artech House, Norwood, MA, 2012.

[4] Flightradar24 Live Air Traffic,
homepage: https//www.flightradar24.com.

[5] J. D. Pararas
"Decision support systems for automated terminal area air traffic control,"
PhD., Massachusetts Institute of Technology, Flight Transportation Laboratory, 1982.

[6] M. O. Ball, R. L. Hoffman, D. Knorr, J. Wetherly, and M. Wambsganss
"Assessing the Benefits of Collaborative Decision Making in Air Traffic Management,"
USA/Europe Air Traffic Management R & D Seminar, June 2000.

[7] L. Bianco, P. Dell'Olmo, and A. Odoni, (Eds.)
"New Concepts and Methods in Air Traffic Management,"
Spinger, 2001.

[8] T. Prevot
"Exploring the Many Perspectives of Distributed Air Traffic Management: The Multi Aircraft Control System MACS,"
AIAA HCI proceedings, 2002.

[9] International Civil Aviation Organization
Air Traffic Management, Doc 4444
ATM/501, 2007.

[10] M. Salerno, D. Rondinella, M. V. Crispino, G. Costantini, M. Carota, and D. Casali
SKY-Scanner: a new paradigm for air traffic management, International Journal of Circuits, Systems and Signal Processing 2(2) (2008), 131–139.

[11] L. Saviciene
"Operationalization of Norms in Aircraft Approach/Departure Decision Support,"
Paris: IRIT press, 79–88, 2011.

[12] A. Gardi, R. Sabatini, S. Ramasamy and K. de Ridder
"4-Dimensional Trajectory Negotiation and Validation System for the Next Generation Air Traffic Management,"
in proceedings of AIAA Guidance, Navigation, and Control conference 2013 (GNC 2013), Boston, MA, USA, 2013. DOI: 10.2514/6.2013-4893.

[13] J. Rakas, B. B. Niu, J. Tom, C. Santiago, and M. Seelhorst
"Analysis of Air Traffic Control Command Entries and the Impact on Decision Support Tool Performance,"
Air Traffic Control Quarterly, Vol. 22(1) pp. 1–20 (2014).

[14] R. Koelle, and W. Strijland
"Semantic Driven Security Assurance for System Engineering In SESAR/NextGen,"
ICNS, 2013.

[15] R. M. Keller
"Ontologies for Aviation Data Management,"
IEEE/AIAA 35th Digital Avionics Systems Conference (DASC), Sacramento, 2016.

[16] The NASA Air Traffic Management Ontology (atmonto), https://data.nasa.gov/ontologies/atmonto/NAS.

[17] J. R. Boyd
The Essence of Winning and Losing,
1996.

[18] E. P. Blasch et al
"High Level Information Fusion (HLIF) Survey of Models, Issues, and Grand Challenges,"
IEEE Aerospace and Electronic Systems Mag., Vol. 27, No. 9, Sept. 2012.

[19] E. Blasch, I. Kadar, J. Salerno, M. M. Kokar, S. Das, G. M. Powell, D. D. Corkill, and E. H. Ruspini
"Issues and Challenges in Situation Assessment (Level 2 Fusion),"
J. of Advances in Information Fusion, Vol. 1, No. 2, pp. 122–139, Dec. 2006.

[20] R. So, L. Sonenberg
"Situation Awareness in Intelligent Agents: Foundations for a Theory of Proactive Agent Behavior,"
Int. Conf. on Intelligent Agent Technology, pp. 86–92, 2004.

[21] E. Blasch, A. Steinberg, S. Das, J. Llinas, C.-Y. Chong, O. Kessler, E. Waltz, and F. White
"Revisiting the JDL model for information Exploitation,"
Int'l Conf. on Info Fusion, 2013.

[22] R. So, L. Sonenberg
"Situation Awareness in Intelligent Agents: Foundations for a Theory of Proactive Agent Behavior,"
Int. Conf. on Intelligent Agent Technology, pp. 86–92, 2004.

[23] M. Dorneich et al.
"Evaluation of information quality and automation visibility in Information automation on the flight deck,"
Human Factors and Ergonomics Society Annual Meeting, 284–288, 2015.

[24] J. Rushby
"Trustworthy Self-Integrated Systems,"
International Conference on Distributed Computing and Internet Technology (ICDCIT), Springer LNCS Vol 9581, pp. 19–29, 2016.

[25] E. P. Blasch
"Ontological Issues in Higher Levels of Information Fusion: User Refinement of the Fusion Process,"
Int. Conf. on Info Fusion, 2003.

[26] P. C. G. Costa, K. B. Laskey, et al.
"Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology,"
Int. Conf. on Info Fusion, 2012.

[27] P. H. Foo, G. W. Ng
"High-level Information Fusion: An Overview,"
J. Of Adv. In Information Fusion, Vol 8, No. 1, June 2013.

[28] E. Blasch, P. C. G. Costa, et al.
"The URREF Ontology for Semantic Wide Area Motion Imagery Exploitation,"
Int'l Conf. on Semantic Technologies for Intelligence, Defense, and Security (STIDS), 2012.

[29] E. Blasch, K. B. Laskey, A-L. Joussselme, V. Dragos, P. C. G. Costa, and J. Dezert
"URREF Reliability versus Credibility in Information Fusion (STANAG 2511),"
Int'l Conf. on Info Fusion, 2013.

[30] V. Dragos
"An otological analysis of uncertainty in soft data,"
Int'l Conf on Information Fusion, 2013.

[31] J. Ziegler, and F. Detje
"Application of empirical methodology to evaluate information fusion approraches,"
Int'l Conf on Info. Fusion, 2013.

[32] E. Blasch, A. Jøsang, J. Dezert, P. C. G. Costa, K. B. Laskey, and A-L. Jousselme
"URREF Self-Confidence in Information Fusion Trust,"
Int'l. Conf. on Information Fusion, 2014.

[33] E. Blasch, X. Li, G. Chen, and W. Li
"Image Quality Assessment for Performance Evaluation of Image fusion,"
Conf. on Info Fusion, 2008.

[34] S. Alsing, et al.
"Three-Dimensional Receiver Operating Characteristic (ROC) trajectory concepts for the Evaluation of Target Recognition algorithms faced with the Unknown target detection problem,"
Proc. SPIE, Vol. 3718, 1999.

[35] R. I. Hammoud, C. S. Sahin, et al.
"Automatic Association of Chats and Video Tracks for Activity Learning and Recognition in Aerial Video Surveillance,"
Sensors, 14, 19843–19860, 2014.

[36] H. Koen, J. P. de Villiers, and G. Pavlin
"A framework for inferring predictive distibutions of rhino poaching events through causual modelling,"
Int'l Conf on Information Fusion, 2014.

[37] A-L. Jousselme, and G. Pallotta
"Dissecting uncertainty-based fusion techniques for maritime anomaly detection,"
Int'l Conf on Information Fusion, 2015.

[38] A. Oltramari, J. H. Cho
"ComTrustO: Composite trust-based ontology framework for information and decision fusion,"
Int'l Conf on Information Fusion, 2015.

[39] J. P. de Villiers, A-L. Jousselme, A. de Waal, et al.
"Uncertainty evaluation of data and infromation fusion withing the context of the decision loop,"
Int'l Conf on Information Fusion, 2016.

[40] L. Snidaro, and J. Garcia Herrero, et al.
Context-Enhanced Information Fusion: Boosting Real-World Performance with Domain Knowledge,
Springer, 2016.

[41] K. Wyffels, and M. Campbell
"Priority-Based Tracking of Extended Objects,"
J. of Advances in Information Fusion, Vol. 12, No. 2., pg 211–227, Dec. 2017.

[42] A-L. Jousselme, J. P. De Villiers, P. Costa, et al.
"Subjects under evaluation with the urref ontology,"
Int'l Conf on Information Fusion, 2017.

[43] E. Gringinger, D. Eier, and D. Merkl
"Ontology-based CNS software development,"
ICNS, 2010.

[44] E. Gringinger D. Eier, and D. Merkl
"NextGen and SESAR Moving towards Ontology-based Software developments,"
IEEE ICNS, 2011.

[45] R. Koelle, and A. Tarter
"Towards a distributed situation Management Capability for SESAR and NEXTEGN,"
Integrated Communications Navigation and Surveillance (ICNS), 2012.

[46] R. Koelle, and W. Strijland
"Semantic Driven Security Assurance for System Engineering In SESAR/NextGen,"
ICNS, 2013.

[47] E. Blasch
"Ontologies for NextGen Avionics Systems,"
IEEE/AIAA Digital Avionics Systems Conference, 2015.

[48] E. Blasch, and M. Bélanger
"Agile battle management efficiency for command, control, communications, computers and intelligence (C4I),"
Proc. SPIE, Vol. 9842, 2016.

[49] Federal Aviation Association Next Generation (NextGen), homepage: https://www.faa.gov/nextgen/.

[50] Single European Sky ATM Research (SESAR), homepage: http://ec.europa.eu/transport/modes/air/sesar/index_en.htm.

[51] Single European Sky ATM Research (SESAR) Programme
"AIRM Primer, edition 03.03.01,"
Sesar Joint Undertaking, Brussels, 2015.

[52] C. Flachberger et al
"Collaboration in Crisis Management—Learning from the transportation Domain,"
Future Security, 2015.

[53] M. Zhao, X. Lou, M. Qi, and F. Zeng
"The Research Synopsis about SWIM in China,"
IEEE 12th Int. Symp. on Autonomous Decent. System, 2015.

[54] System Wide Information Management (SWIM), available at https://www.eurocontrol.int/swim.

[55] F. Burgstaller, D. Steiner, M. Schreff, E. Gringinger, et al.
"AIRM-based, fine-grained semantic filtering of notices to airmen,"
Integrated Communication, Navigation and Surveillance Conference, 2015.

