# Categories of belief fusion

AUDUN JØSANG

Belief Fusion consists of merging beliefs about a domain of interest from multiple separate sources. No single belief fusion method is adequate for all categories of situations, hence the challenge is to determine which belief fusion method is the most appropriate for a given situation. The conclusion to be drawn from this discussion is that the analyst must first understand the dynamics of the situation at hand in order to find the best fusion method for analysing it. The aim of this article is first to demonstrate that there are appropriate situations to use belief fusion, and that different mathematical fusion operators are required for the different situations. Secondly we propose criteria than can be applied to identify the various categories of fusion situations, and describe specific belief fusion operators that are suitable for modeling the fusion situations in each category.

Author's address: University of Oslo, Norway (E-mail: josang@mn.uio.no).

1. INTRODUCTION

When analyzing hypotheses about specific domains of interest there is often a need to combine evidence from multiple sources. This principle belongs to information fusion in general, and is called belief fusion when the evidence is represented as belief. It is important to realise that there is no single fusion method that is suitable for analyzing all situations of belief fusion. It is also quite challenging to determine the best belief fusion method for a specific situation, and there has been considerable confusion around this issue in the literature. It is therefore crucial to have a consistent method for categorising different situations of belief fusion, and to apply this method for selecting the most suitable belief-fusion operator for each category of situations.

Beliefs are represented as subjective opinions throughout this article. A subjective opinion generalises the traditional representations of belief functions by including a base rate distribution over the values of the domain variable. A domain of interest contains the possible hypotheses or states that the analyst is interested in, e.g. for identifying the hypothesis which correspond best with reality. A subjective opinion is denoted ω_X^C, where C represents the source of the opinion and X represents the variable of the opinion's object/target domain.

In general, the source of an opinion can be a human, or it can be a sensor which produces data which in turn can form the basis an opinion. Multiple separate sources, e.g. denoted $C_1, C_2, \ldots C_N$, can produce different and possibly conflicting opinions $\omega_X^{C_1}, \omega_X^{C_2}, \ldots \omega_X^{C_N}$ about the same variable X. In this situation, source fusion consists of merging the different sources into a single source that can be denoted $\diamond(C_1, C_2, \ldots C_N)$, and mathematically fusing their opinions into a single opinion denoted $\omega_X^{\diamond(C_1, C_2, \ldots C_N)}$ which then represents the opinion of the merged sources. The source merger function is here denoted by the symbol '$\diamond$', and the general belief-fusion principle is illustrated in Figure 1.
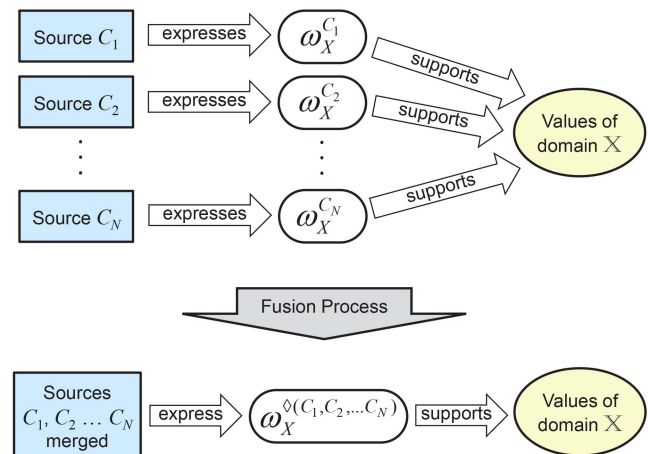
Fig. 1. Belief-fusion principle

Different belief fusion situations can vary significantly and semantically depending on the purpose and nature of the fusion process, and hence require different fusion operators. However, it can be challenging to identify the correct or most suitable fusion operator for a specific situation. In general, a given fusion operator is unsuitable when it produces wrong results in some instances of a situation, even if it produces correct results in most instances of the situation. A fusion operator should produce sound and intuitive results in all realistic instances of the situation to be analysed.

In order to see the importance of using the correct belief fusion method in a given fusion situation it is instructive to consider other situation types where the effect of applying the correct or incorrect formal model and method is more obvious. First, consider the situation of predicting the physical strength of a steel chain, where the classical and correct model is that of the weakest link. Then, consider the situation of determining the competitive strength of a relay swimming team, for which an adequate model is the average strength of each swimmer on the team, in terms of how fast each swimmer can swim.

Applying the weakest-link model (i.e. the slowest swimmer) to predict the overall speed of the relay swimming team is an approximation which might give a relatively good prediction in most instances of high-level swimming championships. However, it is obviously an incorrect model and would produce rather unreliable predictions if there are large variations in speed between the swimmers in a relay swimming team.

Similarly, applying the average strength model for assessing the physical strength of the steel chain represents an approximation which would produce relatively good strength predictions in most instances of high-quality steel chains where the link strength is highly uniform. However, it is obviously a very poor model which would be unreliable in general, and which could have fatal consequences if life depended on it.

These examples illustrate the inadequacy of anecdotal examples for determining whether the weakest-link model is suitable for predicting the strength of relay swimming teams. Similarly it is insufficient to simply use a few anecdotal examples to test whether the averaging principle is adequate for modelling the strength of steel chains. Without a clear understanding of the situation to be modelled, the analyst does not have a basis for selecting the correct and appropriate model. The selection of appropriate models might be obvious for the simple examples above, but it can be challenging to judge whether a fusion operator is suitable for a specific situation of belief fusion [1].

The conclusion to be drawn from this discussion is that the analyst must first understand the dynamics of the situation at hand in order to find the best model for analysing it.

The aim of this article is first to demonstrate that there can be many different categories of situations of belief fusion, and that different mathematical fusion operators are required for the computation of belief fusion in the different categories of belief-fusion situations. Secondly we propose criteria for identifying the various categories of fusion situations, and describe specific belief fusion operators that are suitable for belief fusion in each category.

This work forms part of the effort to define "Evaluation of Techniques for Uncertainty Representation" under the ETUR Working Group [2] where the UR-REF ontology is one of the reference documents [3], [4]. Previous work on defining categories for belief fusion is described in [1], [5]. The contribution of the current work is to generalise and define new operators for belief fusion, and to clarify the understanding of fusion categories. Belief fusion belongs to the domain of high-level fusion [6] in contrast to other types of low-level data fusion.

Section 2 describes a set of belief-fusion categories The criteria defined in Section 3 then describe how a given fusion situation can be understood and categorised. Section 6 describes corresponding fusion operators for the respective categories. Section 7 provides numerical examples to compare the different fusion operators, and Section 8 discusses the implications of the categories of belief fusion presented in this article.

2. CATEGORIES OF FUSION SITUATIONS

Situations of belief fusion take belief arguments from multiple sources through a fusion process to produce a single belief argument. More specifically, a fusion situation is characterised by a domain of two or more state values, and the various sources' different belief arguments about these values. The domain of state values can be interpreted as a set of competing hypotheses, where it is assumed that only one value/hypothesis is TRUE at any one time. Each belief argument can assign belief mass to one or several state values, which thereby represents support for those values in terms of which values are believed to be TRUE. The purpose of belief fusion is to produce a new belief argument that reflect the sources' collective set of belief arguments in the most fair or correct way. It is then assumed that the fused belief argument supports the most correct, acceptable or most preferred value, when seen from the perspective of the collective set of sources.

It is often challenging to determine the correct or the most appropriate fusion operator for a specific situation. Our approach of addressing this challenge is to define categories of similar situations according to their typical characteristics, which then allows to determine a suitable belief fusion operator for each category. Four
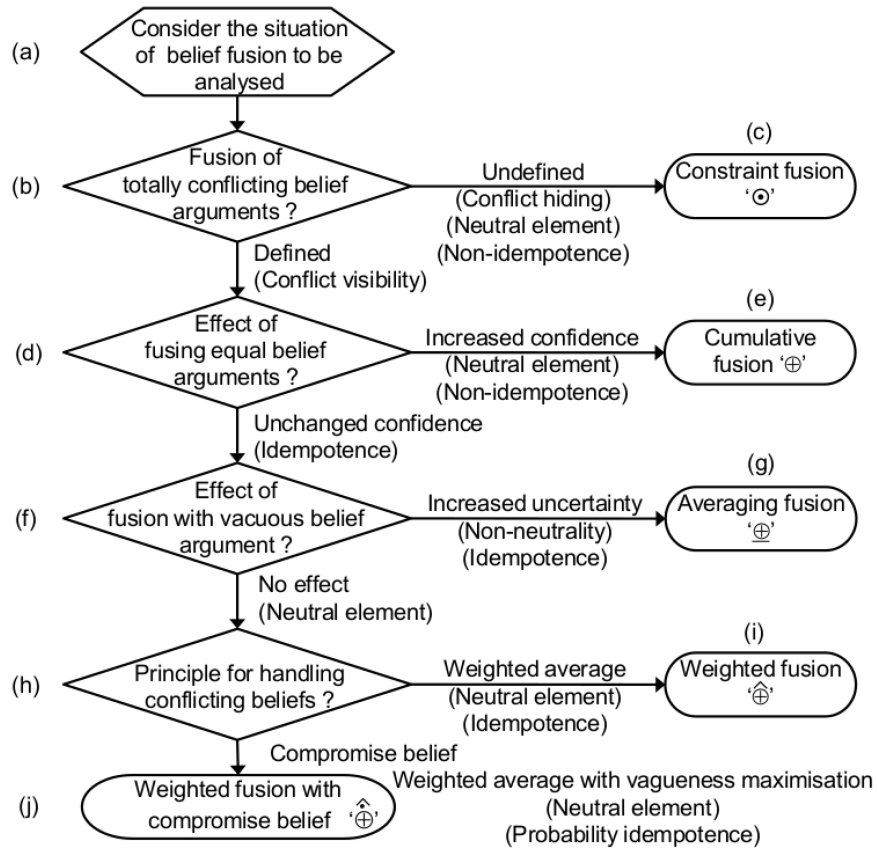
Fig. 2.   Procedure for selecting a suitable belief-fusion operator for each category

distinct categories as well as one hybrid category of fusion situations are described below.

- **Belief Constraint Fusion** (BCF) is suitable when assuming that: 1) belief arguments must not be wrong (sources are totally reliable), and 2) there is no compromise in case of totally conflicting arguments, hence the fusion result is not defined in that case. In some situations these properties are desirable. An example is when two persons try to agree on seeing a movie at the cinema. If their preferences share common movies they can decide to watch one of them. Yet, if their preferences have no movies in common then there is no solution, with the rational consequence they will not watch any movie together. BCF is described in Section 6.1.

- **Cumulative Belief Fusion** (CBF) is suitable when assuming that the amount of independent evidence increases by including more and more sources. For example, when different independent biometric sensors (e.g. fingerprint, voice, face) are being used to authenticate a person, the results from each sensor can be fused with CBF, which produces an opinion with decreasing uncertainty (increased assurance) about the identity of the person. CBF has the vacuous opinion as neutral element, but is not idempotent. CBF is described in Section 6.2. A modification of CBF is when it is assumed or desired that the fusion process produces uncertainty maximised opinions. It is then

possible to apply uncertainty maximisation after CBF, which is called CBF-UM for short. This could e.g. be when witnesses express their opinions about whether Oswald shot Kennedy, which when fused with CBF-UM produces an epistemic opinion about who shot him. CBF-UM is described in Section 6.5.

- **Averaging Belief Fusion** (ABF) is suitable when dependence between sources is assumed, so that including more sources does not necessarily add more evidence behind the fused belief, it just changes the average distribution of evidence. In case of equal belief arguments, the fused result should be the same, which means that idempotence is assumed. An example of this type of situation is when a jury tries to reach a verdict after having observed the court proceedings. It is also assumed that a vacuous belief argument does have an influence on the fused result, which means that ABF does not have a neutral element. This is interpreted in the sense that the source of the vacuous belief argument says: *"I do not see any evidence and therefore do not have any belief about this, and I want my vacuous argument belief to be reflected in the fused output belief"*. ABF is described in Section 6.3.

- **Weighted Belief Fusion** (WBF) is also suitable when dependence between sources is assumed, so that adding more sources does not necessarily add more evidence in total. Equal belief arguments should produce equal fused belief, meaning that idempotence

is assumed. However, it is assumed that a vacuous belief argument has no influence on the fused result, meaning that WBF does have a neutral element in the form of vacuous belief. This is interpreted in the sense that the source of a vacuous belief argument says: *"I do not see any evidence and therefore do not have any belief about this, and I will let the sources that do have evidence and belief about this determine the fused belief without me"*. An example of this type of situation is when experts (e.g. medical doctors) express multinomial opinions about a set of hypothesis (e.g. diagnoses). WBF is described in Section 6.4. In case of hyper-opinions WBF does not identify shared (vague) belief on overlapping (composite) values in the domain, and simply computes the weighted average.

- **Weighted Belief Fusion with Vagueness Maximisation** (WBF-VM) can be used when the analyst naturally wants to preserve shared beliefs from different sources, and to transform conflicting beliefs into vague belief. In this way shared belief is preserved when it exists, and compromise vague belief is formed when necessary. In the case of totally conflicting beliefs, then the resulting fused belief becomes vague. WBF-VM is probability-idempotent, commutative and has the vacuous belief argument as neutral element. Probability-idempotence means that the projected probability distribution is preserved when fusing equal opinions, but the fused opinion will in general have different vague belief. A situations where WBF-VM is suitable is when experts (e.g. medical doctors) express hyper-opinions about a set of hypothesis (e.g. diagnoses). WBF-VM takes into account shared (vague) belief on overlapping (composite) values, and is therefore suitable for preserving shared beliefs when fusing hyper-opinions. WBF-VM is described in Section 6.6.

