Subjective Logic Extensions for the Semantic Web

Davide Ceolin, Archana Nottamkandath, and Wan Fokkink
{d.ceolin, a.nottamkandath, w.j.fokkink}@vu.nl

VU University, Amsterdam, The Netherlands

Abstract. Subjective logic is a powerful probabilistic logic which is useful to handle data in case of uncertainty. Subjective logic and the Semantic Web can mutually benefit from each other, since subjective logic is useful to handle the inner noisiness of the Semantic Web data, while the Semantic Web offers a mean to obtain evidence useful for performing evidential reasoning based on subjective logic. In this paper we propose three extensions and applications of subjective logic in the Semantic Web, namely: the use of semantic similarity measures for weighing subjective opinions, a way for accounting for partial observations, and the new concept of "open world opinion", i.e. subjective opinions based on Dirichlet Processes, which extend multinomial opinions. For each of these extensions, we provide examples and applications to prove their validity.

Keywords: Subjective Logic, Semantic Similarity, Dirichlet Process, Partial Observations

1 Introduction

Subjective logic [7] is a probabilistic logic widely adopted in the trust management domain, based on evidential reasoning and statistical principles. This logic focuses on the representation and the reasoning on assertions of which truth value is not fully determined, but estimated on the basis of the observed evidence. The logic comes with a variety of operators that allow to combine such assertions and to derive the truth values of the consequences.

Subjective logic is well-suited for the management of uncertainty within the Semantic Web. For instance, the incremental access to these data (as a consequence of crawling) can give rise to uncertainty issues which can be dealt with using this logic. Furthermore, the fact that the fulcrum of this logic is the concept of "subjective opinion" (which represent an assertion, its corresponding evidence and the source of this evidence), allows to correctly represent how the estimated truth value of an assertion is bound to the source of the corresponding evidence and allows to easily keep lightweight provenance information. Finally, evidential reasoning allows to limit the typical noisiness of Semantic Web data. On the other hand, we also believe that the Semantic Web can be beneficial to this logic, as an immeasurably important source of information: since the truth value of assertions is based on availability of observations, the more data is available (hopefully of high enough quality), the closer we can get to the correct truth value for our assertions. We believe that this mutual relationship can be improved. This paper proposes extensions and applications of subjective logic that aim at the Semantic Web.

The rest of the paper is organized as follows: Section 2 describes related work, Section 3 proposes a combination of subjective logic and semantic similarity measures, Section 4 introduces a method for dealing with partial observations of evidence, Section 5 introduces the concept of Open World Opinion. Section 6 provides a final discussion about the work presented.

2 Related work

The development of subjective logic's operators has been investigated. Remarkably, the averaging and cumulative fusion [8,9] and the discounting [11] operators are among the most generic and useful operators for this logic. These operators provide the foundations for the work proposed in this paper. The connections among subjective logic and the (Semantic) Web are increasing. Ceolin et al. [4] adopt this logic for computing trust values of annotations provided by experts, using DBpedia and other Web sources as evidence. Unlike this work, they do not use semantic similarity measures. Ceolin et al. [3] and Bellenger et al. [1] provide applications of the combination of evidential reasoning with semantic similarity measures and Semantic Web technologies. In the current paper we provide the theoretical foundations for this kind of approaches, and we generalize them. Sensoy et al. [15] use semantic similarity in combination with subjective logic to import knowledge from one context to another. They use the semantic similarity measure to compute a prior value for the imported data, while we use it to weigh all the available evidence. Kaplan et al. [12] focus on the exploration of uncertain partial observations used for building subjective opinions. Unlike their work, we restrict our focus on partial observations of Web-like data and evaluations, which comprise the number of "likes", links and other similar indicators related to a given Web item. The weighing and discounting based on semantic similarity measures can resemble the work of Jøsang et al. [8], although the additional information that we include in our reasoning (which is semantic similarity) is related only to the frame of discernment in subjective logic, and not to the belief assignment function.

