

Elimination of Subjectivity from Trust Recommendation

Omar Hasan¹, Lionel Brunie¹, Jean-Marc Pierson², and Elisa Bertino³

¹ INSA Lyon, France

{omar.hasan,lionel.brunie}@insa-lyon.fr

² IRIT, France

pierson@irit.fr

³ Purdue University, IN, USA

bertino@cs.purdue.edu

Abstract. In many distributed applications, a party who wishes to make a transaction requires that it has a certain level of trust in the other party. It is frequently the case that the parties are unknown to each other and thus share no pre-existing trust. Trust-based systems enable users to establish trust in unknown users through trust recommendation from known users. For example, Bob may choose to trust an unknown user Carol when he receives a recommendation from his friend Alice that Carol's trustworthiness is 0.8 on the interval [0,1].

In this paper we highlight the problem that when a trust value is recommended by one user to another it may lose its real meaning due to subjectivity. Bob may regard 0.8 as a very high value of trust but it is possible that Alice perceived this same value as only average. We present a solution for the elimination of subjectivity from trust recommendation. We run experiments to compare our subjectivity-eliminated trust recommendation method with the unmodified method. In a random graph based web of trust with high subjectivity, it is observed that the novel method can give better results up to 95% of the time.

1 Introduction

Trust is an indispensable requirement for the successful operation of a number of distributed applications. Trust is defined as “the degree to which one party has confidence in another within the context of a given purpose or decision” [1]. On eCommerce websites, a buyer must trust the seller to deliver the services or goods that are promised. In ad hoc networks, a node trusts neighboring nodes to route its messages. In peer-to-peer file sharing networks, a peer trusts others to deliver authentic content. Internet forums and online communities trust members not to post spam. Without a system in place that enables users to establish the trustworthiness of other parties, a distributed application would suffer from exploitation and eventually fail to provide adequate service.

A variety of trust-based systems [2], [3], [4], [5] have been developed that enable agents (any entity capable of making trust related decisions) to determine

if the party they wish to transact with is trustworthy. Trust recommendation is a key technique that is utilized in trust-based systems for an agent to determine the trustworthiness of an unknown party. A trust recommendation is an attestation of the trustworthiness of an agent Carol by Alice to Bob, where Bob is an agent who is not acquainted with Carol but maintains a trust relationship with Alice.

We present the argument that trust evaluation by each individual is subjective and thus when two individuals exchange a trust value its meaning is distorted due to differences in their perception. For example, Alice may have suggested to Bob that the trustworthiness of Carol is 0.8 on the interval $[0,1]$, which according to her subjective opinion may have been average trustworthiness. However, it is possible that Bob has a different perspective on trust values and regards 0.8 as a very high value. Thus subjectivity prevents the true meaning of Alice's recommendation from being conveyed to Bob.

We subscribe to the definition of subjectivity given by the Merriam-Webster online dictionary (merriam-webster.com) as a judgment that is "modified or affected by personal views, experience, or background" and is "peculiar to a particular individual". Several works [5], [6], [7] propose trust models that aim to capture the subjectivity aspect of human trust. However, the focus is on enabling agents to form trust opinions that are uniquely their own in contrast to delegating trust formation to some external authority. None of the cited works address subjectivity as it affects trust recommendation. In this paper we focus specifically on the problem of subjectivity in trust recommendation.

The remainder of the paper is organized as follows: Section 2 further describes the problem and discusses the notion of disposition to trust. Section 3 presents a basic trust model that serves as a framework for the development of the solution and experiments. In Section 4 we introduce and build the method for elimination of subjectivity from trust recommendation. Experiments in Section 5 that evaluate the effectiveness of the method are followed by a discussion and proposals for future work in Section 6. In Section 7, we present concluding remarks.