[56] The Protégé wiki,
homepage: http://protegewiki.stanford.edu/wiki/Main_Page.

[57] T. Lukoianova, and V. Rubin
"Veracity Roadmap: Is Big Data Objective, Truthful and Credible?,"
Advances in Classification Research Online, 24(1), 2014.

[58] Z. Ding, et al.
"BayesOWL: Uncertainty Modelling in Semantic Web Ontologies,"
Springer, vol. 204, pp. 3–29, 2006.

[59] K. Sambhoos, J. Llinas, and E. Little
"Graphical Methods for Real-Time Fusion and Estimation with Soft Message Data,"
Int'l Conf. on Information Fusion, 2008.

[60] G. Levchuk, et al.
"Probabilistic graphical models for multi-source fusion from text sources,"
IEEE Comp. Intelligence for Security and Defense App., 2015.

[61] Recommended Key Performance Indicators for Measuring ANSP Operational Performance,
homepage: https://www.canso.org/sites/default/files/RecommendedKPIforMeasuringANSOOperationalPerformance.pdf, 2015.

[62] European ATM Master Plan—The Roadmap for Sustainable Air Traffic Management, 2nd Ed.,
homepage: https://ec.europa.eu/transport/sites/transport/files/modes/air/sesar/doc/2012_10_23_atm_master_plan_ed2oct2012.pdf, 2012.

[63] L. K. Asante, and F. J. S. Nieto
"Complexity in the optimisation of ATM performance metrics,"
2nd International Conference on Application and Theory of Automation in Command and Control Systems, may 2012.

Carlos C. Insaurralde is a Senior Lecturer in Electronic Engineering in the Department of Engineering Design and Mathematics, University of the West of England, UK. He is the Programme Leader for BEng(Hons) Robotics, and Module Leader of four Programme Modules. He received the MEng degree in Electronics (recognition of Argentinean qualification as equivalent to the Spanish one) from the National University of Cordoba, Argentina, in 1999, the MAst and PhD degrees in Computer Engineering (Mention "Doctor Europaeus" accredited by the European University Association) from the Universidad Complutense de Madrid, Spain, in 2005 and 2007 respectively, and the MPhil degree in Electrical Engineering from Heriot-Watt University, UK in 2014. He also received a PgCert in Learning and Teaching in Higher Education from Teesside University, UK in 2017. He is a Fellow of the Higher Education Academy (FHEA), UK.

Dr Insaurralde was a Senior Lecturer in Control and Instrumentation Engineering at Teesside University, UK (2015–2018), and an Assistant Professor (Lecturer) at Heriot-Watt University, UK in 2014. He was also a Research Associate in the Institute of Sensors, Signals and Systems, Heriot-Watt University (2010–2013), and in the Software Systems Engineering Initiative, Loughborough University, UK (2008–2010). Additionally, he held post-doc positions in the Laboratory of Applied Research on Software-Intensive Technologies, Atomic Energy Commission (France, 2008), and the Institute of Production Engineering, Tampere University (Finland, 2006–2007). His research interests are architectures of intelligent and autonomous systems, multidisciplinary development of high-integrity systems, and metric assessment of systems performance. He has worked in different application domains (aerospace, defense, industrial automation, and maritime systems). His research has lately been focused on intelligent automation and autonomy (decision-making support) for Air Traffic Management (ATM).

Dr Insaurralde has been involved in several EU-funded projects, MoD-funded projects, and an ANR-funded project (France) under the role of project manager, main researcher, and task leader. He has worked in collaboration with Airbus, EADS, and BAE Systems. He has over twenty years of hands-on experience in software engineering for dependable applications. This includes over ten years of engineering research experience in robotics and autonomous systems. He is author of over eighty international publications, including a book and five book chapters. He is author of fifteen technical project reports. He received the Best Session Paper award from the Digital Avionics Systems Conference (DASC) in 2004, 2008, 2013, and 2014 as well as an award for significant contribution as section chair for the Navigation section at the DASC in 2016. He is also an instructor for conference tutorials on intelligent control architecture for autonomous unmanned vehicles (IEEE International Systems Conference and DASC).

Erik Blasch is a program officer at United States Air Force Research Laboratory (AFRL)—Air Force Office of Scientific Research (AFOSR). Previously, he was a principal scientist at AFRL in Rome, NY (2012–2017); exchange scientist to Defence Research and Development Canada (DRDC) at Valcartier, Quebec (2009–2012); and Information Fusion Evaluation Tech Lead at AFRL in Dayton, OH (2000–2009). Dr. Blasch has been an Adjunct Electrical Engineering Professor at Wright State University teaching signal processing, target tracking, and information fusion.

Dr. Blasch was a founding member of the International Society of Information Fusion (ISIF) (www.isif.org), 2007 President, and Board of Governors (2000–2010). He served on the IEEE Aerospace and Electronics Systems Society (AESS) Board of Governors (2011–2016), distinguished lecturer (2012–2018), co-chair of 5 conferences, with 9 plenary talks. He has focused on information fusion, target tracking, pattern recognition, and robotics research compiling 750+ scientific papers and book chapters. He holds 22 patents, received 33 team-robotics awards, presented over 60 tutorials, and is an associate editor of three academic journals. His co-authored books include *High-Level Information Fusion Management and Systems Design* (Artech House, 2012), *Advances and Applications of DSmT for Information Fusion* (American Research Press, 2015), *Context-Enhanced Information Fusion* (Spriner, 2016), *Multispectral Image Fusion and Colorization* (SPIE Press, 2018), and *Handbook of Dynamic Data Driven Applications Systems* (Springer, 2018).

Dr. Blasch received his B.S. in Mechanical Engineering from the Massachusetts Institute of Technology in 1992 and Master's Degrees in Mechanical ('94), Health Science ('95), and Industrial Engineering (Human Factors) ('95) from Georgia Tech and attended the University of Wisconsin for a MD/PhD in Neurosciences/Mech. Eng until being called to military service in 1996 to the United States Air Force. He completed an MBA ('98), MSEE ('98), MS Econ ('99), and a PhD ('99) in Electrical Engineering from Wright State University and is a graduate of Air War College ('08). He is the recipeint of the Military Sensing Symposium (MSS) Mignogna Leadership in Data Fusion Award, IEEE AESS MImno Best Paper Award, and IEEE Russ Bio-Engineering Award. He is a Fellow of the Society of Photo-Optical and Instrumentation Engineers (SPIE), Associate Fellow of American Institute of Aeronautics and Astronautics (AIAA), and Fellow of IEEE.

Temporal Bayes net information & knowledge entropy

KENNETH J. HINTZ
STEVEN DARCY

Various information measures have been defined on Bayes Nets (BN) with the assumption that the Bayes Net is stationary. Our interest is in the utilization of a BN as a component of a real-time, information-based sensor management system wherein the dynamics of the situation cause changes both in the structure and underlying probabilities of the nodes in the BN. If a BN is used to represent the situation assessment (SA) of an environment as a result of our observations of that environment, we can say that the BN represents our knowledge about the situation in the form of a temporal Bayes net (TBN). If one were to not observe the processes in an environment with additional sensor observations, then the underlying probabilities of at least some of the BN nodes diffuse at a rate dependent on the dynamics of the process whose uncertainty is represented by that node, hence the use of the modifier *temporal*. This loss of knowledge in the form of increasing uncertainty results in information flow from the TBN, or, as we refer to it here, temporal information loss. In order to compensate for this temporal information loss and maintain or improve our knowledge of an environment, the environment needs to be observed by obtaining data. We focus in this paper on choosing a global TBN information measure In doing so, we differentiate between aleatory nodes with stationary uncertainties and epistemic nodes with temporal uncertainties, as well as formulate a dynamic representation of these temporal uncertainties. We provide several examples of temporal information loss under different dynamic assumptions.

Manuscript received February 19, 2018; revised July 27, 2018, February 2, 2019; released for publication February 19, 2019.

Refereeing of this contribution was handled by Paulo Costa.

Authors' addresses: K. J. Hintz, Department of Electrical and Computer Engineering, George Mason University, Fairfax, VA 22030 USA (E-mail: khintz@gmu.edu). S. Darcy, Systems Engineering and Operations Research Department, George Mason University, Fairfax,VA 22030 USA (E-mail: sdarcy1@gmu.edu).

This paper results from research supported by the Naval Postgraduate School Contract No. N62271-17-P-1154 awarded by the Naval Postgraduate School, Director Contracting and Logistics. The views expressed in written materials or publications, and/or made by speakers, moderators, and presenters, do not necessarily reflect the official policies of the Naval Postgraduate School nor does mention of trade names, commercial practices, or organizations imply endorsement by the U.S. Government.

1. INTRODUCTION

This document introduces the concept of information loss over time from Bayes nets (BN) due to the dynamics associated with epistemic nodes in a BN. Epistemic and aleatory uncertainties will be defined in the sequel and adapted to BN. We call this phenomenon *temporal Bayes information (TBI)* loss to distinguish it from the term *dynamic* used in dynamic Bayes nets (DBN) by Kjaerulff [1] and Chang & Sun [2]. DBN can also have a temporal component by incorporating changes in the network structure itself over time rather than just changes in the uncertainty. Furthermore, it will be shown that the decision as to which epistemic node to update in order to maximize the information rate can change with time. This is due to the fact that the uncertainties associated with the processes represented by different epistemic nodes do not change at the same rate. We will relate the concept of TBI to two different interpretations of Bayes nets, one of which is purely aleatory in that it provides a graphical representation of the joint probability distribution of random variables, and one used in target detection and tracking problems which is composed of both epistemic and aleatory nodes. There are two examples presented later in this paper which differentiate between aleatory (yellow) and epistemic (blue) nodes as shown in Figure 2 and Figure 4. The latter BN is representative of a causal Bayes net as introduced by Pearl in his fundamental book [3].