The subtle differences between the fusion situations above illustrate the challenge of modelling them correctly. For instance, consider the task of determining the location of a mobile phone subscriber at a specific point in time by collecting location evidence from a base station, in which case it seems natural to use belief constraint fusion. If two adjacent base stations detect the subscriber, then the belief constraint operator can be used to locate the subscriber within the overlapping region of the respective radio cells. However, if two base stations far apart detect the subscriber at the same time, then the result of belief constraint fusion is not defined so there is no conclusion. With additional assumptions, it would still be reasonable to think that the subscriber is probably located in one of the two cells, but not which one in particular, and that the case needs further investigation because the inconsistent signals might be caused by an error in the system. Some method of trust revision [7] can be applied in this situation.

3. CRITERIA FOR IDENTIFYING FUSION CATEGORIES

While having multiple fusion categories can help in scoping the solution space, there is still the issue of determining which category a specific situation belongs to. In order to select the correct or most adequate fusion method the analyst must consider a set of assumptions about the fusion situation to be analysed and for each assumption judge whether it is applicable. The most adequate fusion method is then identified as a function of the set of assumptions that applies to the situation to be analysed. This procedure for identifying and selecting the most appropriate fusion operator is illustrated in Figure 2. The steps in the selection procedure are further described below.

(a) The analyst first needs a good understanding of the situation to be modelled in order to select the most suitable fusion operator. This includes being able to make the binary choices of (b), (d), (f) and (h) below.

(b) Shall it be possible to fuse totally conflicting beliefs?

(c) In case it is assumed that two totally conflicting belief arguments should leave no room for compromise, then BCF (Belief Constraint Fusion) is probably the most suitable operator. BCF is not defined in case of totally conflicting belief or preference arguments, which reflects the assumption that there is no compromise solution in case of total conflict.

(d) Is idempotence assumed, i.e. should two equal belief arguments produce the same output belief?

(e) In case idempotence is not assumed, then CBF (Cumulative Belief Fusion) is probably the most suitable operator. CBF is suitable when non-idempotent is assumed, meaning that equal belief arguments represent independent support for specific values of the variable, which thereby contribute to reducing the uncertainty in the output belief. In addition to being non-idempotent, CBF can handle totally conflicting opinions, as required for this category.

(f) Should a vacuous belief argument have any influence on the output fusion result?

(g) In case it is assumed that a vacuous belief arguments shall influence the output, then no neutral element exists, which indicates that ABF (Averaging Belief Fusion) is a suitable operator. ABF can be meaningful e.g. for making a survey of opinions where vacuity (lack of belief) in a belief argument shall be reflected as less confidence in the output fused belief.

(h) How should conflicting belief be handled?

(i) The simplest belief conflict management principle is to compute the weighted average of conflicting

belief mass. WBF (Weighted Belief Fusion) is suitable for fusing multinomial opinions, but less so for fusing hyper-opinions because the operator is blind to common belief between two vague belief arguments which assign belief mass to partially overlapping composite values.

(j) In case it is assumed that conflicting belief mass should be transformed into compromise (vague) belief then WBF-VM is suitable, i.e. it would be adequate to apply vagueness maximisation (VM) after the weighted belief fusion (WBF). In contrast to simple WBF, the post-processing with vagueness maximisation takes into account and reflects common belief aspects between different opinion arguments, which often better reflects human intuition.

It can be difficult to tell which category a specific situation belongs to. In addition, the choice of fusion operator can also be influenced by the type of fusion result the analyst wants to obtain, which e.g. could be to have an uncertainty-maximised or vagueness-maximised fused opinion.

The various belief fusion operators corresponding to each category in Figure 2 are described in Section 6 below. Before delving into the the formalism of belief fusion operators it is necessary to first describe the representation of subjective opinions and the corresponding Dirichlet PDF (Probability Density Function).

4. SUBJECTIVE OPINIONS

This section describes subjective opinions which represent beliefs over random variables in subjective logic.

In the formalism of subjective logic, a *domain* is a state space of values which can represent e.g. observable or hidden states, events, hypotheses or propositions [5]. A variable X associated with a domain $\mathbb{X}$ can take values $x \in \mathbb{X}$. A variable with an associated probability distribution over its domain is called a *random variable*.

The different values of the domain are assumed to be mutually exclusive and exhaustive, which means that the variable can take only one value at any time, and that all possible values of interest are included in the domain.

Available evidence may indicate that the variable takes a value in a given subset of values, but it is unclear which specific value in particular. For this reason it is meaningful to consider subsets as composite values, where the *hyperdomain* contains all the singletons as well as composites values. It is then possible to have a belief mass distribution over all these values, instead of only having a probability distributions over singleton values.

A subjective opinion distributes a *belief mass* over the values of the hyperdomain. The sum of the belief masses is less than or equal to 1, and is complemented with an *uncertainty mass* which reflects the opinion's confidence level. Subjective opinions also contain a *base rate* probability distribution expressing prior knowledge about the specific class of random variables, so that in case of significant uncertainty about a specific variable, the base rates provide a basis for default likelihoods. We give formal definitions of these concepts in what follows.

Let X be a variable over a domain $\mathbb{X} = \{x_1, x_2, \ldots, x_k\}$ of cardinality k, where x_i $(1 \leq i \leq k)$ represents a specific value from the domain. Let $\mathcal{P}(\mathbb{X})$ be the powerset of $\mathbb{X}$. The *hyperdomain* is the reduced powerset of $\mathbb{X}$, denoted by $\mathcal{R}(\mathbb{X})$, and defined as:

$$\mathcal{R}(\mathbb{X}) = \mathcal{P}(\mathbb{X}) \setminus \{\mathbb{X}, \emptyset\}. \tag{1}$$

All proper subsets of $\mathbb{X}$ are values of $\mathcal{R}(\mathbb{X})$, but $\mathbb{X}$ and $\emptyset$ are not, because they are not considered as possible observations to which belief mass can be assigned. Since $\mathbb{X}$ and $\emptyset$ are excluded the hyperdomain has cardinality $2^k - 2$. We use the same notation for the values of a domain and its hyperdomain, and say that X is a *hyper-variable* when it takes values from the hyperdomain.

Let A denote a source which can be a human, a sensor, etc. A *subjective opinion* ω_X^A of the source A on the variable X is a tuple

$$\omega_X^A = (\mathbf{b}_X^A, u_X^A, \mathbf{a}_X^A), \tag{2}$$

where $\mathbf{b}_X^A : \mathcal{R}(\mathbb{X}) \to [0,1]$ is a belief mass distribution, the parameter $u_X^A \in [0,1]$ is an uncertainty mass, and $\mathbf{a}_X^A : \mathbb{X} \to [0,1]$ is a base rate probability distribution satisfying the following additivity constrains:

$$u_X^A + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X^A(x) = 1, \tag{3}$$

$$\sum_{x \in \mathbb{X}} \mathbf{a}_X^A(x) = 1. \tag{4}$$

In the notation of the subjective opinion ω_X^A, the superscript is the source A, while the subscript is the object target variable X. An explicit source notation makes is possible to express the fact that different sources produce different opinions on the same variable. The source can be omitted in the opinion notation whenever the source is implicit or irrelevant, for example when there is only one source in the modelled situation.

The belief mass distribution $\mathbf{b}_X^A$ has $2^k - 2$ parameters, whereas the base rate distribution $\mathbf{a}_X^A$ only has k parameters. The uncertainty parameter u_X^A is a simple scalar. A general opinion thus contains $2^k + k - 1$ parameters. However, given that Eq. (3) and Eq. (4) remove one degree of freedom each, an opinion over a domain of cardinality k only has $2^k + k - 3$ degrees of freedom. Note that it is possible to express base rates over composite values as expressed by Eq. (5) below.

$$\mathbf{a}_X(x_i) = \sum_{\substack{x_j \in \mathbb{X} \\ x_j \subseteq x_i}} \mathbf{a}_X(x_j), \quad \forall x_i \in \mathcal{R}(\mathbb{X}). \tag{5}$$

A subjective opinion in which $u_X = 0$, i.e. an opinion without uncertainty, is called a *dogmatic opinion*. A
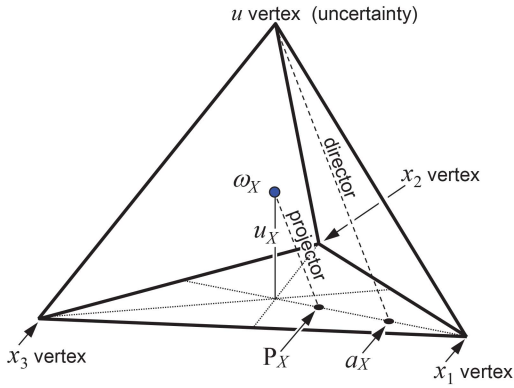
Fig. 3. Example trinomial opinion

dogmatic opinion for which $b_X(x) = 1$, for some x, is called an *absolute opinion*. In contrast, an opinion for which $u_X = 1$, and consequently, $b_X(x) = 0$, for every $x \in \mathcal{R}(\mathbb{X})$, i.e. an opinion with total uncertainty, is called a *vacuous opinion*.

Every subjective opinion 'projects' to a probability distribution $\mathbf{P}_X$ over $\mathbb{X}$ defined through the following function:

$$\mathbf{P}_X(x_i) = \sum_{x_j \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x_i \mid x_j)\mathbf{b}_X(x_j) + \mathbf{a}_X(x_i)u_X, \quad (6)$$

where $a_X(x_i \mid x_j)$ is the *relative base rate* of $x_i \in \mathbb{X}$ with respect to $x_j \in \mathcal{R}(\mathbb{X})$ defined as follows:

$$\mathbf{a}_X(x_i \mid x_j) = \frac{\mathbf{a}_X(x_i \cap x_j)}{\mathbf{a}_X(x_j)}, \quad (7)$$

where a_X is extended on $\mathcal{R}(\mathbb{X})$ additively. For the relative base rate to be always defined, it is enough to assume $a_X^A(x_i) > 0$, for every $x_i \in \mathbb{X}$. This means that everything we include in the domain has a non-zero probability of occurrence in general.

Binomial opinions apply to binary random variables where the belief mass is distributed over the two values in a binary domain. Multinomial opinions apply to random variables in n-ary domains, and where the belief mass is distributed over the values of the domain. Figure 3 visualises a ternary multinomial opinion as a point inside a tetrahedron.

General opinions, also called *hyper-opinions*, apply to hypervariables where belief mass is distributed over values in a hyperdomain which is the reduced powerset of an n-ary domain. Given a hyper-opinion, it is possible to project it onto a multinomial opinion. Assume a hyper opinion ω_X and let $\mathbf{b}_X$ be the belief mass distribution defined by the sum in Eq. (6), i.e.

$$\mathbf{b}_X(x) = \sum_{x' \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x \mid x')\mathbf{b}_X(x'), \quad (8)$$

then it is easy to check that $\mathbf{b}_X : \mathbb{X} \to [0,1]$, and that $\mathbf{b}_X$ together with u_X satisfies the additivity property in Eq. (3). The multinomial opinion denoted $\omega_X =$

$(\mathbf{b}_X, u_X, \mathbf{a}_X)$ is the projected opinion from the hyper-opinion of ω_X. By defining the unary operator $\downarrow$ to represent hyper-to-multinomial projection we can write:

Hyper-to-Multinomial Projection: $\omega_X = \downarrow(\omega_X)$.   (9)

From Eq. (6) and Eq. (8) we obtain $\mathbf{P}(\omega_X) = \mathbf{P}(\omega_X)$. This means that every hyper-opinion can be approximated with its projected multinomial opinion which by definition has the same projected probability distribution as the initial hyper-opinion.

A binomial opinion is equivalent to a Beta probability density function, a multinomial opinion is equivalent to a Dirichlet probability density function, and a hyper-opinion is equivalent to a Dirichlet hyper-probability density function [8]. Binomial opinions thus represent the simplest opinion type, which can be generalised to multinomial opinions, which in turn can be generalised to hyper-opinions. Simple visualisations for binomial and trinomial opinions are based on barycentric coordinate systems as illustrated in Figures 3 and 4.

Consider a domain $\mathbb{X}$ with its hyperdomain $\mathcal{R}(\mathbb{X})$ and powerset $\mathcal{P}(\mathbb{X})$. Recall that $\{\mathbb{X}\} \in \mathcal{P}(\mathbb{X})$. Let x denote a specific value of $\mathcal{R}(\mathbb{X})$ or of $\mathcal{P}(\mathbb{X})$.

