

Continuous Ratings in Discrete Bayesian Reputation Systems

Audun Jøsang, Xixi Luo and Xiaowu Chen

Abstract Reputation systems take as input ratings from members in a community, and can produce measures of reputation, trustworthiness or reliability of entities in the same community. Binomial and multinomial Bayesian reputation systems are discrete in nature meaning that they normally take discrete ratings such as “average” or “good” as input. However, in many situations it is natural to provide input ratings to reputation systems based on continuous measures. This paper describes the principles of discrete Bayesian reputation systems, and how continuous measures can provide input ratings to such systems. The method is based on fuzzy set membership functions.

1 Introduction

Online reputation systems have emerged as important decision support tools that can help reduce the risk of engaging in transactions and interactions on the Internet. Reputation systems stimulate higher quality online services, and are also being investigated as a general method of social control in the online environment.

The same basic principles for creation and propagation of reputation in the physical world used by online reputation systems. The main difference is that online reputation systems are supported by extremely efficient network and computer systems. While reputation formation in the physical world is mostly limited to local communities, online reputation systems have no geographical limits.

Audun Jøsang
University of Oslo, UNIK Graduate Center, Norway e-mail: josang@unik.no

Xixi Luo
Beihang University e-mail: xixiluo@cse.buaa.edu.cn

Xiaowu Chen
Beihang University e-mail: chen@vrlab.buaa.edu.cn

Please use the following format when citing this chapter:

Jøsang, A., Luo, X. and Chen, X., 2008, in IFIP International Federation for Information Processing, Volume 263; *Trust Management II*; Yücel Karabulut, John Mitchell, Peter Herrmann, Christian Damsgaard Jensen; (Boston: Springer), pp. 151–166.

Reputation systems collect information about the performance of a given entity as ratings from other community participants who have had direct experience with that entity. In the typical case of centralised reputation systems, the reputation centre collects all the ratings and derives a reputation score for every party. The reputation scores are published online so that they represent the public reputation of every party in the community. Participants can then use each other's scores, for example, when deciding whether or not to transact with a particular party. The idea is that transactions with reputable parties are likely to result in more favourable outcomes than transactions with disreputable parties.

Fig.1 shows a typical centralised reputation system architecture, where *A* and *B* denote parties with a history of transactions in the past, and who consider transacting with each other in the present.

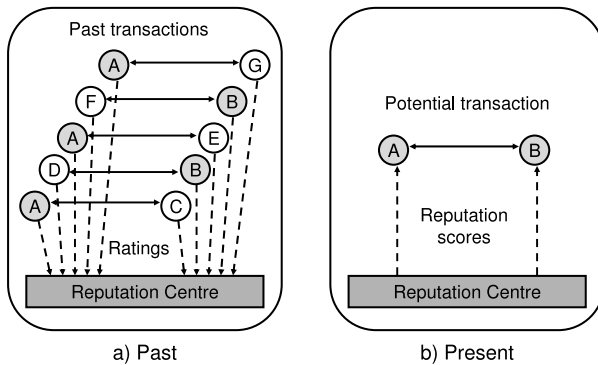


Fig. 1 General reputation system architecture

Fig.1.a shows that the parties provide ratings about each other's performance after each transaction. The reputation centre collects ratings from all the agents, and continuously updates each agent's reputation score as a function of the received ratings.

Fig.1.b shows that updated reputation scores are provided online for all the parties to see. These are used by party *A* and *B* to decide whether or not to transact with each other.

Two fundamental elements of reputation systems are:

1. *Communication protocols* that allow participants to provide ratings about transaction partners to the reputation centre, as well as to obtain reputation scores of potential transaction partners from the reputation centre.
2. A *reputation computation engine* used by the reputation centre to derive reputation scores for each participant, based on received ratings, and possibly also on other information.