3 Combining Subjective Logic with Semantic Similarity

3.1 Preliminaries

Subjective Logic In subjective logic, so-called "subjective opinions" express the belief that source x owns with respect to the value of assertion y chosen among the elements of the set Θ ("frame of discernment"). The belief is assigned to the elements of the set $X = 2^{\Theta} \setminus \Theta$ ("frame"), according to the evidence. In symbols, this is represented as $\omega(b, d, u, a)$ when $|\Theta| = 2$ (binomial opinion) or as $\omega(\overrightarrow{b}, u, \overrightarrow{a})$ when $|\Theta| > 2$ (multinomial opinion). The positive and negative evidence is represented as p and n respectively. The belief (b), disbelief (d), uncertainty (u), and a priori values (a) for binomial opinions are computed as:

$$b = \frac{p}{p+n+2} \quad d = \frac{n}{p+n+2} \quad u = \frac{2}{p+n+2} \quad a = \frac{1}{2} \tag{1}$$

A subjective opinion is equivalent to a Beta probability distribution (binomial opinion) or to a Dirichlet distribution (multinomial opinion). The expected value (E) for the Beta distribution is computed as in equation (2).

Opinions are computed based on contexts. For example source x provides an observation about assertion y in context c (e.g. about an agent's expertise). The trustworthiness of assertion y in context c, represented as $t(x, y : c)$, is the expected value of the Beta distribution corresponding to the opinion and computed as:

$$E = t(x, y : c) = b + a \cdot u \tag{2}$$

Base Rate Discounting Operator in Subjective Logic In subjective logic, the base rate sensitive discounting of opinion of source B on y by opinion of source A on B ω_B^A, $\omega_y^B = (b_y^B, d_y^B, u_y^B, a_y^B)$ by opinion $\omega_B^A = (b_B^A, d_B^A, u_B^A, a_B^A)$ of source A produces transitive belief $\omega_y^{A:B} = (b_y^{A:B}, d_y^{A:B}, u_y^{A:B}, a_y^{A:B})$ where

$$\begin{aligned} b_y^{A:B} &= E(\omega_B^A) b_y^B & d_y^{A:B} &= E(\omega_B^A) d_y^B \\ u_y^{A:B} &= 1 - E(\omega_B^A)(b_y^B + d_y^B) & a_y^{A:B} &= a_y^B \end{aligned} \tag{3}$$

Wu & Palmer Semantic Similarity Measure Many semantic similarity measures have been developed (see the work of Budanitsky and Hirst [2]). We focus on those computed from *WordNet*. *WordNet* groups words into sets of synonyms called synsets that describes semantic relationships between them. It is a directed and acyclic graph with each vertex v, an integer that represents a synset, and each directed edge from v to w represents that w is a hypernym of v. We focus on the Wu & Palmer metric [18], which calculates semantic relatedness in a deterministic way by considering the depths between two synsets in the WordNet taxonomies, along with the depth of the Least Common Subsumer (lcs) as follows:

$$score(s1, s2) = \frac{2 \cdot depth(lcs)}{depth(s1) + depth(s2)} \tag{4}$$

This means that $score \in]0 \ldots 1]$. For deriving the opinions about a concept where no evidence is available, we incorporate $score$, which represents the semantic similarity ($sim(c, c')$) in our trust assessment, where c and c' are concepts belonging to synset $s1$ and $s2$ respectively which represent two contexts.

3.2 Using Semantic Similarity Measures within Subjective Logic

Deriving Opinion about a New or Unknown Context Since we compute opinions based on contexts, it is possible that evidence required to compute the opinion for a particular context is unavailable. For example, suppose that source x owns observations about an assertion in a certain context (e.g. the expertise of an agent about tulips), but needs to evaluate them in a new context (e.g. the agent's expertise about sunflowers), of which it owns no observations. The semantic similarity measure between two contexts, $sim(c, c')$ can be used for obtaining the opinion about an agent y on an unknown or new context through two different methods. In order to derive an opinion about a new or unknown context we can use either the weighing (on the evidence) or the discounting operation (on the opinion) and both the approaches are described below. We will show that the discounting and the weighing are theoretically but not statistically different.

Weighing the Evidence We weigh the positive and negative evidence belonging to a certain context (e.g. *Tulips*) on the corresponding semantic similarity to the new context (e.g. *Sunflowers*), *sim(Tulips, Sunflowers)*. We then perform this for all the contexts for which source x has already provided an opinion, $\forall c' \in C$, by weighing all the positive (p) and negative (n) evidence of c' with the similarity measure $sim(c, c')$ to obtain an opinion about y in c (see the work of Ceolin et al. [3]).

