

Opinion Update in a Subjective Logic Model for Social Networks

Mário S. Alvim¹, Sophia Knight², and José C. Oliveira²

¹ Department of Computer Science, UFMG, Brazil
msalvim@dcc.ufmg.br

² Department of Computer Science, University of Minnesota Duluth, USA
sophia.knight@gmail.com, josecarlosdeoliveirajr@gmail.com

Abstract. *Subjective Logic* (SL) is a logic incorporating uncertainty and opinions for agents in dynamic systems. In this work, we investigate the use of subjective logic to model opinions and belief change in social networks. In particular, we work toward the development of a subjective logic belief/opinion update function appropriate for modeling belief change as communication occurs in social networks. We found through experiments that an update function with belief fusion from SL does not have ideal properties to represent a rational update. Even without these properties, we found that an update function with cumulative belief fusion can describe behaviors not explored by the social network model defined by Alvim, Knight, and Valencia [3].

Keywords: Social networks · Subjective Logic · Belief change · Dynamic opinions · Multi-agent systems

1 Introduction

Recently, social networks have begun to influence every aspect of our lives, with unprecedented, unanticipated consequences, especially in politics and public opinion. Research on social networks has studied opinions and their change over time, but to accurately model real people and their opinions and beliefs, we must include representations of *uncertainty* in formal models of social networks.

To achieve this goal, we use *subjective logic* (SL), which includes information about agents' uncertainty, to develop a more nuanced model of social networks and their changes over time. This work builds upon the model developed by Alvim, Knight, and Valencia (AKV) [3]. Their social network model incorporates quantitative opinions and influence on each agent but only addresses binary opinions and uncertainty is not represented.

The contributions of this paper are the following:

- We propose a model for social networks using elements of the subjective logic model such as multinomial opinions, trust opinions, and belief fusion operators.

- We propose a belief/opinion update function using SL’s trust discount and belief fusion. We use examples to show that our update function, with cumulative, averaging, and weighted belief fusions, does not satisfy proprieties that are useful to model a basic social network with rational update.
- We analyze the update function using trust discount and cumulative belief fusion from subjective logic and how it can represent a different scenario not described in the model proposed by Alvim, Knight, and Valencia [3].

2 Related Work

There is a great deal of work concerned with logical models of social networks and formalizing belief change in social networks, but we are unaware of other work modeling quantitative uncertainty in social networks. More specifically, other work allows uncertainty only in the sense of multiple states or worlds that an agent considers possible, whereas in this work we use subjective logic, considering uncertainty as a *total lack of information*, on top of the multiple states an agent may consider possible. Much of the literature treats opinions as binary, non-quantitative phenomena, whereas we investigate opinions that take on a spectrum of values between 0 and 1, and which also include possible uncertainty.

Before the advent of online social networks, Degroot et al. [7] proposed a model of learning and consensus in multi-agent systems, in which quantitative beliefs are updated by a constant stochastic matrix at each time step. The Degroot model does not include uncertainty but otherwise is close to our work in spirit. The models in [4] are also similar to the models we use, but the focus of that work is on payoffs and optimal decision-making, whereas our focus is purely on the changes in information and uncertainty over time. Incorporating goals and decisions into our models provides an interesting problem for future work, particularly in the context of our focus on uncertainty. In [8], Holliday develops a logic with *ordered* but non-quantitative trust and certainty about propositions: agent a may trust agent b more than they trust agent c , leading them to believe certain propositions more strongly than others. Liu et al. [12] use ideas from doxastic and dynamic epistemic logics to qualitatively model influence and belief changes in social networks. Christoff [6] has developed several non-quantitative logics for social networks, and Young Pedersen [13,15] develops a non-quantitative logic concerned specifically with polarization. In [14], Xiong and Ågotnes develop a logic to analyze signed social networks where agents can have “friends” and “enemies,” a different approach to some of the same questions that concern us, such as polarization and influence.

Hunter [9] introduces a logic of belief updates over social networks with varying levels of influence and trust. Using dynamic epistemic logic, Baltag et al. [5] created a threshold model where agents’ behavior changes when the proportion of supporters changes, but with binary belief and no uncertainty.

This work is a continuation of Alvim et al.’s work [1,2,3]. Alvim et al. develop a formal model for social networks where agents have quantitative opinions and

quantitative influence on each other, with a function for agents' belief update over time. The goal of the current paper is to extend this model by adding the possibility of *uncertainty* to the agents' quantitative opinions.

3 Background: Subjective Logic

This section provides background on the elements of subjective logic that we use in our model. *Subjective Logic* is a logic developed by Jøsang [10] that extends probabilistic logic by adding *uncertainty* and *subjectivity*. In probabilistic logic, a uniform distribution does not express “*we don't know*” because a uniform distribution says that we know that the distribution over the domain is uniform. Subjective logic can distinguish between the situation where the distribution over outcomes is unknown and the situation where the distribution is known and, for example, uniform. In subjective logic, it is also possible to have a situation where some information about the distribution is known and there is some uncertainty. The *subjectivity* comes from the fact that we can assign an opinion, or information, about a proposition to an agent.

Opinion representation The main object of subjective logic is the *opinion*. We represent an opinion by ω_X^A , where A is an agent, X a random variable, and ω_X^A is A 's opinion about X . An opinion expresses support for none, one, or many states of a domain. This section presents the elementary definitions that compose an opinion. A *domain* is a state space consisting of a finite set of values called states, events, outcomes, hypotheses, or propositions. The values are assumed to be exclusive and exhaustive.

Belief mass is a distribution over a domain \mathbb{X} representing an agent's confidence in each value in the domain. The belief mass assigned to a value $x \in \mathbb{X}$ expresses support for x being TRUE. Belief mass is sub-additive, i.e. $\sum_{x \in \mathbb{X}} \mathbf{b}_X(x) \leq$

1. The sub-additivity is complemented by *uncertainty mass* u_X and it represents the lack of support or evidence for the variable X having any specific value.