This paper focuses on the reputation computation engine. Bayesian reputation systems represent a type of mathematically sound and well studied computation en-

gines. We have previously proposed and studied binomial and multinomial Bayesian reputation systems [3, 4, 5, 8]. Binomial reputation systems allow ratings to be expressed with two values, as either positive (e.g. *good*) or negative (e.g. *bad*). Multinomial reputation systems allow the possibility of providing ratings with graded levels such as e.g. *mediocre - bad - average - good - excellent*. In addition, multinomial models are able to distinguish between the case of polarised ratings (i.e. a combination of strictly good and bad ratings) and the case of only average ratings. The ability to indicate when ratings are polarised can provide valuable clues to the user in many situations. Multinomial reputation systems therefore provide great flexibility when collecting ratings and providing reputation scores.

However, it is common that the subject matter to be rated is measured on a continuous scale, such as time, throughput or relative ranking, to name a few examples. Even when it is natural to provide discrete ratings, it may be difficult to express that something is strictly good or average, so that combinations of discrete ratings, such as “*average-to-good*” would better reflect the rater’s opinion. Such ratings can then be considered continuous. To handle this, it is important to have a sound and consistent method for including continuous measures as normal ratings in reputation systems. This paper investigates principles for including ratings based on continuous measures in reputation systems, and combining them with traditional discrete measures. We show that this can be done through membership functions in the same way as fuzzy set membership is computed in traditional fuzzy set theory.

The rest of the paper is structured as follows. Sec.2 briefly reviews the Bayesian multinomial model, and Sec.3 describes how to design reputation systems based on this model. Sec.4 describes how continuous measures can be taken as input ratings in Bayesian reputation systems, and Sec.5 describes an example of using this method. Sec.6 concludes.

2 The Multinomial Bayesian Model

This section briefly reviews the principles of the multinomial Bayesian model which forms the basis for Bayesian reputation systems. For details, see [5, 1].

2.1 The Dirichlet Distribution

Multinomial Bayesian reputation systems allow ratings to be provided over k different levels which can be considered as a set of k disjoint elements. Let this set be denoted as $\Lambda = \{L_1, \dots, L_k\}$, and assume that ratings are provided as votes on the elements of Λ . This leads to a Dirichlet probability density function over the k -component random probability variable $\mathbf{p}(L_i)$, $i = 1 \dots k$ with sample space $[0, 1]^k$, subject to the simple additivity requirement $\sum_{i=1}^k \mathbf{p}(L_i) = 1$.

The Dirichlet distribution with prior captures a sequence of observations of the k possible outcomes with k positive real rating parameters $\mathbf{r}(L_i)$, $i = 1 \dots k$, each corresponding to one of the possible levels. In order to have a compact notation we define a vector $\mathbf{p} = \{\mathbf{p}(L_i) \mid 1 \leq i \leq k\}$ to denote the k -component probability variable, and a vector $\mathbf{r} = \{r_i \mid 1 \leq i \leq k\}$ to denote the k -component rating variable.

In order to distinguish between the *a priori* default base rate, and the *a posteriori* ratings, the Dirichlet distribution must be expressed with prior information represented as a base rate vector \mathbf{a} over the state space. This will be called the Dirichlet Distribution with Prior.

Definition 1 (Dirichlet Distribution with Prior).

Let $\Lambda = \{L_1, \dots, L_k\}$ be a state space consisting of k mutually disjoint elements. Let \mathbf{r} represent the rating vector over the elements of Λ and let \mathbf{a} represent the base rate vector over the same elements. Then the multinomial probability density function over Λ is expressed as:

$$f(\mathbf{p} \mid \mathbf{r}, \mathbf{a}) = \frac{\Gamma(\sum_{i=1}^k (\mathbf{r}(L_i) + \mathbf{Ca}(L_i)))}{\prod_{i=1}^k \Gamma(\mathbf{r}(L_i) + \mathbf{Ca}(L_i))} \prod_{i=1}^k \mathbf{p}(L_i)^{(\mathbf{r}(L_i) + \mathbf{Ca}(L_i) - 1)},$$

$$\text{where } \begin{cases} \sum_{i=1}^k \mathbf{p}(L_i) = 1 \\ \mathbf{p}(L_i) \geq 0, \forall i \end{cases} \quad \text{and} \quad \begin{cases} \sum_{i=1}^k \mathbf{a}(L_i) = 1 \\ \mathbf{a}(L_i) > 0, \forall i. \end{cases} \tag{1}$$

