

Evidence-based discounting rule in Subjective Logic (extended abstract)

Boris Škorić¹, Sebastiaan J.A. de Hoogh², Nicola Zannone¹

1) TU Eindhoven. {b.skoric, n.zannone}@tue.nl

2) Philips Research. sebastiaan.de.hoogh@philips.com

Abstract

We identify an inconsistency in Subjective Logic caused by the discounting operator ‘ \otimes ’. We propose a new operator, ‘ \boxtimes ’, which resolves all the consistency problems. The new algebra makes it possible to compute Subjective Logic trust values (reputations) in arbitrarily connected trust networks. The material presented here is an excerpt of [3].

1 Subjective Logic

Subjective Logic (SL) [1] is a kind of ‘fuzzy’ logic that explicitly keeps track of uncertainties. The central concepts in SL are *evidence* and *opinions*. Let P be a proposition. Evidence about P is denoted as a vector (p, n) , where p is the amount of evidence supporting P , and n the amount of evidence supporting $\neg P$. An *opinion* is a triplet $(b, d, u) \in [0, 1]^3$ satisfying $b + d + u = 1$. The b component is the ‘belief’ in proposition P , and it can be interpreted as the probability that P is provably true given the evidence. Likewise, the d is disbelief (belief in $\neg P$). The u is the uncertainty, the probability that nothing can be proven about P . There is a simple bijection between the evidence vector and the opinion based on it,

$$(b, d, u) = \frac{(p, n, 2)}{p + n + 2}; \quad (p, n) = 2 \frac{(b, d)}{u}. \quad (1)$$

This relation is based on an analysis of a posteriori probability distributions (beta distributions) [1]. Special points are Belief $B = (1, 0, 0)$, Disbelief $D = (0, 1, 0)$ and Uncertainty $U = (0, 0, 1)$. Triplets with $u = 0$ can only be reached with infinite amounts of evidence and are therefore often excluded from opinion space.

There are two important operations for combining opinions about the same proposition: *consensus* and *discounting*. The consensus operation simply adds up evidence vectors. Let $x = (x_b, x_d, x_u)$ be an opinion based on evidence (p_x, n_x) and $y = (y_b, y_d, y_u)$ an opinion based on (p_y, n_y) . Then the combined evidence is $(p_x + p_y, n_x + n_y)$ and the corresponding opinion is given by

$$x \oplus y \stackrel{\text{def}}{=} \frac{(x_u y_b + y_u x_b, x_u y_d + y_u x_d, x_u y_u)}{x_u + y_u - x_u y_u}. \quad (2)$$

The consensus operation \oplus is allowed only if the evidence in x and y is independent, otherwise ‘double counting’ of evidence occurs.

Discounting describes trust transitivity. Let Bob publish opinion y about proposition P . Let Alice have opinion x about Bob’s trustworthiness. Then Alice’s opinion about P is ‘ y discounted through x ’, which is denoted as $x \otimes y$ and defined as

$$x \otimes y \stackrel{\text{def}}{=} (x_b y_b, x_b y_d, x_d + x_u + x_b y_u). \quad (3)$$

2 Problems with the \otimes operator

The definition of \otimes lacks a natural interpretation in evidence space. Let $z = x \otimes y$ in the Alice & Bob example above. The evidence vector (p_z, n_z) obtained using (1) is a messy function of p_x, p_y, n_x and n_y which under certain circumstances yields downright weird results. For instance, if $n_x = 0, n_y = 0$ and $p_y \gg p_x$, then $p_z \approx p_x/4$, which seems to imply that Alice's opinion about P is fully determined by x (which is not even an opinion about P), independent of y .

Furthermore, consider the following case. Alice has trust x in Bob. Bob gathers two independent evidence vectors, (p_1, n_1) and (p_2, n_2) , about proposition P .

Scenario I: Bob forms two independent opinions, y_1 and y_2 , based on the evidence. He publishes first y_1 and later y_2 . Alice forms opinion $x \otimes y_1$ about P and later updates this to $(x \otimes y_1) \oplus (x \otimes y_2)$.

Scenario II: Bob combines his evidence and forms opinion $y_1 \oplus y_2$, which he publishes. Alice forms opinion $x \otimes (y_1 \oplus y_2)$ about P .

It is obvious that these scenarios should yield the same result for Alice. Yet the traditional discounting rule gives $x \otimes (y_1 \oplus y_2) \neq (x \otimes y_1) \oplus (x \otimes y_2)$. In SL the only correct expression is $x \otimes (y_1 \oplus y_2)$. We consider this to be a grave inconsistency in SL.