In DST (Dempster-Shafer Theory) [9], the belief mass on value x is denoted $\mathbf{m}(x)$, and the belief mass distribution is called a *basic belief assignment* (bba). It is possible to define a direct bijective mapping between the bba of DST and the belief mass distribution and uncertainty mass of subjective opinions, as expressed by Eq. (10):

Mapping between the bba of DST and the belief/uncertainty masses of subjective opinions:
$$\begin{cases} \mathbf{m}(x) = \mathbf{b}_X(x), & \forall x \in \mathcal{R}(\mathbb{X}), \\ \mathbf{m}(\mathbb{X}) = u_X. \end{cases}$$
(10)

Technically, the bba of DST and the belief/uncertainty representation of subjective opinions are thus equivalent. Their interpretations however are different. Subjective opinions can not assign belief mass to the domain $\mathbb{X}$ itself. This interpretation corresponds to the (hyper-) Dirichlet model, where only observations of values of $\mathbb{X}$ (or $\mathcal{R}(\mathbb{X})$) are counted as evidence. The domain $\mathbb{X}$ itself can not be an observation in the (hyper-) Dirichlet model, and hence can not be counted as evidence. The difference between the belief representation in DST and the opinion representation in SL is that the DST belief representation does not take base rates into account. As a result the projected (called 'pignistic') probability in DST [9] can only be computed with default base rates equal to the relative cardinalities of (hyper) values in the domain, whereas the projected probability of subjective opinions can be computed with any base rate distribution.

5.   DIRICHLET REPRESENTATION OF BELIEFS

A hyper-opinion is equivalent to a Dirichlet HPDF (hyper probability density function) over a hyperdomain

$\mathcal{R}(\mathbb{X})$, according to the bijective mapping described in Section 5.2. For self-containment, we briefly outline the Dirichlet hypernomial model below, and refer to [10] for details about the Dirichlet model, and to [5] for details about the Dirichlet HPDF. The Dirichlet HPDF can be projected to a Hyper-Dirichlet PDF [11] which is useful for visualisation, but the Hyper-Dirichlet PDF is out of the scope of this presentation.

5.1. The Dirichlet Hypernomial Model

Multinomial probability density over a domain $\mathbb{X}$ of cardinality k is expressed by the k-dimensional Dirichlet PDF, where the special case of a probability density over a binary domain (where $k = 2$) is expressed by the Beta PDF. As a generalisation, hypernomial probability over the hyperdomain $\mathcal{R}(\mathbb{X})$ of cardinality $\kappa = 2^k - 2$ is expressed by the κ-dimensional Dirichlet HPDF [11].

The set of input arguments to the Dirichlet HPDF over $\mathcal{R}(\mathbb{X})$ then becomes a sequence of strength parameters of the κ possible (composite) values $x \in \mathcal{R}(\mathbb{X})$ represented as κ positive real numbers $\alpha_X(x_i)$, $i = 1 \ldots \kappa$, each corresponding to one of the possible values $x \in \mathcal{R}(\mathbb{X})$. Because this is a Dirichlet PDF over a hypervariable, it is called a Dirichlet Hyper-PDF, or Dirichlet HPDF for short.

DEFINITION 1 (**Dirichlet HPDF**). Let $\mathbb{X}$ be a domain consisting of k mutually disjoint values, where the corresponding hyperdomain $\mathcal{R}(\mathbb{X})$ has cardinality $\kappa = (2^k - 2)$. Let α_X represent the strength vector over the κ values $x \in \mathcal{R}(\mathbb{X})$. The hyper-probability distribution $\mathbf{p}_X^H$ and the strength vector α_X are both κ-dimensional. The Dirichlet hyper-probability density function over $\mathbf{p}_X^H$, called Dirichlet HPDF for short, is denoted $\text{Dir}_X^H(\mathbf{p}_X^H; \alpha_X)$, and is expressed as

$$\text{Dir}_X^H(\mathbf{p}_X^H; \alpha_X) = \frac{\Gamma\left(\sum_{x \in \mathcal{R}(\mathbb{X})} \alpha_X(x)\right)}{\prod_{x \in \mathcal{R}(\mathbb{X})} \Gamma(\alpha_X(x))} \prod_{x \in \mathcal{R}(\mathbb{X})} \mathbf{p}_X^H(x)^{(\alpha_X(x) - 1)},$$
(11)

where $\alpha_X(x) \geq 0$, with the restrictions that $p_X^H(x) \neq 0$ if $\alpha_X(x) < 1$.

The strength vector α_X represents the prior as well as the observation evidence, now assumed applicable to values $x \in \mathcal{R}(\mathbb{X})$.

Since the values of $\mathcal{R}(\mathbb{X})$ can contain multiple singletons from $\mathbb{X}$, a value of $\mathcal{R}(\mathbb{X})$ has a base rate equal to the sum of the base rates of the singletons it contains, as expressed by Eq. (5). The strength $\alpha_X(x)$ for each value $x \in \mathcal{R}(\mathbb{X})$ can then be expressed as

$$
\begin{array}{ll}
\forall x \in \mathcal{R}(\mathbb{X}), & \text{where } \begin{cases} \mathbf{r}_X(x) \geq 0, \\ \mathbf{a}_X(x) = \sum_{\substack{x_j \subseteq x \\ x_j \in X}} \mathbf{a}(x_j), \\ W = 2. \end{cases}
\end{array}
$$
(12)

The Dirichlet HPDF over a set of κ possible states $x_i \in \mathcal{R}(\mathbb{X})$ can thus be expressed as a function of the observation evidence $\mathbf{r}_X$ and the base rate distribution $\mathbf{a}_X(x)$, where $x \in \mathcal{R}(\mathbb{X})$. The constant W represents the non-informative prior weight which as a convention is set to $W = 2$ [5] (p.33). The superscript 'eH' in the notation Dir_X^{eH} indicates that it is expressed as a function of the evidence parameter vector $\mathbf{r}_X$ (not the strength parameter vector α_X), and that it is a Dirichlet HPDF (not a traditional Dirichlet PDF). The evidence-based Dirichlet HPDF is expressed as

$$\text{Dir}_X^{eH}(\mathbf{p}_X^H; \mathbf{r}_X, \mathbf{a}_X) = \frac{\Gamma\left(\sum_{x \in \mathcal{R}(\mathbb{X})} (\mathbf{r}_X(x) + \mathbf{a}_X(x)W)\right)}{\prod_{x \in \mathcal{R}(\mathbb{X})} \Gamma(\mathbf{r}_X(x) + \mathbf{a}_X(x)W)}$$
$$\times \prod_{x \in \mathcal{R}(\mathbb{X})} \mathbf{p}_X^H(x)^{(\mathbf{r}_X(x) + \mathbf{a}_X(x)W - 1)},$$
(13)

where $(\mathbf{r}_X(x) + \mathbf{a}_X(x)W) \geq 0$, with the restriction that $\mathbf{p}_X^H(x) \neq 0$ if $(\mathbf{r}_X(x) + \mathbf{a}_X(x)W) < 1$.

Dir_X^{eH} in Eq. (13) is the expression for probability density over hyper-probability distributions p_X^H, where each value $x \in \mathcal{R}(\mathbb{X})$ has a base rate according to Eq. (7).

Because a value $x_j \in \mathcal{R}(\mathbb{X})$ can be composite, the expected probability of any value $x \in \mathbb{X}$ is not only a function of the direct probability density on x, but also of the probability density of all other values $x_j \in \mathcal{R}(\mathbb{X})$ that contain x. More formally, the expected probability of $x \in \mathbb{X}$ results from the probability density of each $x_j \in \mathcal{R}(\mathbb{X})$ where $x \cap x_j \neq \emptyset$.

Given the Dirichlet HPDF of Eq. (13), the expected probability of any of the k values $x \in \mathbb{X}$ can be written as

$$\mathbf{E}_X(x) = \frac{\sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{a}_X(x \mid x_i) \mathbf{r}(x_i) + W \mathbf{a}_X(x)}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}(x_i)} \quad \forall x \in \mathbb{X}.$$
(14)

The mapping between the hyper-opinion and the Dirichlet HPDF is based on defining the expected probability distribution of a Dirichlet HPDF expressed by Eq. (14) to be equal to the projected probability of hyper-opinions expressed by Eq. (6), i.e. $\mathbf{E}_X = \mathbf{P}_X$.

5.2. Mapping Between a Hyper-opinion and a Dirichlet HPDF

Figure 4 is a screenshot of the visualisation of the mapping between binomial opinions $\omega_X^{C_1}$ and $\omega_X^{C_2}$ on the left and the corresponding Beta PDFs on the right.

In general, a hyper-opinion is equivalent to a Dirichlet HPDF according to the mapping defined below.

DEFINITION 2 (**Mapping: Hyper-opinion ↔ Dirichlet HPDF**). Let $\mathbb{X}$ be a domain consisting of k mutually disjoint values, where the corresponding hyperdomain $\mathcal{R}(\mathbb{X})$ has cardinality $\kappa = (2^k - 2)$, and let X be a hypervariable in $\mathcal{R}(\mathbb{X})$. Let ω_X be a hyper-opinion on X, and let $\text{Dir}_X^{eH}(\mathbf{p}_X^H; \mathbf{r}_X, \mathbf{a}_X)$ be a Dirichlet HPDF over the hyper-probability distribution $\mathbf{p}_X^H$. The hyper-opinion ω_X
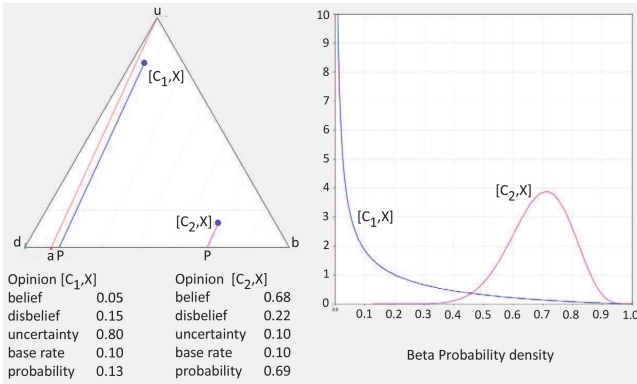
Fig. 4. Mapping opinions $\omega_X^{C_1}$ and $\omega_X^{C_2}$ to Beta PDFs

and the Dirichlet HPDF $\mathrm{Dir}_X^{\mathrm{eH}}(\mathbf{p}_X^{\mathrm{H}}; \mathbf{r}_X, \mathbf{a}_X)$ are equivalent through the following mapping:

$$\forall x \in \mathcal{R}(\mathbb{X}) \quad \begin{cases} \mathbf{b}_X(x) = \dfrac{\mathbf{r}_X(x)}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X(x_i)}, \\[2mm] u_X = \dfrac{W}{W + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X(x_i)}, \end{cases} \Leftrightarrow$$

$$\left(\begin{array}{cc} \text{For } u_X \neq 0: & \text{For } u_X = 0: \\[2mm] \begin{cases} \mathbf{r}_X(x) = \dfrac{W \mathbf{b}_X(x)}{u_X}, \\[2mm] 1 = u_X + \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X(x_i), \end{cases} & \begin{cases} \mathbf{r}_X(x) = \mathbf{b}_X(x) \cdot \infty, \\[2mm] 1 = \sum_{x_i \in \mathcal{R}(\mathbb{X})} \mathbf{b}_X(x_i). \end{cases} \end{array} \right)$$

(15)

The advantage of the Dirichlet HPDF is to provide an interpretation and equivalent representation of hyper-opinions.

This equivalence is very powerful because tools and methods used in Bayesian statistics can be applied to subjective opinions. In addition, the operators of subjective logic, such as conditional deduction, the subjective Bayes' theorem [12] and abduction, can be applied to statistical representations of data based on the Dirichlet model.

6. BELIEF FUSION OPERATORS

There are different categories of belief fusion situations, and each category requires its own operator for the computation of belief fusion [1]. In this article we focus on five different fusion categories, namely *constraint fusion*, *cumulative fusion*, *averaging fusion*, *weighted fusion* and *weighted fusion with vagueness* which are described below.

6.1. Belief Constraint Fusion

A typical application of belief theory in the literature is belief fusion with the classical Dempster's rule [9]. There has been considerable confusion and controversy around the adequacy of belief fusion operators, especially regarding Dempster's rule [13]. The confusion started with Zadeh's example from 1984 [14] where Dempster's rule is applied to a situation for which it is unsuitable and therefore produces erratic results. The

controversy followed when authors failed to realise that it is not a question of whether Dempster's rule is correct or wrong, but of recognising the type of situations for which Dempster's rule is suitable.

As an analogy of the controversy around Dempster's rule, imagine a world where the swim vest (analogy of Dempster's rule) has been invented as a safety device (analogy of a belief fusion operator). Then somebody demonstrates with an example that swim vests provide very poor protection in a car crash (analogy of Zadeh's example). Some researchers explain this by saying that swim vests perform poorly only in the case of high speed (analogy of high conflict) car crashes, and suggest to reduce the driving speed to make swim vests perform better. Other researchers propose the seat belt as an alternative safety device because it works well in car crashes, but this proposal is met with criticism by people who claim that seat belts provide poor protection in a sinking boat, in which case swim vests provide good protection. Many other safety devices are invented, and each device is promoted with an anecdotal example where it provides relatively good protection. In this confusing discussion nobody seems to understand that different safety hazards require different safety devices for protection, and that there is no single safety device that can provide adequate protection in all situations.

In an analogous fashion, the fact that different belief fusion situations require different belief fusion operators has often been ignored in the belief theory literature, and has been a significant source of confusion for many years [13]. There is nothing wrong with Dempster's rule *per se*; there are situations where it is perfectly appropriate, and there are situations where it is clearly inappropriate. No single belief fusion operator is suitable in every situation.

Dempster's rule is traditionally presented as a method for (cumulative) fusion of beliefs from different (independent) sources [9] with the purpose of identifying the most 'correct' hypothesis value from the domain. However, many authors have demonstrated that Dempster's rule is not an appropriate operator for this type of fusion [14]. Motivated by the apparent inconsistency of results produced by Dempster's rule numerous authors have proposed alternative belief fusion operators [15], [16], [17], [18], [19], [20], [21], [22], but the authors often fail to specify which type of situations they model.

