

Determining Model Correctness for Situations of Belief Fusion¹

Audun Jøsang
University of Oslo²
Norway
josang@mn.uio.no

Paulo C.G. Costa
George Mason University
USA
pcosta@c4i.gmu.edu

Erik Blasch
Air Force Research Lab
USA
eric.blash@gmail.com

Abstract—When analyzing hypotheses about specific situations of interest there is often a need to combine information from multiple sources. This principle belongs to information fusion in general, and is called belief fusion when the evidence is represented as belief functions. Because different situations can involve different forms of belief fusion, there is no single formal model that is suitable for analyzing every situation. It is therefore crucial to identify the most adequate fusion operator for modeling each class of situations to be analyzed. It can be challenging to determine the best belief fusion model for a specific situation, and there has been considerable confusion around this issue in the literature. In this paper we illustrate the importance of selecting a belief fusion model that adequately matches the situation to be analyzed, and propose a classification method for this purpose. A set of formal fusion operators is described to provide examples of specific models for classes of fusion situations.

I. INTRODUCTION

Mathematical modeling of practical situations enables us to draw more precise conclusions and to make more accurate predictions than what is possible by mere intuitionistic reasoning. Basic arithmetic operators such as addition, subtraction, multiplication and division represent simple mathematical models that have been applied for thousands of years, e.g. for accounting and engineering. In many cases it is trivial to determine which is the most correct model for analyzing a specific situation. However, there are situations where it is challenging to select an adequate formal model. Even if an operator produces correct or satisfactory results in some instances of a situation, it might not be an adequate operator for that situation in general.

For example, an adequate model for predicting the strength of a chain is the classic principle of the weakest link, meaning that the chain is only as strong as the weakest of all its links. A different situation is e.g. to determine the strength of a relay swimming team, for which an adequate model could be the average strength of each member of the team. Applying the weakest link model to assess the overall strength of the relay swimming team is an approximation that might give the correct result in some instances, but would produce unreliable predictions in general. Similarly, applying the average strength model for assessing the overall strength of the chain represents

an approximation that might produce satisfactory results in some instances, but would often produce a totally wrong result which could be fatal if life depended on it. These simple examples tell us that it is crucial to properly understand the situation at hand in order to find the most correct model for analyzing it.

In the domain of belief reasoning there has been considerable controversy around operators for belief fusion, especially related to Dempster's rule of combination [1]. The traditional interpretation of Dempster's rule is that it fuses separate argument beliefs from independent sources into a single belief. There are well known examples where Dempster's rule produces counterintuitive and clearly wrong results when interpreted in this way, especially in case of strong conflict between the input argument beliefs [2], but also in case of harmony between the input argument beliefs [3]. Motivated by this observation, numerous authors have proposed alternative methods for fusing beliefs [4]–[12]. These operators are not only formally different, they also model very different situations, but the authors often do not specify the type of situations they model. The fact that different situations require different rules and modeling assumptions is often ignored in the literature, and therefore represents a significant source of confusion.

The aim of this paper is to demonstrate that there are several different classes of belief fusion, and that a separate mathematical fusion operator is required for the formal modeling of each class. We describe characteristics that can be used to identify classes of fusion situations, and indicate operators that are suitable for modeling each type of situation. Finally, we argue that the main research question in this area is not about which operator is the correct belief fusion operator, but on how to define which operator is suitable for modeling which class of situations. To address these issues, we initially provide an overview of the belief fusion process, which is briefly described in Sec.II. Sec.III focuses on defining whether a given fusion model is correct, which is an important aspect enabling the definition of classes of fusion situations described in Sec.IV. Section V gives examples of formal fusion models. Sec.VI precedes our conclusions with a discussion on how we plan to validate the contributions of this paper and the future work needed to incorporate them into fusion systems.

¹In the proceedings of the 16th International Conference on Information Fusion (FUSION 2013), Istanbul, July, 2013.

²The work reported in this paper has been partially funded by UNIK.

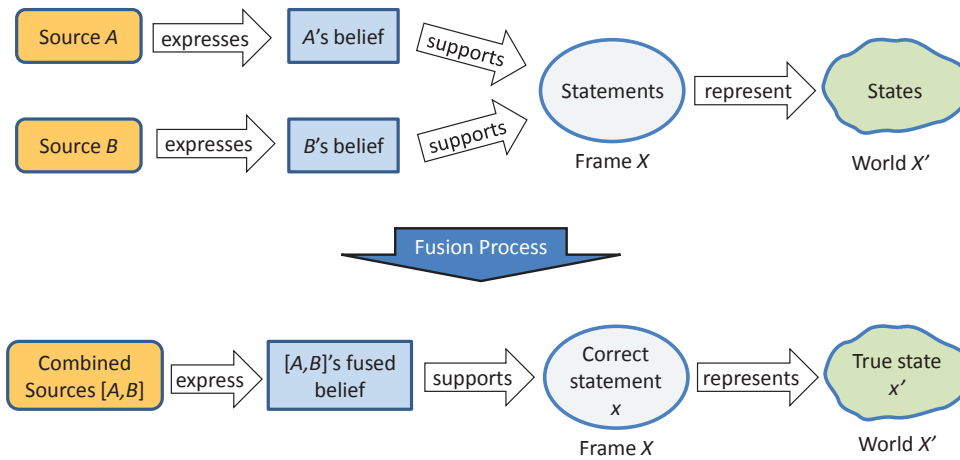


Figure 1. A conceptual model of the belief fusion process.

