Evaluating trust from past assessments with imprecise probabilities: comparing two approaches

Sebastien Destercke

Abstract. In this paper, we consider a trust system where the trust in an agent is evaluated from past assessments made by other agents. We consider that trust is evaluated by values given on a finite scale. To model the agent trustworthiness, we propose to build imprecise probabilistic models from these assessments. More precisely, we propose to derive probability intervals (i.e., bounds on singletons) using two different approaches: Goodman's multinomial confidence regions and the imprecise Dirichlet model (IDM). We then use these models for two purposes: (1) evaluating the chances that a future assessments will take particular values, and (2) computing an interval summarizing the agent trustworthiness, eventually fuzzyfying this interval by letting the confidence value vary over the unit interval. We also give some elements of comparison between the two approaches.
Keywords: trustworthiness, probability intervals, expectations bounds

1 Introduction

The notion of trust and how to evaluate it has taken more and more importance in computer science with the emergence of the semantic web (particularly in the field of e-commerce or security) and multi-agent systems. Once done, trust evaluation can be used to compare agents or to make an absolute judgement whether an agent can be trusted. To perform such an evaluation, many trust systems have been developed in the past years (see Sabater and Sierra [1] for a review).

Note that the notion of trust as well as the information used to evaluate it can take many forms [2]. One can differentiate between individual-level and system-level trusts, the former concerning the trust one has in a particular agent, while the latter concerns the overall system and the way it ensures that no one will be able to use the system in a selfish way (i.e., to its own profit). The collected information about the trustworthiness of an agent may be direct (coming from past transactions one has done with this agent) or indirect (provided by third-party agents), and when it is indirect, it may be a direct evaluation of the agent reputation or information concerning some of its characteristics.

In this paper, we consider that the information whether an agent (called here the *trustee*) can be trusted or not is given in the form of past evaluations provided by other agents on a numerical scale $\mathcal{X} = \{-n, \dots, -1, 0, 1, \dots, n\}$ ranging from $-n$ to n. In this bipolar scale, a rate of $-n$ means that the trustee is totally untrusted, while n means that it is totally trusted, 0 standing for neutral. For sake of clarity, we will also refer to the elements of $\mathcal{X}$ as $\mathcal{X} = \{x_1, \ldots, x_{2n+1}\}$. Using the classification proposed in Ramchurn *et al.* [2], we are working here with indirect information concerning the individual-level trust of the system and the reputation of an agent.

In [3], Ben Naim and Prade discuss the interest of summarising past evaluations by intervals, as they are more informative than mere mean values and precise points (as interval imprecision is a valuable information reflecting our quantity of knowledge), and are far easier to read than the whole set of evaluations. Eventually, the summarising interval can be reduced to a single evaluation, should such precision be needed.

Consider a counting vector $\Theta = \{\theta_1, \ldots, \theta_{2n+1}\}$ where θ_i is the number of times an agent has given x_i as its evaluation of the *trustee* truthfulness, and $\hat{\theta} = \sum_i \theta_i$ the total number of evaluations. The problem we consider here is how to summarise the information provided by this counting vector in an interval representation describing the past behaviour of the trustee. To do so, we propose to use an imprecise probabilistic model well fit to represent uncertainty on multinomial data (here, the ratings), namely probability intervals [4], and to use the notion of lower and upper expectations to compute the summarising interval. As we shall see, this simple model allows for efficient computations of a summarising interval.

Section 2 recalls some basics of probability intervals and presents the two uncertainty models built from the counting vector Θ. Section 3 then details how a summarising interval can be built from these models. It also provides some elements of comparison by exploring the properties of these intervals with respect to Θ and the possibilities of building fuzzy interval as summary rather than a single one.

2 The model

Let us first recall some elements about probability intervals, before studying how they can be derived from Θ by using confidence regions.