Discounting the Opinion In the second approach, every opinion source x has about other related contexts c', where $c' \in C$ is discounted with the corresponding semantic similarity measure $sim(c, c')$ using the Discounting operator in subjective logic. The discounted opinions are then aggregated to form the final opinion of x about y in the new context c.

Discounting Operator and Semantic Similarity Subjective logic offers a variety of operators for "discounting", i.e. for smoothing opinions given by third parties, provided that we have at disposal an opinion about the source itself. "Smoothing" is meant as reducing the belief provided by the third party, depending on the opinion on the source (the worse the opinion, the higher the reduction). Moreover, since the components of the opinion always sum to one, reducing the belief implies an increase of (one) of the other components: hence there exists a discounting operator favoring uncertainty and one favoring disbelief. Finally, there exists a discounting operator that makes use of the expected value E of the opinion. Following this line of thought, we can use the semantic similarity as a discount factor for opinions imported from contexts related to the one of interest, in case of a lack of opinions in it, to handle possible variations in the validity of the statements due to the change of context.

Choosing the Appropriate Discounting Operator We need to choose the appropriate discounting operator that allows us to use the semantic similarity value as a discounting factor for opinions. The disbelief favoring discounting is an operator that is employed whenever one believes that the source considered might be malicious. This is not our case, since the discounting is used to import opinions own by ourselves but computed in different contexts than the one of interest. Hence we do not make use of the disbelief favoring operator.

In principle, we would have no specific reason to choose one between the uncertainty favoring discounting and the base rate discounting. Basically, having that only rarely the belief (and hence the expected value) is equal to 1, the two discounting operators decrease the belief of the provided opinion, one by multiplying it by the belief in the source, the other one by the expected value of the opinion about the source. In practice, we will see that, thanks to Theorem 1 these two operators are equivalent in this context.

Theorem 1 (Semantic Relatedness Measure is a Dogmatic Opinion). *Let $sim(c, c')$ be the semantic similarity between two contexts c and c' obtained by computing the semantic distance between the contexts in a graph through deterministic measurements (e.g. [18]). Then, $\forall$ $sim(c, c') \in [0,1]$, $\omega^{measure}_{c=c'} = (b^{measure}_{c=c'}, d^{measure}_{c=c'}, u^{measure}_{c=c'}, a^{measure}_{c=c'})$ is equivalent to a dogmatic opinion in subjective logic.*

Proof. A binomial opinion is a dogmatic opinion if the value of *uncertainty is 0.* The semantic similarity measure can be represented as an opinion about the similarity of two contexts c and c'. However, since we restrict our focus on *WordNet*-based measures, the similarity is inferred by graph measurements, and not by probabilistic means. This means that, according to the source, this is a "dogmatic" opinion, since it does not provide any indication of uncertainty: $u^{measure}_{c=c'} = 0$. The opinion is not based on evidence observation, rather on actual deterministic measurements.

$$E(\omega^{measure}_{c=c'}) = b^{measure}_{c=c'} + u^{measure}_{c=c'} \cdot a = sim(c, c') \tag{5}$$

where *measure* indicates the procedure used to obtain the semantic distance, e.g. Wu and Palmer Measure. The values of belief and disbelief are obtained as:

$$b^{measure}_{c=c'} = sim(c, c') \qquad d^{measure}_{c=c'} = 1 - b^{measure}_{c=c'} \quad \square \tag{6}$$

Corollary 1 (Discounting an Opinion with a Dogmatic Opinion). *Let A be a source who has an opinion about y in context c' expressed as $\omega^A_{y:c'} = (b^A_{y:c'}, d^A_{y:c'}, u^A_{y:c'}, a^A_{y:c'})$ and let the semantic similarity between the contexts c and c' be represented as a dogmatic opinion $\omega^{measure}_{c=c'} = (b^{measure}_{c=c'}, d^{measure}_{c=c'}, 0, a^{c'}_{c=c'})$. Since, the source A does not have any prior opinion about the context c, we derive the opinion of A about c represented as $\omega^{A:c'}_c = (b^{A:c'}_c, d^{A:c'}_c, u^{A:c'}_c, a^{A:c'}_c)$ using the base rate discounting operator on the dogmatic opinion.*

$$\begin{array}{cc} a^{A:B}_y = a^B_y & b^{A:B}_y = sim(c, c') \cdot b^B_y \\ u^{A:B}_y = 1 - sim(c, c') \cdot (b^B_y + d^B_y) & d^{A:B}_y = sim(c, c') \cdot d^B_y \end{array} \tag{7}$$