II. MODELING THE BELIEF FUSION PROCESS

The belief fusion process involves several related abstractions as illustrated in Fig.1. The belief fusion process focuses on a (part of the) perceived real world denoted X' where various possible states can be identified, and where it is assumed that the world is in one specific true state x' at any given time. A real world situation could e.g. be the possible medical diagnoses of a patient in a hospital, where it is assumed that the patient can only have a single specific diagnosis. The perceived world and its various states are formally represented as a frame X consisting of statements, where each statement is assumed to represent a specific state in the real world, and where there is a single correct statement x which represents the single true state x' . In the literature the term "state" is often used in the meaning of the elements in a frame, but we use the term "statement" in order to make the philosophical distinction between the perceived real world X' and its formal representation X . If the real world is inaccurately perceived by the analyst, e.g. when a patient in reality can have multiple medical diagnoses at the same time, then a frame assuming only a single possible diagnosis is clearly inadequate. There is thus no guarantee that a frame provides an adequate representation of the real world. Instead, it is up to the analyst to design the most appropriate frame for each situation to be analyzed. If the analyst thinks that multiple simultaneous medical diagnoses are possible, then it is more adequate to design a frame which reflects precisely that.

Belief masses are assigned to statements in the frame, and are assumed to provide support for assessing the correctness of those statements. There can be multiple sources of belief assignments to the same frame, so it is necessary to name each specific source as e.g. A , B , etc. The belief source can e.g. be a human analyst, a sensor, or the output from a previous analytic process. It is also possible to consider the same agent as multiple source when the agent expresses multiple beliefs e.g. in the form of multiple hints about a specific frame. As an example consider the case when a person gives hints about the value of a dice with the two propositions

i): "the value is a prime number" and ii): "the value is an even number" which can be expressed as beliefs and fused with an appropriate belief fusion operator to conclude that the correct value must be 2. In case of multiple belief sources it is assumed that each source represents partial evidence about the correct statement, so that the most complete evidence can be obtained by fusing the beliefs with an appropriate fusion operator. As a consequence it is assumed that fused belief from multiple sources provides stronger support for the correct statement x in the frame than the support that can be provided by beliefs from single sources alone.

So far we have used the concepts of true state and correct statement without clearly defining what they mean. The next section defines what we mean by truth of states and correctness of fusion results and fusion models.

III. CORRECTNESS AND CONSISTENCY CRITERIA FOR FUSION MODELS

We argue that meaningful and reliable belief fusion depends on the fusion operator's ability to produce correct results for the practical or hypothetical situation that is being analyzed. This calls for a definition of what it means for results to be correct.

Definition 1 (Correctness of Results.): In general the correctness of results produced by a model is the degree to which the results represent the true state of the real situation that is being modeled.

To clarify this definition it is useful to distinguish between three types of truth: 1) ground truth, 2) consensus truth and 3) subjective truth, as described below, where ground truth is the strongest and subjective truth is the weakest form of truth.

- **Ground truth** about a situation is the objectively observable state of the situation.
- **Consensus truth** about a situation is the state that is identified to be the actual state by a commonly shared opinion about the situation, or the state that is identified to be the actual state according to commonly accepted norms or standards.

- **Subjective truth** about a situation is the state identified to be actual state by the analyst's own opinion about the situation.

The term "*true state*" can also be used in the sense that the state is satisfactory or preferred. For example when a group of people wants to select a movie to watch at the cinema together it would seem strange to say that one specific movie is more true than another. However, in this case the term truth can be interpreted in the sense that one specific movie (the true state) can be considered to be the most satisfactory for all the group members to watch together.

Fusion models produce output results when used to analyze fusion situations. Three different types of result correctness emerge from the three different types of truth, where objective correctness is the strongest, and subjective correctness is the weakest.

- **Objective result correctness** is the degree to which the result represents the ground truth of a situation.
- **Consensus result correctness** is the degree to which the result represents the consensus truth of a situation.
- **Subjective result correctness** is the degree to which the result represents the subjective truth of a situation.

Depending on whether ground truth, consensus truth or subjective truth is available, the strongest form of correctness should be required for assessing the results. For example assume a weather forecast model with all its various input parameters and their complex relationships. Weather forecasts can be compared with the actual weather when the time of the forecast arrives a day or two later, so that it is reasonable to require objective correctness when assessing weather forecasting models.