2.1 Probability intervals

Probability intervals as uncertainty models have been studied extensively by De Campos et al. [4]. Probability intervals on a space $\mathcal{X} = \{x_1, \ldots, x_{2n+1}\}$ are defined as a set $L = \{[l_i, u_i] | i = 1, \ldots, 2n+1\}$ of intervals such that $l_i \leq p(x_i) \leq u_i$, where $p(x_i)$ is the unknown probability of element x_i. In this paper, built probability intervals satisfy a number of reasonable conditions usually required to work with this uncertainty representation, namely

$$\sum_{i=1}^{2n+1} l_i \leq 1 \leq \sum_{i=1}^{2n+1} u_i, \tag{1}$$

and for $i = 1, \dots, 2n+1$

$$u_i + \sum_{\substack{j \in \{1,\dots,2n+1\} \\ j \neq i}} l_j \leq 1 \quad ; \quad l_i + \sum_{\substack{j \in \{1,\dots,2n+1\} \\ j \neq i}} u_j \geq 1. \tag{2}$$

If probability intervals satisfy these conditions, then they induce a set of probability measures $\mathcal{P}_L$ such that

$$\mathcal{P}_L = \{p \in \mathbb{P}_{\mathcal{X}} | i = 1, \dots, 2n+1, \; l_i \leq p(x_i) \leq u_i\},$$

with $\mathbb{P}_{\mathcal{X}}$ the set of all probability measures over $\mathcal{X}$. From $\mathcal{P}_L$can be computed lower and upper probabilities on any event A, respectively as $\underline{P}(A) = \inf_{p \in \mathcal{P}_L} P(A)$ and $\overline{P}(A) = \sup_{p \in \mathcal{P}_L} P(A)$. In the case of probability intervals, their computations are facilitated, since we have [4]

$$\underline{P}(A) = \max(\sum_{x_i \in A} l_i, 1 - \sum_{x_i \notin A} u_i) \quad ; \quad \overline{P}(A) = \min(\sum_{x_i \in A} u_i, 1 - \sum_{x_i \notin A} l_i)$$

The question is now how probability intervals can be derived from the counting vector Θ of past evaluations. Had we an infinite number of evaluations at our disposal, it would be reasonable to adopt as a model of the trustee trustworthiness the probability distribution p_∞ corresponding to limiting frequencies. Therefore, we should ask probability intervals to tend towards such frequencies, i.e.,

$$l_i \xrightarrow[\hat{\theta} \to \infty]{} p_\infty(x_i) \quad ; \quad u_i \xrightarrow[\hat{\theta} \to \infty]{} p_\infty(x_i).$$

In practice there may only be a few evaluations available, and in any case a finite quantity of them. Therefore, the chosen uncertainty representation should both tend towards the limiting frequencies and reflect our potential lack of information.

We propose two approaches to build such representations. The first use Goodman's multinomial confidence intervals [5], while the second use the popular Imprecise Dirichlet Model (IDM for short) [6]. The two approaches as a basis of trustee evaluation are then compared in Section 3.

2.2 Building intervals from Θ: first approach

In this first approach, we propose to use Goodman's multinomial confidence intervals [5] as our representation. Given a space $\mathcal{X}$ and a counting vector Θ, Goodmans intervals $[l_i^{G,\alpha}, u_i^{G,\alpha}]$ with confidence level α read, for $i = 1, \dots, 2n+1$,

$$l_i^{G,\alpha} = \frac{b + 2\theta_i - \sqrt{\Delta_i^\alpha}}{2(\hat{\theta} + b)}, \qquad u_i^{G,\alpha} = \frac{b + 2\theta_i + \sqrt{\Delta_i^\alpha}}{2(\hat{\theta} + b)}, \tag{3}$$

where b is the quantile of order $1 - {}^{(1-\alpha)}/_{(2n+1)}$ of the chi-square distribution with one degree of freedom and where

$$\Delta_i^\alpha = b \left(b + \frac{4\theta_i(\hat{\theta} - \theta_i)}{\hat{\theta}} \right).$$

Note that b is an increasing function of α and n, meaning that confidence interval imprecision increases as α increases and as n (the number of possibilities) increases. These probability intervals satisfy Conditions (1) and (2). They tend towards limiting frequencies and the distance between l_i and u_i decreases as more information is collected (i.e. $u_i - l_i$ is a decreasing function of θ_i and $\hat{\theta}$). Also, they are very simple to compute, since only Θ is needed to estimate them. We will denote by $L^{G,\alpha}$ the obtained probability intervals and by $\mathcal{P}_L^{G,\alpha}$ the induced probability set.