Definition 1. (Belief Mass Distribution) *Let \mathbb{X} be a domain of size $k \geq 2$, and let X be a variable over that domain. A belief mass distribution denoted $\mathbf{b}_X : \mathbb{X} \rightarrow [0, 1]$ assigns belief mass to possible values of the variable X . Belief mass and uncertainty mass sum to one, i.e., $u_X + \sum_{x \in \mathbb{X}} \mathbf{b}_X(x) = 1$.*

Opinions can be semantically different, depending on the situation they apply to. An *aleatory* opinion applies to a variable governed by a random or frequentist process, and represents the likelihood of values of the variable in any unknown past or future instance of the process. “*The (biased) coin will land heads with $p = 0.6$* ” is an aleatory opinion. An *epistemic* opinion applies to a variable that is assumed to be non-frequentist, and that represents the likelihood of the variables in a specific unknown instance. “*Beatriz killed Evandro*” is an epistemic opinion. In an epistemic opinion, opposite/different pieces of evidence should cancel each other out. Therefore, it must be uncertainty-maximized.

Base rate distribution represents a *prior* probability distribution over a domain: the probability distribution before considering evidence about the domain.

Definition 2. (Opinion) *Let \mathbb{X} be a domain of size $k \geq 2$, and X a random variable in \mathbb{X} . An opinion over the random variable X is the ordered triple $\omega_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ where*

- \mathbf{b}_X is a belief mass distribution over X ,
- u_X is the uncertainty mass which represents a lack of evidence,
- \mathbf{a}_X is a base rate distribution (a probability distribution) over \mathbb{X} .

The *projected probability distribution* of an opinion is the *posterior* probability distribution after updating the base rate distribution with the belief mass distribution. The more an opinion depends on the belief mass, the less it depends on the base rate. The projected probability distribution is defined by $\mathbf{P}_X(x) = \mathbf{b}_X(x) + \mathbf{a}_X(x)u_X$, $\forall x \in \mathbb{X}$

This representation is equivalent to representing an opinion as a Beta PDF (or Dirichlet PDF if $k > 2$), where the posterior probability is obtained by updating the parameters α and β (or a vector of parameters α for the Dirichlet PDF) given the observations. The equivalence is defined as the projected probability from SL’s opinion being equivalent to the expected probability of the Beta PDF. More details about the equivalence between opinions and Beta PDFs are presented in Appendix A.

The definition of opinion is useful for our model since it is more expressive than the belief state of an agent about a proposition in [3], which is similarly an opinion with domain $\mathbb{X} = \{true, false\}$, with no uncertainty mass. The agent must commit all of their mass to the values of the domain with no uncertainty.

Example 1. Let $\mathbb{X} = \{x, \bar{x}\}$ be a domain where x is “global warming is happening” and \bar{x} is “global warming is *not* happening”. Let X be a random variable in \mathbb{X} . An opinion about X must be epistemic because it is about a fact in the present instance that is true or false. Let the base rate be uniform. With no evidence, an agent A will hold the opinion $\omega_X^A = ((0, 0), 1, (0.5, 0.5))$ with $\mathbf{P}_X^A(x) = 0.5$, meaning that A is 50% sure that the global warming is happening, but their opinion is relying only on the base rate, with no evidence supporting either of the values.

After gathering evidence from newspapers, scientific studies, and other people, A assigns a belief mass to x . If agent A holds the opinion $\omega_X^A = ((0.6, 0), 0.4, (0.5, 0.5))$, then $\mathbf{P}_X^A(x) = 0.8$. In this case, A is 80% sure that global warming is happening, and has evidence that corresponds to 60% of their mass. The uncertainty mass means that A is relying 40% on the base rate.

Trust discount To model the influence that one agent has on another, subjective logic has *trust opinion*, an opinion an agent has about another agent as a source of information.

Definition 3. (Trust opinion) Let $\mathbb{T}_B = \{t_B, \bar{t}_B\}$ be a trust domain, where t_B means “ B is a good source of information” and \bar{t}_B means “ B is not a good source of information”. Then $\omega_{t_B}^A$, or ω_B^A for short, is the (trust) opinion that A has about the trustworthiness of B as a source of information.

We use trust opinions to model an agent’s updated opinion after communication: $\omega_X^{[A;B]}$ is a new opinion generated by taking belief mass ω_X^B proportional to the belief mass of the trust opinion ω_B^A . And $\omega_X^{[A;B]}$ represents A ’s opinion about X after communicating with B , ω_B^A represents A ’s opinion about B ’s trustworthiness, and ω_X^B represents B ’s opinion about X . The operation is denoted $\omega_X^{[A;B]} = \omega_B^A \otimes \omega_X^B$.

There are several options for computing this value, depending on the situation being modeled. Developing an accurate function for opinion updates in social networks is a focus of the current paper.

Example 2. Let $\omega_B^A = ((1, 0), 0, \mathbf{a}_B^A)$ with $\mathbf{P}_B^A(t_B) = 1$ and $\omega_X^B = ((0.6, 0), 0.4, (0.5, 0.5))$ with $\mathbf{P}_X^B(x) = 0.8$. Here, A completely trusts B and B is 80% sure that x is true with 60% of their mass assigned to x .

$\mathbf{P}_B^A(t_B) = 1$, i.e. A completely trusts B . Then, A by trusting B (in short $[A; B]$) will hold the same opinion as B about X . Therefore, $\omega_X^{[A;B]} = \omega_X^B$. By the opinion that A has about X by trusting B , A is 60% sure that x is true with 80% of their mass assigned to x .

Example 3. Let $\omega_B^A = ((0.5, 0.5), 0, \mathbf{a}_B^A)$ with $\mathbf{P}_B^A(t_B) = 0.5$ and $\omega_X^B = ((0.8, 0), 0.2, (0.5, 0.5))$ with $\mathbf{P}_X^B(x) = 0.9$. Here, A trusts B by 50% and B is 80% sure that x is true with 60% of their mass assigned to x .

$\mathbf{P}_B^A(t_B) = 0.5$. Then, $[A; B]$ will hold 50% of the belief mass of each value from B . Therefore, $\omega_X^{[A;B]} = ((0.4, 0), 0.6, (0.5, 0.5))$ with $\mathbf{P}_X^{[A;B]}(x) = 0.7$. By the opinion that A has about X by trusting B , A is 70% sure that x is true with 40% of their mass assigned to x .