The case of predicting global warming might seem similar to that of forecasting the weather, because models for global warming are also based on many different input parameters with complex relationships. Although predicting global warming to occur over the next 100 years can be objectively verified or refuted, the time scale makes it impossible to require objective correctness in the short term. Instead, practical assessment of model correctness must be based on consensus among experts. So with no ground truth as bench mark it is only possible to require consensus correctness in the short term. An paradoxical observation is that in 100 years (e.g. after year 2100) when ground truth about global warming predicted for the next 100 years finally becomes available there will probably no longer be any interest in assessing the correctness of the models used to make the predictions, and the individuals who designed the models will be long gone. Designers of global warming models will thus never be confronted with the ground truth about their models and predictions. Despite the lack of objective basis, consensus correctness as a criterion is often used for selecting specific models and for determining whether or not the results they produce shall be used in planning and decision making.

In situations where ground truth can not be observed and consensus truth is impossible to obtain, only subjective criteria

for truth can be used. Models for which subjective correctness criteria can be used are e.g. models for making personal decisions about which career path to follow or which partner to live with. In theory such decision are made based on multiple forms of evidence which must be fused to form an opinion. People normally do not use formal models for analyzing such situations, and instead use their intuition. Models assessed under subjective correctness criteria are often only used for practical decision making by an individual a small number of times during a lifetime, so not even statistical evidence can be obtained. However there are expert systems for e.g. career path choice and partner matching, in which case it is possible to determine statistically whether a particular model predicts "good" career choices and "happy" unions in the long term.

With regard to Definition 1 it is necessary to examine the case when it has only been observed once or a small number of times that results from a model represent the true state of a situation. Although a model produces correct results in some instances there might be other instances where the results are clearly wrong, in which case the model can not be considered to be correct in general. As long as only instances with correct results have been observed the analyst might erroneously think that the model is correct in general.

For example, assume a rather naïve analyst who misinterprets the situation of adding apples from two baskets, and erroneously thinks that it should be modeled with the product rule. Assume that the analyst tries a specific example with two apples in each basket, and computes the sum with the product rule, which results in 4 apples. When observing a real example of two baskets of two apples each it turns out that adding them together also produces 4 apples. This result could mistakenly be interpreted so that the model is correct simply because the result represents ground truth in this particular instance. It is questionable whether a model for analyzing a situation can be characterized as a correct model just because it produces results that by coincidence correspond to the truth of the situation. In order for a model to be correct it is natural to require that results produced by it are generally correct and not just by coincidence in specific instances of a situation. In order to distinguish between coincidentally correct results and generally correct results, it is necessary to also consider consistency which we define as follows.

Definition 2 (Model Correctness): A model is correct for a specific situation when it consistently produces correct results in all instances of the situation.

On a high level of abstraction a correct reasoning model according to Definition 2 must faithfully reflect the (class of) situations that are being modeled. A precise way of expressing this principle is that for a given a class of situations there is one correct model. Note that is possible to articulate three types of model correctness according to the three types of result correctness.

- **Objective model correctness** for a specific class of situations is the model's ability to consistently produce objectively correct results for all possible situations in the class.

- **Consensus model correctness** for a specific class of situations is the model's ability to consistently produce consensus correct results for all possible situations in the class.
- **Subjective model correctness** for a specific class of situations is the model's ability to consistently produce subjectively correct results for all possible situations in the class.

Depending on whether ground truth, consensus truth or subjective truth is available, the strongest form of model correctness should be required for practical analysis. Observing result correctness in one instance is not sufficient to conclude that a model is correct. It can be theoretically impossible to verify that all possible results are consistently correct, so proving that a model is correct in general can be challenging. On the other hand, if a single false result is observed it can be concluded that the model is incorrect for the situation. In such cases it might be meaningful to indicate the range of validity of the model which limits the range of input arguments or possibly the range of output results.

IV. CLASSES OF FUSION SITUATIONS

Situations of belief fusion involve belief arguments from multiple sources that must be fused in some way to produce a single belief argument. More specifically, the situation is characterized by a frame consisting of two or more statements, and a set of different belief arguments about these statements. It is assumed that each belief argument supports one or several statements. The purpose of belief fusion is to produce a new belief that identifies the most "correct" statement(s) in the frame. The meaning of most correct statement can also be that it is the most acceptable or most preferred statement.

Different beliefs can be fused in various ways, each having an impact on how the specific situation in evidence fusion is modeled. It is often challenging to determine the correct or the most appropriate fusion operator for a specific situation. One way of addressing this challenge is to categorize these specific situations according to their typical characteristics, which would then allow for determining which fusion operators are more adequate to each category. One of us (Josang), has studied these characteristics and was able to establish four distinct classes of fusion situations, which are described below.

- **Cumulative Belief Fusion** is when it is assumed that it is possible to collect an increasing amount of independent evidence by including more and more arguments, and that certainty about the most correct state increases with the amount of evidence accumulated. A typical case depicting this type of fusion is when one makes statistical observations about possible outcomes, i.e. the more observations the stronger the analyst's belief about the likelihood of each outcome. For example, a mobile network operator could observe the location of a subscriber over time, which will produce increasing certainty about the most frequent locations of that subscriber. However, the result would not necessarily be suitable for indicating the exact location of the subscriber at a specific time.

