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

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

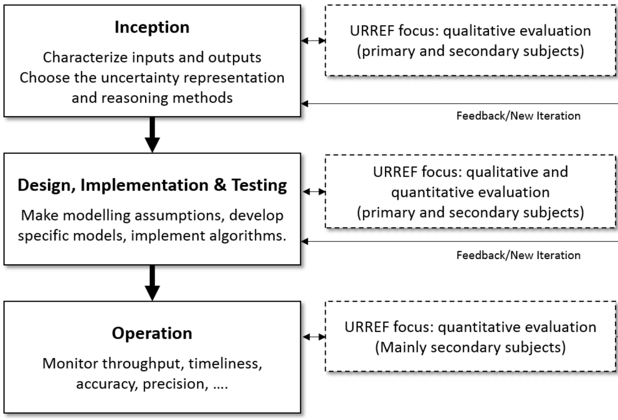

Fig. 1. URREF roles in a development life cycle [3] depicting the inception phase, the design implementation and testing phase, and the operation phase.

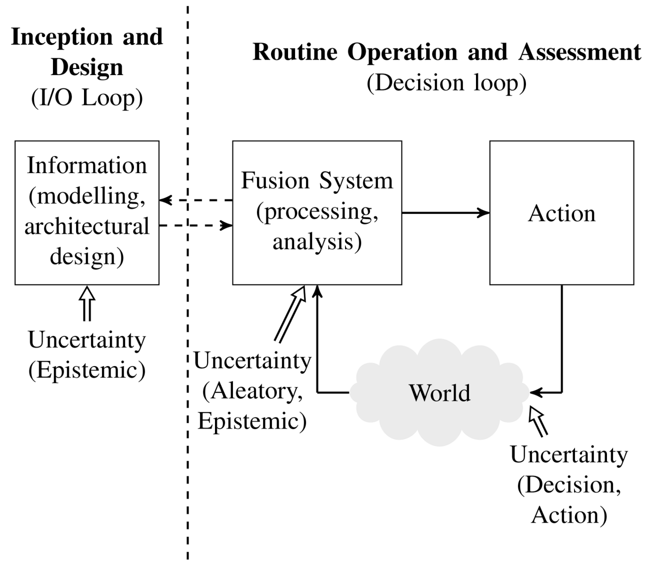

Fig. 2. The two main phases of a fusion system, namely the inception and design phase (input/output loop), and the routine operation phase (decision loop) are depicted. The double arrows depict where uncertainty enters the two phases, and the dashed arrows depict implementation and design refinement. 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-