Example 1. Consider a space $\mathcal{X} = \{-2, -1, 0, 1, 2\}$ containing 5 possible values. The following counting vector $\Theta = (0, 9, 13, 11, 17)$ summarises the various evaluations given by different agents. The probability intervals obtained with a confidence level $\alpha = 0.95$ are summarised in Table 1

	x_1	x_2	x_3	x_4	x_5
$u_i^{G,0.95}$	0.117	0.354	0.441	0.398	0.522
$l_i^{G,0.95}$	0	0.081	0.135	0.107	0.196

Table 1. Example 1 probability intervals

2.3 Building intervals from Θ: second approach

The second approach we propose consists in using the IDM [6] to build the confidence intervals. The IDM basically extends the classical multinomial Dirichlet model by considering all Dirichlet distributions as the initial set of prior distributions. Intervals extracted from the IDM depend on a hyperparameter $s \geq 0$ that determines the influence of prior information on the posterior information. In the IDM, the value s can be seen as a way to settle the speed of convergence of probability intervals to limiting frequencies p_∞, this speed decreasing when s value increases. An often suggested interpretation for s is that it represents the number of "unseen" observations, and on most applications, $s \in \{1, 2\}$.

Given a space $\mathcal{X}$, a counting vector Θ and a positive value s, intervals $[l_i^{I,s}, u_i^{I,s}]$ resulting from the use of the IDM read, for $i = 1, \dots, 2n+1$,

$$l_i^{I,s} = \frac{\theta_i}{\hat{\theta} + s}, \qquad u_i^{I,s} = \frac{\theta_i + s}{\hat{\theta} + s}. \tag{4}$$

As for Goodman's interval, their computation only requires to know Θ, and the distance $u_i - l_i$ decreases as more information is collected. We will denote by $L^{I,s}$ the obtained probability intervals and by $\mathcal{P}_L^{I,s}$ the induced probability set.

Example 2. Consider the space and counting vector of Example 1. The probability intervals obtained with the IDM and a value $s = 2$ are summarised in Table 2.

	x_1	x_2	x_3	x_4	x_5
$u_i^{I,2}$	0.038	0.212	0288	0.25	0.365
$l_i^{I,2}$	0	0.173	0.25	0.212	0.327

Table 2. Example 2 probability intervals

As we can see, these intervals are much narrower than the ones obtained in Example 1. This indicates that small values of s may be unwarranted in the current application (as an agent would be most of the time unwilling to make precise inference from a small number of evaluations). Note that it may be difficult to obtain a general result relating the interval imprecision obtained by the two approaches, since the difference $u_i^{I,s} - l_i^{I,s}$ does not depend on the θ_i, while the difference $u_i^{G,\alpha} - l_i^{G,\alpha}$ does.

Note that in both approaches, one can interpret the built intervals L, and the associated probability set $\mathcal{P}_L$ as a predictive model providing information about the next possible evaluations. In particular, lower and upper probabilities of an event A gives an interval $[\underline{P}(A), \overline{P}(A)]$ characterising our uncertainty about wether the next evaluation will fall in the set A.

3 Summarising Interval

In the first part of this section, we consider that we work with a fixed confidence level α (in the first approach) or with a fixed hyper-parameter s (in the second approach), for sake of clarity. These assumptions will be relaxed in the last subsection.

3.1 Lower and upper expectations

Let us first recall some elements about the notions of lower and upper expectations. Given a probability set $\mathcal{P}$ defined over a domain $\mathcal{X}$ and a real-valued bounded function $f : \mathcal{X} \to \mathbb{R}$, one can compute the lower and upper expectations of f, $\underline{E}_{\mathcal{P}}(f)$ and $\overline{E}_{\mathcal{P}}(f)$ as

$$\underline{E}_{\mathcal{P}}(f) = \inf_{p \in \mathcal{P}} E_p(f), \qquad \overline{E}_{\mathcal{P}}(f) = \sup_{p \in \mathcal{P}} E_p(f),$$

with $E_p(f)$ the expected value of f with respect to probability distribution p. Lower and upper expectations are dual, in the sense that $\underline{E}(f) = -\overline{E}(-f)$, and have the property that if a constant value μ is added to f, $\underline{E}(f + \mu) = \underline{E}(f) + \mu$ and $\overline{E}(f + \mu) = \overline{E}(f) + \mu$.