The vector \mathbf{p} represents probability variables, so that for a given \mathbf{p} the probability density $f(\mathbf{p} \mid \mathbf{r}, \mathbf{a})$ represents their second order probability. The first-order variables of \mathbf{p} represent probabilities of rating levels, whereas the density $f(\mathbf{p} \mid \mathbf{r}, \mathbf{a})$ represents the probability of specific values for the first-order variables. Since the first-order variables \mathbf{p} are continuous, the second-order probability $f(\mathbf{p} \mid \mathbf{r}, \mathbf{a})$ for any given value of $\mathbf{p}(L_i) \in [0, 1]$ is vanishingly small and therefore meaningless as such. It is only meaningful to compute $\int_{p_1}^{p_2} f(\mathbf{p}(L_i) \mid \mathbf{r}, \mathbf{a})$ for a given interval $[p_1, p_2]$ and level L_i , or simply to compute the expectation value of $\mathbf{p}(L_i)$. The most natural is to define the reputation score as a function of the expectation value. This provides a sound mathematical basis for combining ratings and for expressing reputation scores. The probability expectation of any of the k random probability variables can be written as:

$$E(\mathbf{p}(L_i) \mid \mathbf{r}, \mathbf{a}) = \frac{\mathbf{r}(L_i) + \mathbf{Ca}(L_i)}{C + \sum_{i=1}^k \mathbf{r}(L_i)}. \tag{2}$$

The *a priori* constant C will normally be set to $C = 2$ when a uniform distribution over binary state spaces is assumed. Selecting a larger value for C will result in new observations having less influence over the Dirichlet distribution, and can in fact represent specific *a priori* information provided by a domain expert or by another reputation system. It can be noted that it would be unnatural to require a uniform distribution over arbitrary large state spaces because it would make the sensitivity to new evidence arbitrarily small.

For example, requiring a uniform *a priori* distribution over a state space of cardinality 100, would force $C = 100$. In case an event of interest has been observed 100 times, and no other event has been observed, the derived probability expectation value of the event of interest will still only be about $\frac{1}{2}$, which would seem totally counterintuitive. In contrast, when a uniform *a priori* distribution is assumed in the binary case, and the same 100 observations are taken as input, the derived probability expectation of the event of interest would be close to 1, as intuition would dictate.

The value of C determines the approximate number of votes needed for a particular level to influence the probability expectation value of that level from 0 to 0.5

2.2 Visualising Dirichlet Distributions

Visualising Dirichlet distributions is challenging because it is a density function over $k - 1$ dimensions, where k is the state space cardinality. For this reason, Dirichlet distributions over ternary state spaces are the largest that can be easily visualised.

With $k = 3$, the probability distribution has 2 degrees of freedom, and the equation $\mathbf{p}(L_1) + \mathbf{p}(L_2) + \mathbf{p}(L_3) = 1$ defines a triangular plane as illustrated in Fig.2.




Fig. 2 Triangular plane

In order to visualise probability density over the triangular plane, it is convenient to lay the triangular plane horizontally in the x - y plane, and visualise the density dimension along the z -axis.

Let us consider the example of a reputation system with three discrete rating levels: L_1 , L_2 and L_3 (i.e. $k = 3$). Let us first assume that no other information than the cardinality is available, meaning that the default base rate is $\mathbf{a}(L_i) = 1/3$ for all states, and $\mathbf{r}(L_1) = \mathbf{r}(L_2) = \mathbf{r}(L_3) = 0$. Then Eq.(2) dictates that the expected

a priori probability of picking a ball of any specific colour is the default base rate probability, which is $\frac{1}{3}$. The *a priori* Dirichlet density function is illustrated in Fig.3.




Fig. 3 Prior Dirichlet distribution in case of three rating levels

Let us now assume that ratings have been given as $r(L_1) = 6$, $r(L_2) = 1$, and $r(L_3) = 1$. Then the *a posteriori* expected probability of level L_1 can be computed as $E(p(L_1)) = \frac{2}{3}$. The *a posteriori* Dirichlet density function is illustrated in Fig.4.