Definition 1 (Weighing Operator). *Let C be the set of contexts c' of which a source A has an opinion derived from the positive and negative evidence in the* *past. Let c be a new context for which A has no opinion yet. We can derive the opinion of A about facts in c, by weighing the relevant evidences in set C with the semantic similarity measure $sim(c, c')$ $\forall c' \in C$. The belief, disbelief, uncertainty and a priori obtained through the weighing operation are expressed below.*

$$\begin{array}{ll} b_c^A = \frac{sim(c,c')\cdot p_{c'}^A}{sim(c,c')(p_{c'}^A+n_{c'}^A)+2} & d_c^A = \frac{sim(c,c')\cdot n_{c'}^A}{sim(c,c')(p_{c'}^A+n_{c'}^A)+2} \\ u_c^A = 1 - \frac{sim(c,c')\cdot(p_{c'}^A+n_{c'}^A)}{sim(c,c')(p_{c'}^A+n_{c'}^A)+2} & a_c^A = a_{c'}^A \end{array} \tag{8}$$

Theorem 2 (Approximation of the Weighing and Discounting Operators). *Let $\omega_{y:c}^{A:c'} = (b_{y:c}^{A:c'}, d_{y:c}^{A:c'}, u_{y:c}^{A:c'}, a_{y:c}^{A:c'})$ be a discounted opinion which source A has about y in a new or unknown context c, derived by discounting A's opinion on known contexts c' $\in C$ represented as $\omega_{c'}^A = (b_{c'}^A, d_{c'}^A, u_{c'}^A, a_{c'}^A)$ with the corresponding dogmatic opinions (e.g. sim(c,c')). Let source A also obtain an opinion about the unknown context c based on the evidence available from the earlier contexts c', by weighing the evidence (positive and negative) with semantic similarity between c and c', sim(c,c') $\forall c' \in C$. Then the difference between the results from the weighing and from the discount operator in subjective logic are statistically insignificant.*

Proof. We substitute the values of belief, disbelief, uncertainty values in equation (9) for Base Rate Discounting with the values from equation (1) and expectation value from equation (5). We obtain the new value of the discounted base rate opinion as follows:

$$\begin{array}{ll} b_c^{A:c'} = \frac{sim(c,c')\cdot p_{c'}^A}{(p_{c'}^A+n_{c'}^A+2)} & d_c^{A:c'} = \frac{sim(c,c')\cdot n_{c'}^A}{(p_{c'}^A+n_{c'}^A+2)} \\ u_c^{A:c'} = 1 - \frac{sim(c,c')\cdot(p_{c'}^A+n_{c'}^A)}{(p_{c'}^A+n_{c'}^A+2)} & a_c^{A:c'} = a_{c'}^A \end{array} \tag{9}$$

Equation (9) and (8) are pretty similar, except for the $sim(c, c').(p_{c'}^A + n_{c'}^A)$ factor in the weighing operator. In the following section we use a 95% t-student and Wilcoxon signed-rank statistical test to prove that the difference due to that factor is not statistically significant for large values of $sim(c, c')$ (at least 0.5).

3.3 Evaluations

We prove statistically the similarity between the weighing and the discounting.[1]

First Validation: Discounting and Weighing in a Real-Life Case

Steve Social Tagging Project Dataset For the purpose of our evaluations, we use the "Steve Social Tagging Project" [16] data (in particular, the "Researching social tagging and folksonomy in the ArtMuseum"), which is a collaboration of museum professionals and others aimed at enhancing social

[1] Complete results are available at `http://tinyurl.com/bp43k5d`

tagging. In our experiments, we used a sample of tags which the users of the system provided for the 1784 images of the museum available online. Most of the tags were evaluated by the museum professionals to assess their trustworthiness. We used only the evaluated tags for our experiments. The tags can be single words or a string of words provided by the user regarding any objective aspect of the image displayed to them for the tagging.

Gathering Evidence for Evaluation We select a set of tags highly semantically related, by using a Web-based *WordNet* interface [14]. We then gather the list of users who provided the tags regarding the chosen words and count the number of positive and the negative evidence.
The opinions are calculated using two different methods. First by weighing the evidence with the semantic distance using equation (8) and the second method is by discounting the evidence with the semantic distance using equation (9). We consider the *Chinese-Asian* pair (semantic similarity 0.933) and the *Chinese-Buddhist* pair (semantic similarity 0.6667).