- **Averaging Belief Fusion** is when dependence between arguments is assumed. In other words, including more arguments does not mean that more evidence is supporting the conclusion. An example of this type of situation is when a jury tries to reach a verdict after having observed the court proceedings. Because the evidence is limited to what was presented to the court, the certainty about the verdict does not increase by having more jury members expressing their beliefs, since they were all exposed to the same evidence.
- **Constraining Belief Fusion** is when it is assumed that (a) each belief argument can dictate which states of the frame are the most correct, and (b) conflicting belief between the two sources is assigned to common states considered correct by both sources. In this fusion class, if two belief arguments express totally conflicting beliefs, i.e. no common state is considered correct by both sources, then they effectively veto each other's beliefs - which means that no state is correct. An example is when two persons try to agree on seeing a movie at the cinema. If their preferences include some common movies they can decide to see one of them. Yet, if their preferences do not have any movies in common then there is no solution, so the rational consequence is that they will not watch any movie together.
- **Consensus & Compromise Fusion** is when no single belief argument alone can dictate that specific states of the frame are the most correct. In this fusion class the analyst naturally wants to preserve shared beliefs from each argument, and in addition transform conflicting beliefs into new shared beliefs on union subsets. In this way consensus belief is preserved when it exists and compromise belief is formed when necessary. In case of totally conflicting beliefs on a binary frame, then the resulting fused belief is totally uncertain. An example is when analysing evidence about the Kennedy murder case, where the analyst collects statements from two witnesses. Assuming that both witnesses claim to know with some certainty that Oswald killed Kennedy, the consensus & compromise fusion would say the same, because there is a consensus. However, when assuming that witness 1 claims to know with certainty that Oswald killed Kennedy, and that witness 2 claims to know with certainty that Oswald did not kill Kennedy, then consensus & compromise fusion would return the result that it is totally uncertain whether Oswald killed Kennedy, because uncertainty is the best compromise in case of totally conflicting beliefs.

These examples illustrate the variety of fusion situations and the challenge of modeling 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 base station, in which case it seems natural to use constraining belief 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 constraining belief 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.

While defining classes of fusion situations helps in scoping the solution space, there is still the issue of determining which class a specific situation belongs to. The approach we propose for this classification problem is to specify a set of assumptions about a fusion situation, where each assumption can be judged to be either valid or invalid for the situation. In other words, we decompose the classification problem so it now becomes a matter of defining whether specific assumptions apply to the situation. The set of assumptions below can be used to determine which class a situation belongs to.

In order to select the correct fusion model the analyst must consider each assumption and determine whether it applies to the situation to be analyzed. The correct fusion model is identified as a function of the assumption that applies to the situation to be analyzed.

- 1) It is assumed that in case two belief arguments are in total conflict, then no statement can be correct.
 \implies Constraining Fusion
- 2) It is assumed that in case two belief arguments are in total conflict, then all statements supported by the arguments can be correct in a statistical sense.
 \implies Cumulative Fusion or Averaging Fusion
- 3) It is assumed that in case two belief arguments are in total conflict, then one of the supported statements is correct, but it is uncertain which of them is correct.
 \implies Consensus & Compromise Fusion
- 4) Idempotent belief fusion is assumed, i.e. a belief argument fused with itself should always produce the same belief argument.
 \implies Averaging Fusion or Consensus & Compromise Fusion
- 5) Idempotent belief fusion is not assumed, i.e. a partially uncertain belief argument fused with itself should produce a fusion result with less uncertainty.
 \implies Cumulative Fusion or Constraining Fusion

The set of assumptions above is not exhaustive, additional assumptions can be specified and used to identify the correct fusion model. There are also other classes of fusion situations than those described above, for which different assumptions should be used. By associating specific assumptions to each class of fusion situation it is easier for analysts to identify and

select the correct fusion model to be applied.

In the end, classifying a specific situation is a matter of assessing whether specific assumptions are valid or not valid for that situation.

V. FORMAL BELIEF FUSION MODELS

Subjective opinions [13] generalize traditional belief functions and lend themselves to simple mathematical expressions of fusion models. We therefore use the opinion representation for describing the various formal fusion models, but the expressions can easily be mapped to traditional belief functions.

A subjective opinion expresses belief about statements in a frame. Let X be a frame of cardinality k . An opinion distributes belief mass over the reduced powerset $\mathcal{R}(X)$ of cardinality κ . The reduced powerset $\mathcal{R}(X)$ is defined as:

$$\mathcal{R}(X) = \mathcal{P}(X) \setminus \{X, \emptyset\}, \quad (1)$$

where $\mathcal{P}(X) = 2^X$ denotes the powerset of X . All proper subsets of X are elements of $\mathcal{R}(X)$, but the frame $\{X\}$ and empty set $\{\emptyset\}$ are not elements of $\mathcal{R}(X)$.