3 The Dirichlet Reputation System

Multinomial Bayesian systems are based on computing reputation scores by statistical updating of Dirichlet Probability Density Function (PDF). This can be called Dirichlet reputation system [5]. The *a posteriori* (i.e. the updated) reputation score is computed by combining the *a priori* (i.e. previous) reputation score with the new rating. The same principle is also used for binomial Bayesian reputation systems based on the Beta distribution [2, 4, 6, 7].

In Dirichlet reputation systems, an agent is allowed to rate another agent or service, with any level from a set of predefined rating levels, and the reputation scores are not static but will gradually change with time as a function of the received ratings. Initially, each agent’s reputation is defined by the base rate reputation which is distributed evenly among all agents. After evidence about a particular agent is gathered, its reputation will change accordingly. Moreover, the reputation score can be represented on different forms.




Fig. 4 A posteriori Dirichlet distribution after 6 L_1 -ratings 1 L_2 -rating and 1 L_3 -rating

3.1 Collecting Ratings

A general reputation system allows for an agent to rate another agent or service, with any level from a set of predefined rating levels. Some form of control over what and when ratings can be given is normally required, such as e.g. after a transaction has taken place, but this issue will not be discussed here. Let there be k different discrete rating levels. This translates into having a state space of cardinality k for the Dirichlet distribution. Let the rating level be indexed by i . The aggregate ratings for a particular agent y are stored as a cumulative vector, expressed as:

$$\mathbf{R}_y = (\mathbf{R}_y(L_i) \mid i = 1 \dots k) . \tag{3}$$

The simplest way of updating a rating vector as a result of a new rating is by adding the newly received rating vector \mathbf{r} to the previously stored vector \mathbf{R} . The case when old ratings are aged is described in Sec.3.2.

Each new discrete rating of agent y by an agent x takes the form of a trivial vector \mathbf{r}_y^x where only one element has value 1, and all other vector elements have value 0. The index i of the vector element with value 1 refers to the specific rating level.

3.2 Aggregating Ratings with Aging

Ratings may be aggregated by simple addition of the components (vector addition).

Agents (and in particular human agents) may change their behaviour over time, so it is desirable to give relatively greater weight to more recent ratings. This can

be achieved by introducing a longevity factor $\lambda \in [0, 1]$, which controls the rapidity with which old ratings are aged and discounted as a function of time. With $\lambda = 0$, ratings are completely forgotten after a single time period. With $\lambda = 1$, ratings are never forgotten.

Let new ratings be collected in discrete time periods. Let the sum of the ratings of a particular agent y in period t be denoted by the vector $\mathbf{r}_{y,t}$. More specifically, it is the sum of all ratings \mathbf{r}_y^x of agent y by other agents x during that period, expressed by:

$$\mathbf{r}_{y,t} = \sum_{x \in M_{y,t}} \mathbf{r}_y^x \quad (4)$$

where $M_{y,t}$ is the set of all agents who rated agent y during period t .

Let the total accumulated ratings (with aging) of agent y after the time period t be denoted by $\mathbf{R}_{y,t}$. Then the new accumulated rating after time period $t + 1$ can be expressed as:

$$\mathbf{R}_{y,(t+1)} = \lambda \cdot \mathbf{R}_{y,t} + \mathbf{r}_{y,(t+1)}, \text{ where } 0 \leq \lambda \leq 1. \quad (5)$$

Eq.(5) represents a recursive updating algorithm that can be executed once every period for all agents. Assuming that new ratings are received between time t and time $t + n$, then the new rating can be computed as:

$$\mathbf{R}_{y,(t+n)} = \lambda^n \cdot \mathbf{R}_{y,t} + \mathbf{r}_{y,(t+n)}, \quad 0 \leq \lambda \leq 1. \quad (6)$$

3.3 Convergence Values for Reputation Scores

The recursive algorithm of Eq.(5) makes it possible to compute convergence values for the rating vectors, as well as for reputation scores. Assuming that a particular agent receives the same ratings every period, the Eq.(5) defines a geometric series. We use the well known result of geometric series:

$$\sum_{j=0}^{\infty} \lambda^j = \frac{1}{1-\lambda} \text{ for } -1 < \lambda < 1. \quad (7)$$

Let \mathbf{r}_y represent the rating vector of agent y for each period. The Total accumulated rating vector after an infinite number of periods is then expressed as:

$$\mathbf{R}_{y,\infty} = \frac{\mathbf{r}_y}{1-\lambda}, \text{ where } 0 \leq \lambda < 1. \quad (8)$$

Eq.(8) shows that the longevity factor determines the convergence values for the accumulated rating vectors.