2 Background

2.1 Trust Representation and Subjectivity

How does one represent the amount of trust that one individual associates with another? A common approach is to represent the spectrum of trust quantitatively as a numerical range. Marsh's formalism [8] represents trust as a continuous variable over an interval of $[-1,1]$. Golbeck's FilmTrust [9] defines an integer range of 1 to 10. Gambetta [10], Griffiths [6], and Toivonen [11] utilize an interval of $[0,1]$ for the purpose.

An alternate approach is to divide the span of trust into strata and assign them qualitative labels. The stratification used by Abdul-Rahman and Hailes [2] is given as the set {Very Trustworthy, Trustworthy, Untrustworthy, Very Untrustworthy}. Jonker and Treur [12] use a similar stratification defined as the ordering: Unconditional Distrust < Conditional Distrust < Conditional Trust

< Unconditional Trust. Levien's Advogato [13] allows users to rate each other as an Apprentice (minimum trust), Journeyer (medium trust), or as a Master (maximum trust).

Let's consider a scenario where Alice assigns a trust value of 0.8 to Carol on an interval of $[0,1]$ with 1 representing maximum trust. Let's assume that 0.8 is an average trust value if it is viewed in the context of trust values that Alice has assigned to other entities in the past. Thus Alice perceives Carol as someone being moderately trustworthy. With whatever skew Alice assigns trust values to other entities, it presents no problem inside her local environment since all those values lie in the same context.

The problem of subjectivity arises when Alice conveys to Bob that her trust in Carol is represented by the value 0.8. It is likely that a value of 0.8 signifies something very different to Bob. Is 0.8 an average value of trust for Bob as was the case for Alice? Or is 0.8 a very high value of trust for Bob? Given the context of Bob's history of trust value assignments, we may discover that Bob rarely ever assigns a value of 0.8 to any entity and thus associates very high trust with such a value. In Alice's position Bob might have assigned a value such as 0.6 to Carol. Bob may make a misjudgment of Carol's trustworthiness if he bases his decision on his own perception of the trust value conveyed to him by Alice. We observe that due to subjectivity, the meaning of a trust value is distorted when it is propagated from one individual to another. Subjectivity occurs due to differences in the dispositions to trust of individuals. Disposition to trust is defined and discussed in the next section.

The use of strata with qualitative labels may initially be considered as a solution to the problem of subjectivity. We may argue that a stratified trust representation model, such as the four distinct strata defined by Abdul-Rahman and Hailes [2], provides clear semantics and avoids the ambiguity associated with numerical values. The reasoning being that a qualitative label such as "trustworthy" should hold the same meaning for one entity as it does for another.

However, we concur with Griffiths [6] and Marsh [8] that the stratification approach also suffers from the problem of subjectivity. Different entities may associate the same experiences with different strata. For example, based on their own perception of trust, what is viewed by Alice as "very trustworthy" may be judged as merely "trustworthy" by Bob.

We note that subjectivity, as we describe it, is not an issue for the trust representation model used by some popular commercial websites, such as Epinions (epinions.com). This is due to the fact that the resolution they provide for evaluating users is minimal. Epinions allows users to only either "Block" (not trust) or "Trust" other users. This model relies more on the quantity of ratings received per user rather than the degree of trustworthiness specified in an individual rating. On eBay (ebay.com), which uses a somewhat similar model, users value each other's trustworthiness in the same stratum (that is "positive") over 99% of the time [14]. Our work addresses systems that employ broader ranges for the expression of trust.

2.2 Disposition to Trust

Disposition to trust is the inherent propensity of an individual to trust or distrust others. An individual's disposition to trust does not vary for specific entities but is a stable characteristic of their personality that governs how they view the trustworthiness of every other entity that they encounter.

McKnight et al [15] define disposition to trust as the "extent to which a person displays a tendency to be willing to depend on others across a broad spectrum of situations and persons".

Rotter [16], [17] notes that an individual's "generalized attitude" towards trust is a product of life experiences, such as interactions with parents, peers, and authorities. Boone and Holmes [18] suggest that good experiences lead to a greater disposition to trust and vice versa.