Our interest in TBI is intimately tied with our use of a BN as an underlying component of our method of information based sensor management (IBSM). IBSM will not be discussed further here as it is has been presented in previous papers by Hintz & McVey [4], and Hintz & Kadar [5]. Briefly, the IBSM situation information expected value net (SIEV-net) takes an information measure defined on a situation assessment Bayes net and combines it with mission values and the probability of obtaining information to compute the expected situation information value rate. We use the resulting expected situation information value rate ($\mathrm{EIVR}_{\mathrm{sit}}$) to choose from among the several situation information needs that information request which will yield the highest value of $\mathrm{EIVR}_{\mathrm{sit}}$. Our interest in TBI stems from the fact that the predictable loss of information from a BN will yield different values of the maximum $\mathrm{EIVR}_{\mathrm{sit}}$ depending on the delay in fulfilling that information request. The different values of $\mathrm{EIVR}_{\mathrm{sit}}$ at different times results in different choices of which information to request.

As a brief preliminary example of how TBI can affect the amount of information which could be obtained from mutually exclusive sensing actions which could be taken at two different times in the future, let's assume that we are tracking 2 targets, the state of each one being represented as individual nodes in a BN. Let's further hypothesize that one has highly dynamic kinematics, e.g., a fighter aircraft, with a large process noise, and

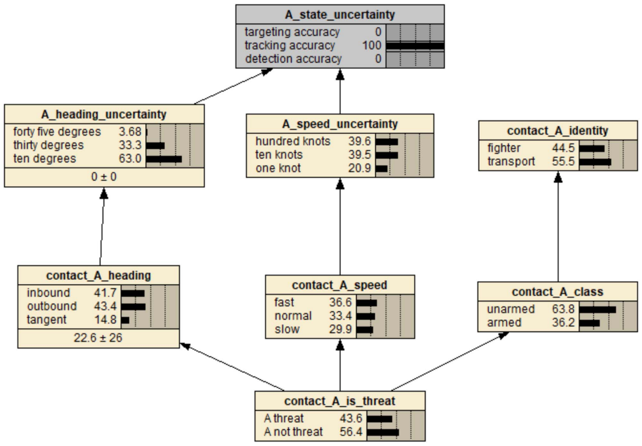

Fig. 1. BN showing the use of the target kinematic state produced by an external K-filter fusion process to populate nodes of a BN in order to estimate whether a contact is a threat.

a second target with slow dynamics, e.g., a helicopter, and smaller process noise. If we extrapolate the Kalman filter (K-filter) error covariance matrix of each of these to some proximate time in the future, it may be that we will obtain more situation information if we choose to observe the helicopter rather than the fighter, as in the case where we may have just detected the helicopter and started tracking it. However, if we were to wait to make an observation to some later time, then the high kinematic dynamics of the fighter may, through the extrapolation of both the helicopter and fighter error covariance matrices to this later time, result in the fact that more situation information will be gained by observing the fighter.

This loss of information over time has been recognized, but not explicitly evaluated by Ciftcioglu et al. [6] in dealing with maximizing information from multiple sensors. They say that "[t]he main property of QoI [Quality of Information] is that it is a composite metric which deteriorates with age and increases with time due to additional information gathered. The amount of information that sensors collect varies randomly throughout time, which leads to uncertainty in the QoI utility evolution."

1.1. Aleatory vs epistemic definitions

Winkler [7] states that "[a]t a fundamental level, uncertainty is uncertainty, yet the distinctions [aleatory and epistemic, reducible and irreducible, stochastic and subjective] are related to very important practical aspects of modelling [sic] and obtaining information." Costa, et al. [8], state that "Uncertainty Type is a concept that focuses on underlying characteristics of the information that make it uncertain. Its subclasses are Ambiguity, Incompleteness, Vagueness, Randomness, and Inconsistency…" Shafer [9], in discussing the distinction between belief and chance, provides a simple example by writing that "[c]hances arise only when one describes an aleatory (or random) experiment, like the throw of a die or the toss of a coin." We focus on two particular uncertainties, aleatory and epistemic, as they apply to BN in order to differentiate between those nodes that participate in an information measure and those that don't.

Aleatory and epistemic are terms used in seismic hazard analysis, reliability engineering, system safety, structural reliability, and risk analysis, but are not common in the information fusion literature. The general meaning of aleatory and epistemic can be taken from the Oxford English Dictionary as:

aleatory: Dependent on uncertain events or occurrences; haphazard, random [10]

epistemic: Of or relating to knowledge, or to its extent, linguistic expression, or degree of validation [11]

Unfortunately these are not very satisfying definitions for our intended use in information fusion and, in particular, situation assessment utilizing BN.

To facilitate the discussion we first make clear what we mean by aleatory and epistemic by quoting from Der Kiureghian and Ditlevsen [12] in the field of structural reliability or risk analysis:

> The word aleatory derives from the Latin *alea*, which means the rolling of dice. Thus, an aleatoric uncertainty is one that is presumed to be the intrinsic randomness of a phenomenon. Interestingly, the word is also used in the context of music, film and other arts, where a randomness or improvisation in the performance is implied. The word epistemic derives from the Greek $\epsilon\pi\iota\sigma\tau\epsilon\mu\epsilon$ (episteme), which means knowledge. Thus, an epistemic uncertainty is one that is presumed as being caused by lack of knowledge (or data).

In Abrahamson's paper related to seismic hazard analysis we find [13]

> Aleatory variability is the natural randomness in a process. For discrete variables, the randomness is parameterized by the probability of each possible value. For continuous variables, the randomness is parameterized by the probability density function.
>
> Epistemic uncertainty is the scientific uncertainty in the model of the process. It is due to limited data and knowledge. The epistemic uncertainty is characterized by alternative models. For discrete random variables, the epistemic uncertainty is modelled [sic] by alternative probability distributions. For continuous random variables, the epistemic uncertainty is modelled [sic] by alternative probability density functions. In addition, there is epistemic uncertainty in parameters that are not random by have [sic] only a single correct (but unknown) value.

While other authors have presented alternative views of these two terms, we believe that Abrahamson's meaning best serves our purposes in the field of information fusion and situation assessment.

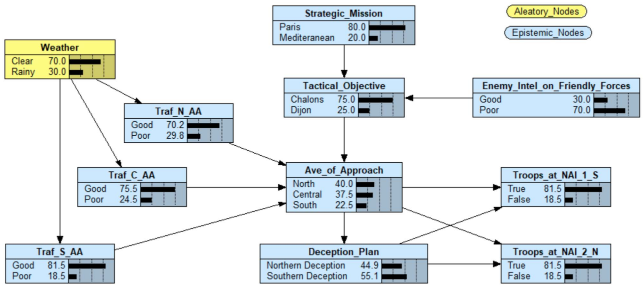

Fig. 2. Bayesian network representing enemy intent showing both aleatory (shown in yellow) and epistemic (shown in blue) nodes. After Buede, et al. [22].

According to DK&D [12], "[a]ny discussion on the nature and character of uncertainties should be stated within the confines of the model universe." They further suggest this determination should be a pragmatic choice based on how the modeler intends to use the uncertainty in the model. Since BN situation models in the information fusion world are quite diverse, it would seem that the allowance for both aleatory and epistemic chance nodes is appropriate. Furthermore, there is the possibility that the nodes change from aleatory to epistemic over time as the model is used. For example, DK&D [12] use the concept of the strength of concrete in a building having a known statistical uncertainty before the building is built (aleatory); however, after the building is built, measurements of the strength can be taken over the time the concrete is curing leading to an epistemic statistical uncertainty. Note that the process is the curing of the concrete with an associated uncertainty which can be reduced if measured, but remains the same or increases in uncertainty if not measured. In the case of situation assessment, an example is converting from the probability that a target is going to enter a volume of space (aleatory) to the probability that a target has been detected (epistemic) once a detection is made. The fact that there has been a detection does not mean that there is a target in that volume with absolute certainty since each detection has associated with it a probability of detection less than 1 ($P_d < 1$) and a probability of false alarm of greater than zero ($P_{fa} > 0$). The uncertainty about whether a target is actually in the volume is reduced with repeated measurements. Another example is shown in Figure 2 wherein the weather is an explicit aleatory node during mission planning, but becomes an epistemic node during its use when particular values of the weather can be acquired as evidence.

If we relate the above to Pearl's causal networks [14] [3] which require a directed relation between nodes, epistemic is a straightforward uncertainty which is added to either the linear or nonlinear functional relation between nodes as in the linear relationship below with the additive aleatory random variable u_i:

$$x_i = \sum_{k \neq i} \alpha_{ik} x_k + u_i \quad i = 1, \ldots, n \tag{1}$$

The nodal value of interest, x_i, is epistemic as its uncertainty can be refined with repeated measurements thereby reducing the uncertainty introduced by the random additive component, u_i.

1.2. Aleatory or epistemic: stationary or nonstationary?

Aleatory uncertainties may change, but cannot be improved with repeated measurements as they are associated with a naturally occurring randomness. A counter argument to this, which we will ignore without loss of generality, can be best exemplified by the probabilities associated with the roll of a die. Mathematically, a fair die has equal probability of the single event comprised of a face of a die. No amount of experiments on this mathematically fair die will change that. In reality, no physical die is perfect and hence, not fair. That is, repeated rolling of the die will show that some faces will occur more than others due to the imperfections in the physical die. This is not to be confused with the typical gambler's mistake most easily associated with the flip of a coin. If the coin toss results in an unusually long run of heads or tails, one wants to think that the next toss will be the opposite even though we know that there is equal probability of the two faces of the coin occurring as a result of the next toss.