Let \vec{b}_X be a belief vector over the elements of $\mathcal{R}(X)$, let u_X be the complementary uncertainty mass, and let \vec{a} be a base rate vector over X . Whenever relevant, a superscript such as A denotes the opinion owner. Then a subjective opinion ω_X^A is the composite function expressed as:

$$\omega_X^A = (\vec{b}_X, u_X, \vec{a}_X). \quad (2)$$

The attribute A is thus the belief source, and X is the target frame. The belief, uncertainty and base rate parameters satisfy the following additivity constraints.

$$\text{Belief additivity: } u_X + \sum_{x_i \in \mathcal{R}(X)} \vec{b}_X(x_i) = 1, \text{ where } x \in \mathcal{R}(X) \quad (3)$$

$$\text{Base rate additivity: } \sum_{i=1}^k \vec{a}_X(x_i) = 1, \text{ where } x \in X \quad (4)$$

The belief vector \vec{b}_X has $\kappa = (2^k - 2)$ parameters, whereas the base rate vector \vec{a}_X only has k parameters. The uncertainty parameter u_X 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, opinions over a frame of cardinality k only have $(2^k + k - 3)$ degrees of freedom. The probability projection of hyper opinions is the vector denoted as \vec{E}_X in Eq.(5).

$$\vec{E}_X(x_i) = \sum_{x_j \in \mathcal{R}(X)} \vec{a}_X(x_i/x_j) \vec{b}_X(x_j) + \vec{a}_X(x_i) u_X, \forall x_i \in \mathcal{R}(X) \quad (5)$$

where $\vec{a}_X(x_i/x_j)$ defined in Eq.(6) denotes relative base rate, i.e. the base rate of subset x_i relative to the base rate of (partially) overlapping subset x_j .

$$\vec{a}_X(x_i/x_j) = \frac{\vec{a}_X(x_i \cap x_j)}{\vec{a}_X(x_j)}, \quad \forall x_i, x_j \subset X. \quad (6)$$

General opinions are also called *hyper opinions*. A *multinomial opinion* is when belief mass only applies to singleton

statements in the frame. A *binomial opinion* is when the frame is binary. A *dogmatic opinion* is an opinion without uncertainty, i.e. where $u = 0$. A *vacuous opinion* is an opinion that only contains uncertainty, i.e. where $u = 1$.

Equivalent probabilistic representations of opinions, e.g. as a Beta pdf (probability density function) in case of binomial opinions, as a Dirichlet pdf in case of multinomial opinions, or as a hyper Dirichlet pdf in case of hyper opinions offer an alternative interpretation of subjective opinions in terms of traditional statistics [13].

A. Constraining Fusion

The belief constraint operator described here is an extension of Dempster's rule which in Dempster-Shafer belief theory is often presented as a method for fusing evidence from different sources [1] in order to identify the most likely hypothesis from the frame. Many authors have however demonstrated that Dempster's rule is not an appropriate operator for evidence fusion [2], and that it is better suited as a method for combining constraints [14], [15], which is also our view.

Assume two opinions ω_X^A and ω_X^B over the frame X of cardinality k with reduced powerset $\mathcal{R}(X)$ of cardinality κ . The superscripts A and B are attributes that identify the respective belief sources or belief owners. These two opinions can be mathematically merged using the belief constraint operator denoted as " \odot " which can be expressed as:

$$\text{Constraining Belief Fusion: } \omega_X^{A\&B} = \omega_X^A \odot \omega_X^B. \quad (7)$$

Belief source combination denoted with "&" thus corresponds to opinion fusion with " \odot ". The algebraic expression of the belief constraint operator for subjective opinions is described next.

$$\omega_X^{A\&B} = \quad (8)$$

$$\begin{cases} \vec{b}^{A\&B}(x_i) = \frac{\text{Har}(x_i)}{(1-\text{Con})}, & \forall x_i \in \mathcal{R}(X), x_i \neq \emptyset \\ u_X^{A\&B} = \frac{u_X^A u_X^B}{(1-\text{Con})} \\ \vec{a}^{A\&B}(x_i) = \frac{\vec{a}^A(x_i)(1-u_X^A) + \vec{a}^B(x_i)(1-u_X^B)}{2-u_X^A-u_X^B}, & \forall x_i \in X, x_i \neq \emptyset \end{cases}$$

The term $\text{Har}(x_i)$ represents the degree of *Harmony* (overlapping belief mass) on x_i . The term Con represents the degree of *Conflict* (non-overlapping belief mass) between ω_X^A and ω_X^B . These are defined below:

$$\text{Har}(x_i) = \vec{b}^A(x_i)u_X^B + \vec{b}^B(x_i)u_X^A + \sum_{y \cap z = x_i} \vec{b}^A(y)\vec{b}^B(z) \quad (9)$$

$$\text{Con} = \sum_{y \cap z = \emptyset} \vec{b}^A(y)\vec{b}^B(z) \quad (10)$$

The divisor $(1 - \text{Con})$ in Eq.(8) normalizes the derived belief mass; it ensures belief mass and uncertainty mass additivity.

The use of the belief constraint operator is mathematically possible only if ω^A and ω^B are not totally conflicting, i.e., if $\text{Con} \neq 1$.

The belief constraint 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 would require a generalization of Eq.(8) for multiple agents, i.e. for multiple input arguments, which is relatively trivial.

B. Cumulative Fusion

The cumulative fusion rule is equivalent to *a posteriori* updating of Dirichlet distributions. Its derivation is based on the bijective mapping between the belief and evidence notations described in [13].