We argue that Dempster's rule is better suited as a method for *belief constraint fusion* [13], [23], as shown in Figure 2. Situations of this type are e.g. when agents express different preferences with regard to a common decision that the agents must agree on [23] or when the analyst is presented with specific hints that are guaranteed to be valid [24], which is expressed by saying that the sources are 'reliable'.

It is common to see situations where people with different preferences try to agree on a single choice, or situations where evidence is presented as factual hints.

This must not be confused with fusion of belief from different agents to determine the most likely correct hypothesis or actual event, because the beliefs can not be taken as factual. Multi-agent preference combination assumes that each agent has already made up her mind, and then that they together want to determine the most acceptable decision or choice for all. Similarly, the fusion of hints assumes that the truth is known to the sources, but that they only reveal parts of the truth in the form of hints. Preferences and hints over a variable can be expressed in the form of subjective opinions. The constraint fusion operator of subjective logic can be applied as a method for merging preferences and hints from multiple sources into a single conclusion for the group of sources. This operator is expressive and flexible, and produces perfectly intuitive results. Preference can be represented as belief mass, and indifference can be represented as uncertainty mass. Positive and negative beliefs are considered as symmetric concepts, so they can be represented in the same way and combined using the same operator. Vacuous belief has no influence on the conclusion, and thereby represents the neutral element.

6.1.1. Method of Belief Constraint Fusion

The BCF (Belief Constraint Fusion) operator described next is an extension of Dempster's rule. The notation is also generalised to cover multiple sources, not only two sources.

DEFINITION 3 (**The Constraint Fusion Operator**). Assume the domain $\mathbb{X}$ and its hyperdomain $\mathcal{R}(\mathbb{X})$, and assume the hypervariable X which takes its values from $\mathcal{R}(\mathbb{X})$. Let $\mathbb{C} = \{C_1, C_2, \dots C_N\}$ denote a set of N independent sources. Let $C \in \mathbb{C}$ denote a specific source, and let ω_X^C denote its opinion about the variable X.

The respective opinions can be mathematically merged using the BCF (Belief Constraint Fusion) operator denoted '$\odot$' which can be expressed as

$$\omega_X^{\&(\mathbb{C})} = \underset{C \in \mathbb{C}}{\odot} (\omega_X^C)$$
$$= \omega_X^{C_1} \odot \omega_X^{C_2} \odot \cdots \omega_X^{C_N}. \qquad (16)$$

Source combination denoted '&' thus corresponds to belief fusion with '$\odot$'. The multi-source expression for BCF is given by Eq. (17):

$$\forall x \in \mathcal{R}(\mathbb{X}), \quad \omega_X^{\&(\mathbb{C})}:$$

$$\begin{cases} \mathbf{b}_X^{\&(\mathbb{C})}(x) & = \dfrac{\text{Har}(x)}{(1 - \text{Con})}, \\[2ex] u_X^{\&(\mathbb{C})} & = \dfrac{\prod_{C \in \mathbb{C}} u_X^C}{(1 - \text{Con})}, \\[2ex] \mathbf{a}_X^{\&(\mathbb{C})}(x) & = \dfrac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)(1 - u_X^C)}{N - \sum_{C \in \mathbb{C}} u_X^C}, \quad \exists u_X^C < 1, \\[2ex] \mathbf{a}_X^{\&(\mathbb{C})}(x) & = \dfrac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)}{N}, \qquad \forall u_X^C = 1. \end{cases}$$
$$(17)$$

The term $\text{Har}(x)$ represents the relative *harmony* between the constraint opinion ω_X^C (in terms of overlapping belief mass) on x. The term Con represents the relative *conflict* between constraints (in terms of non-overlapping belief mass) between the constraint opinions ω_X^C. DST's notation $\mathbf{m}(x)$ for belief-mass of $x \in \mathcal{P}(\mathbb{X})$ given by Eq. (10) gives the most compact notation for computing 'Har' and 'Con':

$$\text{Har}(x) = \sum_{\substack{\cap x^C = x \\ x^C \in \mathcal{P}(\mathbb{X})}} \prod_{C \in \mathbb{C}} \mathbf{m}_X^C(x^C), \qquad (18)$$

$$\text{Con} = \sum_{\substack{\cap x^C = \varnothing \\ x^C \in \mathcal{P}(\mathbb{X})}} \prod_{C \in \mathbb{C}} \mathbf{m}_X^C(x^C). \qquad (19)$$

The divisor $(1 - \text{Con})$ in Eq. (17) normalises the belief mass and uncertainty mass; i.e. it ensures their additivity. The application of the BCF operator is mathematically possible only if the constraint opinions ω_X^C are not totally conflicting, i.e., if $\text{Con} \neq 1$.

The BCF operator is commutative and non-idempotent. Associativity is preserved when the base rate is equal for all agents. Associativity in case of different base rates requires that all preference opinions be combined in a single operation which requires that Eq. (17) is applied for all input arguments in a single operation, which then represents semi-associativity.

The base rates of the two arguments are normally equal, but different base rates can be used in case of base rate disagreement between the sources, in which case the fused base rate distribution is the confidence-weighted average base rate.

Associativity in case of different base rates requires that all arguments opinions be combined in a single operation according to Definition 3. A totally indifferent opinion acts as the neutral element for constraint fusion, formally expressed as

$$\text{IF } (\omega_X^A \text{ is indifferent, i.e. } u_X^A = 1)$$
$$\text{THEN } (\omega_X^A \odot \omega_X^B = \omega_X^B). \qquad (20)$$

Having a neutral element in the form of the totally indifferent (i.e. vacuous) opinion can be useful when modelling situations of preference combination.

The rich format of subjective opinions makes it simple to express positive and negative preferences within the same framework, as well as indifference/uncertainty. Because preferences can be expressed over arbitrary subsets of the domain, this is in fact a multi-polar model for expressing and combining preferences. Even in the case of totally conflicting dogmatic opinions the belief constraint fusion operator produces meaningful results, namely that the preferences are incompatible. Examples in Sections 6.1.2–6.1.5 demonstrates the usefulness of this property.

TABLE 1
Example preferences and corresponding subjective opinions

Example Type	Domain & Opinion	Expression
"Ingredient x is mandatory"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Hard positive	Binomial opinion	$\omega_x : (1, 0, 0, \frac{1}{2})$
"Ingredient x is totally out of the question"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Hard negative	Binomial opinion	$\omega_x : (0, 1, 0, \frac{1}{2})$
"I prefer x with rating 0.3"	Binary domain	$\mathbb{X} = \{x, \bar{x}\}$
Quantitative	Binomial opinion	$\omega_x : (0.3, 0.7, 0.0, \frac{1}{2})$
"I prefer x or y, but z is also acceptable"	Ternary domain	$\Theta = \{x, y, z\}$
Qualitative	Trinomial opinion	$\omega_\Theta : (b(\{x, y\}) = 0.6,$ $b(z) = 0.3, u = 0.1,$ $a(x_1), a(x_2), a(x_3) = \frac{1}{3})$
"I like x, but I like y even more"	Binary domains	$\mathbb{X} = \{x, \bar{x}\}$ and $\mathbb{Y} = \{y, \bar{y}\}$
Positive rank	Binomial opinions	$\omega_x : (0.6, 0.3, 0.1, \frac{1}{2}),$ $\omega_y : (0.7, 0.2, 0.1, \frac{1}{2})$
"I don't like x, and I dislike y even more"	Binary domains	$\mathbb{X} = \{x, \bar{x}\}$ and $\mathbb{Y} = \{y, \bar{y}\}$
Negative rank	Binomial opinions	$\omega_x : (0.3, 0.6, 0.1, \frac{1}{2}),$ $\omega_y : (0.2, 0.7, 0.1, \frac{1}{2})$
"I'm indifferent about x, y and z"	Ternary domain	$\Theta = \{x, y, z\}$
Neutral	Trinomial opinion	$\omega_\Theta : (u_\Theta = 1.0,$ $a(x_1), a(x_2), a(x_3) = \frac{1}{3})$
"I'm indifferent but most people prefer x"	Ternary domain	$\Theta = \{x, y, z\}$
Neutral with bias	Trinomial opinion	$\omega_\Theta : (u_\Theta = 1.0, a(x) = 0.6,$ $a(y), a(z) = 0.2)$

TABLE 2
Fusion of film preferences

| | Belief preferences of: | | | Fusion results: | |
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	*A&B* $\omega_X^{A\&B}$	*A&B&C* $\omega_X^{A\&B\&C}$
$b(x_1)$	0.99	0.00	0.00	0.00	0.00
$b(x_2)$	0.01	0.01	0.00	1.00	1.00
$b(x_3)$	0.00	0.99	0.00	0.00	0.00
$b(\{x_2, x_3\})$	0.00	0.00	1.00	0.00	0.00

Default base rates are specified in all but the last example, which indicates total indifference, but with a bias that expresses the average preference in the population. Base rates are useful in many situations, such as for default reasoning. Base rates influence the computed results only in case of significant indifference or uncertainty.

6.1.3. Example: Going to the Cinema, First Attempt

Assume three friends, Alice, Bob and Clark, who want to see a film together at the cinema one evening, and that the only films showing are *Black Dust* (x_1), *Grey Matter* (x_2) and *White Powder* (x_3), represented as the ternary domain $X = \{x_1, x_2, x_3\}$. Assume that the friends express their preferences in the form of the opinions of Table 2.

Alice and Bob have strong and conflicting preferences. Clark, who strictly does not want to watch *Black Dust* (x_1), and who is indifferent about the two other films, is not sure whether he wants to come along, so Table 2 shows the results of applying the belief/preference constraint fusion operator, first without him, and then when including him in the party.

By applying belief constraint fusion, Alice and Bob conclude that the only film they are both interested in seeing is *Grey Matter* (x_2). Including Clark in the party does not change that result because he is indifferent to *Grey Matter* (x_2) and *White Powder* (x_3) anyway, he just does not want to watch *Black Dust* (x_1).

The belief mass values of Alice and Bob in the above example are in fact equal to those that Zadeh [14] used to demonstrate the unsuitability of Dempster's rule for fusing beliefs by showing how they produce counter-intuitive results. Zadeh's example describes a medical case where two medical doctors express their expert opinions about possible diagnoses, which typically should not have been modelled with Dempster's rule (BCF), but with the weighted belief fusion (WBF) operator [1], and possibly followed by vagueness maximisation (WBF-VM). In order to select the appropriate operator, it is crucial to fully understand the nature of the situation to be modelled. The failure to understand that Dempster's rule does not represent an operator for cumulative or averaging belief fusion, combined with the unavailability of the general cumulative, averaging and weighted fusion operators during that period

6.1.2. Expressing Preferences with Subjective Opinions

Preferences can be expressed as soft or hard constraints, qualitative or quantitative, ordered or partially ordered, etc. It is possible to specify a mapping between qualitative verbal tags and subjective opinions, which enables easy solicitation of preferences [25]. Table 1 describes examples of how preferences can be expressed.

All the preference types of Table 1 can be interpreted in terms of subjective opinions, and further combined by considering them as constraints expressed by different sources/agents. The examples which comprise two binary domains could equally well have been modelled with a quaternary product domain with a corresponding quatronomial product opinion. In fact, to compute product opinions over product domains is an alternative approach of simultaneously considering preferences over multiple variables.

	Belief preferences of:			Fusion results:	
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	$A\&B$ $\omega_X^{A\&B}$	$A\&B\&C$ $\omega_X^{A\&B\&C}$
$\mathbf{b}(X_1)$	0.98	0.00	0.00	0.490	0.000
$\mathbf{b}(x_2)$	0.01	0.01	0.00	0.015	0.029
$\mathbf{b}(x_3)$	0.00	0.98	0.00	0.490	0.961
$\mathbf{b}(\{x_2,x_3\})$	0.00	0.00	1.00	0.000	0.010
u	0.01	0.01	0.00	0.005	0.000
$\mathbf{a}(x_1)$	0.6	0.6	0.6	0.6	0.6
$\mathbf{a}(x_2) = a(x_3)$	0.2	0.2	0.2	0.2	0.2

TABLE 3
Fusion of film preferences with indifference and non-default base rates

	Belief preferences of:			Fusion results:	
	Alice ω_X^A	Bob ω_X^B	Clark ω_X^C	$A\&B$ $\omega_X^{A\&B}$	$A\&B\&C$ $\omega_X^{A\&B\&C}$
$b(x_1)$	1.00	0.00	0.00	Undefined	Undefined
$b(x_2)$	0.00	0.00	0.00	Undefined	Undefined
$b(x_3)$	0.00	1.00	0.00	Undefined	Undefined
$b(\{x_2,x_3\})$	0.00	0.00	1.00	Undefined	Undefined

TABLE 4
Combination of film preferences with hard and conflicting preferences

(1976 [9]–2013 [1]), has often led to inappropriate applications of Dempster's rule to cases of belief fusion [13]. However, when specifying the same numerical values as in [14] in a case of preference constraints such as in the example above, the belief constraint fusion operator (which is a simple extension of Dempster's rule) is the correct fusion operator which produces perfectly intuitive results.

6.1.4. Example: Going to the Cinema, Second Attempt

In this example Alice and Bob soften their preference with some indifference in the form of $u = 0.01$, as specified by Table 3. Clark has the same opinion as in the previous example, and is still not sure whether he wants to come along, so Table 3 shows both the results without him, and with his preference included.