When the lower (resp. upper) probabilities of a credal set $\mathcal{P}$ satisfies the property of 2-monotonicity (resp. 2-alternance), that is when, for any two events $A, B \subseteq \mathcal{X}$, we have $\underline{P}(A) + \underline{P}(B) \leq \underline{P}(A \cup B) + \underline{P}(A \cap B)$ (resp. $\overline{P}(A) + \overline{P}(B) \geq \overline{P}(A \cup B) + \overline{P}(A \cap B)$), one can use the Choquet integral [7] to evaluate the lower and upper expectations. Consider a positive bounded function[1] f. If we denote by $()$ a reordering of elements of $\mathcal{X}$ such that $f(x_{(1)}) \leq \ldots \leq f(x_{(2n+1)})$, Choquet integrals giving lower and upper expectations are given by

$$\underline{E}(f) = \sum_{i=1}^{N} (f(x_{(i)}) - f(x_{(i-1)}) \underline{P}(A_{(i)}), \tag{5}$$

$$\overline{E}(f) = \sum_{i=1}^{N} (f(x_{(i)}) - f(x_{(i-1)}) \overline{P}(A_{(i)}), \tag{6}$$

with $f(x_{(0)}) = 0$ and $A_{(i)} = \{x_{(i)}, \ldots, x_{(N)}\}$. In Walley's [8] behavioural interpretation of lower and upper expectations, $\underline{E}(f)$ represents the maximum buying price an agent would pay for a gamble whose gains are represented by f, and $\overline{E}(f)$ the minimum selling price an agent would be ready to accept for the gamble f.

3.2 Expectation bounds as a summarising interval

Let us now come back to our trust evaluation problem, and consider the first approach. Information about the trustee are given by probability intervals $L^{G,\alpha}$ resulting from the counting vector Θ and inducing a probability set $\mathcal{P}_L^{G,\alpha}$. It is known [4] that probability intervals induce 2-monotone and 2-alternating lower and upper probabilities.

Given this information $L^{G,\alpha}$, we propose to summarise the trustworthiness of the trustee as the interval given by lower and upper expectations of a function f such that $f(x_1) = -n, f(x_2) = -n+1, \ldots, f(x_{n+1}) = 0, \ldots, f(x_{2n+1}) = n$ with respect to the probability set $\mathcal{P}_L^{G,\alpha}$. $\underline{E}^{G,\alpha}(f)$ can then be interpreted as the maximal price an agent would be ready to pay to be in interaction with the trustee, while $\overline{E}^{G,\alpha}(f)$ can be interpreted as the minimal price an agent would be ready to accept for being forbidden to interact with the trustee. Algorithm 1 provides an easy way to compute lower and upper expectations.

Algorithm 1 uses the facts that function f values are always rank-ordered in the same way and that the difference of two consecutive values of f is 1. Therefore, Equations (5) and (6) reduce to sums of lower and upper probabilities in this particular case. The adaptation of Algorithm 1 to the second approach is straightforward, since it consists of replacing $l_i^{G,\alpha}$, $u_i^{G,\alpha}$ with $l_i^{I,s}, u_i^{I,s}$. In this latter case, the resulting interval will be denoted by $I^{I,s} := [\underline{E}^{I,s}(f), \overline{E}^{I,s}(f)]$.

Example 3. The summarising intervals corresponding to interval probabilities of Examples 1 and 2 are

$$I^{G,0.95} = [\underline{E}^{G,0.95}, \overline{E}^{G,0.95}] = [-0.09, 1.225]$$

$$I^{I,2} = [\underline{E}^{I,2}, \overline{E}^{I,2}] = [0.615, 0.769]$$

[1] Note that any bounded function f can be made positive by adding a suitable constant to it.

Algorithm 1: Algorithm giving summarising interval

Input: Θ,α

Output: $I^{G,\alpha} = [\underline{E}^{G,\alpha}(f), \overline{E}^{G,\alpha}(f)]$

$\underline{E}^{G,\alpha}(f) = 0, \overline{E}^{G,\alpha}(f) = 0$;

Evaluate $\hat{\theta} = \sum_{i=1}^{2n+1} \theta_i$;

for $i = 1, \ldots, 2n+1$ **do**

  Evaluate $l_i^{G,\alpha}$ (Eq.(3));

  Evaluate $u_i^{G,\alpha}$ (Eq.(3)) ;

for $i = 1, \ldots, 2n+1$ **do**

  if $i == 1$ **then**

    $\underline{E}^{G,\alpha}(f) = \underline{E}^{G,\alpha}(f) + 1$;

    $\overline{E}^{G,\alpha}(f) = \overline{E}^{G,\alpha}(f) + 1$;

  else

    $\underline{E}^{G,\alpha}(f) = \underline{E}^{G,\alpha}(f) + \max(\sum_{k=i}^{2n+1} l_k^{G,\alpha}, 1 - \sum_{k=1}^{i-1} u_k^{G,\alpha})$;