Aleatory uncertainties may change over time, but not due to measurements. In the case of weather and whether or not it is going to rain, the aleatory uncertainty changes if there are observed clouds, but repeated measurements to determine if clouds are present do not change the uncertainty about whether it is going to rain or not.

Epistemic uncertainties can be non-stationary or changing over time due to observations of, or changes in, the process dynamics. An interesting example is a situation assessment node which represents the kinematic state of a target in track. If this nodal estimate is derived from a Kalman filter, then we can see both aleatory and epistemic statistics depending on how the modeler utilizes the K-filter. The equations from the discrete K-filter are [15]

system model,

$$\vec{x}_k = \mathbf{\Phi}_{k-1}\vec{x}_{k-1} + w_{k-1}, \quad w_k \sim N(0, \mathbf{Q}_k) \tag{2}$$

in which Φ_{k-1} is the state transition matrix and w_{k-1} is the process noise, both having subscripts indicating that they are non-stationary and may change over time,

measurement model,

$$\vec{z}_k = \mathbf{H}_k\vec{x}_k + v_k, \quad v_k \sim N(0, \mathbf{R}_k) \tag{3}$$

in which $\vec{z}_k$ is the observation vector, $\mathbf{H}_k$ is the observation matrix, and v_k is the additive, white, Gaussian measurement noise, all having subscripts indicating that they are non-stationary and may change over time, state estimate extrapolation,

$$\hat{\vec{x}}_k^- = \mathbf{\Phi}_{k-1}\hat{\vec{x}}_{k-1}^+ \quad (4)$$

error covariance matrix extrapolation,

$$\mathbf{P}_k^- = \mathbf{\Phi}_{k-1}\mathbf{P}_{k-1}^+\mathbf{\Phi}_{k-1}^T\mathbf{Q}_{k-1} \quad (5)$$

Kalman gain matrix,

$$\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}_k^T[\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \mathbf{R}_k]^{-1} \quad (6)$$

state estimate update,

$$\hat{\vec{x}}_k^+ = \hat{\vec{x}}_k^- + \mathbf{K}_k[\vec{z}_k - \mathbf{H}_k\hat{\vec{x}}_k^-] \quad (7)$$

and, error covariance matrix update.

$$\mathbf{P}_k^+ = [I - \mathbf{K}_k\mathbf{H}_k]\mathbf{P}_k^- \quad (8)$$

Notice the similarity in form of the K-filter equations (2) and (3) to (1) of Pearl in that there is an additive random component to both the system model (2), and the measurement (3) which ripples through the other state estimator equations. The random variable in the system equation, (2), $\vec{w}(t)$, generally called the process noise, represents the unmodeled uncertainties (Pearl's latent variables) associated with the process dynamics including the random maneuvers of the target. The random variable in the measurement equation, $\vec{v}(t)$, represents the additive noise due to the fact that no observation is perfect and there are uncertainties associated with it.

As an example of using the state estimate produced by a K-filter to populate or update the parameters of a node in a BN, we present the simple BN of Figure 1. This network shows how the various uncertainties in the components of the kinematic state vector can affect the uncertainty in a situation assessment node which is not directly determined by the kinematic state.

In these most general K-filter equations (2) through (8) Φ_{k-1}, w_{k-1}, $\vec{z}_k$, $\mathbf{H}_k$, and v_k all have subscripts indicating that they are non-stationary and may change over time, indicating that the process and the resulting state estimates are not stationary. The system model propagates based on the previous state with a time-dependent random component added to it and a reduction in uncertainty based on noisy observations. It is important for our purpose here to note that if an observation of the system is not taken, then the uncertainty of the extrapolated state variable, $\hat{\vec{x}}_k^-$, as represented by the extrapolated error covariance matrix, $\mathbf{P}_k^-$, grows. The uncertainty in the extrapolated state estimate is represented by some norm of the error covariance matrix. The trace will not do as a norm as it is dimensionally nonconformal; the determinant, while dimensionally conformal, and monotonically related to information, does not have meaningful dimensions but may still be a useful norm. Alternatively, the error covariance matrix can be normalized to meaningful spatial units by pre- and post-multiplying by a dimension conforming matrix.

It can be seen from the error covariance extrapolation (5) that the state estimate (7) depends on the propagation of the previous state estimate (4) plus the Kalman gain (6), $\mathbf{K}(t)$, multiplying the difference between the previous estimate and the observation. If there is no observation (3), $\vec{z}(t)$, then the uncertainty continues to grow. This is our first hint that without continual observations of the state of a process corrupted by random process (latent variable) noise (2), our uncertainty about its state (5) grows, and hence its entropy.

Whether the elements of the K-filter are treated as stationary or non-stationary depends on the modeler's understanding of the process and how the model is to be used. One might consider the state propagation matrix in (2), $\mathbf{\Phi}_{k-1}$, representing the physics of the target's trajectory, to be stationary and unchanging over time yielding a constant $\mathbf{\Phi}$. If the same sensor is used to obtain a measurement of the target, then the observation matrix of (3) also becomes a constant, $\mathbf{H}_k = \mathbf{H}$. Even with this simplifying assumption, we still need to deal with the additive random components in (2) and (3), the process noise, w_k, and the measurement noise, v_k which are characterized in (2) and (3) by their covariance matrices, $\mathbf{Q}_k$ and $\mathbf{R}_k$, respectively. Typically target trackers include multiple model (e.g., IMM, Interacting Multiple Model [16]) methods with different process noise covariances $\mathbf{Q}_k$ at different stages in the target tracking process. Since $\mathbf{Q}_k$ is directly involved in the computation of the extrapolated error covariance matrix of (5), $\mathbf{P}_k^-$, which is used to compute the error covariance matrix update of (8), $\mathbf{P}_k^+$, the amount of information change associated with target observations under different model assumptions will change independently of the observation noise.

The measurement noise covariance matrix in (3), $\mathbf{R}_k$, also directly affects the amount of information associated with the computation of the state estimate update as well as the Kalman gain of (6), $\mathbf{K}_k$, which is used to compute the error covariance matrix update of (8), $\mathbf{P}_k^+$. Of course, the measurement covariance may change from observation to observation, but let's assume it is constant for the sake of discussion. The point here is to show that while the process and measurement noises can be considered stationary and hence aleatory, the resulting uncertainties in the target state updates as measured by the updated error covariance matrix are epistemic.

1.3. Aleatory and epistemic BN nodes

Concerning the modeler's view of the K-filter state estimate of a target in track as part of a BN representing the situation assessment, the random components can be viewed as either aleatory or epistemic. Continuing with our K-filter example which uses observation data of (3), $\vec{z}_k$, to reduce our uncertainty about the kinematic state of a target in track, we can look at the sources of randomness and see that they are, in the general model, non-stationary as they are all functions of time.

We do not discuss here the choices made by the modeler, but will, in our example later, show how evolving process dynamics of a target model affect the amount of information that one can extract from a measurement. We do recognize that in a situation assessment there will be a combination of both aleatory and epistemic nodes. Our concern with respect to a global entropy of a BN is limited to the epistemic nodes since there is no change in uncertainty in the aleatory nodes.

As we will see in the sequel, entropy changes reflect a gain or loss of information. We will extend this epistemic entropy change to a BN and see that there is a global gain or loss of information over time but only due to epistemic nodes.

Section 1 is the introduction which provides some necessary background information. Section 2 differentiates and makes clear the distinction among data, information, and knowledge. Section 3 investigates properties of hard (i.e., physics based) and soft (i.e., human-derived) data to conclude that there exist probabilistic measures which can be used to compute entropy of both hard and soft data. In section 4 we present sample computations which may help to clarify some of the newly introduced concepts.

2. DATA, INFORMATION, AND KNOWLEDGE

Waltz [17] distinguishes among three levels of abstraction of knowledge: data, information, and knowledge.

- Data are "individual observations, measurements, and primitive messages [which] form the lowest level. Human communication, text messages, electronic queries, or scientific instruments that sense phenomena are the major sources of data."
- Information is "organized sets of data...The organization process may include sorting, classifying, or indexing and linking data to place data elements in relational context for subsequent searching and analysis."
- Knowledge or foreknowledge (predictions or forecasts) is "information once analyzed, understood, and explained..."

For our purposes we take a slightly different approach by considering *information* to be a change in our uncertainty about processes in the environment which result from temporal changes, an observation, or the acquisition of relevant data (evidence). *Knowledge* in our context is expressed in the form of a Bayes net because the BN contains both the causal processes in the environment as well as the uncertainties associated with them. Furthermore, the fact that BN uncertainty increases over time is already known as Singhal & Brown [18] note in their discussion of dynamic Bayes nets, "[a] decay function is associated with the PDFs that increases the variance of the beliefs when they have not been updated for a period of time." They also recognize that observations do not have to be regular or synchronous. "To relax synchronization issues and constraints, we employ an asynchronous update policy that uses dynamic Bayesian networks to create new probability density functions (PDF)." [18]

If we consider the BN as the repository of knowledge about the situation, then the changes in uncertainty associated with the BN can be considered as information gain or loss. It is common to think about information gain as a result of obtaining data, but it is less common to think about the information loss associated with increases in our uncertainty about a process as the process evolves over time. Generally the BN is considered to be stationary until more data, i.e., evidence, are obtained to decrease the remaining uncertainty, but that is not the case when we are dealing with processes. As we saw in the K-filter target tracking example previously presented in Figure 1, our uncertainty about the kinematic state estimate grows as time advances if we do not observe the process due to the additive process noise. In the case of K-filters, the loss of information can be computed as a change in the entropy of our state estimate over time [4]. If we make an observation and obtain data, then the difference in uncertainties represented by the entropies of a norm of the extrapolated error covariance matrix (5), $\mathbf{P}_k^-$, and the norm of the error covariance update (8), $\mathbf{P}_k^+$, is a measure of the amount of information gain. This information gain represents the increase in knowledge as a result of obtaining data and decreases the uncertainty in the BN.