Having some indifference in the preferences would mean that Alice and Bob should pick film *Black Dust* (x_1) or *White Powder* (x_3), because in both cases, one of them actually prefers one of the films, and the other finds it acceptable. Neither Alice nor Bob prefers *Grey Matter* (x_2), they only find it acceptable, so it would be a bad choice for both of them. When taking into consideration the base rates $a(x_1) = 0.6$ for *Black Dust* and $a(x_3) = 0.2$ for *White Powder*, the expected preference levels according to Eq. (6) are such that

$$\mathbf{P}_X^{A\&B}(x_1) > \mathbf{P}_X^{A\&B}(x_3). \qquad (21)$$

More precisely, the preference probabilities from Eq. (6) are

$$\mathbf{P}_X^{A\&B}(x_1) = 0.493, \qquad \mathbf{P}_X^{A\&B}(x_3) = 0.491. \qquad (22)$$

Because of the higher base rate, *Black Dust* (x_1) also has a higher expected preference than *White Powder* (x_3), so the rational choice would be to watch *Black Dust* (x_1).

However, when including Clark, who does not want to watch *Black Dust* (x_1), the base rates no longer dictate the result. In this case constraint fusion with Eq. (6) produces $\mathbf{P}^{A\&B\&C}(x_3) = 0.966$ so the obvious choice is to watch *White Powder* (x_3).

6.1.5. Example: Not Going to the Cinema

Assume now that Alice and Bob have totally conflicting preferences as specified in Table 4, i.e. Alice has a hard preference for *Black Dust* (x_1) and Bob has a hard preference for *White Powder* (x_3). As before, Clark still does not want to watch *Black Dust* (x_1), and is indifferent about the other two films.

In this case, the belief constraint fusion operator can not be applied because Eq. (17) involves a division by zero. The conclusion is that the friends will not go to the cinema to see a film together that evening. The test for detecting this situation is to observe $Con = 1$ in Eq. (19). It makes no difference to include Clark in the party, because a conflict can not be resolved by including additional preferences. However it would have been possible for Bob and Clark to watch *White Powder* (x_3) together without Alice.

6.2. Cumulative Belief Fusion

Cumulative Belief Fusion (CBF) is when it is assumed that the amount of evidence increases by including additional sources of independent evidence. An example of this type of situation is when different witnesses express their opinions about whether they saw the accused at the crime scene, and where their independent testimonies can be fused to produce an opinion about whether the accused really was there.

Assume a hyperdomain $\mathcal{R}(\mathbb{X})$ and a process where the outcome variable X takes values from $\mathcal{R}(\mathbb{X})$. Assume further that the outcome can be observed by different independent sources which can be expressed as $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$. Let $C \in \mathbb{C}$ denote a specific source, and let ω_X^C denote its opinion about the variable X. Assume that the sources in $\mathbb{C}$ produce independent opinions about the same variable X.

Observations can be vague, meaning that sometimes the sources observe an outcome which might be one of multiple possible singletons in $\mathbb{X}$, but the sources are unable to identify the observed outcome uniquely.

For example, assume that sources C_1 and C_2 observe coloured balls being picked from an urn, where the balls can have one of four colours: black, white, red or green. Assume further that the observer C_2 is colour-blind, which means that in poor light conditions he is unable see the difference between red and green balls, although

he can always tell the other colour combinations apart. As a result, his observations can be vague, meaning that sometimes he perceives a specific ball to be either red or green, but is unable to identify the ball's colour precisely. This corresponds to the situation where X is a hypervariable which can take composite values from $\mathcal{R}(\mathbb{X})$.

The symbol '$\diamond$' denotes the fusion of independent sources $C \in \mathbb{C}$ into a single cumulative merged source denoted $\diamond(\mathbb{C})$.

Let $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$ be a frame of N sources with the respective opinions $\omega_X^{C_1}, \omega_X^{C_2}, \ldots \omega_X^{C_N}$ over the same variable X. Let C denote a specific source $C \in \mathbb{C}$. The cumulative merger of all the sources in the source frame $\mathbb{C}$ is denoted $\diamond(\mathbb{C})$. The opinion $\omega_X^{\diamond(\mathbb{C})} \equiv (\mathbf{b}_X^{\diamond(\mathbb{C})}, u_X^{\diamond(\mathbb{C})}, \mathbf{a}_X^{\diamond(\mathbb{C})})$ is the cumulative fused opinion expressed as:

Case I: $u_X^C \neq 0, \forall C \in \mathbb{C}$:

$$\mathbf{b}_X^{\diamond(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}}(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j})}{\sum_{C \in \mathbb{C}}\left(\prod_{C_j \neq C} u_X^{C_j}\right) - (N-1)\prod_{C \in \mathbb{C}} u_X^C},$$

$$u_X^{\diamond(\mathbb{C})} = \frac{\prod_{C \in \mathbb{C}} u_X^C}{\sum_{C \in \mathbb{C}}\left(\prod_{C_j \neq C} u_X^{C_j}\right) - (N-1)\prod_{C \in \mathbb{C}} u_X^C},$$

$$\mathbf{a}_X^{\diamond(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}}\left(\mathbf{a}_X^C \prod_{C_j \neq C} u_X^{C_j}\right) - \sum_{C \in \mathbb{C}} \mathbf{a}_X^C \cdot \prod_{C \in \mathbb{C}} u_X^C}{\sum_{C \in \mathbb{C}}\left(\prod_{C_j \neq C} u_X^{C_j}\right) - N \prod_{C \in \mathbb{C}} u_X^C},$$

$$\mathbf{a}_X^{\diamond(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C}{N}, \quad \forall u_X^C = 1, \tag{23}$$

Case II: $\exists u_X^C = 0$, define $\mathbb{C}^{\text{dog}} = \{C \text{ where } u_X^C = 0\}$:

$$\mathbf{b}_X^{\diamond(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{b}_X^C(x),$$

$$u_X^{\diamond(\mathbb{C})} = 0,$$

$$\mathbf{a}_X^{\diamond(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{a}_X^C(x), \tag{24}$$

where

$$\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\text{dog}}} \to 0} \frac{u_X^C}{\sum_{C_j \in \mathbb{C}^{\text{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\text{dog}}. \tag{25}$$

The notation $u_X^{\mathbb{C}^{\text{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\text{dog}}$. The cumulative fused opinion $\omega_X^{\diamond(\mathbb{C})}$ results from fusing the respective opinions ω_X^C of the sources $C \in \mathbb{C}$. The symbol '$\oplus$' denotes the cumulative belief fusion operator, hence we define

$$\omega_X^{\diamond(\mathbb{C})} \equiv \bigoplus_{C \in \mathbb{C}} (\omega_X^C) \tag{26}$$

$$\equiv \omega_X^{C_1} \oplus \omega_X^{C_2} \oplus \cdots \omega_X^{C_N}. \tag{27}$$

It can be verified that the cumulative fusion operator is commutative, associative and non-idempotent. In Case II of Eq. (24), the associativity depends on preserving the relative weights of intermediate results with

the additional weight parameter γ. In this case, the cumulative fusion operator is equivalent to the weighted average of probabilities.

The argument base rate distributions are normally equal. When that is not the case the fused base rate distribution over X is specified to be the evidence-weighted average base rate.

In case of N dogmatic arguments ω_X^C where $C \in \mathbb{C}$ it can be assumed that the limits in Eq. (24) are defined as $\gamma_X^C = 1/N$.

6.2.1. Justification for the Cumulative Fusion Operator

The cumulative belief fusion operator of Eq. (23) is derived by mapping the argument belief opinions to evidence parameters through the bijective mapping of Eq. (15). Cumulative fusion of evidence opinions simply consists of summing up the evidence parameters, where the sum is mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the justification of the cumulative fusion operator of Eq. (23). A more detailed explanation is provided below.

Let the sources $C \in \mathbb{C}$ have respective belief opinions expressed as ω_X^C. The corresponding Dirichlet PDFs $\text{Dir}_X^e(\mathbf{p}_X; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence vectors $\mathbf{r}_X^C$.

The cumulative fusion of these evidence vectors consists of vector summation of $\mathbf{r}_X^C$ where $C \in \mathbb{C}$, expressed as

$$\mathbf{r}_X^{\diamond(\mathbb{C})} = \sum_{C \in \mathbb{C}} \mathbf{r}_X^C. \tag{28}$$

For each value $x \in \mathcal{R}(\mathbb{X})$ the evidence sum $\mathbf{r}_X^{\diamond(\mathbb{C})}(x)$ is

$$\mathbf{r}_X^{\diamond(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}} \mathbf{r}_X^C(x) \tag{29}$$

$$= \sum_{C \in \mathbb{C}} \frac{W \mathbf{b}_X^C(x)}{u_X^C} \tag{30}$$

$$= \frac{W \sum_{C \in \mathbb{C}}(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j})}{\prod_{C \in \mathbb{C}} u_X^C}. \tag{31}$$

The cumulative fused belief opinion $\omega_X^{\diamond(\mathbb{C})}$ of Eq. (23) results from mapping the fused evidence belief mass of Eq. (28) back to a belief opinion by applying the bijective mapping of Eq. (15).

$$\mathbf{b}_X^{\diamond(\mathbb{C})}(x) = \frac{\mathbf{r}_X^{\diamond(\mathbb{C})}(x)}{W + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\diamond(\mathbb{C})}(x)} \tag{32}$$

$$= \frac{\sum_{C \in \mathbb{C}}(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j})}{\prod_{C \in \mathbb{C}} u_X^C + \sum_{x \in \mathcal{R}(\mathbb{X})}\left(\sum_{C \in \mathbb{C}}(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j})\right)} \tag{33}$$

$$= \frac{\sum_{C \in \mathbb{C}} (\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j})}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right) - (N-1) \prod_{C \in \mathbb{C}} u_X^C},$$
$$\exists u_X^C \neq 0. \tag{34}$$

The transition from Eq. (32) to Eq. (33) results from inserting Eq. (31) into Eq. (32). The transition from Eq. (33) to Eq. (34) results from applying Eq. (3).

$$u_X^{\Diamond(\mathbb{C})} = \frac{W}{W + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\Diamond(\mathbb{C})}(x)} \tag{35}$$

$$= \frac{\prod_{C \in \mathbb{C}} u_X^C}{\prod_{C \in \mathbb{C}} u_X^C + \sum_{x \in \mathcal{R}(\mathbb{X})} \left(\sum_{C \in \mathbb{C}} (\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j}) \right)} \tag{36}$$

$$= \frac{\prod_{C \in \mathbb{C}} u_X^C}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right) - (N-1) \prod_{C \in \mathbb{C}} u_X^C},$$
$$\text{where } \exists u_X^C \neq 0. \tag{37}$$

The transition from Eq. (35) to Eq. (36) results from inserting Eq. (31) into Eq. (35). The transition from Eq. (36) to Eq. (37) results from applying Eq. (3).

6.3. Averaging Belief Fusion

Averaging Belief Fusion (ABF) is when dependence between sources is assumed. In other words, including more sources does not mean that more evidence is supporting the conclusion. An example of this type of situations is when a jury tries to reach a verdict after having observed the court proceedings. The assumption is that the correctness of the verdict does not increase as a function of the number of jury members, because the amount of evidence is fixed by what was presented in court.

Let $\mathbb{C}$ denote a group of N separate sources which can be expressed as $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$. Assume that the sources in $\mathbb{C}$ produce separate opinions based on the same evidence about the same variable, so their opinions are necessarily dependent. Still, their perceptions might be different, e.g. because their cognitive capabilities are different. For example, assume that sources C_1 and C_2 together observe the picking of coloured balls from an urn, where the balls can have one of four colours: black, white, red or green. Assume that observer C_2 is colour-blind, which means that sometimes he has trouble distinguishing between red and green balls, although he can always distinguish between the other colour combinations. Observer C_1 has perfect colour vision, and normally can always tell the correct colour when a ball is picked. As a result, when a red ball is picked, observer C_1 almost always identifies it as red, but observer C_2 identifies it as green relatively frequently. This can lead to C_1 and C_2 having different and conflicting opinions about the same variable, although their observations and opinions are totally dependent. The averaging belief fusion operator is perfectly suitable for this fusion situation.