Belief fusion To model A ’s concurrent interactions with multiple other agents, we use belief fusion [10,11]. Belief fusion combines a set of opinions into a single opinion which then represents the opinion of the collection of sources. There is more than one possible definition for the belief fusion operator. The possible definitions differ in their properties and applications. Below, we omit the details of the calculation of the fusion operators and instead, explain the intuition. For more details, see [10]. For our model, we consider the following operators from [10,11]:

- *Cumulative belief fusion* (denoted $\omega_X^{(A \diamond B)} = \omega_X^A \oplus \omega_X^B$): Used when it is assumed that the amount of independent evidence increases by including more sources. The idea is to sum the amount of evidence of the opinions. It is non-idempotent. E.g. a set of agents flips a coin several times and produces an opinion about the bias of the coin. An opinion produced by cumulative belief fusion represents all the experiments made by the agents.

- *Averaging belief fusion* (denoted $\omega_X^{(A \circ B)} = \omega_X^A \oplus \omega_X^B$): It is used when including more sources does not mean that more evidence supports the conclusion. The idea is to take the average of the amount of evidence of the opinions. It is idempotent, but it has no neutral element. E.g. After observing the court proceedings, each member of a jury produces an opinion about the same evidence. The verdict is the fusion between those opinions.
- *Weighted belief fusion* (denoted $\omega_X^{(A \hat{\circ} B)} = \omega_X^A \hat{\oplus} \omega_X^B$): It is used when we take the average of the amount of evidence of the opinions weighted by their lack of uncertainty. In particular, opinions with no belief mass are rejected. It is idempotent and it has a neutral element $u_X = 1$. E.g. a group of medical doctors needs to decide on a diagnosis. Each of them has an opinion, but some of them are more certain (assigned more belief mass) than others. Those opinions must have more weight than the others upon fusion.

Each of these operators is defined in the Appendix B.

4 A Subjective Logic model for social networks

We describe how to use subjective logic to model a social network. Existing notions from subjective logic work well for modeling the relevant aspects of agents' opinions about an issue in a social network, except for the update function. Our goal is to expand subjective logic with an appropriate update function.

4.1 Static elements of the model

The *static elements* represent a snapshot of a social network at a given point in time. They include the following components:

- A (finite) set $\mathcal{A} = \{A_0, A_1, \dots, A_{n-1}\}$ of $n \geq 1$ *agents*. An agent is a user in a social network
- A *domain* of k disjoint events $\mathbb{X} = \{x_0, x_1, \dots, x_{k-1}\}$ and a random variable X over \mathbb{X} . A domain is a generalization of a proposition from binary logic by having multiple truth values for a proposition. A proposition from binary logic would have a domain of size 2. It can represent topics that cannot be answered with just a YES or NO. In the examples in this work, we use a domain of size two for simplicity, and the random variable X represents opinions about a single issue.
- A set of *opinions* $\{\omega_X^{A_i}\}_{A_i \in \mathcal{A}}$, one for each agent, about a single random variable X .

4.2 Dynamic elements of the model

The *dynamic elements* of the model formalize the evolution of agents' beliefs as they interact and communicate about their opinions over time. They include:

- A set of *trust opinions* $\{\omega_{A_j}^{A_i}\}_{A_i, A_j \in \mathcal{A}}$ over the domain $\mathbb{T}_{A_j} = \{t_{A_j}, \bar{t}_{A_j}\}$ where $A_i, A_j \in \mathcal{A}$ and $i \neq j$. Each trust opinion represents how much A_i trusts A_j as a source of evidence. We consider trust opinion as dogmatic, i.e. the uncertainty mass is $i_{T_j} = 0$ for all trust opinions. Therefore, $\mathbf{P}_{A_j}^{A_i}(t_{A_j}) = \mathbf{b}_{A_j}^{A_i}(t_{A_j})$. Throughout this paper, we represent a trust opinion only by $\mathbf{P}_{A_j}^{A_i}(t_{A_j})$.
- A *time frame* $\mathcal{T} = \{0, 1, 2, \dots, t_{max}\}$ representing the discrete passage of time.
- An *update function* f such that

$$\omega_X^{A_i[t+1]} = f(\omega_X^{A_i[t]}, \{\omega_{A_j}^{A_i}\}_{A_j \in \mathcal{A}}, \{\omega_X^{A_j[t]}\}_{A_j \in \mathcal{A}}). \quad (1)$$

An update function f takes A_i 's opinions at time t and updates it using all other agents' opinions and trust opinions to them. Here we consider that agents have complete trust in themselves. Otherwise, we can define this function with A_i 's opinion included in the set of opinions $\{\omega_{A_j}^{A_i}\}_{A_j \in \mathcal{A}}$.

5 Applying trust discount and belief fusion

The intuition behind our planned update function is to fuse agent A 's current opinion with all the opinions that A can gather by trusting other agents. Define a *dogmatic opinion* as an opinion with no uncertainty, i.e. $u_X = 0$. For this update function, we are not considering situations with dogmatic opinions, because it means the agent has an infinite amount of evidence and the belief fusion operators remove non-dogmatic opinions when at least one is present.

Definition 4. (Update function with Belief Fusion and Trust) Let $\omega_X^{A_0[t]}, \dots, \omega_X^{A_{n-1}[t]}$ be non-dogmatic opinions. Let \oplus be a belief fusion operator. Define the update function for $\omega_X^{A_n[t+1]}$ as

$$\omega_X^{A_n[t+1]} = \omega_X^{A[t]} \oplus \bigoplus_{\substack{A_m \in \mathcal{A} \\ m \neq n}} (\omega_{A_m}^{A_n} \otimes \omega_X^{A_n[t]}). \quad (2)$$

Our goal in this section is to understand the effects and properties of the proposed update function applied to only two agents. Doing so will properly set a foundation to understand the same effects and properties on a large-scale model. Now we define the update function for two agents with examples and experiments.

Definition 5. (Update function for two agents with Belief Fusion and Trust) Let $\omega_X^{A[t]}$ and $\omega_X^{B[t]}$ be non-dogmatic opinions. Let \oplus be a belief fusion operator. Define the update function for $\omega_X^{A[t+1]}$ as

$$\omega_X^{A[t+1]} = \omega_X^{A[t]} \oplus (\omega_B^A \otimes \omega_X^{B[t]}). \quad (3)$$

We call $(\omega_B^A \otimes \omega_X^{B[t]})$ the opinion that A will learn by interacting with B . $\omega_X^{A[t+1]}$ is the opinion that A holds after merging their previous opinion $(\omega_X^{A[t]})$ with the opinion that A learned $(\omega_B^A \otimes \omega_X^{B[t]})$.