Results We employ the Student's t-test and the Wilcoxon signed-rank test to assess the statistical significance of the difference between two sample means. At 95% confidence level, both tests show a statistically significant difference between the two means. This difference, for the *Chinese-Asian* pair is 0.025, while for the *Chinese-Buddhist* pair is 0.11, thanks also to the high similarity (higher than 0.5) between the considered topics. Having removed the average difference from the results obtained from discounting (which, on average, are higher than those from weighing), both the tests assure that the results of the two methods distribute equally.

Second Validation: Discounting and Weighing on a Large Simulated Dataset In order to validate our hypothesis that weighing with semantic distance produces results that are highly similar to those obtained with the discounting operator of subjective logic, we perform the Student's t-test and the Wilcoxon signed-rank test on a larger dataset consisting of 1000 samples. For semantic distance values $sim(c, c') > 0.7$, the mean difference between the belief values obtained by weighing and discounting is 0.092. Thus with 95% confidence interval, both tests assure that both the weighing operator and the discounting operator produce similar results. The semantic similarity threshold $sim(c, c') > 0.7$ is relevant and reasonable, because it becomes more meaningful to compute opinions for a new context based on the opinions provided earlier for the most semantically related contexts, while also in case of lack of evidence for a given context, evidence about a very diverse context can not be much significant.

4 Partial Evidence Observation

The Web and the Semantic Web are pervaded of data that can be used as evidence for a given purpose, but that constitute partially positive/negative evidence for others. Think about the *Waisda?* tagging game [13]. Here, users challenge each other about video tagging. The more users insert the same tag about the same video within the same time frame, the more the tag is believed to be correct. Matching tags can be seen as positive observations for a specific tag to be correct. However, consider the orthogonal issue of the user reputation. User reputation is based on past behavior, hence on the trustworthiness of the tags previously inserted by him/her. Now, the trustworthiness of each tag is not deterministically computed, since it is roughly estimated from the number of matching tags for each tag inserted by the user. The expected value of each tag, which is less than one, can be considered as a partial observation of the trustworthiness of the tag itself. Vice-versa, the remainder can be seen as a negative partial observation. After having considered tag trustworthiness, one can use each evaluation as partial evidence with respect to the user reliability: no tag (or other kind of observation) is used as a fully positive or fully negative evidence, unless its correctness has been proven by an authority or by another source of validation. However, since only rarely the belief (and therefore, the expected value) is equal to one, these observations almost never count as a fully positive or fully negative evidence. We propose an operator for building opinions based on indirect observations, i.e., on observations used to build these opinions, each of which counts as an evidence.

Theorem 3 (Partial Evidence-Based Opinions). *Let p be a vector of positive observations (e.g. a list of "like" counts) about distinct facts related to a given subject s. Let l be the length of p. Let each opinion based on each entry of p have an a priori value of $\frac{1}{2}$. Then we can derive an opinion about the reliability of the subject in one of these two manners.*

- *By cumulating the expected values (counted as partial positive evidence) of each opinion based on each element of p:*

$$b = \frac{1}{l+2}\Sigma_{i=1}^{l}\frac{p_i+1}{p_i+2} \quad d = \frac{1}{l+2}\Sigma_{i=1}^{l}\frac{1}{p_i+2} \quad u = \frac{2}{l+2} \tag{10}$$

- *By averaging the expected values of the opinions computed on each of the elements of p:*

$$b = \frac{1}{l(l+2)}\Sigma_{i=1}^{l}\frac{p_i+1}{p_i+2} \quad d = \frac{1}{l(l+2)}\Sigma_{i=1}^{l}\frac{1}{p_i+2} \quad u = \frac{2}{l(l+2)} \tag{11}$$

Proof. The expected value of each opinion is computed as:

$$E = b + a \cdot u = \frac{p}{p+2} + \frac{1}{2} \cdot \frac{2}{p+2} = \frac{p+1}{p+2} \tag{12}$$