2.1. Bayes net information

We can extend this concept from the kinematic state of a single target in track to the global knowledge of a situation as represented in a BN [18]. Situation information and sensor information are differentiated by the authors [19], and we only focus here on situation information as represented by the global change in uncertainties in a BN, namely entropy changes among the situation assessment nodes. The Shannon entropy [20], H, of a discrete random variable, X, with possible values $\{x_1, x_2, \ldots x_n\}$ and probability mass function $P(X)$ is computed in the usual manner as

$$H(X) = \mathcal{E}[I(X)] \tag{9}$$

$$H(X) = \mathcal{E}[-\ln(P(X))] \tag{10}$$

with $\mathcal{E}$ being the expectation operator and I being the information content of the RV. Alternatively, and letting the RV, X, be a BN node, N_j, the node entropy can be computed as follows

$$H(X) = \sum_{i=1}^{n} P(x_i) I(x_i) \tag{11}$$

$$H(N_j) = -\sum_{i=1}^{n} P(x_i) \log_b P(x_i) \tag{12}$$

here n is the number possible values that can be taken on by a single node of a BN, N_j, and b is the radix of the logarithm used. Utilizing the radix 2 yields the entropy measured in bits.

After an observation which changes the probability distribution associated with node N_j, the single node BN information, I^{+j}, [19] is, therefore,

$$I^+(N_j) = H^-(N_j) - H^+(N_j) \tag{13}$$

where j is the index number of a single node of the BN, the superscript "+" indicates the values associated with the jth node after the effects of the observation have changed the probability within the jth node (the *a posteriori* value), and the superscript "−" reflects the probability of the jth node before the observation (the *a priori* value).

The global BN information gain or loss is the sum of the information gain/loss of all the nodes. Since we assume a mixture of aleatory and epistemic nodes in the BN, and furthermore that the aleatory nodes are stationary, there is no information gain/loss associated with them. That is, we only have to sum the information gain/loss of epistemic nodes. Without observations, there will be a net increase in our uncertainty of each process node with an associated loss of information over time. Assuming mutually exclusive sensor observation opportunities, there may be either a net global gain or net global loss of information in the situation assessment BN with the observation of a single node. For the observed node, there may be either a gain or a loss of information based on whether the observation decreases the uncertainty more than it had increased since the last observation. For the non-observed process nodes there may be a loss of information since they are not being observed and their uncertainty may have grown since their last observation.

The global temporal Bayes net information at the kth observation is

$$I_k^{\mathrm{TBN}} = \sum_{\substack{\text{all epistemic} \\ \text{nodes}}} [H_k - H_{k-1}] \tag{14}$$

which can be reduced to

$$I_k^{\mathrm{TBN}} = \sum_{j-1}^{m} I^+(N_j) \tag{15}$$

where m is the number of epistemic nodes in the BN and $I^+(N_j)$ is defined in (13).

2.2. Hard/soft knowledge and information

The implication until now in this paper is that the situation assessment in the form of a BN represents only kinematic uncertainty of processes in the domain of concern. We take a more egalitarian view of situation assessment in that nodes can represent kinematic uncertainties as well as *intentional* (not purposeful, but rather motivational as in the node *contact_A_is_threat* of Figure

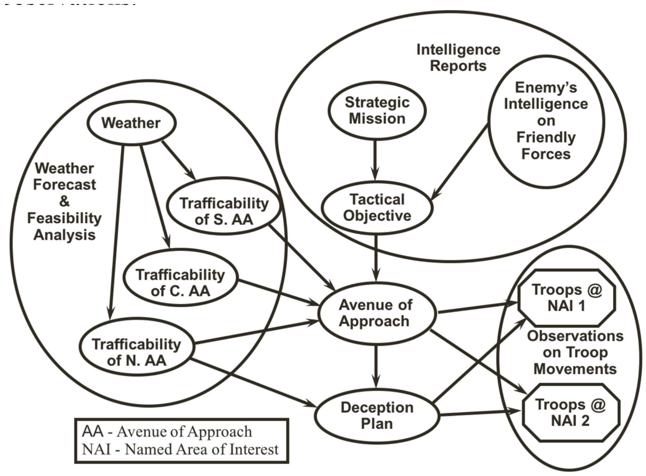

Fig. 3. Relational diagram representing enemy intent. After Buede, et al. [22].

1) uncertainties about hypotheses such as whether the enemy is going to attack or not. Hypothesis nodes like this can be partially resolved if there is overt physical action or observable preparatory action on the part of the enemy; however, it is more likely that actionable intelligence is derived from the interception and analysis of communications intelligence (COMINT) or other automatically processed natural language communications, i.e., soft data. According to Dragos, [21] "[s]oft data are a mix of both facts and opinions" the difference between the two being the source and the probability associated with each.

Yet it doesn't matter whether the source of data is hard or soft, but rather whether the acquisition of data changes our uncertainty about a particular aspect of the situation being assessed as reflected by a changed probability in one or more nodes of the BN.

2.3. Changing BN structure, information gain/loss?

Since we consider a BN as a knowledge representation structure with uncertainties associated with each node, the addition of another node, be it aleatory or epistemic, does not inherently add any *information* unless the addition of the node connects to other nodes which are conditioned on it. The addition of a node may affect the amount of temporal information that is gained or lost as time progresses or observations are made since the BN information is the sum of the information lost over time and likely regained with observations.

But the question arises as to how much *knowledge* is contained in a BN and whether adding or deleting a node changes that amount of knowledge. If we take the entropic view of uncertainty and the information theoretic view of information being a change in entropy, then we can consider the maximum uncertainty, the total entropy, in a BN to be the sum of the entropies of all the nodes *as if* each of these entropies were at its maximum uncertainty. If we measure the entropy of a node, we can view this as its *potential information (PI)* since

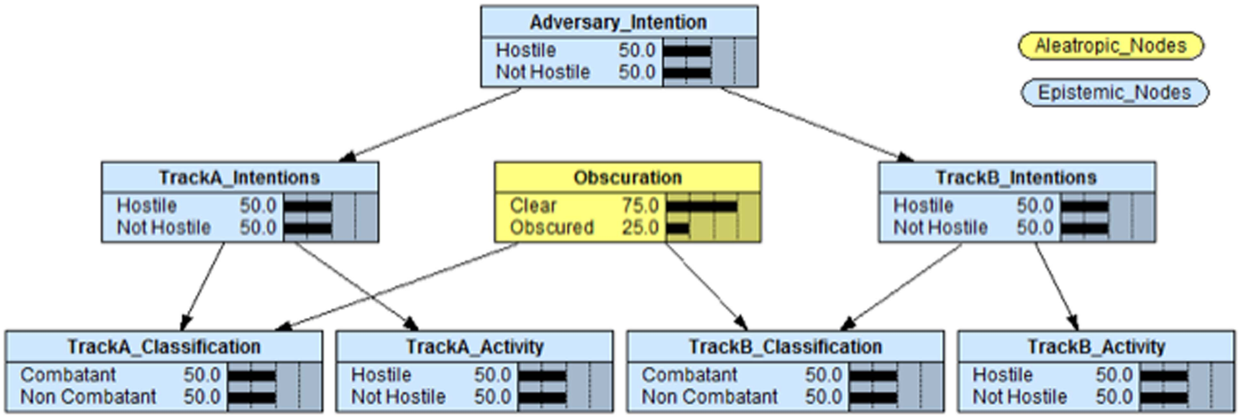

Fig. 4. Example BN containing uniform probabilities (no evidence) and both aleatory (yellow) and epistemic (blue) nodes, $KEn = 7$.

it represents the amount of uncertainty that can be resolved through measurements if it is an epistemic node. The maximum PI of a BN node occurs when all values are equally probable and represents the maximum information that can be obtained from a BN node resulting in perfect knowledge of that node since the entropy of a node with no uncertainty is zero. If we extend this to the BN itself, then we can talk about the maximum potential information of a BN as the sum of the maximum PIs of the individual nodes. Note, however, that the probabilities of a node may not be at their maximum uncertainty since we may have some *a priori* knowledge which skews the probabilities. In this case, the PI is the entropy of that skewed distribution associated with the node.

Since we have made a distinction in this paper between aleatory and epistemic nodes, we need to define which nodes need to be included in our definition of potential information. We assert that only epistemic nodes should be included as, by definition, the uncertainty in aleatory nodes cannot be reduced by measurements and those in epistemic nodes can. The decision between which ones are epistemic nodes and which ones are aleatory nodes is a modeling decision and can change with point of view and over time and there is no general rule that can be applied other than whether observations of any other node in the network changes a node's probabilities.

We also can view the observing of a node to reduce its uncertainty as gaining information about the node. We call this *kinetic information (KI)* because it results from a physical or cyber action and a change in the BN as well as a reduction in the potential information yet available to be gained. We can actively observe the process associated with a node to obtain kinetic information. Alternatively, by not observing a random variable related to a dynamic process represented by a node, the BN can leak KI over time which increases its PI.

Currently, there is no unit for the uncertainty of knowledge in a BN. We propose to use units of *Knowledge Entropy (KEn)* to represent the uncertainty in a BN. Zero KEn results when there is no uncertainty in any of the epistemic nodes. There is precedence for this new use of an old (if not archaic) word if one examines definitions found in the Oxford English Dictionary (Oxford English Dictionary, 2017) which defines *ken* as:

- *ken*, v.1, 11. a. To know (a thing); to have knowledge of or about (a thing, place, person, etc.), to be acquainted with; † to understand. Now chiefly *Sc*.
- *ken*, v.1, 12. a. *intr.* or *absol*. To have knowledge (*of* or *about* something). † Also with inf.: To know how to, to be able to (*obs.*).