Let $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$ be a frame of N sources with the respective opinions $\omega_X^{C_1}, \omega_X^{C_2}, \ldots \omega_X^{C_N}$ over the same variable X. Let C denote a specific source $C \in \mathbb{C}$. The averaging merger of all the sources in the source frame $\mathbb{C}$ is denoted $\underline{\Diamond}(\mathbb{C})$. The opinion $\omega_X^{\underline{\Diamond}(\mathbb{C})} \equiv (\mathbf{b}_X^{\underline{\Diamond}(\mathbb{C})}, u_X^{\underline{\Diamond}(\mathbb{C})}, \mathbf{a}_X^{\underline{\Diamond}(\mathbb{C})})$ is the averaging-fused opinion expressed as:

Case I: $u_X^C \neq 0, \forall C \in \mathbb{C}$:

$$\mathbf{b}_X^{\underline{\Diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right)}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right)},$$

$$u_X^{\underline{\Diamond}(\mathbb{C})} = \frac{N \prod_{C \in \mathbb{C}} u_X^C}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right)},$$

$$\mathbf{a}_X^{A \Diamond B}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)}{N}, \tag{38}$$

Case II: $\exists u_X^C = 0$, define $\mathbb{C}^{\text{dog}} = \{C \text{ where } u_X^C = 0\}$:

$$\mathbf{b}_X^{\underline{\Diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{b}_X^C(x),$$

$$u_X^{\underline{\Diamond}(\mathbb{C})} = 0,$$

$$\mathbf{a}_X^{\underline{\Diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{a}_X^C(x), \tag{39}$$

where

$$\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\text{dog}}} \to 0} \frac{u_X^C}{\sum_{C_j \in \mathbb{C}^{\text{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\text{dog}}. \tag{40}$$

The notation $u_X^{\mathbb{C}^{\text{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\text{dog}}$. The averaging-fused opinion $\omega_X^{\underline{\Diamond}(\mathbb{C})}$ results from averaging fusion of the respective opinions ω_X^C of the sources $C \in \mathbb{C}$. By using the symbol '$\underline{\oplus}$' to designate the averaging belief fusion operator, we define

$$\omega_X^{\underline{\Diamond}(\mathbb{C})} \equiv \underset{C \in \mathbb{C}}{\underline{\oplus}} (\omega_X^C). \tag{41}$$

It can be verified that the averaging belief fusion operator is commutative, idempotent, and non-associative. The non-associativity means that

$$(\omega_X^{C_1} \underline{\oplus} \omega_X^{C_2}) \underline{\oplus} \omega_X^{C_3} \neq \omega_X^{C_1} \underline{\oplus} (\omega_X^{C_2} \underline{\oplus} \omega_X^{C_3}). \tag{42}$$

However, semi-associativity exists as expressed by Eq. (41) where the argument order is irrelevant because all the arguments are fused in one single operation. The only way to apply averaging fusion to more than two arguments is thus by fusing all arguments in one operation as described in Eq. (38) and expressed by the notation of Eq. (41). For three argument sources, this is expressed as:

$$\omega_X^{\underline{\Diamond}(C_1, C_2, C_3)} \equiv \underline{\oplus}(\omega_X^{C_1}, \omega_X^{C_2}, \omega_X^{C_3}). \tag{43}$$

The argument base rate distributions are normally equal. When that is not the case the fused base rate

distribution is specified to be the average base rate distribution. In case the opinions of the N sources in $\mathbb{C}$ are all dogmatic opinions, then the limits in Eq. (39) can be set to $\gamma_X^C = 1/N$.

6.3.1. Justification for the Averaging Fusion Operator

The averaging belief fusion operator of Eq. (38) is derived by mapping the argument belief opinions to evidence opinions through the bijective mapping of Eq. (15). Averaging fusion of evidence opinions simply consists of computing the average of the evidence parameters. The fused evidence opinion is then mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the justification of the averaging fusion operator of Eq. (38). A more detailed explanation is provided below.

Let the sources $C \in \mathbb{C}$ have respective belief opinions expressed as ω_X^C. The corresponding Dirichlet PDFs $\mathrm{Dir}_X^e(\mathbf{p}_X; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence vectors $\mathbf{r}_X^C$.

The averaging fusion of these evidence vectors consists of vector averaging of $\mathbf{r}_X^C$ where $C \in \mathbb{C}$, expressed as

$$\mathbf{r}_X^{\diamond(\mathbb{C})} = \frac{\sum_{C \in \mathbb{C}} \mathbf{r}_X^C}{N}. \tag{44}$$

For each value $x \in \mathcal{R}(\mathbb{X})$ the average evidence $\mathbf{r}_X^{\diamond(\mathbb{C})}(x)$ is

$$\mathbf{r}_X^{\diamond(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{r}_X^C(x)}{N} = \frac{\sum_{C \in \mathbb{C}} W \mathbf{b}_X^C(x)/u_X^C}{N} \tag{45}$$

$$= \frac{W \sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right)}{N \prod_{C \in \mathbb{C}} u_X^C}. \tag{46}$$

The averaging-fused belief opinion $\omega_X^{\diamond(\mathbb{C})}$ of Eq. (38) results from mapping the fused evidence belief mass of Eq. (44) back to a belief opinion by applying the bijective mapping of Eq. (15).

$$\mathbf{b}_X^{\diamond(\mathbb{C})}(x) = \frac{\mathbf{r}_X^{\diamond(\mathbb{C})}(x)}{W + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\diamond(\mathbb{C})}(x)} \tag{47}$$

$$= \frac{\sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right)}{N \prod_{C \in \mathbb{C}} u_X^C + \sum_{x \in \mathcal{R}(\mathbb{X})} \left(\sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right) \right)} \tag{48}$$

$$= \frac{\sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right)}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right)}, \tag{49}$$

where $\exists u_X^C \neq 0$.

The transition from Eq. (47) to Eq. (48) results from inserting Eq. (46) into Eq. (47). The transition from

Eq. (48) to Eq. (49) results from applying Eq. (3).

$$u_X^{\diamond(\mathbb{C})} = \frac{W}{W + \sum_{x \in \mathcal{R}(\mathbb{X})} \mathbf{r}_X^{\diamond(\mathbb{C})}(x)} \tag{50}$$

$$= \frac{N \prod_{C \in \mathbb{C}} u_X^C}{N \prod_{C \in \mathbb{C}} u_X^C + \sum_{x \in \mathcal{R}(\mathbb{X})} \left(\sum_{C \in \mathbb{C}} \left(\mathbf{b}_X^C(x) \prod_{C_j \neq C} u_X^{C_j} \right) \right)} \tag{51}$$

$$= \frac{N \prod_{C \in \mathbb{C}} u_X^C}{\sum_{C \in \mathbb{C}} \left(\prod_{C_j \neq C} u_X^{C_j} \right)}, \tag{52}$$

where $\exists u_X^C \neq 0$.

The transition from Eq. (50) to Eq. (51) results from inserting Eq. (46) into Eq. (50). The transition from Eq. (51) to Eq. (52) results from applying Eq. (3).

6.4. Weighted Belief Fusion

The weighted belief fusion (WBF) operator produces averaging beliefs weighted by the opinion confidences.

The confidence c_X of an opinion ω_X is computed as:

$$c_X = 1 - u_X. \tag{53}$$

WBF is suitable for fusing source opinions in situations where the confidence should determine the opinion weight in the fusion process, which e.g. means that a vacuous opinion (i.e. an without confidence) has no effect on the fusion result.

When the arguments are conflicting multinomial opinions the fused result will be a dissonant multinomial opinion. This property could be seen as counter-intuitive when fusing opinions from human expert sources, because humans would tend to leverage belief on overlapping values and prefer vagueness over dissonance [26]. WBF is therefore best suited for frequentist situations where dissonance is preferred over vagueness. When vagueness is preferred the WBF-VM operator described in Section 6.6 can be used because it transforms dissonance into vagueness.

The definition of 2-source WBF specified in [5] was extended to multi-source WBF in [27] which is expressed below.

DEFINITION 4 (**The Weighted Belief Fusion Operator**). Assume a hyperdomain $\mathcal{R}(\mathbb{X})$ and a situation where the variable X takes values from the domain $\mathcal{R}(\mathbb{X})$. Assume further that the different sources from a frame of N sources $\mathbb{C} = \{C_1, C_2, \ldots C_N\}$ have their respective independent opinions on X. A specific source is denoted by $C \in \mathbb{C}$, and its opinion about the variable X is denoted ω_X^C.

Let $\omega_X^{\hat{\diamond}(\mathbb{C})}$ be the opinion such that

$$\omega_X^{\hat{\diamond}(\mathbb{C})} = (\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}, u_X^{\hat{\diamond}(\mathbb{C})}, \mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}), \quad \text{where} \tag{54}$$

Case I: $(\forall C \in \mathbb{C} : u_X^C \neq 0) \wedge (\exists C \in \mathbb{C} : u_X^C \neq 1)$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{b}_X^C(x)(1 - u_X^C) \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i}}{\left(\sum_{C \in \mathbb{C}} \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i} \right) - N \prod_{C \in \mathbb{C}} u_X^C},$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = \frac{(N - \sum_{C \in \mathbb{C}} u_X^C) \prod_{C \in \mathbb{C}} u_X^C}{\left(\sum_{C \in \mathbb{C}} \prod_{\substack{C_i \in \mathbb{C} \\ C_i \neq C}} u_X^{C_i} \right) - N \prod_{C \in \mathbb{C}} u_X^C},$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)(1 - u_X^C)}{N - \sum_{C \in \mathbb{C}} u_X^C}, \tag{55}$$

Case II: $\exists C \in \mathbb{C} : u_X^C = 0$. Let $\mathbb{C}^{\text{dog}} = \{C \in \mathbb{C} : u_X^C = 0\}$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{b}_X^C(x),$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = 0,$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \sum_{C \in \mathbb{C}^{\text{dog}}} \gamma_X^C \mathbf{a}_X^C(x), \tag{56}$$

where

$$\gamma_X^C = \lim_{u_X^{\mathbb{C}^{\text{dog}}} \to 0} \frac{u_X^C}{\sum_{C_j \in \mathbb{C}^{\text{dog}}} u_X^{C_j}}, \quad \forall C \in \mathbb{C}^{\text{dog}}.$$

Case III: $\forall C \in \mathbb{C} : u_X^C = 1$:

$$\mathbf{b}_X^{\hat{\diamond}(\mathbb{C})}(x) = 0,$$

$$u_X^{\hat{\diamond}(\mathbb{C})} = 1,$$

$$\mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{a}_X^C(x)}{N}. \tag{57}$$

The notation $u_X^{\mathbb{C}^{\text{dog}}} \to 0$ means that $u_X^C \to 0$ for each $C \in \mathbb{C}^{\text{dog}}$. $\omega_X^{\hat{\diamond}(\mathbb{C})}$ denotes the WBF (Weighted Belief Fusion) opinion resulting from the opinions ω_X^C provided by the sources $C \in \mathbb{C}$. By using the symbol '$\hat{\oplus}$' to denote this belief operator, we define

$$\omega_X^{\hat{\diamond}(\mathbb{C})} \equiv \underset{C \in \mathbb{C}}{\hat{\oplus}} (\omega_X^C). \tag{58}$$

It can be verified that WBF is commutative, idempotent and has the vacuous opinion as neutral element. Semi-associativity requires that three or more arguments must first be combined together in the same operation.

The argument base rate distributions are normally equal among the sources. When that is not the case the fused base rate distribution over X is specified to be the confidence-weighted average base rate distribution. In case of dogmatic arguments assume the limits in Eq. (56) to be $\gamma_X^C = 1/N$ where $N = |\mathbb{C}|$.

The WBF operator is equivalent to updating Dirichlet PDFs as the confidence-weighted average of source agents' evidence to produce posterior Dirichlet PDFs. The derivation of the confidence-weighted fusion operator is based on the bijective mapping between the belief and evidence notations described in Eq. (15).

THEOREM 1 *The weighted belief fusion operator of Definition 4 is equivalent to confidence-weighted averaging of the evidence parameters of the Dirichlet HPDF in Eq. (14).*

PROOF 1. The weighted belief fusion operator of Definition 4 is derived by mapping the argument belief opinions to evidence opinions through the bijective mapping of Eq. (15). Weighted belief fusion of evidence opinions simply consists of computing the confidence-weighted average of the evidence parameters. The fused evidence opinion is then mapped back to a belief opinion through the bijective mapping of Eq. (15). This explanation is in essence the proof of Theorem 1. A more detailed explanation is provided below.

Let the N sources $C \in \mathbb{C}$ have the respective belief opinions ω_X^C. The corresponding evidence opinions $\text{Dir}_X^{\text{eH}}(\mathbf{p}_X^{\text{H}}; \mathbf{r}_X^C, \mathbf{a}_X^C)$ contain the respective evidence parameters $\mathbf{r}_X^C$.

The weighted fusion of these bodies of evidence simply consists of weighted vector averaging of the parameters in the evidence opinions $\text{Dir}_X^{\text{eH}}(\mathbf{p}_X^{\text{H}}; \mathbf{r}_X^C, \mathbf{a}_X^C)$:

$$\text{Dir}_X^{\text{eH}}(\mathbf{p}_X^{\text{H}}; \mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}, \mathbf{a}_X^{\hat{\diamond}(\mathbb{C})}) = \underset{C \in \mathbb{C}}{\hat{\oplus}} \text{Dir}_X^{\text{eH}}(\mathbf{p}_X^{\text{H}}; \mathbf{r}_X^C, \mathbf{a}_X^A). \tag{59}$$

More specifically, for each value $x \in \mathcal{R}(\mathbb{X})$ the confidence-weighted fusion evidence $\mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}(x)$ is computed as

$$\mathbf{r}_X^{\hat{\diamond}(\mathbb{C})}(x) = \frac{\sum_{C \in \mathbb{C}} \mathbf{r}_X^C(x)(1 - u_X^C)}{N - \sum_{C \in \mathbb{C}} u_X^C}. \tag{60}$$

The weighted fusion opinion $\omega_X^{\hat{\diamond}(\mathbb{C})}$ of Definition 4 results from mapping the fused evidence belief mass of Eq. (59) back to a belief opinion as defined in Definition 4 by applying the bijective mapping of Eq. (15).

6.5. Uncertainty Maximisation

Uncertainty maximisation consists of transforming belief mass of an opinion ω_X into uncertainty mass while preserving the projected probability distribution $\mathbf{P}_X$.

Given a specific multinomial opinion ω_X, the corresponding uncertainty-maximised opinion is denoted $\ddot{\omega}_X = (\ddot{\mathbf{b}}_X, \ddot{u}_X, \mathbf{a}_X)$. Obviously, the base rate distribution $\mathbf{a}_X$ is not affected by uncertainty-maximisation.