    $\overline{E}^{G,\alpha}(f) = \overline{E}^{G,\alpha}(f) + \min(\sum_{k=i}^{2n+1} u_k^{G,\alpha}, 1 - \sum_{k=1}^{i-1} l_k^{G,\alpha})$;

$\underline{E}^{G,\alpha}(f) = \underline{E}^{G,\alpha}(f) - (n+1)$;

$\overline{E}^{G,\alpha}(f) = \overline{E}^{G,\alpha}(f) - (n+1)$;

3.3 Some Properties

Let us now study some of the properties of each summarising intervals. The first property, satisfied by the two approaches, show that two similar evaluation profiles (in the sense that empirical frequencies are equal) with different amount of information (quantity of evaluations) give coherent summarising intervals, in the sense that the interval obtained with a greater amount of evaluations is included in the one obtained with less evaluations.

Proposition 1. *Let Θ and Θ' be two counting vectors with $\Theta' = \beta\Theta$, $\beta > 1$. Then, given a confidence value α or a hyper-parameter s, we have*

$$I^{G',\alpha} \subset I^{G,\alpha} \quad \textit{and} \quad I^{I',s} \subset I^{I,s},$$

with $I^{G',\alpha}$, $I^{I',s}$ the summarising intervals obtained from Θ', and $I^{I,s}$, $I^{G,\alpha}$ the summarising intervals obtained from Θ.

Proof. We will only prove the inclusion for the first approach, the proof for the second being similar. $\Theta' = \beta\Theta$ implies that for $i = 1, \ldots, 2n+1$, $\theta'_i = \beta\theta_i$. By Eq. (3), we have that $l_i^{G,\alpha} < l_i^{G',\alpha}$ and $u_i^{G',\alpha} < u_i^{G,\alpha}$, hence $[l_i^{G',\alpha}, u_i^{G',\alpha}] \subset [l_i^{G,\alpha}, u_i^{G,\alpha}]$. This means that $\mathcal{P}_L^{G',\alpha} \subset \mathcal{P}_L^{G,\alpha}$, and that infinimum and supremum of expectations over these two sets are such that $I^{G',\alpha} \subset I^{G,\alpha}$.

Let us now demonstrate a proposition that only holds for the IDM approach, and that basically says that better evaluations should provide a better global score (both higher lower and upper expectations) for the trustee.

Proposition 2. *Let Θ and Θ' be two counting vectors, with $\hat{\theta} = \hat{\theta}'$ and for which there is an index i such that $\forall j \geq i$, $\theta'_j \geq \theta_j$ and $\forall j < i$, $\theta'_j \leq \theta_j$. Then, given a hyper-paramater s, we have*

$$\underline{E}^{I,s} \leq \underline{E}^{I',s} \quad and \quad \overline{E}^{I,s} \leq \overline{E}^{I',s}, \tag{7}$$

with $\underline{E}^{I,s}$ and $\underline{E}^{I',s}$ the lower expectation resp. obtained from Θ and Θ', and likewise for the upper expectations.

Proof. Let us consider the initial counting vector Θ. As $\hat{\theta} = \hat{\theta}'$, going from Θ to Θ' can be done by transferring some evaluations, e.g. of index $k < i$ to better ones e.g. of index $i \geq m$ one at a time. Therefore, all we have to do is to consider the counting vector Θ'' such that $\theta''_k = \theta_k - 1$, $\theta''_m = \theta_m + 1$ and $\theta''_i = \theta_i$ for all other indices, and to prove that Eq. (7) holds in this case.

By Eq (4), we have that $l_i^{I'',s} = l_i^{I,s}$ and $u_i^{I'',s} = u_i^{I,s}$ for any i different of k, m. We also have that $l_k^{I'',s} \leq l_k^{I,s}$, $u_k^{I'',s} \leq u_k^{I,s}$, $l_m^{I'',s} \geq l_m^{I,s}$ and $u_m^{I'',s} \geq u_m^{I,s}$. Now, concentrating on the lower expectation and using Eq (5), to prove Eq. (7), we need to prove $\sum_{i=1}^N \underline{P}(A_{(i)}) \leq \sum_{i=1}^N \underline{P}''(A_{(i)})$, with $\underline{P}$ and $\underline{P}''$ the lower probabilities induced by the probability intervals obtained from Θ and Θ''. The two sums read:

$$\sum_{i=1}^{N} \underline{P}(A_{(i)}) = \sum_{i=1}^{N} \max\{\sum_{j=i}^{2n+1} l_j, 1 - \sum_{j=1}^{i-1} u_j\}$$

$$\sum_{i=1}^{N} \underline{P}''(A_{(i)}) = \sum_{i=1}^{N} \max\{\sum_{j=i}^{2n+1} l''_j, 1 - \sum_{j=1}^{i-1} u''_j\}$$