So we can refer to the KEn of a BN at any time as measured in bits of uncertainty in our situation knowledge. The KEn can change over time due to the leakage of KI or the acquisition of KI through observations, and can be computed as the sum of the entropies of all epistemic nodes in the BN. Formally, the knowledge entropy of a BN, KEn, is

$$KEn(t) = \sum_{\substack{\text{all epistemic} \\ \text{nodes}}} H(t) \tag{16}$$

and the amount of temporal Bayes information, TBI, which results from a change in nodal probabilities or network structure from time t_0 to t_1, is

$$TBI(t_1) = KEn(t_0) - KEn(t_1) \tag{17}$$

As previously mentioned, TBI may be zero, positive, or negative. Zero TBI means that no network information was gained or lost over the time period although there could be individual nodal information changes whose net sum is zero. Positive or negative TBI indicates a gain or loss of information respectively with a concomitant change in our situation knowledge as represented by the temporal BN.

One would like to hypothesize that there is a conservation of knowledge law associated with BN, i.e., the *KNowledge Entropy (KEn)* is conservative, and that there is a one-for-one exchange between KI and PI, but this does not appear to be the case due to the conditional probabilities relating nodes. Increasing or decreasing the

TABLE 1.
Entropy of uniformly distributed probabilities based on the number of bins.

# of bins in X	Uniform probability	$\log_2(p_i)$	$PI = H(X)$ (bits)
2	0.500	−1.000	1.000
4	0.250	−2.000	2.000
5	0.200	−2.322	2.322
8	0.125	−3.000	3.000

uncertainty in one node may increase or decrease the entropy in other nodes, but our preliminary investigations lead us to conjecture that the gain in KI is not offset by an equal loss in PI. This is a topic that bears further investigation but is not the main point of this paper so we leave it for now.

$$KEn_{\text{total}} \neq \sum PI + \sum KI \tag{18}$$

We can now answer the question about what to do about a node which is added to, or deleted from, a dynamic BN and that is to simply consider it to be adding or deleting potential information to or from the knowledge represented in the BN. Adding a node will increase the KEn, but by less that the PI of the node if its connections affect the other probabilities. Furthermore, the loss of information, KI, results in an increase in the PI although not on an equivalent basis. The acquisition of information, KI, by observing a relevant RV decreases the PI. The total knowledge, the KEn, in a structurally stationary BN is knowable and computable. As with the question of conservation of information in a BN, we defer this topic to a later paper.

3. BN HARD/SOFT INFORMATION

We have already described the uncertainty associated with kinematic state estimates utilizing the K-filter formulation. This epistemic uncertainty is fully described by the increase in a norm of the error covariance matrix as it propagates over time or decreases with an observation. Other similar physical state estimates have continuous or discretized uncertainties that are straightforward to work with. Soft knowledge in BN, on the other hand, requires additional explanation since it includes other forms of uncertainty as noted by Dragos [21] namely

- Intrinsic uncertainties such as ambiguity, vagueness, and precision
- Source related uncertainties which are a mixture of facts and opinions
- Relational uncertainties which are concerned with inaccuracies, overlappings, and contradictions in analysis

TABLE 2
Conditional Probability Table (CPT) of the TrackA_Classification node.

TrackA_Int	Obscured	Combatant	NonCombatant
Hostile	Clear	75	25
Hostile	Obscured	50	50
Not Hostile	Clear	25	75
Not Hostile	Obscured	50	50

Dragos [21] continues with methods for estimating all of these uncertainties which will not be repeated here. For our purposes, we will assume that soft uncertainties can be estimated allowing us to compute entropies of soft data.

The suitability of uncertainty in any form…hard or soft, social or physical, quantitative or fuzzy…has been shown by Kjaerulff [23] to be applicable to formulation as a BN. "Note, that the method in this paper can be applied to other evidential frameworks where independent pieces of evidence are combined into a joint evidence e.g., Bayesian combination. For more high-dimensional problems, i.e., when it could be more suitable to utilize a graph structure for modeling dependencies e.g. Bayesian Networks,…" That is, relations among aleatory and epistemic processes such as the example of Buede et al., [24] as shown in Figure 2, can be represented in a causal BN, also from Buede, et al., as shown in Figure 3.

4. EXAMPLE GLOBAL TEMPORAL BN COMPUTATIONS

The following simplified examples will demonstrate some of the concepts introduced in this paper. First, looking at an individual node and computing the entropy of a uniform distribution of discrete values in accordance with (11), we see as exemplified in Table 1 that the potential information of the jth node is simply the $\log_2$(*number of bins in* jth *node*) and the potential information of the BN consisting of m nodes each having k_j values associated with each nodes.

$$PI(N_j) = \sum_{j=1}^{m} -\log_2 k_j \text{ bits} \tag{19}$$

If each node were a true/false hypothesis node, then there would be one bit of PI/node resulting in an m-node BN containing an upper bound of m-bits of PI since the connectivity of the BN reduces the actual amount of information that is available. Clearly as one changes the number of uniformly distributed bins/node, the summation of PI is easily calculated as well as the amount of PI if a node is added or deleted.

4.1. Simple BN PI and KI example for epistemic nodes

In order to instantiate some of the concepts introduced here, we perform PI, KI, and KEn com-

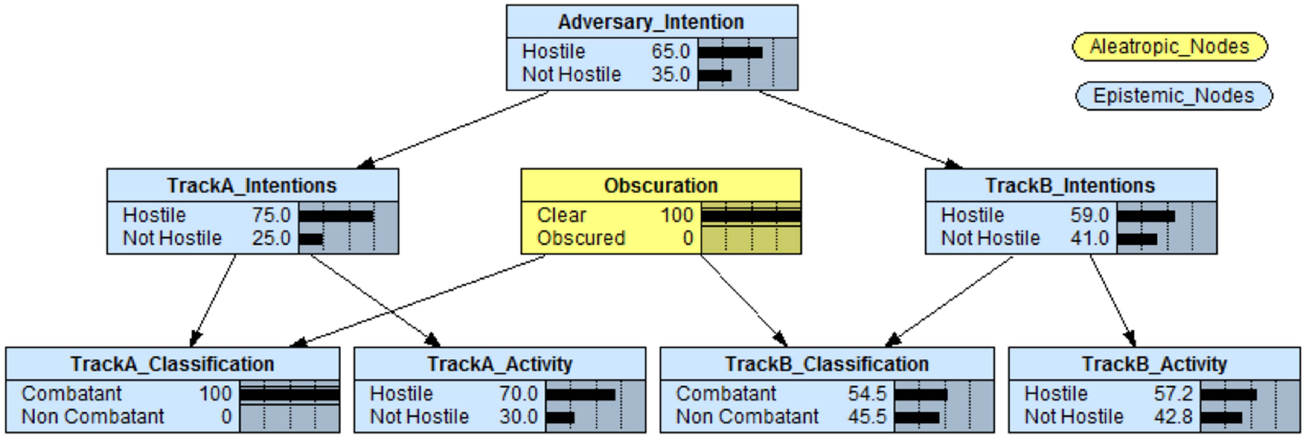

Fig. 5. Classification evidence at time T_0.

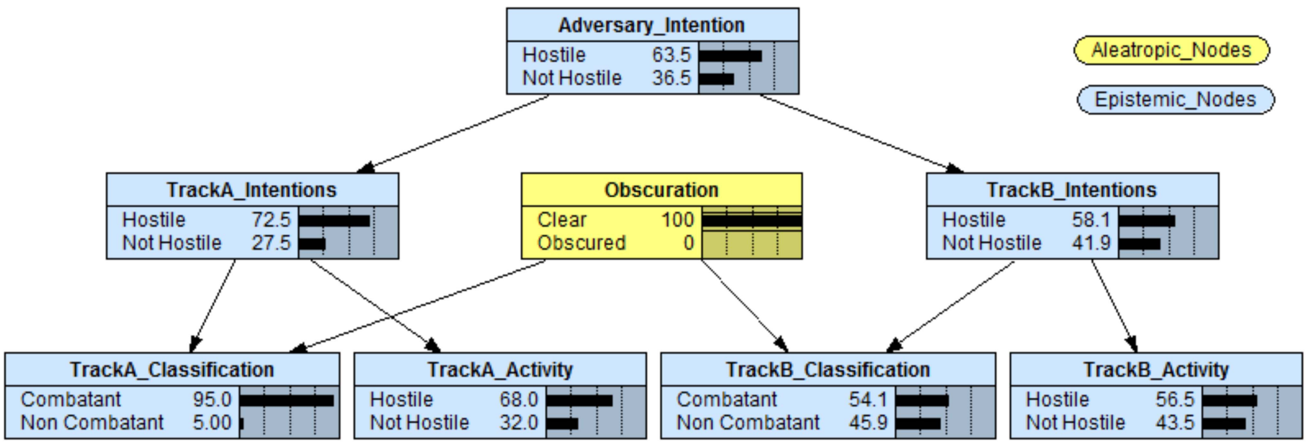

Fig. 6. Classification evidence at time T_{30} showing a change in the BN's initial knowledge with a KEn = 5.58 and the result of a loss of information in the TrackA_Classification node resulting in an increase in KEn to 5.95.

putations on a simple 8-node BN as shown in Figure 4, with all the nodes being epistemic and uniform probability nodes, except one (Obscuration) which is aleatory. Associated with the TrackA_Classification and TrackB_Classification nodes is a Conditional Probability Table (CPT) as shown in Table 3.

For our numerical example, we use the sensitivity as computed in Norsys Netica [25] BN program. The mathematical formulations utilized by Netica are documented in their on-line documentation [26]) and are the same as (19) above. Furthermore, we have done sample calculations outside of Netica utilizing the net of Figure 4 without the aleatory "obscurations" node and the results of our calculations match those produced by the Netica sensitivity analysis.