A study in the context of ecommerce by McCord and Ratnasingam [19] has demonstrated that there is a strong relationship between an individual's disposition to trust and the trust related decisions that they make.

A thorough treatment of the literature on disposition to trust is provided by Kaluscha [20].

We now revisit Alice, Bob and Carol from our previous example. Alice and Bob are two individuals with different dispositions to trust. Alice has a high disposition to trust and thus assigns a high trust value of 0.8 to Carol. In contrast, Bob who has a lower disposition to trust, rates Carol's trustworthiness as only 0.6. This subjectivity occurs despite the fact that Carol exhibits the same behavior in her interactions with both Alice and Bob.

3 Trust Model

In this section we define a trust model. An important constituent of the model is the provision for trust recommendation and propagation. The objective is not to define a novel trust model but to establish a basic one that will serve as a framework within which we will develop and test our method for elimination of subjectivity from trust recommendation.

We define A as a set of agents. $A = \{a_0, a_1, \dots, a_n\}$. We define a binary relation T on the set A . T is a subset of $A \times A$. $T = \{(u, v) : u, v \in A\}$. The relation T represents the *trusts* relation between two agents. We will use the notation uTv , u trusts v , and (u, v) interchangeably. In our model, the properties of the *trusts* relation are as follows:

Property 1. *The relation T is reflexive. uTu . An agent trusts itself.*

Property 2. *The relation T is not symmetric. $uTv \not\Rightarrow vTu$. If agent u trusts agent v then this does not imply that v also trusts u .*

Property 3. *The relation T is not transitive. $a_0Ta_1 \wedge a_1Ta_2 \not\Rightarrow a_0Ta_2$. If agent a_0 trusts agent a_1 who in turn trusts agent a_2 , then this does not imply that a_0 also trusts a_2 . a_0 may trust a_2 or it may not.*

We define a Web of Trust as a weighted directed graph $G = (A, T)$. The agents in the set A form the vertices of the graph. The trust relations between agents given as ordered pairs in the set T are the edges of the graph. Since G is a directed graph, an edge (u, v) is incident from u and incident to v .

A weight is associated with every edge (u, v) in the graph, which represents the amount of trust that agent u holds for agent v . The weight associated with an edge (u, v) is given as the function $t(u, v)$. $t : T \rightarrow X$. The set X is defined as $X = [0, 1]$.

The range of $t(u, v)$ is real numbers bounded by 0 and 1. 0 implies “minimum trust” and 1 implies “maximum trust”. Real numbers between 0 and 1 give us infinite resolution for expressing trust.

The absence of (u, v) in T implies that no trust relationship exists between agents u and v . We do not address distrust in this model.

A path $p(a_0, a_k)$ of length k from an agent a_0 to an agent a_k is a sequence $\langle a_0, a_1, a_2, \dots, a_k \rangle$ of agents such that $(a_{i-1}, a_i) \in T$ for $i = 1, 2, \dots, k$.

3.1 Trust Recommendation and Propagation

If $(a_0, a_1) \in T \wedge (a_1, a_2) \in T$, then $t(a_1, a_2)$ may be considered as a recommendation from a_1 to a_0 . That is, taking into consideration $t(a_0, a_1)$ and $t(a_1, a_2)$, a_0 may choose to establish (a_0, a_2) and $t(a_0, a_2)$. We say that the trust of a_1 in a_2 is propagated to a_0 .

To facilitate the discussion we establish the following terminology:

Source agent – the agent from whom the path originates; the agent that may establish trust in a previously unknown agent based on the given recommendations

Recommender agent – an agent that recommends another agent

Target agent – the agent at whom the path terminates; the agent whom the source agent may choose to trust

In the preceding case, a_0 is the source agent, a_1 a recommender agent, and a_2 the target agent.

We stress that since trust is not transitive in our model, the propagated trust is only a suggestion to the source agent regarding the trustworthiness of the target agent. The source agent may or may not choose to establish a trust belief based on this suggestion.