Next consider the trust network in Fig. 1. Due to the complicated mixup of evidence components in expressions of the form $x \otimes y$, combined with the prohibition on combining dependent evidence in \oplus operations, it is impossible to write down a consistent SL result ('canonical expression' [2]) expressing the trust that node 1 has in node 6.




Figure 1: *Example of a trust network that is problematic for Subjective Logic.*

3 Bijection between evidence and opinion: Simplified derivation

We have found a simple way to obtain a bijection between evidence (p, n) and opinion $x = (b, d, u)$. Instead of looking at a posteriori probability distributions, we ask ourselves which natural constraints should be satisfied by such a bijection. If we impose the following conditions,

1. $b/d = p/n$
2. $b + d + u = 1$
3. $p + n = 0 \Rightarrow u = 1$
4. $p + n \rightarrow \infty \Rightarrow u \rightarrow 0$

then the relation between x and (p, n) can only be

$$x = (b, d, u) = \frac{(p, n, c)}{p + n + c} \quad ; \quad (p, n) = c \frac{(b, d)}{u} \quad (4)$$

where $c > 0$ is a constant. Eq. (4) is precisely of the form (1), except for the constant '2' versus c . We make two important remarks: (i) The more generic mapping (4) is consistent with the \oplus definition (2), i.e. the value of c does not matter, as long as all entities use the same c . (ii) The analysis of [1] can be re-done using a general constant c , and then still yields a consistent result. We see no reason to set $c = 2$.

4 New discounting operator: \boxtimes

We first define a new operation in Subjective Logic, *multiplication of a scalar and an opinion*. Let $x = (b, d, u)$ be an opinion based on evidence (p, n) . Let $\lambda \geq 0$ be a scalar. In evidence space the product $\lambda \cdot x$ is defined as $(\lambda p, \lambda n)$. In opinion space this corresponds to the definition

$$\lambda \cdot x \stackrel{\text{def}}{=} \frac{(\lambda b, \lambda d, u)}{\lambda(b+d) + u}. \quad (5)$$

Next we define our new discounting operator ‘ \boxtimes ’. Let g be a function that maps opinions to $[0, 1]$, satisfying $g(B) = 1$ and $g(D) = 0$. We define

$$x \boxtimes y \stackrel{\text{def}}{=} g(x) \cdot y. \quad (6)$$

The function g can be chosen at will, depending on the context.

We refer to {SL with the new discounting operator} as Evidence-Based Subjective Logic (EBSL). EBSL avoids all the inconsistencies of the \otimes operation,

- The expression $x \boxtimes y$ has a very simple interpretation in evidence space: Due to the disbelief and uncertainty present in x , only a fraction $g(x)$ of the evidence in y is accepted by the recipient.
- It holds that $x \boxtimes (y_1 \oplus y_2) = (x \boxtimes y_1) \oplus (x \boxtimes y_2)$, which is what is intuitively expected of a discounting operation.
- Due to the cleanness of the \boxtimes operation, there is a strict separation between evidence on the one hand and the way it is carried over trust links on the other hand. Consequently EBSL can handle any trust network, no matter how complicated the graph. (See the next section.)

5 Arbitrary trust networks

Let opinion A_{ij} be the amount of trust that a node i has in node j , based on *direct* evidence, e.g. past interaction between i and j . We set the diagonal to $A_{ii} = U$. All nodes publish these direct opinions. Every node wants to know how much the other nodes can be trusted, and is willing to make use of the opinions published by others (‘indirect evidence’). The mathematical problem is now to compute a meaningful reputation matrix R from A , giving proper weights to all the direct and indirect evidence. The diagonal of R is undefined, so we are free to set it arbitrarily. We set it to $B\mathbf{1}$, where $\mathbf{1}$ is the unit matrix. The following relation must be satisfied,

$$R = B\mathbf{1} \oplus (R \boxtimes A), \quad (7)$$

where the ‘matrix multiplication’ $R \boxtimes A$ is defined as $(R \boxtimes A)_{ij} = \oplus_k (R_{ik} \boxtimes A_{kj})$. Eq. (7) says that a reputation R_{ij} consists of a weighted sum of direct opinions A_{kj} , where the weights are determined by the reputations R_{ik} . Eq. (7) is a fixed-point equation. It can be solved e.g. by iterative methods such as repeatedly substituting (7) into itself. Experiments on synthetic as well as real data show fast convergence.

References

- [1] A. Jøsang. A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–311, 2001.
- [2] A. Jøsang, E. Gray, and M. Kinatader. Simplification and Analysis of Transitive Trust Networks. *Web Intelligence and Agent Systems*, 4:139–161, 2006.
- [3] B. Škorić, S.J.A. de Hoogh, and N. Zannone. Flow-based reputation with uncertainty: Evidence-Based Subjective Logic, 2014. <http://arxiv.org/abs/1402.3319>.