Referring to Figure 4, the initial entropy of the Adversary_Intention node at time t_0 is the expected 1.0 bits with uniform distributions in the other two nodes. If the TrackA_Classification node of Figure 4 is set to 100% as shown in Figure 5 and Figure 6, the Adversary_Intention changes to 65% Hostile/35% Non-Hostile and the BN KEn changes from 7 to 5.58 indicating a global network decrease in uncertainty (or increase in KI) of 1.42 bits as a result of the sensing action which provided the Combatant classification with 100% certainty.

We demonstrate the temporal increase in uncertainty in the TBN by changing the probabilities of the TrackA_Classification node. At some later time, t_{30}, we assume the uncertainty has decreased from 100% Combatant to 95% Combatant/5% Non-Combatant. This temporal loss of information results in the Adversary_Intention changing to 63.5% Hostile/36.5% Non-Hostile and the BN KEn increasing from its t_0 of 5.58 to its t_{30} value of 5.95 of a KI loss of 0.07 bits.

As an example of an alternative type of information loss related to a different sensing node, TrackA_Activity, Figure 7 and Figure 8 shows the BN of Figure 4 with initial, non-uniform uncertainties in the TrackA_Activity node of 100% Hostile/0% non-hostile. With this initial condition, the Adversary_Intention becomes Hostile 74%/Non-Hostile 26% for an initial KEn at t_0 of 5.06 bits.

We model the temporal change in our certainty of the TrackA_Activity node by decreasing the uncertainty at some later time, t_{30}, from 100% Hostile to 80%/20% as shown in Figure 7 and Figure 8. This temporal loss of information results in the Adversary_Intention changing

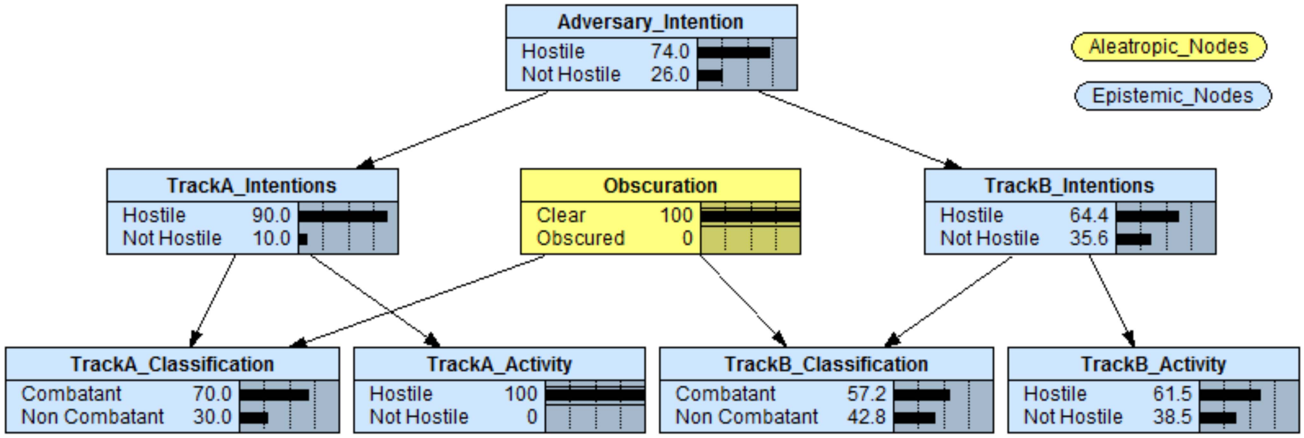

Fig. 7. Activity evidence at time T_0.

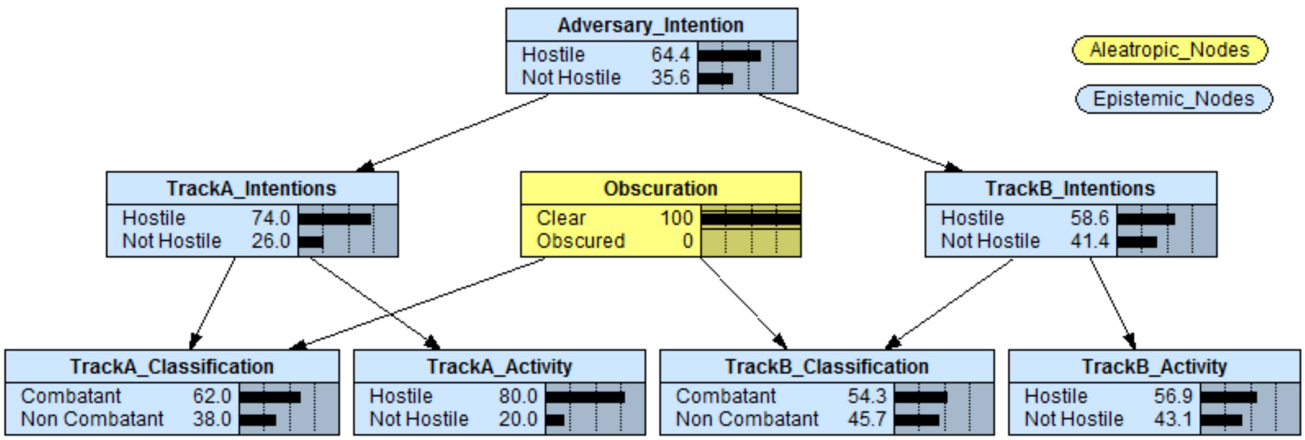

Fig. 8. Activity evidence at time T_{30} showing a change in the BN's initial knowledge with a KEn = 5.06 and the result of a loss of information in the TrackA_Activity node resulting in an increase in KEn to 6.41.

TABLE 3
Table summarizing the results of information loss over time due to decreased uncertainty in classification and, alternatively, identity. No obscuration in the aleatory node.

Scenario with no obscuration	Total Knowledge Entropy (epistemic only)
No evidence	7.00
Classify TrackA, t_0	5.58
Classify TrackA, t_{30}	5.95
Identify Activity TrackA, t_0	5.06
Identify Activity TrackA, t_{30}	6.41

to 64.4% Hostile/35.6% Non-Hostile and the BN KEn increasing from its t_0 of 5.06 to its t_{30} value of 6.41 of a KI loss of 1.35 bits. The loss in KEn with time is summarized in the table of Table 3.

This example shows that if we were to use a BN with no evidence and a KEn of 7.0 with the expected information gains at t_0 of 1.42 bits if we choose to classify as opposed to 1.94 bits if we choose to identify, we would choose to classify since it yields the maximum information. If, on the other hand, if we chose to wait until t_{30} the expected information gain from the initial KEn of 7.0 would yield 1.05 bits for classify and 0.59 bits for identify showing that accounting for the temporal loss of information from t_0 to t_{30} results in a different choice of which sensor function to use.

Other findings have been computed which result in higher losses of information while most result in a positive flow of KI into the BN.

4.2. Information in the presence of aleatory node

Remembering that the computation of KEn only includes epistemic nodes, the question arises as to the effect of an aleatory node on the amount of information gain and choice of sensor function if one makes different assumptions about the probabilities in an aleatory node. If the aleatory Obscuration node is changed from its 100% clear as used for the previous example to 75% clear/25% obscured, the following results. The results are shown in Table 4.

Referring to Figure 4, the initial entropy of the Adversary_Intention node at time t_0 is still the expected 1.0 bits with uniform distributions in the other two nodes. If the TrackA_Classification node is set to 100%, the Adversary_Intention changes to 61.3% Hostile/38.8% Non-Hostile and the BN KEn changes from 7 to 5.77

TABLE 4
Table summarizing the results of information loss over time due to decreased uncertainty in classification and, alternatively, identity. Twenty-five percent obscuration in the aleatory node.

Scenario with 25% obscuration	Total Knowledge Entropy (epistemic only)
No evidence	7.00
Classify TrackA, t_0	5.77
Classify TrackA, t_{30}	6.10
Identify Activity TrackA, t_0	5.12
Identify Activity TrackA, t_{30}	6.43

indicating a global network decrease in uncertainty (or increase in KI) of 1.23 bits as a result of the sensing action which provided the Combatant classification with 100% certainty.

As before, we model the temporal change in our certainty of the TrackA_Classification node by decreasing the uncertainty at some later time, t_{30}, from 100% Combatant to 95%/5%. This temporal loss of information results in the Adversary_Intention changing to 60.1% Hostile/39.9% Non-Hostile and the BN KEn increasing from its t_0 of 5.77 to its t_{30} value of 6.10 of a KI loss of 0.33 bits.

As an example of an alternative type of information loss under aleatory uncertainty related to a different sensing node, the TrackA_Activity, node is changed to 100% Hostile/0% non-hostile. With this initial condition, the Adversary_Intention becomes Hostile 74%/Non-Hostile 26% for an initial KEn at t_0 of 5.12 bits.

Again, modeling the temporal change in our certainty of the TrackA_Activity node by decreasing the uncertainty at some later time, t_{30}, from 100% Hostile to 80%/20%. This temporal loss of information results in the Adversary_Intention changing to 64.4% Hostile/35.6% Non-Hostile and the BN KEn increasing from its t_0 of 5.12 to its t_{30} value of 6.42 of a KI loss of 1.30 bits. The loss in KEn with time is summarized in Table 4.

This aleatory example, Table 4, shows that if we were to use a BN with no evidence and a KEn of 7.0, choosing to classify would yield an expected information gains at t_0 of 1.23 $(7.0 - 5.77)$. If, instead, choosing to identify would yield 1.88 bits $(7.0 - 5.12)$. Because of this expected differential information gain, would choose to identify since it yields the maximum information. If, on the other hand, we choose to wait until t_{30} the expected information gain from the initial KEn of 7.0 would yield 0.9 bits $(7.0 - 6.1)$ for classify and 0.57 bits $(7.0 - 6.43)$ for identify, leading us to choose to classify as the maximum information choice. That is, accounting for the temporal loss of information from t_0 to t_{30} results in a different choice of which sensor function to use in order to maximize the information gain for a single observation.