For $i \leq k$ or $i > m$, we have $\max\{\sum_{j=i}^{2n+1} l_j, 1 - \sum_{j=1}^{i-1} u_j\} = \max\{\sum_{j=i}^{2n+1} l''_j, 1 - \sum_{j=1}^{i-1} u''_j\}$, because $l_k^{I,s} - l_k^{I'',s} = l_m^{I'',s} - l_m^{I,s}$ $(i \leq k)$ and $u_m^{I'',s} - u_m^{I,s} = u_k^{I,s} - u_k^{I'',s}$ $(i < m)$. Now, consider the case where $k < i \leq m$, we do have $l_m^{I'',s} \geq l_m^{I,s}$ and $u_k^{I,s} \geq u_k^{I'',s}$, therefore $\max\{\sum_{j=i}^{2n+1} l_j, 1 - \sum_{j=1}^{i-1} u_j\} \leq \max\{\sum_{j=i}^{2n+1} l''_j, 1 - \sum_{j=1}^{i-1} u''_j\}$. Hence, we have $\underline{E}^{I,s} \leq \underline{E}^{I',s}$. The proof concerning the upper expectation is similar.

As shows the next example, the approach using Goodman's confidence intervals does not satisfy this property, that may seem intuitive at first sight. This is mainly due to the fact that differences between upper and lower probability bounds u_i, l_i derived from Goodman's confidence intervals depend on the number of evaluations θ_i, i.e. more evaluations θ_i will provide a narrower interval $[l_i, u_i]$. This means that the model precision depends on how evaluations are distributed, while it can be argued that it is not the case for the IDM (where differences $u_i - l_i$ depend solely on parameter s and $\hat{\theta}$).

Example 4. Consider a space $\mathcal{X} = \{-2, -1, 0, 1, 2\}$ containing 5 possible values and the two following counting vectors $\Theta = (0, 0, 10, 0, 0)$ and $\Theta' = (0, 0, 8, 2, 0)$. With a confidence degree $\alpha = 0.95$, we have

$$I^{G,0.95} = [-0.8, 0.8] \qquad I^{G',0.95} = [-0.92, 1]$$

From the example, it can be seen that Goodman's intervals somewhat reflect the dispersion of evaluations, i.e. the model imprecision depends on how concentrated evaluations are. Indeed, more dispersed evaluations may improve the upper score, while providing a more imprecise interval $[\underline{E}, \overline{E}]$ (as in Example 4).

It would be interesting to relate this kind of behaviour (interval imprecision increase) with some dispersion measures of the empirical frequencies distributions (e.g., entropy, Gini index, ...). Also, it could be checked whether Goodman's intervals approach satisfy a weaker condition than Proposition 2, namely that for two counting vectors Θ and Θ' satisfying condition of Proposition 2 and a given confidence value α, we have $\overline{E}^{G,\alpha} \leq \overline{E}^{G',\alpha}$.

These two properties may be seen as monotonic properties w.r.t. evaluation quantity and evaluation score, respectively. Other properties, such as adaptation of the ones proposed by Ben-Naim and Prade [3], should be investigated in further studies.

3.4 Towards fuzzy evaluations

In this subsection, we relax some of the previous assumptions (i.e. fixed confidence level α and parameter s) and propose some methods to obtain a fuzzy interval as the evaluation summary rather than a crisp interval.

Recall that a fuzzy set μ is a mapping $\mu : \mathcal{X} \to [0, 1]$ from $\mathcal{X}$ (here, the interval $[-n, n]$) to the unit interval, where $\mu(x)$ is called the membership value of x. The β-cut of a fuzzy set μ is the set $A_\beta := \{x \in \mathcal{X} | \mu(x) \geq \beta\}$.