Example 4. For brevity, we write the value of $\mathbf{P}_B^A(t_B)$ where it should be ω_B^A . If

$$\begin{aligned} \omega_X^{A[t]} &= ((0, 0), 1, (0.5, 0.5)), & \mathbf{P}_X^{A[t]}(x) &= 0.5 \\ \omega_B^A &= ((0.5, 0.5), 0, \mathbf{a}_B^A), & \mathbf{P}_B^A(t_B) &= 0.5 \\ \omega_X^{B[t]} &= ((0.8, 0), 0.2, (0.5, 0.5)) & \mathbf{P}_X^{B[t]}(x) &= 0.9, \end{aligned} \quad (4)$$

then

$$\begin{aligned} \omega_X^{A[t+1]} &= ((0, 0), 1, (0.5, 0.5)) \oplus (0.5 \otimes ((0.8, 0), 0.2, (0.5, 0.5))) \\ &= ((0, 0), 1, (0.5, 0.5)) \oplus ((0.4, 0), 0.6, (0.5, 0.5)). \end{aligned} \quad (5)$$

Agent A is 50% sure about x and B is 90% about x . Both trust each other by 50%. As we show below, there are different outcomes depending on the choice of fusion operator.

Example 5. Now consider these two agents iteratively updating their opinion over time.

$$\begin{aligned} \omega_X^{A[0]} &= ((0.2, 0), 0.8, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) &= 0.6 & \mathbf{P}_B^A(t_B) &= 0.5 \\ \omega_X^{B[0]} &= ((0.8, 0), 0.2, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) &= 0.9 & \mathbf{P}_A^B(t_A) &= 0.5 \end{aligned} \quad (6)$$

Here, initially agent A is 60% sure about x and B is 90% about x . Both trust each other by 50%. The iterated evolution of $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ over 20 time steps is shown in Fig. 1.




Fig. 1. $\mathbf{P}_X^A(x)$ (blue) and $\mathbf{P}_X^B(x)$ (orange) updated 20 times as in Example 5.

Say that weak convergence means that A cannot move further from B upon update, i.e. $\mathbf{P}_X^{A[t]} \leq \mathbf{P}_X^{A[t+1]} \leq \mathbf{P}_X^{B[t]}$ or $\mathbf{P}_X^{A[t]} \geq \mathbf{P}_X^{A[t+1]} \geq \mathbf{P}_X^{B[t]}$. We expected that the update function would weakly converge by $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converging to some value between $\mathbf{P}_X^{A[0]}(x) = 0.6$ and $\mathbf{P}_X^{B[0]}(x) = 0.9$. But because evidence

keeps accumulating over time, with cumulative belief fusion, $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 1. Therefore, the update function with cumulative fusion does not weakly converge.

For averaging and weighted belief fusion, $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 0.5, violating weak convergence. With epistemic opinions, increasing uncertainty over time is expected, but the same happens with aleatory opinions because the trust discount removes an amount of evidence from each agent at every step.

Example 6. In this case, A and B have the same opinion and the same trust.

$$\begin{aligned} \omega_X^{A[0]} &= ((0.6, 0), 0.4, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) &= 0.8 & \mathbf{P}_B^A(t_B) &= 0.5 \\ \omega_X^{B[0]} &= ((0.6, 0), 0.4, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) &= 0.8 & \mathbf{P}_A^B(t_A) &= 0.5 \end{aligned} \quad (7)$$

Here, A and B have the same opinion. They are 80% sure about x . Both trust each other by 50%. The evolution of $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ is shown in Fig. 2.




Fig. 2. $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ (both orange) updated 20 times as in Example 6.

Even starting with the same opinion, agents A and B do not keep the same opinion over time. With cumulative belief fusion, A and B keep accumulating evidence and $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 1. With averaging or weighted belief fusion, uncertainty keeps increasing, and $\mathbf{P}_X^A(x)$ and $\mathbf{P}_X^B(x)$ converge to 0.5. The example shows none of the belief fusion operators make the update idempotent with respect to the agents' opinions. None of these operators can represent the case where agents with the same opinion, when interacting, keep the same opinion over time.

In this section, we showed that an update function when belief fusion and trust discount do not have ideal properties that can be useful to model a rational update in a social network. Even without these properties, in the next section, we show that an update function with cumulative belief fusion can represent a different kind of bias not explored in Alvim, Knight, and Valencia's work [1,2,3].

6 Analysis of the update function with cumulative belief

In this section, we analyze the cumulative belief fusion operator, which we believe in practice provides some interesting outcomes in simple simulations of social

networks, even though it does not always weakly converge, as seen above. We will show that the cumulative belief function is promising for modeling belief updates in social networks.

Recalling the definition of cumulative belief fusion, it is assumed that the amount of evidence increases by including more sources. An opinion in subjective logic can be translated to a Beta PDF using the amount of evidence for each state in the domain as parameters. Say that $\mathbf{r}_X^A(x)$ and $\mathbf{r}_X^B(x)$ are the amount of evidence supporting the state x for agents A and B respectively. The cumulative fusion operator gets the sum $\mathbf{r}_X^A(x) + \mathbf{r}_X^B(x)$ to be the amount of evidence of the merged opinion and translates it back to a subjective logic opinion. More details are presented in the appendix. For repeated interactions using the cumulative belief fusion, we interpret repeated evidence contained in an opinion as reinforcement of the belief in the state.

We did many simulations using the two-agent update function with cumulative fusion and we found three cases for initial epistemic opinions with the projected probability $\mathbf{P}_X^A(x)$ ranging from 0 to 1.