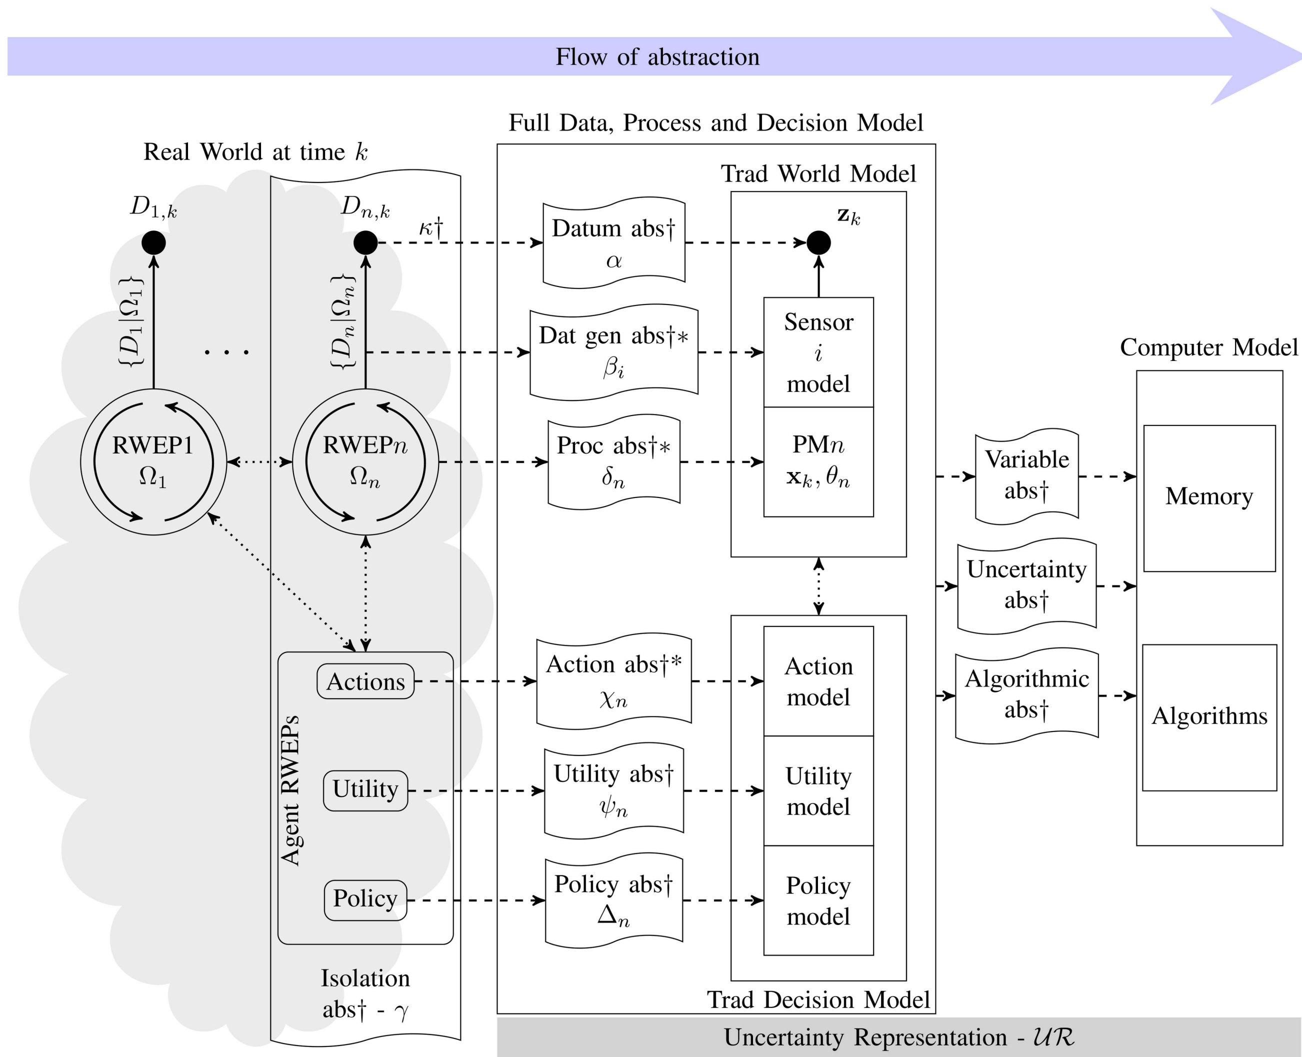

Fig. 3. The modelling (abstraction) of a fusion system making a measurement at time k is depicted. The principal components depicted are a) real world entities and processes (RWEPs), b) agents acting in the world (a specific type of RWEP), and c) models/abstractions of these RWEPs. Solid arrows indicate 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 *n*th 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