We generalize the notion of trust recommendation and propagation for a path of length k :

If $(a_0, a_1), (a_1, a_2), (a_2, a_3), \dots, (a_{k-2}, a_{k-1}), (a_{k-1}, a_k) \in T$, then $t(a_{k-1}, a_k)$ may be considered as a recommendation from a_{k-1} to a_{k-2} , $t(a_{k-2}, a_{k-1})$ as a recommendation from a_{k-2} to a_{k-3}, \dots , and $t(a_1, a_2)$ as a recommendation from a_1 to a_0 . Taking into consideration $t(a_0, a_1), t(a_1, a_2), t(a_2, a_3), \dots, t(a_{k-2}, a_{k-1}), t(a_{k-1}, a_k)$, a_0 may choose to establish (a_0, a_k) and $t(a_0, a_k)$. We say that the trust of a_{k-1} in a_k is propagated to a_0 .

According to the classification introduced by Ziegler and Lausen [21], the trust metric presented in this section may be categorized as local and scalar. The model discussed here shares similarities with those defined by Golbeck et al [22], Chen and Yeager [23], and Abdul-Rahman and Hailes [24].

4 A Method for Elimination of Subjectivity from Trust Recommendation

In this section we introduce our method for the elimination of subjectivity from trust recommendation.

4.1 Quantitative Representation of an Agent's Disposition to Trust

The method requires quantitative representation of the disposition to trust of agents. We discuss three possible alternatives for this purpose.

Manually specified by the agent. The agent may be presented with a scale, for example, 1 to 10 or $[0,1]$ and asked to rate their disposition to trust manually. The approach is simple and straightforward. However, the disadvantage of this approach is that the agent has to be explicitly engaged by the process. Moreover, it is debatable if an agent himself is a true judge of his own disposition to trust.

Assessed through a trust scale. A number of researchers have developed trust scales that help assess the disposition to trust of a person. The subject is required to respond to a series of questions with weighted multiple choice answers. The cumulative score of the subject indicates their disposition to trust.

Rotter's Interpersonal Trust Scale [17] and Christie and Geis's Machiavellianism Scale [25] are examples of this approach. A sample question from Rotter's Interpersonal Trust Scale is as follows:

"In dealing with strangers one is better off to be cautious until they have provided evidence that they are trustworthy."

Answer choices: strongly agree (weight: 1), mildly agree (2), agree and disagree equally (3), mildly disagree (4), strongly disagree (weight: 5).

Rotter's and the Machiavellianism trust scales are likely to assess the disposition to trust of an individual accurately. However, the requirement that each agent make themselves available for a series of questions discounts their practicality.

Inferred from an agent's history of trust value assignments. Several examples from the computer science literature may be cited where historical patterns are used to predict future behavior with considerable success. Instances include Self-Customizing Software [26] or Adaptive User Interfaces [27], and Branch Predictors in Microprocessors [28].

We propose an approach based on similar lines for determining the disposition to trust of an agent. The trust values that an agent has assigned in the past may be considered as an indication of their disposition to trust. For example, given an agent who has a pattern of assigning high values of trust, we may infer that the agent has a high disposition to trust, and vice versa. We thus propose to represent an agent's disposition to trust by the collection of their previous trust value assignments in a system.

A close approximation of an agent's disposition to trust is possible only if they have made a significant number of trust value assignments in the past. The question is what number can be considered as significant. We experiment with multiple values in Section 5.

The primary reason we choose this approach for the representation of disposition to trust is that it does not require additional input from an agent. Given a web of trust, we can test our method without requiring each agent to explicitly establish their disposition to trust.

4.2 The Method

As we have discussed earlier, the trust values assigned by an agent are subjective to their disposition to trust. When a recommender agent recommends a target, the meaning of the associated trust value is distorted due to the different disposition to trust of the source agent.