Suppose that $\mathbb{X} = \{x, \bar{x}\}$ and all opinions are epistemic. Recall that the update function for two agents with trust discount and belief fusion is

$$\omega_X^{A[t+1]} = \omega_X^{A[t]} \oplus (\omega_B^A \otimes \omega_X^{B[t]}) \quad (8)$$

where A and B are agents, $\omega_X^{A[t]}$ and $\omega_X^{B[t]}$ are non-dogmatic epistemic opinions, and ω_B^A is a trust opinion. Also, recall that $\omega_B^A \otimes \omega_X^{B[t]}$ is the opinion that A learns after trusting B about X . The update function behaves like these three cases:

1. *Consensus*: This case happens when both agents are agreeing or disagreeing at the same time, not necessarily with the same projected probability, i.e.

$$\begin{aligned} &\mathbf{P}_X^{A[0]}(x) < 0.5 \text{ and } \mathbf{P}_B^A(t_b) \otimes \mathbf{P}_X^{B[0]}(x) < 0.5 \text{ or} \\ &\mathbf{P}_X^{A[0]}(x) > 0.5 \text{ and } \mathbf{P}_B^A(t_b) \otimes \mathbf{P}_X^{B[0]}(x) > 0.5 \end{aligned} \quad (9)$$

When these agents interact, they will accumulate evidence about the same outcome in each interaction. The fusion leads $\mathbf{b}_X^A(x)$ and $\mathbf{b}_X^B(x)$ to converge to 0 or 1, depending on what they are agreeing upon, and both with uncertainty mass to 0. Increasing the trust discount or the uncertainty will increase the speed of convergence.

This represents a situation when two agents agree that x is TRUE, with different levels of projected probabilities. That leads them to be completely certain about the proposition.

Example 7. Let A and B be initially in consensus agreeing that x is TRUE as follows.

$$\begin{aligned} \omega_X^{A[0]} &= ((0.2, 0.0), 0.8, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}x &= 0.6 & \mathbf{P}_B^A(t_b) &= 0.5 \\ \omega_X^{B[0]} &= ((0.4, 0.0), 0.6, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}x &= 0.7 & \mathbf{P}_A^B(t_b) &= 0.5 \end{aligned} \quad (10)$$

A is 60% sure about x and B is 70% sure about x . Both trust each other by 50%. On the limit of the interactions between A and B , both will be sure by 100% that x is true.

$$\mathbf{P}_X^{A[t]}(x) = \mathbf{P}_X^{B[t]}(x) \xrightarrow{t \rightarrow \infty} 1 \quad (11)$$

Similarly if both disagree that x is TRUE, i.e. that \bar{x} is TRUE, $\mathbf{P}_X^{A[t]}(x)$ and $\mathbf{P}_X^{B[t]}(x)$ will converge to 0.




Fig. 3. On the left, the evolution of $\mathbf{P}_X^{A[t]}(x)$ and $\mathbf{P}_X^{B[t]}(x)$ for the Example 7. On the right, the similar case for when both agents disagree that x is TRUE.

2. *Balanced opposite*: This case happens when both agents are learning the exact opposite opinion that A already had, i.e.

$$\mathbf{P}_X^{A[0]}(x) = 1 - \mathbf{P}_B^A(t_b) \otimes \mathbf{P}_X^{B[0]}(x), \quad \forall x \in \mathbb{X} \quad (12)$$

Because contrary pieces of evidence cancel each other, the fusion leads $\mathbf{b}_X^{A[0]}(x)$ to be a vacuous opinion. The speed of convergence is defined by the trust opinion. The more an agent trusts another, the faster the convergence. This represents a situation when two agents support opposite views but with the same projected probability. That leads them to be completely indecisive about the proposition.

Example 8. Let A and B have balanced opposite opinions.

$$\begin{aligned} \omega_X^{A[0]} &= ((0, 0.4), 0.6, (0.5, 0.5)) & \mathbf{P}_X^{A[0]}(x) &= 0.3 & \mathbf{P}_B^A(t_B) &= 1 \\ \omega_X^{B[0]} &= ((0.4, 0), 0.6, (0.5, 0.5)) & \mathbf{P}_X^{B[0]}(x) &= 0.7 & \mathbf{P}_A^B(t_A) &= 1 \end{aligned} \quad (13)$$

A is 30% sure about x and B is 70% sure about x . Both trust each other by 100%. On the limit of interactions between A and B , both will be sure by 50% that x is true.

$$\mathbf{P}_X^{A[t]}(x) = \mathbf{P}_X^{B[t]}(x) \xrightarrow{t \rightarrow \infty} 0.5 \quad (14)$$




Fig. 4. Evolution of $\mathbf{P}_X^{A[t]}(x)$ and $\mathbf{P}_X^{B[t]}(x)$ for the Example 8.

3. *Unbalanced opposite*: This happens when there are opposite beliefs, but they don't have the same projected probability in their respective views. In other words, both are in conflict but one agent is more radical than the other.

$$\begin{aligned} & (\mathbf{P}_X^{A[0]}(x) < 0.5 \text{ and } \mathbf{P}_B^A(t_b) \oplus \mathbf{P}_X^B > 0.5 \\ \text{or } & \mathbf{P}_X^{A[0]}(x) > 0.5 \text{ and } \mathbf{P}_B^A(t_b) \oplus \mathbf{P}_X^B < 0.5) \\ & \text{and } \mathbf{P}_X^{A[0]} \neq 1 - \mathbf{P}_B^A(t_b) \mathbf{P}_X^{B[0]}(x) \end{aligned} \quad (15)$$

In this case, contrary pieces of evidence will cancel each other, but one agent will transmit more evidence than the other upon interaction. The agents that transmit more evidence will win in the limit. Increasing the trust or becoming more radical will increase the speed of convergence.