E is considered as partial positive evidence. Hence $1 - E$ is considered as partial negative evidence. Given that we have l pieces of partial evidence (because we have l distinct elements in $\overrightarrow{p}$), we compute the opinion about s following equations (1). Having that p (positive evidence of ω_s) is equal to $\frac{p'+1}{p'+2}$, we obtain equation (10). If we choose to average the evidence (and hence, the expected values) instead of cumulate them, what we obtain is $p = \frac{1}{l}\Sigma\frac{p_i+1}{p_i+2}$, hence $b = \frac{\frac{1}{l}\Sigma\frac{p_i+1}{p_i+2}}{l+2}$ and therefore we obtain equation (11). □

5 Dirichlet Process-Based Opinions: Open World Opinions

5.1 Preliminaries: Dirichlet Process

The Dirichlet Process [6] is a stochastic process representing a probability distribution whose domain is a random probability distribution. As we previously saw, the binomial and multinomial opinions are equivalent to Beta and Dirichlet probability distributions. The Dirichlet distribution represents an extension of the Beta distribution from a two-category situation to a situation where one among n possible categories has to be chosen. A Dirichlet process over a set S is a stochastic process whose sample path (i.e. an infinite-dimensional set of random variables drawn from the process) is a probability distribution on S. The finite dimensional distributions are from the Dirichlet distribution: if H is a finite measure on S, α is a positive real number and X is a sample path drawn from a Dirichlet process, written as

$$X \sim DP(\alpha, H) \tag{13}$$

then for any partition of S of cardinality m, say $\{B_i\}_{i=1}^m$

$$(X(B_1), \ldots, X(B_m)) \sim Dirichlet(\alpha H(B_1), \ldots, \alpha H(B_m)). \tag{14}$$

Moreover, given n draws from X, we can predict the next observation as:

$$obs_{n+1} = \begin{cases} x_i^*(i \in [1 \ldots k]) & \text{with probability } \frac{n(x_i^*)}{n+\alpha} \\ H & \text{with probability } \frac{\alpha}{n+\alpha} \end{cases} \tag{15}$$

where x_i^* is one of the k unique value among the observations gathered.

5.2 Open World Opinions

Having to deal with real data coming from the Web, which are accessed incrementally, the possibility to update the relative probabilities of possible outcomes might not be enough to deal with them. We may need to handle unknown categories of data which should be accounted and manageable anyway. Ceolin et al. [5] show how it is important to account for unseen categories, when dealing with Web data. Here, we propose a particular subjective opinion called "open world opinion" which accounts for partial knowledge about the possible outcomes. A subjective opinion resemble personal opinion provided by sources with respect to facts. Open world opinions represent the case when something about a given fact has been observed, but the evidence allow also for some other (not yet observed) outcome to be considered as plausible. With this extension we allow the frame of discernment to have infinite cardinality. In practice, open world opinions allow to represent situations when the unknown outcome of an event can be equal to one among a list of already observed values (proportionally to the amount of observations for each of them), but it is also possible that (and so some probability mass is reserved to) the outcome is different from what has been observed so far, and is drawn from an infinitely large domain.

Definition 2 (Open World Opinion). *Let: X be a frame of infinite cardinality, $\alpha \in \mathbb{R}^+$, k be the number of categories observed, $\overrightarrow{p}$ be the array of evidence per category, $\overrightarrow{B}$ be a belief function over X. We define the open world opinion ω_x:*

$$\omega_x(\overrightarrow{B}, U, H) \quad B_{x_i} = \frac{p_{x_i}}{\alpha + \Sigma_{x=1}^{k} p_{x_i}} \quad U = \frac{\alpha}{\alpha + \Sigma_{x=1}^{k} p_{x_i}} \quad 1 = U + \Sigma_{x_i} B_{x_i} \tag{16}$$

Definition 3 (Expected Value of Open World Opinion).
The expected value of an open world opinion is computed as follows:

$$E(p(x_i)|r, H) = \frac{p_{x_i} + H(x_i)}{\alpha + \Sigma p_{x_t}} = \frac{p_{x_i}}{\alpha + \Sigma p_{x_t}} \tag{17}$$