The symbol " \diamond " denotes the cumulative fusion of two observers A and B into a single imaginary observer $A \diamond B$.

Let ω^A and ω^B be opinions respectively held by agents A and B over the same frame X of cardinality k with reduced powerset $\mathcal{R}(X)$ of cardinality κ . Let $\omega^{A \diamond B}$ be the opinion where:

$$\text{Case I: For } u^A \neq 0 \vee u^B \neq 0:$$

$$\begin{cases} b^{A \diamond B}(x_i) = \frac{b^A(x_i)u^B + b^B(x_i)u^A}{u^A + u^B - u^A u^B} \\ u^{A \diamond B} = \frac{u^A u^B}{u^A + u^B - u^A u^B} \end{cases} \quad (11)$$

$$\text{Case II: For } u^A = 0 \wedge u^B = 0:$$

$$\begin{cases} b^{A \diamond B}(x_i) = \gamma^A b^A(x_i) + \gamma^B b^B(x_i) \\ u^{A \diamond B} = 0 \end{cases} \quad (12)$$

$$\text{where } \begin{cases} \gamma^A = \lim_{\substack{u^A \rightarrow 0 \\ u^B \rightarrow 0}} \frac{u^B}{u^A + u^B} \\ \gamma^B = \lim_{\substack{u^A \rightarrow 0 \\ u^B \rightarrow 0}} \frac{u^A}{u^A + u^B} \end{cases}$$

Then $\omega^{A \diamond B}$ is the cumulatively fused opinion of ω^A and ω^B , representing the combination of independent opinions of A and B . By using the symbol ' \oplus ' to designate this belief operator, cumulative fusion is expressed as:

$$\text{Cumulative Belief Fusion: } \omega_X^{A \diamond B} = \omega_X^A \oplus \omega_X^B. \quad (13)$$

The cumulative fusion operator is commutative, associative and non-idempotent. In Eq.(12) the associativity depends on the preservation of relative weights of intermediate results through the weight variable γ , in which case the cumulative rule is equivalent to the weighted average of probabilities.

C. Averaging Fusion

Averaging belief fusion is derived from averaging arguments represented as evidence (not belief) through the bijective mapping between evidence and belief in subjective logic [13].

The symbol " \diamond " denotes the averaging fusion of two observers A and B into a single imaginary observer $A \diamond B$.

Let ω^A and ω^B be the respective opinions of agents A and B over the same frame X of cardinality k with reduced powerset $\mathcal{R}(X)$ of cardinality κ . Let $\omega^{A \diamond B}$ be the opinion such that:

Case I: For $u^A \neq 0 \vee u^B \neq 0$:

$$\begin{cases} b^{A \diamond B}(x_i) &= \frac{b^A(x_i)u^B + b^B(x_i)u^A}{u^A + u^B} \\ u^{A \diamond B} &= \frac{2u^A u^B}{u^A + u^B} \end{cases} \quad (14)$$

Case II: For $u^A = 0 \wedge u^B = 0$:

$$\begin{cases} b^{A \diamond B}(x_i) &= \gamma^A b^A(x_i) + \gamma^B b^B(x_i) \\ u^{A \diamond B} &= 0 \end{cases} \quad (15)$$

$$\text{where } \begin{cases} \gamma^A = \lim_{\substack{u^A \rightarrow 0 \\ u^B \rightarrow 0}} \frac{u^B}{u^A + u^B} \\ \gamma^B = \lim_{\substack{u^A \rightarrow 0 \\ u^B \rightarrow 0}} \frac{u^A}{u^A + u^B} \end{cases}$$

Then $\omega^{A \diamond B}$ is the averaging opinion of ω^A and ω^B , representing the combination of the possibly dependent opinions of A and B . By using the symbol ' \oplus ' to designate this belief operator, averaging fusion is expressed as:

$$\text{Averaging Belief Fusion: } \omega_X^{A \diamond B} = \omega_X^A \oplus \omega_X^B. \quad (16)$$

It can be verified that the averaging fusion rule is commutative and idempotent; but it is *not* associative.

D. Consensus & Compromise Fusion

CC-fusion (Consensus & Compromise) is a new fusion model specifically designed to satisfy the requirements of being idempotent, having a neutral element, and where conflicting beliefs result in compromise beliefs. This shows that it is possible to design fusion models to fit particular requirements.

Assume two opinions ω_X^A and ω_X^B over the frame X of cardinality k with reduced powerset $\mathcal{R}(X)$ of cardinality κ . The superscripts A and B are attributes that identify the respective belief sources or belief owners. These two opinions can be mathematically merged using the CC-fusion operator denoted as " \odot " which can be expressed as:

$$\text{Consensus \& Compromise Fusion: } \omega_X^{A \odot B} = \omega_X^A \odot \omega_X^B. \quad (17)$$

Belief source combination denoted with " \heartsuit " thus corresponds to opinion fusion with " \odot ". The CC-operator is

formally described next. It is a two-step operator where the consensus step comes first, and then the compromise step.