- If A is sure that x is TRUE more than B is sure that x is FALSE discounted by the trust, i.e.

$$\mathbf{P}_X^{A[t]}(x) > 1 - \mathbf{P}_B^A(t_b) \otimes \mathbf{P}_X^{B[t]}(x) \quad (16)$$

then, A and B will eventually both agree that x is TRUE in some degree, i.e.

$$\mathbf{P}_X^{A[t]}(x) = \mathbf{P}_X^{B[t]}(x) \xrightarrow{t \rightarrow \infty} p \in (0.5, 1] \quad (17)$$

- If A is sure that x is FALSE more than B is sure that x is TRUE discounted by the trust, i.e.

$$\mathbf{P}_X^{A[t]}(x) < 1 - \mathbf{P}_B^A(t_b) \otimes \mathbf{P}_X^{B[t]}(x) \quad (18)$$

then, A and B will eventually both agree that x is FALSE in some degree, i.e.

$$\mathbf{P}_X^{A[t]}(x) = \mathbf{P}_X^{B[t]}(x) \xrightarrow{t \rightarrow \infty} p \in [0, 0.5) \quad (19)$$

The unbalanced opposite is the most interesting case. It can represent some bias when conflicting agents are interacting. The behavior is different from any update function described by [3]. After experiments, we found that

- If the agents are too far, they will radicalize, i.e. $\mathbf{P}_X^{A[t]}(x) \xrightarrow{t \rightarrow \infty} 0$ or $\mathbf{P}_X^{A[t]}(x) \xrightarrow{t \rightarrow \infty} 1$
- Otherwise, they will converge to some value at the winning outcome. $[0, 0.5)$ if A is sure that x is FALSE more than B is sure that x is TRUE discounted by the trust, or $(0.5, 1]$ if A is sure that x is TRUE more than B is sure that x is FALSE discounted by the trust.

This behavior shows that we have a fixed-point for $\mathbf{P}_X^{A[0]}(x)$ for each $\mathbf{P}_X^{B[t]}(x)$ such that $\mathbf{P}_X^{B[0]}(x) = \mathbf{P}_X^{B[t]}(x)$ when t goes to infinity, in other words, for every initial A 's opinion there is an initial B 's opinion such that A 's and B 's opinions will both eventually converge to B 's initial opinion, when A and B communicate repeatedly. Fig. 5 shows the results from the experiments. We found the fixed-points by a binary search for each $\mathbf{P}_X^{A[0]}(x)$.




Fig. 5. Fixed point for $\mathbf{P}_X^{A[0]}(x)$ for each $\mathbf{P}_X^{B[t]}(x)$

The experiments show that the fixed-point function has a curve similar to the logistic function. Each fixed-point represents the boundary between the region where two agents will radicalize and the region where they will eventually agree on some less radical point. We believe that the radicalization phenomenon happens because the agents are transmitting evidence in an *unstable* way, when the proportion of evidence x vs. \bar{x} increases over time. They do not radicalize when the agents are transmitting evidence in a *stable* way, when the proportion of evidence x vs. \bar{x} is the same over time.

7 Future Work and Conclusions

In this work, we described a subjective logic model for social networks. It is defined by a set of agents representing users in a social network, a domain of

disjoint events that can be used to represent a topic that cannot be answered with simply YES or NO, a set of multinomial opinions to represent the opinion of each user, trust opinions to represent the relationship between users and an update function to represent the interaction between users.

The main focus of this work is to define an update function. Many update functions can be used depending on what kind of interaction is represented. We did experiments with three update functions using cumulative belief fusion, averaging belief fusion, and weighted belief fusion, with trust discount. The experiments showed that none of these functions have useful properties to model the rational update function described in [3] such as idempotency, weak convergence, and non-increasing uncertainty.

Through experiments, we showed that our update function with cumulative belief fusion has the potential to represent a kind of interaction not described in [3], even though it does not meet our initial desired properties for an update function. We also showed that if the agents disagree about a proposition, i.e. they have opposite opinions, there is a close enough distance between them where the update is stable, and they converge to a non-radical point. But if they are far enough, they will still converge, but also radicalize, converging to completely agree or completely disagree about the proposition. This is different from the update function describing confirmation bias from [3], where . The next step for the update function with cumulative fusion is to have a clear definition of this behavior by calculating the convergence of the update function for each case.

References

1. Alvim, M.S., Amorim, B., Knight, S., Quintero, S., Valencia, F.: A multi-agent model for polarization under confirmation bias in social networks. In: Peters, K., Willemse, T.A.C. (eds.) *Formal Techniques for Distributed Objects, Components, and Systems - 41st IFIP WG 6.1 International Conference, FORTE 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings*. Lecture Notes in Computer Science, vol. 12719, pp. 22–41. Springer (2021). https://doi.org/10.1007/978-3-030-78089-0_2
2. Alvim, M.S., Amorim, B., Knight, S., Quintero, S., Valencia, F.: A formal model for polarization under confirmation bias in social networks. *Log. Methods Comput. Sci.* **19**(1) (2023). [https://doi.org/10.46298/lmcs-19\(1:18\)2023](https://doi.org/10.46298/lmcs-19(1:18)2023)
3. Alvim, M.S., Knight, S., Valencia, F.: Toward a formal model for group polarization in social networks. In: *The Art of Modelling Computational Systems*. LNCS, vol. 11760, pp. 419–441. Springer (2019)
4. Bala, V., Goyal, S.: Learning from Neighbours. *The Review of Economic Studies* **65**(3), 595–621 (1998). <https://doi.org/10.1111/1467-937X.00059>, <https://doi.org/10.1111/1467-937X.00059>
5. Baltag, A., Christoff, Z., Rendsvig, R.K., Smets, S.: Dynamic epistemic logics of diffusion and prediction in social networks. *Studia Logica* **107**(3), 489–531 (2019)
6. Christoff, Z.L.: *Dynamic logics of networks: information flow and the spread of opinion*. Ph.D. thesis, Universiteit van Amsterdam (2016)
7. DeGroot, M.H.: Reaching a consensus. *Journal of the American Statistical Association* **69**(345), 118–121 (1974). <https://doi.org/10.2307/2285509>

8. Holliday, W.H.: Trust and the dynamics of testimony. In: Logic and Interactive RAtionality. Seminar's yearbook. pp. 118–142 (2009), <https://events.ilc.uva.nl/dg/wp-content/uploads/2012/07/yearbook.2009.pdf>
9. Hunter, A.: Reasoning about trust and belief change on a social network: A formal approach. In: Liu, J.K., Samarati, P. (eds.) Information Security Practice and Experience 2017. Lecture Notes in Computer Science, vol. 10701, pp. 783–801. Springer (2017)
10. Jøsang, A.: Subjective Logic - A Formalism for Reasoning Under Uncertainty. Artificial Intelligence: Foundations, Theory, and Algorithms, Springer (2016)
11. Jøsang, A.: Categories of belief fusion. Journal of Advances in Information Fusion **13**(2), 235–254 (2018)
12. Liu, F., Seligman, J., Girard, P.: Logical dynamics of belief change in the community. Synthese **191**(11), 2403–2431 (2014)
13. Pedersen, M.Y., Smets, S., Ågotnes, T.: Modal logics and group polarization. J. Log. Comput. **31**(8), 2240–2269 (2021). <https://doi.org/10.1093/logcom/exab062>
14. Xiong, Z., Ågotnes, T.: On the logic of balance in social networks. J. Log. Lang. Inf. **29**(1), 53–75 (2020). <https://doi.org/10.1007/s10849-019-09297-0>
15. Young Pedersen, M.: Polarization and Echo Chambers: A Logical Analysis of Balance and Triadic Closure in Social Networks. Master's thesis, Universiteit van Amsterdam (2019)

A Appendix: Mapping an opinion to a Beta PDF

This appendix shows the mapping of an opinion to a Beta PDF defined by Jøsang [10] We show the evidence notation of the multivariate Beta PDF (or Dirichlet PDF) and show its equivalence with the opinion representation defined at Def. 2.