First approach: Goodman's intervals Extending the first approach to obtain a fuzzy representation is straightforward, since the formalism of fuzzy sets is particularly well suited to the representation of confidence intervals [9]. Indeed, a β-cut can be interpreted as a confidence set or interval with a confidence level $1 - \beta$.

An interval $I^{G,\alpha} = [\underline{E}^{G,\alpha}, \overline{E}^{G,\alpha}]$ for a given α can therefore be directly associated to the $(1-\alpha)$-cut of a fuzzy set giving a global evaluation of the trustee trustworthiness. The resulting fuzzy set μ^G is such that, for any $\alpha \in (0, 1]$

$$\mu^G(\underline{E}^{G,\alpha}) = 1 - \alpha \qquad \mu^G(\overline{E}^{G,\alpha}) = 1 - \alpha.$$

Example 5. Consider the counting vector $\Theta = (0, 9, 13, 11, 17)$ provided in Example 1. Figure 1 illustrates the obtained summarising fuzzy interval. The representation shows that the trustee has a positive score, centred around 0.7. Only intervals given by conservative confidence values (above 0.9) provide summarising intervals that include negative values.

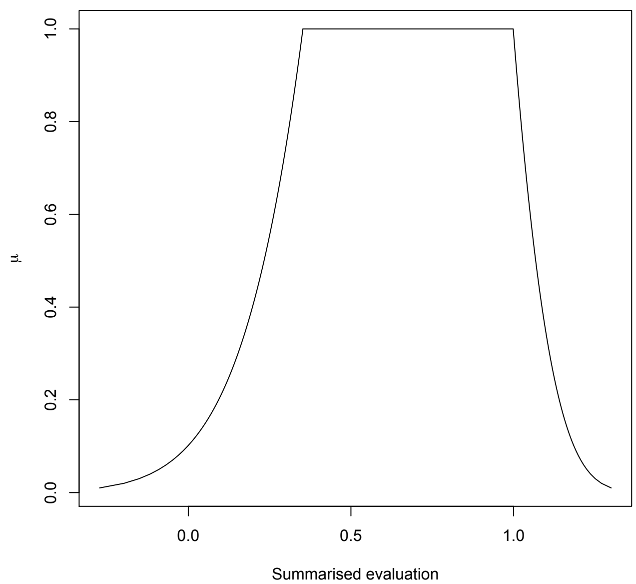

Fig. 1. Fuzzy evaluation with Goodman's intervals (Example 5).

Second approach: IDM intervals How to build a fuzzy evaluation by using the IDM approach is less straightforward. An idea (the one we take here) is to let the parameter s vary within some bounds $[0, \overline{s}]$, and to build a fuzzy set μ^I such that, for any $s \in [0, \overline{s}]$,

$$\mu^I(\underline{E}^{I,s}) = \frac{\overline{s} - s}{\overline{s}} \qquad \mu^I(\overline{E}^{I,s}) = \frac{\overline{s} - s}{\overline{s}}.$$

This is indeed a fuzzy set, since for two $s, s' \in [0, \overline{s}]$ such that $s \leq s'$, we do have $[\underline{E}^{I,s}, \overline{E}^{I,s}] \subset [\underline{E}^{I,s'}, \overline{E}^{I,s'}]$. However, the interpretation in terms of confidence intervals is in this case less clear, and the final fuzzy global evaluation is highly dependent of the value $\overline{s}$ (the lower $\overline{s}$, the more precise will be the fuzzy set). Hence this extension is more ad hoc, as well as questionable.

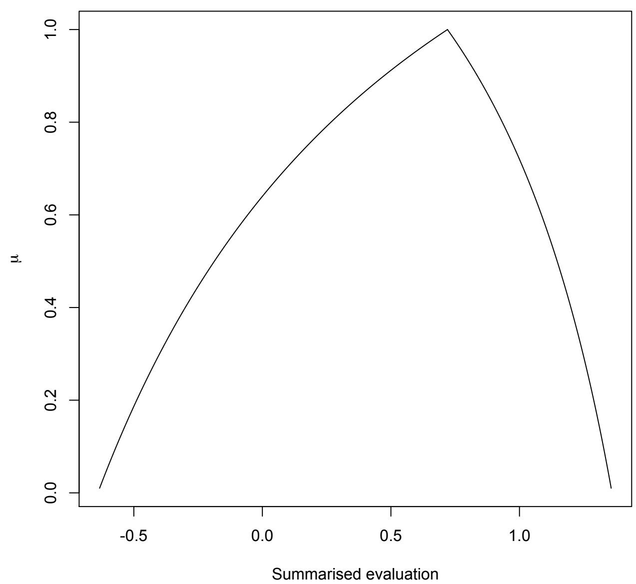

Fig. 2. Fuzzy evaluation with IDM intervals (Example 6).