3.4 Reputation Representation

A reputation score applies to member agents in a community M . Before any evidence is known about a particular agent y , its reputation is defined by the base rate reputation which is the same for all agents. As evidence about a particular agent is gathered, its reputation will change accordingly.

The reputation score of a multinomial system can be represented on different forms, which can be *evidence representation*, *density representation*, *multinomial probability representation*, or *point estimate representation*. Each form will be described in turn below.

3.4.1 Evidence Representation

The most direct form of representation is to simply express the aggregate evidence vector \mathbf{R}_y . The amount of ratings of level i for agent y is denoted by $\mathbf{R}_y(L_i)$.

It is not necessary to express individual base rate vectors, as it will be the same for all agents.

3.4.2 Density Representation

The reputation score of an agent can be expressed as a multinomial probability density function (PDF) in the form of Eq.(1). For ternary state spaces, the PDF can be visualised as in Fig.4. Visualisation of PDFs for state spaces larger than ternary is not practical.

3.4.3 Multinomial Probability Representation

The most natural is to define the reputation score as a function of the probability expectation values of each element in the state space. The expectation value for each rating level can be computed with Eq.(2).

Let \mathbf{R} represent a target agent's aggregate ratings. Then the vector \mathbf{S} defined by:

$$\mathbf{S}_y : \left(\mathbf{S}_y(L_i) = \frac{\mathbf{R}_y(L_i) + \mathbf{Ca}(L_i)}{C + \sum_{j=1}^k \mathbf{R}_y(L_j)}; | i = 1 \dots k \right) . \quad (9)$$

is the corresponding multinomial probability reputation score. As already stated, $C = 2$ is the value of choice, but larger value for the constant C can be chosen if a reduced influence of new evidence over the base rate is required.

The reputation score \mathbf{S} can be interpreted like a multinomial probability measure as an indication of how a particular agent is expected to behave in future transactions. It can easily be verified that

$$\sum_{i=1}^k S(L_i) = 1 . \tag{10}$$

The multinomial reputation score can for example be visualised as columns, which would clearly indicate if ratings are polarised. Assume for example 5 levels:

$$\text{Discrete rating levels: } \begin{cases} L_1 : \text{Mediocre,} \\ L_2 : \text{Bad,} \\ L_3 : \text{Average,} \\ L_4 : \text{Good,} \\ L_5 : \text{Excellent.} \end{cases} \tag{11}$$

We assume a default base rate distribution. Before any ratings have been received, the multinomial probability reputation score will be equal to 1/5 for all levels. Let us assume that 10 ratings are received. In the first case, 10 *average* ratings are received, which translates into the multinomial probability reputation score of Fig.5.a. In the second case, 5 mediocre and 5 excellent ratings are received, which translates into the multinomial probability reputation score of Fig.5.b.




Fig. 5 Illustrating score difference resulting from average and polarised ratings

With a binomial reputation system, the difference between these two rating scenarios would not have been visible.

In case an agent receives the same ratings every period, the reputation scores will converge to specific values. These values emerge by inserting the convergence values of Eq.(8) into Eq.(9). Let \mathbf{r}_y be the constant ratings that agent y receives every period. The convergence score value for each rating level i can then be expressed as:

$$S_{y,\infty}(L_i) = \frac{\lambda \cdot \mathbf{r}_y(L_i) + (1 - \lambda)Ca(L_i)}{(1 - \lambda)C + \lambda \sum_{j=1}^k \mathbf{r}_y(L_j)} \tag{12}$$

In particular it can be seen that when no ratings are received (i.e. \mathbf{r}_y is the null vector), then the convergence score value for each level is simply the base rate for that level.

3.4.4 Point Estimate Representation




Fig. 6 Sliding windows

While informative, the multinomial probability representation can require considerable space to be displayed on a computer screen. A more compact form can be to express the reputation score as a single value in some predefined interval. This can be done by assigning a point value v to each rating level i , and computing the normalised weighted point estimate score σ .