1) *Consensus Step*: The consensus step simply consists of determining shared belief mass between the two arguments, which is stored as the belief vector \vec{b}_X^{cons} expressed by Eq.(18).

$$\vec{b}_X^{\text{cons}}(x_i) = \min(\vec{b}_X^A(x_i), \vec{b}_X^B(x_i)). \quad (18)$$

The sum of consensus belief denoted b_X^{cons} is expressed as:

$$b_X^{\text{cons}} = \sum_{x_i \in \mathcal{R}(X)} \vec{b}_X^{\text{cons}}(x_i) \quad (19)$$

The residue belief masses of the arguments are:

$$\begin{cases} \vec{b}_X^{\text{resA}}(x_i) &= \vec{b}_X^A(x_i) - \vec{b}_X^{\text{cons}}(x_i) \\ \vec{b}_X^{\text{resB}}(x_i) &= \vec{b}_X^B(x_i) - \vec{b}_X^{\text{cons}}(x_i) \end{cases} \quad (20)$$

2) *Compromise Step*: The compromise step redistributes conflicting residue belief mass to produce compromise belief mass, stored in \vec{b}_X^{comp} expressed by Eq.(21).

$$\begin{aligned} \vec{b}_X^{\text{comp}}(x_i) &= \vec{b}_X^{\text{resA}}(x_i)u_X^B + \vec{b}_X^{\text{resB}}(x_i)u_X^A \\ &+ \sum_{\{y \cap z\} = x_i} \vec{a}_X(y/z) \vec{a}_X(z/y) \vec{b}_X^{\text{resA}}(y) \vec{b}_X^{\text{resB}}(z) \\ &+ \sum_{\substack{\{y \cup z\} = x_i \\ \{y \cap z\} \neq \emptyset}} (1 - \vec{a}_X(y/z) \vec{a}_X(z/y)) \vec{b}_X^{\text{resA}}(y) \vec{b}_X^{\text{resB}}(z) \\ &+ \sum_{\substack{\{y \cup z\} = x_i \\ \{y \cap z\} = \emptyset}} \vec{b}_X^{\text{resA}}(y) \vec{b}_X^{\text{resB}}(z), \quad \text{where } x_i \in \mathcal{P}(X). \end{aligned} \quad (21)$$

The preliminary uncertainty u_X^{pre} is computed as:

$$u_X^{\text{pre}} = u_X^A u_X^B. \quad (22)$$

The sum of compromise belief denoted b_X^{comp} is:

$$b_X^{\text{comp}} = \sum_{x_i \in \mathcal{P}(X)} \vec{b}_X^{\text{comp}}(x_i). \quad (23)$$

In general $b_X^{\text{cons}} + b_X^{\text{comp}} + u_X^{\text{pre}} < 1$, so normalisation of \vec{b}_X^{comp} is required. The normalisation factor denoted f_{norm} is:

$$f_{\text{norm}} = \frac{1 - (b_X^{\text{cons}} + u_X^{\text{pre}})}{b_X^{\text{comp}}}. \quad (24)$$

Because belief on X is uncertainty, the fused uncertainty is:

$$u_X^{A \odot B} = u_X^{\text{pre}} + f_{\text{norm}} \vec{b}_X^{\text{comp}}(X). \quad (25)$$

After computing the fused uncertainty the compromise belief mass on X must be set to zero, i.e.

$$\vec{b}_X^{\text{comp}}(X) = 0. \quad (26)$$

After normalisation the resulting CC-fused belief is:

$$\vec{b}_X^{A \odot B}(x_i) = \vec{b}_X^{\text{cons}}(x_i) + f_{\text{norm}} \vec{b}_X^{\text{comp}}(x_i), \quad \forall x_i \in \mathcal{R}(X). \quad (27)$$

The CC-operator is commutative, idempotent and semi-associative, with the vacuous opinion as neutral element. Semi-associativity means that 3 or more arguments must first be combined together in the Consensus Step, and then together again in the Compromise Step before normalisation.

VI. DISCUSSION AND FUTURE WORK

The approach laid out above is aimed at improving the process of modeling belief fusion problems.

As part of our work we participate in ETURWG¹, a working group of the International Society for Information Fusion (ISIF) devoted to the development of URREF, a framework for evaluation of uncertainty representation and reasoning² [16]. URREF includes use cases, associated data sets, and an ontology on uncertainty representation and reasoning [17].

The problem of uncertainty in belief fusion has in the last few years received much attention from the information fusion community [18]. We have designed our approach to be not only consistent with URREF, but also compatible with, and agnostic to, the various approaches for uncertainty representation discussed in the ETURWG forum.

Among the main issues we still need to address is tool support for the classification process, to help the analyst in selecting the most optimal fusion model. At this point, we are considering whether to create a specific ontology for the belief function process, or to simply add the main concepts and relationships to the URREF ontology. The latter is appealing since it would leverage the work already done and the support from ISIF, while the former would provide us with more flexibility to define the specific aspects involved in answering the propositions.