Multinomial probability density over a domain of cardinality k is expressed by the k -dimensional Beta PDF. Assume a domain \mathbb{X} of cardinality k , and a random variable X over \mathbb{X} with probability distribution \mathbf{p}_X . The Beta PDF can be used to represent *probability density* over \mathbf{p}_X .

The multivariate Beta PDF takes as variable the k -dimensional probability distribution \mathbf{p}_X . The strength parameters for the k possible outcomes are represented as k positive real numbers $\alpha_X x$, each corresponding to one of the possible values $x \in \mathbb{X}$. The strength parameters represent evidence/observations of $X = x$ where $x \in \mathbb{X}$.

Definition 6. (Multivariate Beta Probability Density Function). *Let \mathbb{X} be a domain consisting of k mutually disjoint values. Let α_X represent the strength vector over the values of \mathbb{X} , and let \mathbf{p}_X denote the probability distribution over \mathbb{X} . With \mathbf{p}_X as a k -dimensional variable, the multivariate Beta PDF denoted $\text{Beta}(\mathbf{p}_X, \alpha_X)$ is expressed as:*

$$\text{Beta}(\mathbf{p}_X, \alpha_X) = \frac{\Gamma\left(\sum_{x \in \mathbb{X}} \alpha_X(x)\right)}{\prod_{x \in \mathbb{X}} \Gamma(\alpha_X(x))} \prod_{x \in \mathbb{X}} \mathbf{p}_X(x)^{(\alpha_X(x)-1)}, \text{ where } \alpha_X(x) \geq 0, \quad (20)$$

with the restrictions that $\mathbf{p}_X(x) \neq 0$ if $\alpha_X(x) < 1$.

Now assume that $x \in \mathbb{X}$ represents a frequentist event. Let $\mathbf{r}_X(x)$ denote the number of observations for x . The strength vector α_X can be expressed as a function of the observations $\mathbf{r}_X(x)$ and the base rate \mathbf{a}_X :

$$\alpha_X(x) = \mathbf{r}_X(x) + \mathbf{a}_X(x)W, \text{ where } \mathbf{r}_X(x) \geq 0 \ \forall x \in \mathbb{X}. \quad (21)$$

By expressing the strength vector α_X in terms of the evidence observation \mathbf{r}_X , the base rate \mathbf{a}_X , and the non-informative prior weight W , we get the representation of the Beta PDF denoted $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$. The exact definition of $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$ is described at [10].

Given a Beta PDF $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$, the expected distribution over \mathbb{X} can be written as

$$\mathbf{E}_X(x) = \frac{\mathbf{r}_X(x) + \mathbf{a}_X(x)W}{W + \sum_{x_j \in \mathbb{X}} \mathbf{r}_X(x_j)}, \ \forall x \in \mathbb{X}. \quad (22)$$

The Beta model translates observation evidence directly into a PDF over a k -component probability variable. The representation evidence, together with the base rate, can be used to determine subjective opinions. In other words, it is possible to define a bijective mapping between Beta PDFs and opinions.

The bijective mapping between ω_X and $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$ is based on the requirement for equality between the projected probability distribution \mathbf{P}_X derived from ω_X and the expected probability distribution \mathbf{E}_X derived from $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$. This means that the more evidence in favor of a particular outcome x , the greater the belief mass on the outcome. Furthermore, the more total evidence available, the less uncertainty mass.

Definition 7. (Mapping: Opinion \leftrightarrow Beta PDF) *Let $\omega_X = (\mathbf{b}_X, u_X, \mathbf{a}_X)$ be an opinion and let $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X, \mathbf{a}_X)$ be a Beta PDF, both over the same variable $X \in \mathbb{X}$. These are equivalent through the following mapping,*

$$\forall x \in \mathbb{X} \quad \left\{ \begin{array}{l} \mathbf{b}_X(x) = \frac{\mathbf{r}_X(x)}{W + \sum_{x_i \in \mathbb{X}} \mathbf{r}_X(x_i)} \\ u_X = \frac{W}{W + \sum_{x_i \in \mathbb{X}} \mathbf{r}_X(x_i)} \end{array} \right\} \Leftrightarrow \left\{ \begin{array}{l} \left\{ \begin{array}{l} \mathbf{r}_X(x) = \frac{W\mathbf{b}_X(x)}{u_X} \\ 1 = u_X + \sum_{x_i \in \mathbb{X}} \mathbf{b}_X(x_i) \end{array} \right. \quad \text{if } u_X \neq 0 \\ \left\{ \begin{array}{l} \mathbf{r}_X(x) = \mathbf{b}_X(x) \cdot \infty \\ 1 = \sum_{x_i \in \mathbb{X}} \mathbf{b}_X(x_i) \end{array} \right. \quad \text{if } u_X = 0 \end{array} \right. \quad (23)$$

This equivalence between opinions and Beta PDFs is very powerful because it makes it possible to determine opinions from statistical observations.

B Appendix: Belief fusion operators

This appendix shows a more detailed definition of the belief fusion operators intuitively described at Sec. 3. Here, we will define the belief fusions in terms

of Beta PDFs. Since there is a mapping between opinions and Beta PDFs, the direct definition for each operator will be omitted. These definitions can be found at [10,11].

Cumulative belief fusion The cumulative belief fusion is used when it is assumed that the amount of independent evidence increases by including more sources. The idea is to sum the amount of evidence of the opinions.

Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . W.l.o.g., let A and B be agents, and ω_X^A and ω_X^B be opinions. The cumulative belief fusion between A and B is denoted $\omega_X^{(A \diamond B)}$ and it can be represented as a Beta PDF.

Definition 8. (Cumulative Belief Fusion) *Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . Let A and B be agents, and ω_X^A and ω_X^B be opinions. The cumulative belief fusion between ω_X^A and ω_X^B is denoted:*

$$\omega_X^{(A \diamond B)} = \omega_X^A \oplus \omega_X^B \quad (24)$$

Let $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A)$ and $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B)$ be the Beta PDFs equivalent to ω_X^A and ω_X^B , respectively.

The opinion $\omega_X^{(A \diamond B)}$ is the opinion equivalent to $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \diamond B)}, \mathbf{a}_X^{(A \diamond B)})$ defined as:

$$\begin{aligned} \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \diamond B)}, \mathbf{a}_X^{(A \diamond B)}) &= \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A) \oplus \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B) \\ &= \text{Beta}_X^e(\mathbf{p}_X, (\mathbf{r}_X^A + \mathbf{r}_X^B), \mathbf{a}_X^{(A \diamond B)}). \end{aligned} \quad (25)$$

More specifically, for each values $x \in \mathbb{X}$ the accumulated source evidence $\mathbf{r}^{(A \diamond B)}$ is computed as:

$$\mathbf{r}^{(A \diamond B)}(x) = \mathbf{r}_X^A(x) + \mathbf{r}_X^B(x), \quad \forall x \in \mathbb{X}. \quad (26)$$

The fusion of three or more opinions is defined at [11]. It can be verified that the cumulative fusion operator is commutative, associative, and non-idempotent.

Averaging Belief Fusion The averaging belief fusion is used when including more sources does not mean that more evidence supports the conclusion. The idea is to take the average of the amount of evidence of the opinions.

Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . W.l.o.g., let A and B be agents, and ω_X^A and ω_X^B be opinions. The averaging belief fusion between A and B is denoted $\omega_X^{(A \oslash B)} = \omega_X^A$ and it can be represented as a Beta PDF.

Definition 9. (Averaging Belief Fusion) *Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . Let A and B be agents, and ω_X^A and ω_X^B be opinions. The averaging belief fusion between ω_X^A and ω_X^B is denoted:*

$$\omega_X^{(A \oslash B)} = \omega_X^A \oslash \omega_X^B \quad (27)$$

Let $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A)$ and $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B)$ be the Beta PDFs equivalent to ω_X^A and ω_X^B , respectively.

The opinion $\omega_X^{(A \circ B)}$ is the opinion equivalent to $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \circ B)}, \mathbf{a}_X^{(A \circ B)})$ defined as:

$$\begin{aligned} \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \circ B)}, \mathbf{a}_X^{(A \circ B)}) &= \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A) \oplus \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B) \\ &= \text{Beta}_X^e(\mathbf{p}_X, (\mathbf{r}_X^A + \mathbf{r}_X^B)/2, \mathbf{a}_X^{(A \circ B)}). \end{aligned} \quad (28)$$

More specifically, for each values $x \in \mathbb{X}$ the accumulated source evidence $\mathbf{r}^{(A \circ B)}$ is computed as:

$$\mathbf{r}^{(A \circ B)}(x) = \frac{\mathbf{r}_X^A(x) + \mathbf{r}_X^B(x)}{2}, \quad \forall x \in \mathbb{X}. \quad (29)$$

The fusion of three or more opinions is defined at [11]. It can be verified that the cumulative fusion operator is commutative, idempotent, but non-associative.

Weighted Belief Fusion The weighted belief fusion is used when including more sources does not mean that more evidence supports the conclusion. The idea is used when we take the average of the amount of evidence of the opinions weighted by their lack of uncertainty. In particular, opinions with no belief mass are rejected.

Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . W.l.o.g., let A and B be agents, and ω_X^A and ω_X^B be opinions. The averaging belief fusion between A and B is denoted $\omega_X^{(A \hat{\circ} B)} = \omega_X^A$ and it can be represented as a Beta PDF.

Definition 10. (Weighted Belief Fusion) Let \mathbb{X} be a domain, and X be a random variable over \mathbb{X} . Let A and B be agents, and ω_X^A and ω_X^B be opinions. The averaging belief fusion between ω_X^A and ω_X^B is denoted:

$$\omega_X^{(A \hat{\circ} B)} = \omega_X^A \hat{\oplus} \omega_X^B \quad (30)$$

Let $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A)$ and $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B)$ be the Beta PDFs equivalent to ω_X^A and ω_X^B , respectively. Also, let c_X^A be the confidence that A has in their opinion. Formally:

$$c_X^A = 1 - \sum_{x \in \mathbb{X}} \mathbf{b}_X^A(x) \quad (31)$$

The opinion $\omega_X^{(A \hat{\circ} B)}$ is the opinion equivalent to $\text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \hat{\circ} B)}, \mathbf{a}_X^{(A \hat{\circ} B)})$ defined as:

$$\begin{aligned} \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^{(A \hat{\circ} B)}, \mathbf{a}_X^{(A \hat{\circ} B)}) &= \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^A, \mathbf{a}_X^A) \hat{\oplus} \text{Beta}_X^e(\mathbf{p}_X, \mathbf{r}_X^B, \mathbf{a}_X^B) \\ &= \text{Beta}_X^e(\mathbf{p}_X, (c_X^A \mathbf{r}_X^A + c_X^B \mathbf{r}_X^B)/2, \mathbf{a}_X^{(A \hat{\circ} B)}). \end{aligned} \quad (32)$$

More specifically, for each values $x \in \mathbb{X}$ the accumulated source evidence $\mathbf{r}^{(A \hat{\circ} B)}$ is computed as:

$$\mathbf{r}^{(A \hat{\circ} B)}(x) = \frac{c_X^A \mathbf{r}_X^A(x) + c_X^B \mathbf{r}_X^B(x)}{c_X^A + c_X^B}, \quad \forall x \in \mathbb{X}. \quad (33)$$

The fusion of three or more opinions is defined at [11]. It can be verified that the cumulative fusion operator is commutative, idempotent, but non-associative.