Assume e.g. k different rating levels with point values evenly distributed in the range $[0,1]$, so that $v(L_i) = \frac{i-1}{k-1}$. The point estimate reputation score is then computed as:

$$\sigma = \sum_{i=1}^k v(L_i)S(L_i) . \tag{13}$$

However, this point estimate removes information, so that for example the difference between the average ratings and the polarised ratings of Fig.5.a and Fig.5.b is no longer visible. The point estimates of the reputation scores of Fig.5.a and Fig.5.b are both 0.5, although the ratings in fact are quite different. A point estimate in the range $[0,1]$ can be mapped to any range, such as 1-5 stars, a percentage or a probability.

3.5 Dynamic Community Base Rates

Bootstrapping a reputation system to a stable and conservative state is important. In the framework described above, the base rate distribution \mathbf{a} will define initial default reputation for all agents. The base rate can for example be evenly distributed, or biased towards either a negative or a positive reputation. This must be defined by those who set up the reputation system in a specific market or community.

Agents will come and go during the lifetime of a market, and it is important to be able to assign new members a reasonable base rate reputation. In the simplest case, this can be the same as the initial default reputation that was given to all agents during bootstrap.

However, it is possible to track the average reputation score of the whole community, and this can be used to set the base rate for new agents, either directly or with a certain additional bias.

Not only new agents, but also existing agents with a standing track record can get the dynamic base rate. After all, a dynamic community base rate reflects the whole community, and should therefore be applied to all the members of that community.

The aggregate reputation vector for the whole community at time t can be computed as:

$$\mathbf{R}_{M,t} = \sum_{y_j \in M} \mathbf{R}_{y,t} \quad (14)$$

This vector then needs to be normalised to a base rate vector as follows:

Definition 2 (Community Base Rate). Let $\mathbf{R}_{M,t}$ be an aggregate reputation vector for a whole community, and let $\mathbf{S}_{M,t}$ be the corresponding multinomial probability reputation vector which can be computed with Eq.(9). The community base rate as a function of existing reputations at time $t + 1$ is then simply expressed as the community score at time t :

$$\mathbf{a}_{M,(t+1)} = \mathbf{S}_{M,t}. \quad (15)$$

The base rate vector of Eq.(15) can be given to every new agent that joins the community. In addition, the community base rate vector can be used for every agent every time their reputation score is computed. In this way, the base rate will dynamically reflect the quality of the market at any one time.

If desirable, the base rate for new agents can be biased in either negative or positive direction in order to make it harder or easier to enter the market.

When base rates are a function of the community reputation, the expressions for convergence values with constant ratings can no longer be defined with Eq.(8), and will instead converge towards the average score from all the ratings.

4 Taking Continuous Ratings

This section describes a method for taking continuous ratings as a basis for input to multinomial and binomial Bayesian reputation systems.

4.1 The Multinomial Case

For a multinomial reputation system with k discrete levels, the parameters of the Dirichlet distribution are \mathbf{r} . Our method is based on a sliding window for determining the discrete rating as a function of the continuous rating.

In general, when there are k rating levels, the parameters $(\mathbf{r}(L_1), \mathbf{r}(L_2), \dots, \mathbf{r}(L_k))$ can be computed as a function of the continuous rating q according to fuzzy triangular membership functions.

Let each rating level L_i be a fuzzy subset, and each rating q is assigned a membership grade $\mathbf{r}(L_i, q)$ taking values in $[0, 1]$, with $\mathbf{r}(L_i, q) = 0$ corresponding to non-membership in L_i , $0 < \mathbf{r}(L_i, q) < 1$ to partial membership in L_i , and $\mathbf{r}(L_i, q) = 1$ to full membership in L_i . According to the above analysis, the fuzzy set triangular membership functions can be expressed in terms of Eq.(16), Eq.(17), and Eq.(18).