Theorem 4 (Equivalence between the Subjective and Dirichlet Process Notation). *Let $\omega_X^{bn} = (\overrightarrow{B}, U, H)$ be an opinion expressed in belief notation, and $\omega_X^{pn} = (E, \alpha, H)$ be an opinion expressed in probabilistic notation, both over the same frame X. ω_X^{bn} and ω_X^{pn} are equivalent when the following mappings holds:*

$$\begin{cases} B_{x_i} = \frac{p_{x_i}}{\alpha + \Sigma_{x=1}^{k} p_{x_i}} \\ \\ U = \frac{\alpha}{\alpha + \Sigma_{x=1}^{k} p_{x_i}} \end{cases} \Leftrightarrow \begin{cases} p_{x_i} = \frac{\alpha B_{x_i}}{U} \\ \\ 1 = U + \Sigma B_{x_i} \end{cases} \tag{18}$$

Proof. Each step of the Dirichlet Process can be seen as a Dirichlet Distribution. Hence the mapping between Dirichlet Distributions and multinomial opinions [9] holds also here. □

Theorem 5 (Mapping between Open World Opinion and Multinomial Opinion). *Let $\omega 1_y^x(\overrightarrow{B}, U, H)$ be an open world opinion and let $\omega 2_y^x(\overrightarrow{B}, U, \overrightarrow{a})$ be a multinomial opinion. Let X_2 and Θ_2 be the frame and the frame of discernment of $\omega 2_y^x$. Let $\{B_i\}_{i=1}^k$ be the result of the partition of $dom(H)$ such that:*

1. $|\Theta_2| = |\{B_i\}|$
2. $\bigcup \{B_i\}_{i=1}^k = dom(H)$
3. $\forall \{x_i\}[(\{x_i\} \in X_2 \wedge |\{x_i\}| = 1 \wedge x_i \in B_j) \Rightarrow \nexists x_{k \neq j} \in B_i]$
4. *$W = k$, where W is the non-informative constant of multinomial opinions*

Then there exists a function $D : Dom(H) \rightarrow \{B_i\}$ such that $D(\omega 1_y^x) = \omega 2_y^x$.

Proof. The equivalence between the discretized open world opinion and the multinomial opinion is proven by showing that:

- given equation (14), since the partition $\{B_i\}_{i=1}^k$ covers the entire $dom(H)$, then the partition distributes like the corresponding Dirichlet distribution;
- to each category of $\omega 2_y^x$ corresponds one and only one partition of $\{B_i\}$ as per item 2 of Theorem 5. □

In other words, open world opinions extend multinomial opinions by allowing the frame of discernment Θ to be infinite. However, by properly discretizing an open world opinion, what we obtain is an equivalent multinomial opinion.

5.3 Example: Using Open World Opinions

Piracy at sea is a well know problem. Every year, several ships are attacked, hijacked, etc. by pirates. The International Chamber of Commerce has created a repository of reports about ship attacks.[2] Van Hage et al. [17] have created an enriched Semantic Web version of such a repository, the Linked Open Piracy (LOP).[3] On the basis of LOP, one might think to be able to predict the frequency of attacks from one year based on the previously available data. However, a problem arises in this situation, since new attack types appear every year and this makes that frequencies vary. Ceolin et al. [5] have shown how the Dirichlet process can be employed to model such situations. Having the possibility to represent this information by means of an open world opinion adds the power of subjective logic to the Dirichlet process based representation. We can merge contributions from different sources, taking into account their reliability. Moreover, we can combine these facts with others in a logical way and then estimate the opinion (and the corresponding probability to be true) of the consequent facts. By using open world opinions, we can easily apply usual subjective operators to these data and easily represent them in a way that takes into account basic provenance information (e.g. data source) when applying fusing or discounting operators. For instance, if according to LOP, in Asia in 2010 we had 10 hijacking events and 10 attempted boarding, then we would represent this as:

$$\omega^{LOP}_{Attacks\ in\ Asia\ in\ 2010}([0.48, 0.48], 0.04, U(0,1))$$

If our opinion about LOP is that is a reliable but not fully accountable source (e.g. $\omega^{us}_{LOP}(0.8, 0.1, 0.1)$), then we can take this information into account by weighing the opinion given by LOP as follows:

$$\omega^{us}_{LOP}(0.8, 0.1, 0.1) \otimes \omega^{LOP}_{Attacks\ in\ Asia\ in\ 2010}([0.48, 0.48], 0.04, U(0,1)) =$$

$$= \omega^{us:LOP}_{Attacks\ in\ Asia\ in\ 2010}([0.384, 0.384], 0.232, U(0,1))$$

The resulting weighted opinion is more uncertain than the initial one, because, even though the two observed types are more likely to happen, the small uncertainty about the source reliability makes the other probabilities to rise.