The theoretical maximum uncertainty mass $\ddot{u}_X$ is determined by converting as much belief mass as possible into uncertainty mass, while preserving consistent projected probabilities. This process is illustrated in Figure 5 which shows an opinion ω_X as well as the corresponding uncertainty-maximised opinion $\ddot{\omega}_X$.

The projector line defined by the equations

$$\mathbf{P}_X(x_i) = \mathbf{b}_X(x_i) + \mathbf{a}_X(x_i)u_X, \quad i = 1, \dots k, \tag{61}$$

which by definition is parallel to the base rate director line, and which joins $\mathbf{P}_X$ and $\ddot{\omega}_X$ in Figure 5, defines possible opinions ω_X for which the projected probability distribution is constant. As the illustration shows, the opinion $\ddot{\omega}_X$ is the uncertainty-maximised opinion when
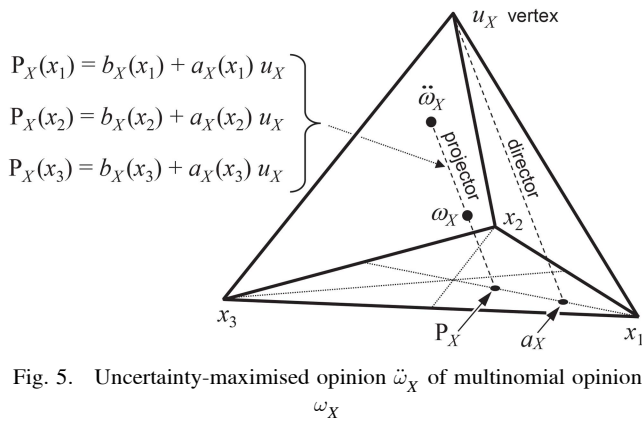
$$P_X(x_1) = b_X(x_1) + a_X(x_1)\, u_X$$
$$P_X(x_2) = b_X(x_2) + a_X(x_2)\, u_X$$
$$P_X(x_3) = b_X(x_3) + a_X(x_3)\, u_X$$

Fig. 5.   Uncertainty-maximised opinion $\ddot{\omega}_X$ of multinomial opinion ω_X

Eq. (61) is satisfied and at least one belief mass of $\ddot{\omega}_X$ is zero, since the corresponding point would lie on a side of the simplex. In general, not all belief masses can be zero simultaneously, except for vacuous opinions. The example of Figure 5 shows the case where $\ddot{\mathbf{b}}_X(x_1) = 0$.

The candidate maximum uncertainty mass $\check{u}_X(x_i)$ at each point where the projector intersects a side plane defined by $\mathbf{b}_X(x_i) = 0$ can be determined by Eq. (62):

$$\check{u}_X(x_i) = \frac{P_X(x_i)}{\mathbf{a}_X(x_i)}. \tag{62}$$

All belief masses determined according to Eq. (65) must be non-negative, which is satisfied through the constraint of Eq. (63):

$$\check{u}_X(x_i) \le \frac{P_X(x)}{\mathbf{a}_X(x)}, \quad \forall x \in \mathbb{X}. \tag{63}$$

Under the constraint of Eq. (63) the maximised uncertainty $\ddot{u}_X$ is the minimum candidate uncertainty from Eq. (62):

$$\ddot{u}_X = \min_{x_i \in \mathbb{X}}[\check{u}_X(x_i)]. \tag{64}$$

The belief masses under uncertainty maximisation emerge from Eq. (65) which is simply a transformation of Eq. (6):

$$\ddot{\mathbf{b}}_X(x) = P_X(x) - \mathbf{a}_X(x)\ddot{u}_X. \tag{65}$$

The uncertainty-maximised opinion consists of the components denoted $\ddot{\omega}_X = (\ddot{\mathbf{b}}_X, \mathbf{a}_X, \ddot{u}_X)$. By defining $\ddot{\uparrow}$ to be the unary operator for uncertainty maximisation we can write:

Uncertainty Maximisation: $\ddot{\omega}_X = \ddot{\uparrow}(\omega_X).$ (66)

A natural application of uncertainty maximisation is to produce epistemic opinions during opinion fusion. For that it is necessary to first generate a fused opinion, and subsequently to apply vagueness maximisation. In the case of e.g. CBF (Cumulative Belief Fusion) the combination with uncertainty maximisation is called CBF-UM (Cumulative Belief Fusion with Uncertainty Maximisation). An situation where it would be natural to apply CBF-UM could be when different witnesses express highly confident and highly conflicting opinions

about whether Oswald shot Kennedy in 1968, which when fused with e.g. CBF would produce an opinion with high confidence. Since the combined testimonies in this case would be inconclusive it could be natural to apply uncertainty maximisation to the result of CBF to produce CBF-UM, as shown in the example of Section 7.

6.6.   Vagueness Maximisation

In situations where people give different hypotheses it is fair to acknowledge that anyone can be wrong, and that a good consensus might be to agree that one of the hypotheses probably is right. This would typically be the situation in Zadeh's example [14] where two medical doctors give different diagnoses to explain a patient's symptoms, so that it would be natural for the doctors to agree that one of the diagnoses is correct, but that they are unable to identify which diagnosis in particular is correct. In this situation the combination of the two doctors result in a *vague* diagnosis.

Composite values $x \in \mathcal{R}(\mathbb{X})$ are state values containing multiple singleton values which e.g. can be different hypotheses such as medical diagnoses. Vague belief is belief mass assigned to a composite value, meaning that the belief mass applies to multiple singletons simultaneously. Vague belief mass thus reflects that the source believes that one of the singletons in the composite value is TRUE, without being able to identify which singleton in particular is TRUE. Vagueness is relevant for belief fusion, especially for WBF because vagueness can express compromise belief between conflicting sources. Vagueness maximisation consists of transforming belief masses on multiple singleton values into belief mass on a composite value, while preserving the projected probability distribution of Eq. (6).

In case the fused opinion ω_X is hypernomial we need to first apply Eq. (8) to compute the projected multinomial opinion ω_X.

Vagueness maximisation consists of transforming belief masses on multiple singleton values into a vague belief mass on the composite value containing the singletons. In case ω_X has belief mass on every singleton $x \in \mathbb{X}$ then a transformation into belief mass on $\mathbb{X}$ would not be meaningful because this is the same as uncertainty mass, and the transformation would break the assumption of preserving the amount of belief mass. We must identify the value(s) $x_i \in \mathbb{X}$ that should not be subject to vague belief mass, which can be done by computing the uncertainty-maximised opinion $\ddot{\omega}_X$ as described in Section 6.5 above.

The method of uncertainty maximisation described above forms the basis for the computation of vagueness-maximised opinions which is described below in the form of 4 consecutive steps. Note that this method of vagueness maximisation applies to multinomial opinions. Hence, if the goal is to apply vagueness maximisation to a hyper-opinion, a necessary preliminary step

is to first project it to a multinomial opinion according to Eq. (9)

Step 1:
Compute $\ddot{u}_X$ according to the procedure for uncertainty-maximisation described in Section 6.5. Let $\mathbb{X}_{\text{cut}}^{[1]}$ be the cut-out set of values x_i for which $\check{u}_X(x_i) = \ddot{u}_X$ with reference to Eq. (62) and Eq. (64). Note that $\mathbb{X}_{\text{cut}}^{[1]}$ may contain a single or multiple values.

Case A: $|\mathbb{X}_{\text{cut}}^{[1]}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ of the singleton value $x \in \mathbb{X}_{\text{cut}}^{[1]}$ and proceed to Step 2.

Case B: $1 < |\mathbb{X}_{\text{cut}}^{[1]}| < |\mathbb{X}|$. The composite value $x_{\text{vag}}^{[1]} = \{x \in \mathbb{X}_{\text{cut}}^{[1]}\}$ gets assigned the vague belief mass $\mathbf{b}_X(x_{\text{vag}}^{[1]})$ according to Eq. (67).

$$\mathbf{b}_X(x_{\text{vag}}^{[1]}) = \sum_{x \in \mathbb{X}_{\text{cut}}^{[1]}} \mathbf{b}_X(x). \tag{67}$$

Then proceed to Step 2.

Case C: $|\mathbb{X}_{\text{cut}}^{[1]}| = |\mathbb{X}|$: Split $\mathbb{X}$ into two exclusive sets $\mathbb{X}_{\text{res}}^{[1]}$ and $\mathbb{X}_{\text{res}}^{[2]}$ for which the respective sums of projected probability $\mathbf{P}(\mathbb{X}_{\text{res}}^{[1]})$ and $\mathbf{P}(\mathbb{X}_{\text{res}}^{[2]})$ are (approximately) equal. While this is a form of the knapsack problem we propose to simply sum up the greatest projected probabilities until the sum is greater than 0.5, and assign the corresponding set of values to $\mathbb{X}_{\text{res}}^{[1]}$, and the remaining values to $\mathbb{X}_{\text{res}}^{[2]}$. Define the composite values $x_{\text{vag}}^{[1]} = \{x \in \mathbb{X}_{\text{res}}^{[1]}\}$ and $x_{\text{vag}}^{[2]} = \{x \in \mathbb{X}_{\text{res}}^{[2]}\}$. Assign the vague belief masses $\mathbf{b}_X(x_{\text{vag}}^{[1]}) = \sum_{x \in \mathbb{X}_{\text{res}}^{[1]}} \mathbf{b}_X(x)$ and $\mathbf{b}_X(x_{\text{vag}}^{[2]}) = \sum_{x \in \mathbb{X}_{\text{res}}^{[2]}} \mathbf{b}_X(x)$. Proceed to the Final Step.

Step 2:
We exclude $\mathbb{X}_{\text{cut}}^{[1]}$ to produce the residual set $\mathbb{X}_{\text{res}}^{[2]}$:

$$\mathbb{X}_{\text{res}}^{[2]} = \mathbb{X} \setminus \mathbb{X}_{\text{cut}}^{[1]}. \tag{68}$$

Case A: $|\mathbb{X}_{\text{res}}^{[2]}| = 0$. Proceed to the Final Step.
Case B: $|\mathbb{X}_{\text{res}}^{[2]}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ on the singleton value $x \in \mathbb{X}_{\text{res}}^{[2]}$. Proceed to the Final Step.
Case C: $|\mathbb{X}_{\text{res}}^{[2]}| \geq 2$. Now we focus exclusively on values $x_i \in \mathbb{X}_{\text{res}}^{[2]}$ when applying the constraint of Eq. (63). The next synthetic maximum uncertainty mass is:

$$\ddot{u}_X^{[2]} = \min_{x_i \in \mathbb{X}_{\text{res}}^{[2]}} [\check{u}_X(x_i)]. \tag{69}$$

Eq. (70) gives the corresponding synthetic belief masses:

$$\ddot{\mathbf{b}}_X^{[2]}(x) = \mathbf{P}_X(x) - \mathbf{a}_X(x)\ddot{u}_X^{[2]}, \quad \forall x \in \mathbb{X}_{\text{res}}^{[2]}. \tag{70}$$

We define the composite value $x_{\text{vag}}^{[2]} = \{x \in \mathbb{X}_{\text{res}}^{[2]}\}$. The vague belief mass $\mathbf{b}_X(x_{\text{vag}}^{[2]})$ can then be assigned according to Eq. (71)

$$\mathbf{b}_X(x_{\text{vag}}^{[2]}) = \sum_{x \in \mathbb{X}_{\text{res}}^{[2]}} (\mathbf{b}_X(x) - \ddot{\mathbf{b}}_X^{[2]}(x)). \tag{71}$$

Let the iterative step index be denoted η. Set $\eta = 3$ and proceed to Step η.

Step η:
Let $\mathbb{X}_{\text{cut}}^{[\eta-1]}$ be the set of values x_i for which $\check{u}_X(x_i) = \ddot{u}_X^{[\eta-1]}$ with reference to Eq. (62) and Eq. (64). We exclude $\mathbb{X}_{\text{cut}}^{[\eta-1]}$ from $\mathbb{X}_{\text{res}}^{[\eta-1]}$ to produce the residual set $\mathbb{X}_{\text{res}}^{[\eta]}$:

$$\mathbb{X}_{\text{res}}^{[\eta]} = \mathbb{X}_{\text{res}}^{[\eta-1]} \setminus \mathbb{X}_{\text{cut}}^{[\eta-1]}. \tag{72}$$

Case A: $|\mathbb{X}_{\text{res}}^{[\eta]}| = 0$. Proceed to the Final Step.
Case B: $|\mathbb{X}_{\text{res}}^{[\eta]}| = 1$. Keep the singular belief mass $\mathbf{b}_X(x)$ on the singleton value $x \in \mathbb{X}_{\text{res}}^{[\eta]}$. Proceed to the Final Step.
Case C: $|\mathbb{X}_{\text{res}}^{[\eta]}| \geq 2$. Now we focus exclusively on values $x_i \in \mathbb{X}_{\text{res}}^{[\eta]}$ when applying the constraint of Eq. (63). The next synthetic maximum uncertainty mass is:

$$\ddot{u}_X^{[\eta]} = \min_{x_i \in \mathbb{X}_{\text{res}}^{[\eta]}} [\check{u}_X(x_i)]. \tag{73}$$

The computation of the belief masses emerges from Eq. (74):

$$\ddot{\mathbf{b}}_X^{[\eta]}(x) = \mathbf{P}_X(x) - \mathbf{a}_X(x)\ddot{u}_X^{[\eta]}. \tag{74}$$

We define the composite value $x_{\text{vag}}^{[\eta]} = \{x \in \mathbb{X}_{\text{res}}^{[\eta]}\}$. The vague belief mass $\mathbf{b}_X(x_{\text{vag}}^{[\eta]})$ can then be assigned according to Eq. (75)

$$\mathbf{b}_X(x_{\text{vag}}^{[\eta]}) = \sum_{x \in \mathbb{X}_{\text{res}}^{[\eta]}} (\ddot{\mathbf{b}}_X^{[\eta-1]}(x) - \ddot{\mathbf{b}}_X^{[\eta]}(x)). \tag{75}$$

Increment the step index η as $\eta := \eta + 1$, then repeat Step η.