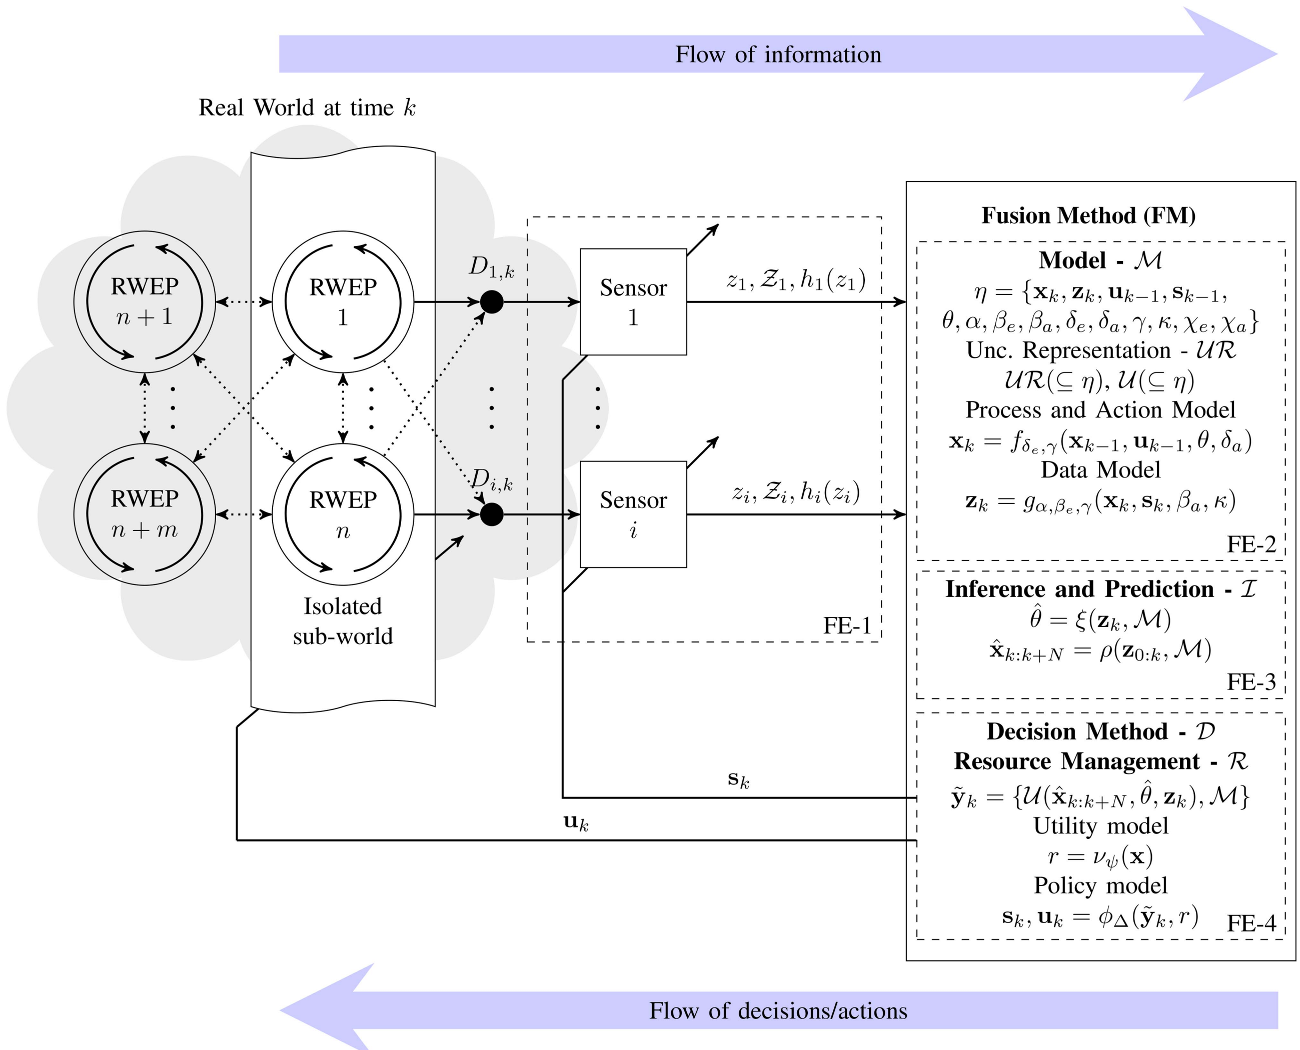

Fig. 4. The fusion decision (e.g., OODA) loop depicting the flow of information through sensors (observe) and the Fusion Method (FM) (orient), and the flow of decisions (decide) and actions (act) out of the decision method and resource management block. These actions in turn influence the real world. 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].

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.

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."

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).

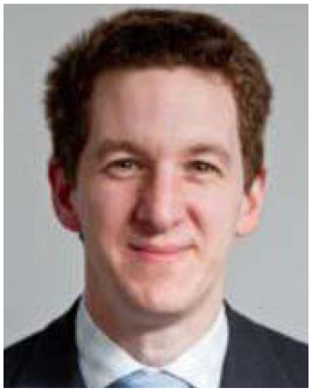

Simon Maskell is a Professor of Autonomous Systems at the University of Liverpool. His research interests centre on using Bayesian inference to address real-world challenges across a range of different applications. Often, doing so ends up demanding the development of advanced numerical Bayesian methods, 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.

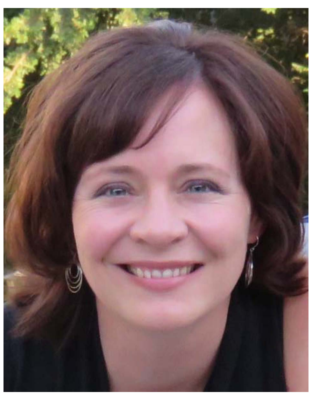

Alta de Waal currently holds a senior lecturer position in the Department of Statistics, University of Pretoria. Previously she was employed as a senior researcher at the Council for Scientific and Industrial Research (CSIR). She received her MSc (Mathematical Statistics) cum laude from the University 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.

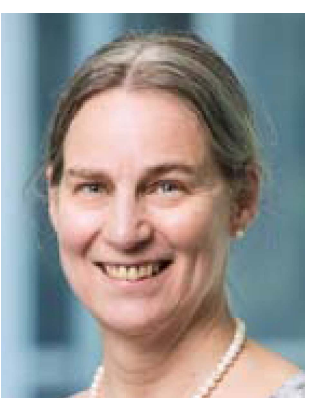

Kathryn Blackmond Laskey, Ph.D., is Professor of Systems Engineering and Operations Research at George Mason University and Associate Director of the Center of Excellence in Command, Control, Communications, Computing and Intelligence (C4I Center). She teaches and performs research on multisource information 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.

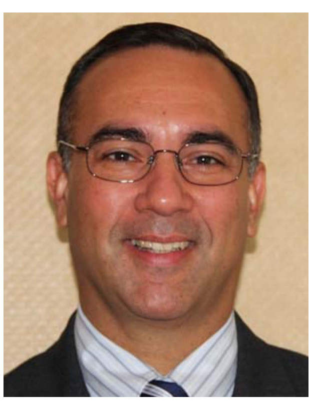

Paulo C. G. Costa is an Associate Professor of Systems Engineering and Operations Research at George Mason University, and Associate Director of the C4I & Cyber Center's Radio and Radar Engineering Laboratory.

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).