Part of our future work also includes defining the formalism behind our ontology. More specifically, most languages for expressing ontologies, such as the most popular variant of the W3C Recommendation OWL [19], are based on Description Logics [20]. Therefore, they do not have standardized support for uncertainty reasoning, a major aspect in our research. At this point, we are considering to work with probabilistic ontologies written in PR-OWL [21], [22], which extend traditional ontologies to capture both domain semantics and associated uncertainty about the domain

VII. CONCLUSIONS

The term belief fusion is vague in the sense that it can mean different things, making it very difficult for an analyst to define how best to address a belief fusion problem. The main objective of this work is to provide a means for analysts to define the most adequate fusion operator for the situation at hand. This paper brings two major contributions towards this objective. The first is a taxonomy of fusion situations based on their main characteristics, which synthesizes the problem to a limited number of classes. The second is a straightforward approach to the complex problem of assigning any specific situation to a given class, based on a set of assumptions the analyst could easily test against the situation to be classified. Taken together, these contributions have the potential to improve the way analysts address the fusion problem, both by standardizing the process and making it less subjective. We are currently leveraging the work on the ISIF ETURWG group for validation of our results.

¹Evaluation of Techniques for Uncertainty Representation

²URREF is available at <http://eturwg.c4i.gmu.edu>

REFERENCES

- [1] G. Shafer, *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [2] L. Zadeh, "Review of Shafer's A Mathematical Theory of Evidence," *AI Magazine*, vol. 5, pp. 81–83, 1984.
- [3] J. Dezert, P. Wang, and A. Tchamova, "On the Validity of Dempster Shafer Theory," in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.
- [4] S. Destercke and D. Dubois, "Idempotent merging of belief functions: Extending the minimum rule of possibility theory," in *Workshop on the Theory on Belief Functions (WTBF 2010)*, T. Denœx, Ed., Brest, April 2010.
- [5] 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.
- [6] A. Jøsang, "The Consensus Operator for Combining Beliefs," *Artificial Intelligence*, vol. 142, no. 1–2, pp. 157–170, October 2002.
- [7] E. Lefevre, O. Colot, and P. Vannooenbergh, "Belief Functions Combination and Conflict Management," *Information Fusion*, vol. 3, no. 2, pp. 149–162, June 2002.
- [8] M. Daniel, "Associativity in Combination of Belief Functions," in *Proceedings of 5th Workshop on Uncertainty Processing. - Praha, Edicni oddeleni VSE 2000*. Springer, 2000, pp. 41–54.
- [9] C. K. Murphy, "Combining belief functions when evidence conflicts," *Decision Support Systems*, vol. 29, pp. 1–9, 2000.
- [10] P. Smets, "The Combination of Evidence in the Transferable Belief Model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 5, pp. 447–458, 1990.
- [11] D. Dubois and H. Prade, "Representation and combination of uncertainty with belief functions and possibility measures," *Comput. Intell.*, vol. 4, pp. 244–264, 1988.
- [12] R. Yager, "On the Dempster-Shafer framework and new combination rules," *Information Sciences*, vol. 41, pp. 93–137, 1987.
- [13] 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)*, Singapore, July 2012.
- [14] A. Jøsang, "Multi-Agent Preference Combination using Subjective Logic," in *International Workshop on Preferences and Soft Constraints (Soft'11)*, Perugia, Italy, 2011.
- [15] A. Jøsang and S. Pope, "Dempster's Rule as Seen by Little Colored Balls," *Computational Intelligence*, vol. 28, no. 4, November 2012.
- [16] P. C. G. Costa, R. N. Carvalho, K. B. Laskey, and C. Y. Park, "Evaluating uncertainty representation and reasoning in HLF systems," in *Proceedings of the 14th International Conference on Information Fusion (FUSION 2011)*, Chicago, July 2011.
- [17] 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.
- [18] 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.
- [19] P. F. Patel-Schneider, P. Hayes, and I. Horrocks, "OWL web ontology language semantics and abstract syntax," <http://www.w3.org/TR/owl-semantics/>, Feb. 2004, W3C Recommendation. [Online]. Available: <http://www.w3.org/TR/owl-semantics/>
- [20] F. Baader, I. Horrocks, and U. Sattler, "Description logics as ontology languages for the semantic web," in *Mechanizing Mathematical Reasoning*. Springer Berlin Heidelberg, 2005, pp. 228–248. [Online]. Available: <http://www.springerlink.com/content/axh20n8134bc3ecb>
- [21] P. C. G. Costa, "Bayesian semantics for the semantic web," PhD Dissertation, George Mason University, Fairfax, VA, USA, Jul. 2005, Brazilian Air Force. [Online]. Available: <http://hdl.handle.net/1920/455>
- [22] P. C. G. Costa and K. B. Laskey, "PR-OWL: a framework for probabilistic ontologies," in *Proceedings of the International Conference on Formal Ontology in Information Systems (FOIS 2006)*, ser. Frontiers in Artificial Intelligence and Applications, B. Bennet and F. Christiane, Eds., vol. 150. Baltimore, MD, USA: IOS Press, Nov. 2006, pp. 237–249. [Online]. Available: <http://www.booksonline.iospress.nl/Content/View.aspx?pid=2210>