$$\begin{array}{l} \text{Membership function for } L_1 : \\ \hline \mathbf{r}(L_1, q) = \begin{cases} 1-q(k-1) & \text{IF } 0 \leq q \leq \frac{1}{(k-1)} \\ 0 & \text{ELSE} \end{cases} \end{array} \quad (16)$$

$$\begin{array}{l} \text{Membership function for } L_i \text{ where } 1 < i < k : \\ \hline \mathbf{r}(L_i, q) = \begin{cases} i-q(k-1) & \text{IF } \frac{(i-1)}{(k-1)} \leq q \leq \frac{i}{(k-1)} \\ 2-i+q(k-1) & \text{IF } \frac{(i-2)}{(k-1)} \leq q \leq \frac{(i-1)}{(k-1)} \\ 0 & \text{ELSE} \end{cases} \end{array} \quad (17)$$

$$\begin{array}{l} \text{Membership function for } L_k : \\ \hline \mathbf{r}(L_k, q) = \begin{cases} 2-k+q(k-1) & \text{IF } \frac{(k-2)}{(k-1)} \leq q \leq 1 \\ 0 & \text{ELSE} \end{cases} \end{array} \quad (18)$$

For example with five rating levels the sliding window function can be illustrated as in Fig.6.

The continuous q -value determines the position of the sliding window. The relative overlap between the window and a specific level determines the r -value for that level.

As an example, Fig.6 indicates the continuous value $q = 3/8$, which causes the sliding window to overlap with rating levels L_2 and L_3 . It can be seen that $q = 3/8$ results in the level rating vector expressed by:

$$\begin{array}{l} \text{Discrete level ratings} \\ \text{resulting from } q = 3/8: \end{array} \begin{cases} \mathbf{r}(L_1) = 0.0 \\ \mathbf{r}(L_2) = 0.5 \\ \mathbf{r}(L_3) = 0.5 \\ \mathbf{r}(L_4) = 0.0 \\ \mathbf{r}(L_5) = 0.0 \end{cases} \quad (19)$$

These r -ratings can then be fed into the reputation system described in Sec.3.

Visualisation of fuzzy membership functions provide an alternative way of intuitively deriving the discrete level ratings. The fuzzy membership functions in the case of 5 discrete rating levels are illustrated in Fig.7.




Fig. 7 Fuzzy triangular membership functions

A discrete rating vector derived from a continuous measure will have either one or two vector elements with positive value, where the sum is always one. This property emerges from the formal expressions of Eqs.(16), (17) and (18). The same property becomes immediately obvious through the visualisation of the fuzzy membership functions in Fig.7.

4.2 The Binomial Case

The binomial case is simply a special case of the multinomial case. Let (r_1, r_2) be the parameters of the Beta distribution. Let q be the continuous rating in the range $[0, 1]$. Then (r_1, r_2) can be determined by the fuzzy set membership function

$$\begin{cases} r_1(q) = 1 - q \\ r_2(q) = q \end{cases} \tag{20}$$

For every continuous rating, we can compute its membership value to each rating level, and then taking this membership value to be the rating of that level. The Eq.(4)···Eq.(13) are the same with continuous ratings, and the parameter r is allowed to be any number between 0 and 1, but not limited to be 0 and 1.

5 Example

In this example, agents can be rated on continuous measures in the range $[0, 1]$, and the reputation system has 5 discrete levels, with base rates evenly distributed. Let an agent be rated over 10 time periods as expressed in Table 1. The longevity factor is set to $\lambda = 0.9$.

Computing the rating levels with Eqs.(16), (17), and (18), we can get the level ratings expressed in the middle row of Table 1.

Then applying Eq.(9), we can get the corresponding multinomial reputation scores in the bottom row of Table 1. The same scores are visualized as a function of the time period in Fig.8.

Time Period	0	1	2	3	4	5	6	7	8	9	10
Continuous ratings:	0.05	0.05	0.05	0.05	0.00	0.10	0.90	0.80	0.80	0.80	0.90
Level ratings		↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
L_1		0.8	0.8	0.8	1.0	0.6					
L_2		0.2	0.2	0.2	0	0.4					
L_3											
L_4							0.4	0.8	0.8	0.8	0.4
L_5							0.6	0.2	0.2	0.2	0.6
Level scores	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
L_1	0.2	0.4	0.4923	0.5452	0.6161	0.5998	0.4982	0.4209	0.3604	0.3121	0.2728
L_2	0.2	0.2	0.2	0.2	0.1632	0.2033	0.1728	0.1496	0.1315	0.1170	0.1052
L_3	0.2	0.1333	0.1026	0.0849	0.0735	0.0656	0.0598	0.0554	0.0520	0.0492	0.0470
L_4	0.2	0.1333	0.1026	0.0849	0.0735	0.0656	0.1197	0.2162	0.2916	0.3519	0.3540
L_5	0.2	0.1333	0.1026	0.0849	0.0735	0.0656	0.1496	0.1580	0.1645	0.1698	0.2210