Example 6. Consider the counting vector $\Theta = (0, 9, 13, 11, 17)$ provided in Example 1 and a value $\overline{s} = 50$. Figure 2 illustrates the obtained summarising fuzzy interval. Although the fuzzy set is centred around the same values as in Figure 1, its shape is quite different. Indeed, in this case the imprecision growth decreases as α-value decreases, while in the case of Goodman's intervals the imprecision growth increases as α-value increases.

Note that, since Propositions 1 and 2 are valid for any confidence level or hyper-parameter values, their conclusions can directly be extended to the proposed fuzzy extensions.

4 Conclusion

In this paper, we have proposed and compared two imprecise probabilistic models to evaluate the trustworthiness of an agent (the trustee) from previous evaluations made by other agents. The two models are based on the estimation of lower and upper expectations induced by probability intervals, themselves induced by the counting vector of evaluations. In the first model, these probability intervals are given by Goodman's statistical confidence intervals, while in the second, probability intervals are provided by the (popular) Imprecise Dirichlet model. Lower and upper expectations summarise the counting vector of evaluations in a richer way than single point values, since the interval they provide also reflects the dispersion of evaluations and their quantity. Both methods are computationally efficient, and we have proposed for both of them extensions to fuzzy evaluations.

From our study, it appears that the two approaches are at odds. Indeed, Goodman's intervals approach does not satisfy some monotonic properties (Property 2) that intuitively one may wish to satisfy, while the IDM approach does. However, it could be argued that the IDM probability intervals and the induced summarising interval only takes account of the quantity of evaluations, as the imprecision in both of them only depends on the number of evaluations (once s is fixed). On the contrary, Goodman's probability intervals imprecision and the induced summarising interval are also influenced by the evaluations distribution across space $\mathcal{X}$. Goodman's confidence intervals also have a clear statistical interpretation, allowing for a very natural extension of the summarising process to fuzzy intervals. Such an extension, although possible, is more tricky to interpret in the case of the IDM, for which the choice of s and its meaning are still discussed among researchers.

In conclusion, our preference would be to let go of Property 2 and to use Goodman's intervals, since they account for evaluations dispersion in $\mathcal{X}$ and have a clear statistical interpretation. However, if one considers that Property 2 have to be satisfied, then the IDM approach should be used.

A possible improvement of the current approach would be to integrate additional features to the evaluations or the way they are taken into account. For instance, it could be desirable to allow for imprecise evaluations or to consider the time at which evaluations were given (recent evaluations being more reliable than old ones). However, such additional information would also mean that the counting vector Θ would no longer be sufficient "statistic" to provide a summary.

Another interesting topic to explore is how trust information coming from past evaluations can be combined with other trust information sources (e.g., direct interactions).

References

1. Sabater, J., Sierra, S.: Review on computational trust and reputation models. Artificial Intelligence Review **24**(33-60) (2005)
2. Ramchurn, S., Huynh, D., Jennings, N.: Trust in multi-agent systems. The Knowledge Engineering Review **19** (2004) 1–25
3. Ben-Naim, J., Prade, H.: Evaluating trustworthiness from past performances: Interval-based approaches. In: SUM 2008. (2008) 33–46
4. de Campos, L., Huete, J., Moral, S.: Probability intervals: a tool for uncertain reasoning. I. J. of Uncertainty, Fuzziness and Knowledge-Based Systems **2** (1994) 167–196
5. Goodman, L.: On simultaneous confidence intervals for multinomial proportions. Technometrics **7** (1964) 247–254
6. Bernard, J.M.: An introduction to the imprecise dirichlet model for multinomial data. I. J. of Approximate Reasoning **39** (2004) 123–150
7. Choquet, G.: Theory of capacities. Annales de l'institut Fourier **5** (1954) 131–295
8. Walley, P.: Statistical reasoning with imprecise Probabilities. Chapman and Hall, New York (1991)
9. Dubois, D., Foulloy, L., Mauris, G., Prade, H.: Probability-possibility transformations, triangular fuzzy sets, and probabilistic inequalities. Reliable Computing **10** (2004) 273–297