For this example of changes in our *a priori* assumption about the probabilities associated with an aleatory node, there is a change in the expected information gain even though the entropy of the aleatory node is not included in the information measure. This shows that our model assumptions about the unmeasurable causal probabilities can affect our choice of sensing actions since they may affect our expected situation information expected value rate ($EIVR_{sit}$).

5. CONCLUSION

The differentiation between aleatory and epistemic nodes in Bayes nets has been defined and illustrated. It is also shown that BN are not limited to hard data as the analogy to Kalman-filter shows, but soft data nodes can be included since there exist soft data entropy measures. The fact that both hard and soft data uncertainty measures can be expressed as entropies allows one to put the two types of knowledge in the same BN and apply information measures based on entropy changes. Potential information and kinetic information are defined and it is conjectured that a conservation of knowledge law exists, but the details of this will be the subject of further research. Finally, a simple example of a temporal BN was presented showing how the leakage of information over time could lead to increases in entropy over time which could affect the choice of expected situation information gain when utilizing an information based sensor measurement (IBSM) approach to sensor management.

REFERENCES

[1] U. Kjaerulff
"A computational scheme for reasoning in dynamic probabilistic networks,"
in *Eighth Conference on Uncertainty in Artificial Intelligence*, 1992.

[2] K. C. Chang and W. Sun
"Performance modeling for dynamic Bayesian networks,"
in *Signal Processing, Sensor Fusion, and Target Recognition XIII*, Orlando, FL, 2004.

[3] J. Pearl
Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference,
San Francisco, CA: Morgan Kaufmann, 1988.

[4] K. J. Hintz and E. S. McVey
"Multi-process constrained estimation,"
IEEE Transactions on Systems, Man, and Cybernetics, vol. 21, no. 1, pp. 237–244, 1991.

[5] K. Hintz and I. Kadar
"Improving ORM utilizing implicit collaboration & context sensitive fusion,"
in *Proceedings International Conference on Information Fusion, Fusion2015*, Washington, DC, 2015.

[6] E. N. Ciftcioglu, A. Yener and M. J. Neely
"Maximizing quality of information from multiple sensor devices: the exploration vs exploitation tradeoff,"
IEEE Journal of Selected Topics in Signal Processing, vol. 7, no. 5, pp. 883–894, 2013.

[7] R. L. Winkler
"Uncertainty in probabilistic risk assessment,"
Reliability Engineering and System Safety, vol. 54, pp. 127–132, 1996.

[8] P. C. G. Costa, K. B. Laskey, E. Blasch and A.-L. Jousselme
"Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology,"
in *15th International Conference on Information Fusion, Fusion2012*, Singapore, Singapore, 2012.

[9] G. Shafer
A Mathematical Theory of Evidence,
Princeton: Princeton University Press, 1976.

[10] Oxford University Press
"OED Online,"
March 2018. [Online]. Available: http://www.oed.com/view/Entry/4795?redirectedFrom=aleatory. [Accessed 04 June 2018].

[11] Oxford University Press
"OED Online,"
March 2018. [Online]. Available: http://www.oed.com/view/Entry/63541?rskey=rMvb1N&result=1&isAdvanced=false. [Accessed 04 June 2018].

[12] A. Der Kiureghian and O. Ditlevsen
"Aleatory or epistemic? Does it matter?,"
Structural Safety, vol. 31, pp. 105–112, 2009.

[13] N. A. Abrahamson
"5: Aleatory Variability and Epistemic Uncertainty,"
[Online]. Available: http://www.ce.memphis.edu/7137/PDFs/Abrahamson/C05.pdf. [Accessed 20 January 2018].

[14] J. Pearl
Causality, Models, Reasoning, and Inference,
2nd Edition, New York: Cambridge University Press, 2009.

[15] A. Gelb
Applied Optimal Estimation,
Cambridge, Massachusetts: MIT Press, 1974.

[16] E. Mazor, A. Averbuch, Y. Bar-Shalom and J. Dayan
"Interacting multiple model methods in target tracking: a survey,"
IEEE Transactions on Aerospace and Electronic Systems, vol. 34, no. 1, pp. 103–123, 1998.

[17] E. Waltz
Knowledge Management in the Intelligence Enterprise,
Boston: Artech House, 2003.

[18] A. Singhal and C. Brown
"Dynamic Bayes Net Approach to Multimodal Sensor Fusion,"
in *Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, Pittsburgh, PA, 1997.

[19] K. J. Hintz and S. Darcy
"Valued situation information in IBSM,"
in *20th International Conference on Information Fusion, Fusion2017*, Xi'an, China.

[20] C. E. Shannon
"A Mathematical Theory of Communication,"
Bell System Technical Journal, vol. 27, pp. 623–657, 1948.

[21] V. Dragos
"An ontological analysis of uncertainty in soft data,"
in *16th International Conference on Information Fusion, Fusion2013*, Istanbul, Turkey, 2013.

[22] Oxford English Dictionary
"OED,"
Oxford University Press, 2017. [Online]. Available: http://www.oed.com. [Accessed 03 February 2018].

[23] H. J. Steinhauer, A. Karlsson and S. F. Andler
"Traceable uncertainty,"
in *16th International Conference on Information Fusion, Fusion 2013*, Istanbul, Turkey, 2013.

[24] D. M. Buede, J. A. Tatman and T. A. Bresnick
Introduction to Bayesian Networks, Monterey,
CA, 1998.

[25] Norsys Software Corp.
"Netica Application,"
2018. [Online]. Available: https://www.norsys.com/netica.html. [Accessed 06 June 2018].

[26] Norsys
"Sensitivity Equations, Entropy Reduction,"
[Online]. Available: https://www.norsys.com/WebHelp/NETICA/X_Sensitivity_Equations.htm#entropy. [Accessed 25 June 2018].

Kenneth J. Hintz (BSEE '67–MSEE '79–PhD '81) Dr. Hintz has been an Associate Professor in the Department of Electrical and Computer Engineering at George Mason University since 1987. He designed and established the Bachelor and Masters in Computer Engineering Degree Programs at GMU and teaches courses in sensor engineering, image processing, and computer engineering.

Dr. Hintz' current research interest is in Orchestrated Resource Management (ORM) and Information Based Sensor Management (IBSM) supported by the Naval Postgraduate School. He also developed a new method for pre-shot detection of barreled weapons based on his discovery of cavity induced modulation (CIM). Additional signal processing developments include syntactic pattern recognition applied to ground penetrating radar images for landmine detection and classification. His other research interests are in information-based real-time sensor management, X-ray, thermal, and visual, image processing, as well as algorithms for real time signal processing.

Before joining GMU, Dr. Hintz was with the Naval Surface Warfare Center, Dahlgren, VA, working in electronic warfare and radar signal processing where he designed and built the original AN/ULQ-16 pulse analyzer. Prior to working at NSWC, Dr. Hintz was with the U. S. Navy as a designated Naval Aviator stationed for 3 years in Rota, Spain flying Electronic Warfare Reconnaissance with Fleet Air Reconnaissance Squadron Two (VQ-2). During that time he became designated Electronic Warfare Aircraft Commander (EWAC) in both the EC-121 and EP-3E aircraft. Dr. Hintz holds 24 patents, 6 patents pending, is a Fellow of SPIE, a Senior Life Member of IEEE, and lead author on a book on Microcontrollers. He received his B.S. degree in Electrical Engineering from Purdue University, West Lafayette, Indiana in 1967 and his M.S. and Ph.D. degrees in Electrical Engineering from the University of Virginia in 1979 and 1981 respectively.

Steven Darcy is an analyst with Innovative Decisions Inc. He has extensive experience as an OR analyst and software developer. He has helped develop over a dozen significant computer models. He is a Ph.D. candidate in the Systems Engineering and Operations Research Department of George Mason University. He has a BA in Physics from Rutgers University and a MS in Operations Research from George Mason University.

INTERNATIONAL SOCIETY OF INFORMATION FUSION

ISIF Website: http://www.isif.org

SOCIETY VISION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied information fusion technologies.

SOCIETY MISSION

Advocate

To advance the profession of fusion technologies, propose approaches for solving real-world problems, recognize emerging technologies, and foster the transfer of information.

Serve

To serve its members and engineering, business, and scientific communities by providing high-quality information, educational products, and services.

Communicate

To create international communication forums and hold international conferences in countries that provide for interaction of members of fusion communities with each other, with those in other disciplines, and with those in industry and academia.

Educate

To promote undergraduate and graduate education related to information fusion technologies at universities around the world. Sponsor educational courses and tutorials at conferences.

Integrate

Integrate ideas from various approaches for information fusion, and look for common threads and themes– look for overall principles, rather than a multitude of point solutions. Serve as the central focus for coordinating the activities of world-wide information fusion related societies or organizations. Serve as a professional liaison to industry, academia, and government.

Disseminate

To propagate the ideas for integrated approaches to information fusion so that others can build on them in both industry and academia.

Call for Papers

The Journal of Advances in Information Fusion (JAIF) seeks original contributions in the technical areas of research related to information fusion. Authors are encouraged to submit their manuscripts for peer review http://isif.org/journal.

Call for Reviewers

The success of JAIF and its value to the research community is strongly dependent on the quality of its peer review process. Researchers in the technical areas related to information fusion are encouraged to register as a reviewer for JAIF at http://jaif.msubmit.net. Potential reviewers should notify via email the appropriate editors of their offer to serve as a reviewer.

Journal of Advances in Information Fusion Volume 13 Number 2 December 2018 ISSN 1557-6418