Table 1 Scalar ratings translated into level ratings that in turn generate level scores

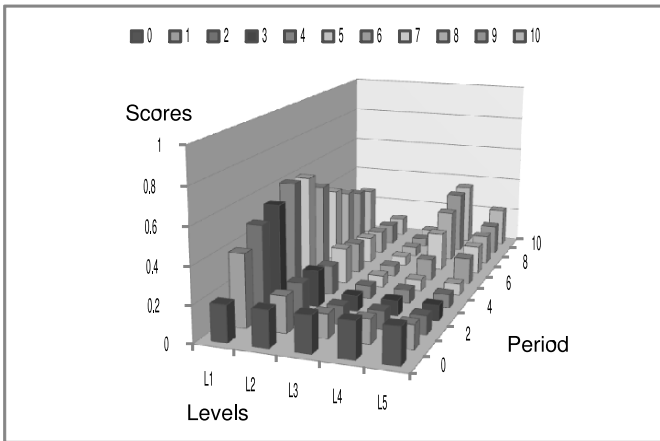


Fig. 8 Evolution of an agent’s reputation scores after the rating sequence of Table 1

It can be seen that the first five periods are characterised by very low continuous ratings, resulting in the score for L_1 increasing rapidly. Then in the five last periods, the continuous rating is relatively high, resulting in increasing scores for L_4 and L_5 and decreasing scores for L_1 , L_2 and L_3 .

6 Conclusion

Bayesian reputation systems normally take discrete ratings as input. This could represent a limitation to the applicability of such reputation systems when the observations to be rated are continuous in nature. This paper focuses on transforming continuous ratings into discrete ratings by using fuzzy set membership function. This work makes the Bayesian reputation systems more practical and generally applicable. The traditional reputation system principles such as aggregating rating with aging, convergence value for reputation scores, methods for reputation representation, and dynamic community base rates are equally applicable both with discrete and continuous ratings.

References

1. A. Gelman et al. *Bayesian Data Analysis, 2nd ed.* Chapman and Hall/CRC, Florida, USA, 2004.
2. A. Jøsang. Trust-Based Decision Making for Electronic Transactions. In L. Yngström and T. Svensson, editors, *Proceedings of the 4th Nordic Workshop on Secure Computer Systems (NORDSEC'99)*. Stockholm University, Sweden, 1999.
3. A. Jøsang, S. Hird, and E. Faccar. Simulating the Effect of Reputation Systems on e-Markets. In P. Nixon and S. Terzis, editors, *Proceedings of the First International Conference on Trust Management (iTrust)*, Crete, May 2003.
4. A. Jøsang and R. Ismail. The Beta Reputation System. In *Proceedings of the 15th Bled Electronic Commerce Conference*, June 2002.
5. A. Jøsang and Haller J. Dirichlet Reputation Systems. In *In the Proceedings of the International Conference on Availability, Reliability and Security (ARES 2007)*, Vienna, Austria, April 2007.
6. L. Mui, M. Mohtashemi, C. Ang, P. Szolovits, and A. Halberstadt. Ratings in Distributed Systems: A Bayesian Approach. In *Proceedings of the Workshop on Information Technologies and Systems (WITS)*, 2001.
7. L. Mui, M. Mohtashemi, and A. Halberstadt. A Computational Model of Trust and Reputation. In *Proceedings of the 35th Hawaii International Conference on System Science (HICSS)*, 2002.
8. A. Withby, A. Jøsang, and J. Indulska. Filtering Out Unfair Ratings in Bayesian Reputation Systems. *The Icfa Journal of Management Research*, 4(2):48–64, 2005.