A difference with respect to multinomial opinions arises in case of fusion, because the fusion operator requires that the *a priori values* have to be merged (averaged). Since the a priori values in the case of the open world opinions are represented by the distribution H (supposedly, H_1 and H_2 for two opinions to be merged). The averaging is still performed, and in this case the averaged distribution corresponds to the distribution Z having $E(Z) = b \cdot E(X_1) + a \cdot E(X_2)$ and $VAR(X) = b^2 \cdot (VAR(X_1)) + a^2 \cdot (VAR(X_2))$, where a, b are the two weights (e.g. u_1 and u_2 in case of cumulative fusion).

[2] http://www.icc-ccs.org

[3] http://semanticweb.cs.vu.nl/lop

6 Discussion

We have shown the potential for employing subjective Logic as a basis for reasoning on Web and Semantic Web data. We have shown how it can be really powerful for handling uncertainty and how little extensions can help in improving the mutual benefit that Semantic Web and subjective logic obtain from cooperating together. Part of this work is based on previously mentioned practical applications that show the usefulness of it, and here we provide theoretical foundations for it. We foresee that other extensions will be possible as well like, for instance, the usage of hyperopinions [10] to handle subsumption reasoning about uncertain data.

References

1. A. Bellenger, S. Gatepaille, H. Abdulrab, and J.-P. Kotowicz. An Evidential Approach for Modeling and Reasoning on Uncertainty in Semantic Applications. In *URSW*, volume 778, pages 27–38. CEUR-WS.org, 2011.
2. A. Budanitsky and G. Hirst. Evaluating WordNet-based Measures of Lexical Semantic Relatedness. *Computational Linguistics*, 32(1):13–47, Mar. 2006.
3. D. Ceolin, A. Nottamkandath, and W. Fokkink. Automated Evaluation of Annotators for Museum Collections using Subjective Logic. In *IFIPTM*, pages 232–239. Springer, May 2012.
4. D. Ceolin, W. Van Hage, and W. Fokkink. A Trust Model to Estimate the Quality of Annotations Using the Web. In *WebSci10*. Online, 2010.
5. D. Ceolin, W. van Hage, and W. Fokkink. Estimating the Uncertainty of Categorical Web Data. In *URSW*, volume 778, pages 15–26. CEUR-WS.org, 2011.
6. T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, 2:209–230, 1973.
7. A. Jøsang. A Logic for Uncertain Probabilities. *Int. Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–212, 2001.
8. A. Jøsang, M. Daniel, and P. Vannoorenberghe. Strategies for combining conflicting dogmatic beliefs. In *FUSION*. IEEE, 2003.
9. A. Jøsang, J. Diaz, and M. Rifqi. Cumulative and averaging fusion of beliefs. *Information Fusion*, 11(2):192–200, 2010.
10. A. Jøsang and R. Hankin. Interpretation and fusion of hyper opinions in subjective logic. In *FUSION*. IEEE, 2012.
11. A. Jøsang, S. Marsh, and S. Pope. Exploring Different Types of Trust Propagation. In *iTrust*, pages 179–192. Springer, 2006.
12. L. Kaplan, S. Chakraborty, and C. Bisdikian. Subjective Logic with Uncertain Partial Observations. In *FUSION*. IEEE, 2012.
13. Netherlands Inst. for Sound and Vision. Waisda? `http://wasida.nl`, Aug. 2012.
14. Princeton University. Wordnet::Similarity. `http://marimba.d.umn.edu/cgi-bin/similarity/similarity.cgi`, Feb. 2012.
15. M. Sensoy, J. Pan, A. Fokoue, M. Srivatsa, and F. Meneguzzi. Using Subjective Logic to Handle Uncertainty and Conflicts. In *TrustCom*. IEEE, 2012.
16. U.S Inst. of Museum and Library Service. Steve social tagging project, Jan. 2012.
17. W. van Hage, V. Malaisé, and M. van Erp. Linked Open Piracy: A story about e-Science, Linked Data, and statistics. *Journal of Data Semantics*, 2012.
18. Z. Wu and M. Palmer. Verbs semantics and lexical selection. In *ACL '94*, pages 133–138. ACL, 1994.