Final Step:
Finally, the components of the vagueness-maximised opinion $\dot{\omega}_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ can be assembled, consisting of the computed vague belief masses $\mathbf{b}_X(x_{\text{vag}}^{[\eta]})$, and whenever applicable the singular belief masses $\mathbf{b}_X(x_i)$, in addition to the original uncertainty mass u_X and base rate distribution $\mathbf{a}_X$. This ends the process of vagueness maximisation.

By defining the unary operator $\dot{\uparrow}$ to represent vagueness maximisation we can write

Vagueness Maximisation: $\quad \dot{\omega}_X = \dot{\uparrow}(\omega_X).$ (76)

A natural application of vagueness maximisation is to produce compromise belief when fusing opinions from multiple (conflicting) sources. To this end it is necessary to first generate a fused opinion with WBF, and subsequently to apply vagueness maximisation. This combination is called WBF-VM (Weighted Belief Fusion with Vagueness Maximisation) and is denoted '$\hat{\oplus}$'.

As an alternative to WBF-VM for belief fusion with compromise, the belief fusion operator CCF (Consensus & Compromise Fusion) has been described with a simple two-source version [5] as well as with a multi-source version [27]. The definition of multi-source CCF is rather complex [27], whereas multi-source WBF-VM is rather simple in comparison. In situations where it

TABLE 5

Zadeh's numerical example applied to belief constraint fusion (BCF), cumulative belief fusion (CBF), cumulative belief fusion with uncertainty maximisation (CBF-UM), averaging belief fusion (ABF), weighted belief fusion (WBF) and weighted belief fusion with vagueness maximisation (WBF-VM)

| | | Source opinions: | | Fused opinions resulting from applying: | | | | | |
		A	B	BCF	CBF	CBF-UM	ABF	WBF	WBF-VM
$b_X(x_1)$	=	0.99	0.00	0.00	0.495	0.485	0.495	0.495	0.000
$b_X(x_2)$	=	0.01	0.01	1.00	0.010	0.000	0.010	0.010	0.010
$b_X(x_3)$	=	0.00	0.99	0.00	0.495	0.485	0.495	0.495	0.000
$b_X(x_1,x_2)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.000
$b_X(x_1,x_3)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.990
$b_X(x_2,x_3)$	=	0.00	0.00	0.00	0.000	0.000	0.000	0.000	0.000
u_X	=	0.00	0.00	0.00	0.000	0.030	0.000	0.000	0.000

TABLE 6

A variation of Zadeh's example applied to belief constraint fusion (BCF), cumulative belief fusion (CBF), cumulative belief fusion with uncertainty maximisation (CBF-UM), averaging belief fusion (ABF), weighted belief fusion (WBF) and weighted belief fusion with vagueness maximisation (WBF-VM)

| | | Source opinions: | | Fused opinions resulting from applying: | | | | | |
		A	B	BCF	CBF	CBF-UM	ABF	WBF	WBF-VM
$b_X(x_1)$	=	0.98	0.00	0.889	0.890	0.880	0.882	0.889	0.806
$b_X(x_2)$	=	0.01	0.01	0.011	0.010	0.000	0.010	0.010	0.010
$b_X(x_3)$	=	0.00	0.90	0.091	0.091	0.081	0.090	0.083	0.000
$b_X(x_1,x_2)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.000
$b_X(x_1,x_3)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.166
$b_X(x_2,x_3)$	=	0.00	0.00	0.000	0.000	0.000	0.000	0.000	0.000
u_X	=	0.01	0.09	0.009	0.009	0.039	0.018	0.018	0.018

is suitable to apply a fusion operator with belief compromise, the most practical choice is therefore to apply WBF-VM which is included in the example of Section 7.

7. COMPARISON OF FUSION OPERATORS

The fusion example in Table 5 takes as input arguments the numerical belief masses from Zadeh's example [14]. In this example, the sources are two medical doctors who each have an opinion about the hypothesis space of three possible diseases, and Dempster's rule (called BCF (Belief Constraint Fusion) in subjective logic) is applied for fusing the two opinions. The counter-intuitive results produced by Dempster's rule (BCF) demonstrate that Dempster's rule is unsuitable for this particular category of situations. A more suitable operator for the situation of the two doctors is WBF-VM (Weighted Belief Fusion with Vagueness Maximisation), because it preserves common belief and produces compromise belief from conflicting belief sources.

Exactly the same pair of argument opinions can of course occur in other fusion situations as well. Table 5 shows the results of fusion with each operator described in the previous sections, where the the fused result opinion produced by a given operator is sound and intuitive according to the corresponding situation category described in Section 2.

On an abstract level, sources A and B provide opinions about the hypothesis space $\mathbb{X} = \{x_1, x_2, x_3\}$ with variable X. The base rate distributions are assumed to be equal and uniform, expressed as $\mathbf{a}_X^A = \mathbf{a}_X^B = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$.

Each operator produces intuitive results given respective relevant situations for which the operators are suitable. For example, in the medical situation of the original Zadeh's example where two medical doctors A and B have conflicting opinions about the diagnosis of a patient, WBF-VM produces vague belief in the form of $\mathbf{b}_X^{A\hat{\diamond}B}(x_1, x_3) = 0.99$ which seems natural until the doctors can agree on a single diagnosis for the patient. The BCF operator produces a sound and intuitive fused opinion with the same argument opinions when e.g. assuming a situation where two friends express preferences for watching a film at the cinema.

Fusion of dogmatic conflicting opinions, i.e. where $u_X = 0$, is defined for all operators except for BCF. If the fusion situation is determined to be in the BCF category the interpretation of fusing dogmatic conflicting opinions is that there is no solution, which is perfectly logic. See Section 6.1.5 for an example of this situation.

Zadeh's example as in Table 5 does not clearly expose the difference between the various belief fusion operator because many fusion operators produce equal results when the sources are dogmatic as in this case. The modified example in Table 6 brings greater differentiation in the fusion results by introducing unbalanced

levels of uncertainty in the argument opinions. The difference between the arguments of Table 5 and Table 6 can be interpreted and explained through the assumptions of the various belief-fusion categories with regard to how conflicting belief arguments are handled in the belief fusion process.

8. DISCUSSION AND CONCLUSION

We argue that the main research question in belief fusion is not about finding the single most correct belief fusion operator, because no single operator is suitable for all situations. Instead, the interesting question and the biggest challenge is how to select the most suitable belief fusion operator for a given situation of belief fusion. For this purpose we propose to classify situations of belief fusion into different categories, where a set of belief-fusion assumptions can be used as criteria for selecting the category to which a specific belief fusion situation belongs.

This article illustrates the importance of selecting a belief fusion operator that adequately matches the situation to be modelled and analyzed. It is scientifically misguided to follow the approach of always applying the favourite belief fusion operator with which the analyst or scientist happens to be familiar, without regard to the nature of the situation to be modelled. By using the selection criteria to categorise a given belief-fusion situation and applying the corresponding belief fusion operator the analyst is able to obtain sound and useful results more consistently than by simply making an uninformed choice when selecting a belief fusion operator for a given application.

ACKNOWLEDGMENT

REFERENCES

[1] A. Jøsang, P. C. Costa, and E. Blash
"Determining model correctness for situations of belief fusion."
in *Proceedings of the 16th International Conference on Information Fusion (FUSION 2013)*. IEEE, Los Alamitos, 2013, pp. 1225–1232.

[2] "ETURWG (Evaluation of Techniques for Uncertainty Representation Working Group),"
Available from http://eturwg.c4i.gmu.edu/, accessed June 2018.

[3] P. C. G. Costa, K. B. Laskey, E. P. Blasch, and A.-L. Jousselme
"Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.

[4] E. P. Blasch, P. C. G. Costa, K. B. Laskey, D. Stampouli, G. W. Ng, J. Schubert, R. Nagi, and P. Valin
"Issues of Uncertainty Analysis in High-Level Information Fusion: Fusion 2012 Panel Discussion."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.

[5] A. Jøsang
Subjective Logic: A Formalism for Reasoning Under Uncertainty.
Springer, Heidelberg, 2016.

[6] P. H. Foo and G. W. Ng
"High-level information fusion: An overview,"
J. Adv. Inf. Fusion, vol. 8, no. 1, pp. 33–72, 2013.

[7] A. Jøsang, M. Ivanovska, and T. Muller
"Trust Revision for Conflicting Sources."
in *Proceedings of the 18th International Conference on Information Fusion (FUSION 2015)*. IEEE, Los Alamitos, 2015.

[8] A. Jøsang and R. Hankin
"Interpretation and fusion of hyper-opinions in subjective logic."
in *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*. IEEE, Los Alamitos, Singapore, July 2012.

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

[10] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin
Bayesian Data Analysis, 3rd ed.
Chapman and Hall/CRC, Boca Raton, 2013.

[11] R. K. Hankin
"A generalization of the Dirichlet distribution,"
Journal of Statistical Software, vol. 33, no. 11, pp. 1–18, February 2010.

[12] A. Jøsang
"Generalising Bayes' Theorem in Subjective Logic."
in *International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2016)*. IEEE, Los Alamitos, 2016.

[13] A. Jøsang and S. Pope
"Dempster's rule as seen by little colored balls,"
Computational Intelligence, vol. 28, no. 4, pp. 453–474, November 2012.

[14] L. A. Zadeh
"Review of Shafer's 'A Mathematical Theory of Evidence',"
AI Magazine, vol. 5, pp. 81–83, 1984.

[15] M. Daniel
"Associativity in combination of belief functions,"
in *Proceedings of 5th Workshop on Uncertainty Processing*, J. Vejnarová, Ed. Ediční oddelění VŠE 2000, Praha, 2000, pp. 41–54.

[16] D. Dubois and H. Prade
"Representation and combination of uncertainty with belief functions and possibility measures,"
Comput. Intell., vol. 4, no. 3, pp. 244–264, 1988.

[17] A. Jøsang
"The consensus operator for combining beliefs,"
Artificial Intelligence, vol. 142, no. 1–2, pp. 157–170, October 2002.

[18] A. Jøsang, J. Diaz, and M. Rifqi
"Cumulative and averaging fusion of beliefs,"
Information Fusion, vol. 11, no. 2, pp. 192–200, 2010, doi:10.1016/j.inffus.2009.05.005.

[19] E. Lefevre, O. Colot, and P. Vannoorenberghe
"Belief functions combination and conflict management,"
Information Fusion, vol. 3, no. 2, pp. 149–162, June 2002.

[20] C. K. Murphy
"Combining belief functions when evidence conflicts,"
Decision Support Systems, vol. 29, pp. 1–9, 2000.

[21] P. Smets
"The Combination of Evidence in the Transferable Belief Model,"
IEEE Transansactions on Pattern Analysis and Machine Intelligence, vol. 12, no. 5, pp. 447–458, 1990.

[22] R. Yager
"On the Dempster-Shafer framework and new combination rules,"
Information Sciences, vol. 41, pp. 93–137, 1987.

[23] A. Jøsang
"Multi-agent preference combination using subjective logic,"
in *International Workshop on Preferences and Soft Constraints (Soft'11)*, Perugia, Italy, 2011, pp. 61–75.

[24] J. Kohlas and P. Monney
A Mathematical Theory of Hints. An Approach to Dempster-Shafer Theory of Evidence,
ser. Lecture Notes in Economics and Mathematical Systems. Springer-Verlag, 1995, vol. 425.

[25] S. Pope and A. Jøsang
"Analsysis of Competing Hypotheses using Subjective Logic,"
in *Proceedings of the 10th International Command and Control Research and Technology Symposium (ICCRTS)*. United States Department of Defense Command and Control Research Program (DoDCCRP), 2005.

[26] A. Jøsang, J.-H. Cho, and F. Chen
"Uncertainty Characteristics of Subjective Opinions,"
in *Proceedings of the 21st International Conference on Information Fusion (FUSION 2018)*. Available at: http://folk.uio.no/josang/papers/JCC2018-FUSION.pdf: IEEE, Los Alamitos, 2018.

[27] R. W. van der Heijden, H. Kopp, and F. Kargl
"Multi-Source Fusion Operations in Subjective Logic,"
in *Proceedings of the 21st International Conference on Information Fusion (FUSION 2018)*. Available at: https://arxiv.org/abs/1805.01388: IEEE, Los Alamitos, 2018.

Audun Jøsang is a Professor in information security at the University of Oslo, and is also an Adjunct Professor at QUT in Australia. Prof. Jøsang is the main author behind subjective logic which is widely applied by researchers and practitioners in the areas of uncertainty representation, belief fusion, subjective Bayesian reasoning, reputation systems, and computational trust. He holds a Master's degree in information security from Royal Holloway College, University of London and a PhD from NTNU in Norway.