The solution we propose is to report trust not as an absolute score but a value that is relative to the disposition to trust of the recommender agent. In other words, we report the relative standing of the recommender agent's trust in the target agent in terms of the trust value assignments that the recommender agent has made in the past.

Two simple options for implementing this idea are reporting trust as either a standard score (z -score), or as a percentile. We opt for a solution based on percentiles and not one based on standard scores since the latter requires that the trust values assigned by agents be normally distributed.

A percentile value indicates the recommender agent's perception of the target agent in relation to the others that the recommender agent has rated in the past.

Going back to the example discussed in Section 2 if Alice conveys to Bob an absolute value such as 0.8, Bob does not know if according to Alice the value 0.8 is an average value or a very high value of trust. However, if the trust is reported as a percentile value, Bob does have this information. For example, if the percentile value is in the vicinity of 50%, Bob would know that according to Alice, Carol has an average trustworthiness. If the percentile value is around 80% or 90%, it is clear that Alice regards Carol as highly trustworthy. The absolute value that Alice locally assigned to Carol becomes irrelevant.

To convert the percentile to a local absolute score the source agent reads the value that is at the given percentile in the collection of trust values that he

himself has assigned to other agents. This absolute score holds perfect meaning for the source agent since it is in the context of his own disposition to trust.

Thus going through a relative value as an intermediary, the subjectivity and misinterpretation associated with an absolute trust value are eliminated.

We note that this method does not require agents to make any modifications to the way they evaluate other agents. Locally, each agent establishes their trust beliefs as usual, in terms of their own disposition to trust. Another positive aspect of this solution is that it does not require the involvement of any third parties and is therefore suitable for decentralized networks.

4.3 Formal Description of the Method

Within the framework of the trust model discussed in Section 3, a formal description of the method follows.

d_u is a collection of the weights associated with the outgoing edges of agent u , that is, all $t(u, v)$ where v is a node adjacent to u . As discussed in Section 4.1, the collection of trust values previously assigned or d_u represents the disposition to trust of agent u .

The values in d_u are arranged in ascending order and indexed $1, 2, \dots, n_u$, where n_u is the number of outgoing edges of agent u (as well as the number of values in d_u). The j^{th} value in d_u is referred to by $d_u[j]$. We define a function $first(x, d_u)$ that returns the index of the first occurrence of a value x present in d_u .

$c(u, v)$ is the percentile of $t(u, v)$ in d_u . The function which calculates $c(u, v)$ is given as:

$$\begin{aligned} c(u, v) &= \text{percentile}(t(u, v), d_u) \\ &= \frac{100 \cdot \text{first}(t(u, v), d_u)}{n_u + 1} \end{aligned}$$

As an example, consider $d_{Alice} = \langle 0.4, 0.4, 0.5, 0.6, 0.8, 0.8, 0.8, 0.8, 0.8, 0.9, 0.9 \rangle$ and $t(Alice, Carol) = 0.8$. Then $n_u = 11$ and $first(t(Alice, Carol), d_{Alice}) = 5$. $c(Alice, Carol)$ is calculated as follows:

$$\begin{aligned} c(Alice, Carol) &= \text{percentile}(t(Alice, Carol), d_{Alice}) \\ &= \frac{100 \cdot \text{first}(t(Alice, Carol), d_{Alice})}{n_{Alice} + 1} \\ &= \frac{100 \cdot 5}{11 + 1} = 41.67\text{percentile} \end{aligned}$$

$t(u, v)_w$ is defined as the value in d_w at the $c(u, v)^{\text{th}}$ percentile. The function which calculates $t(u, v)_w$ is stated as:

$$\begin{aligned}
t(u, v)_w &= \text{trustvalue}(c(u, v), d_w) \\
&= \begin{cases} d_w[i] + f \cdot (d_w[i + 1] - d_w[i]) & \text{if } 0 < i < n_w \\ d_w[1] & \text{if } i = 0 \\ d_w[n_w] & \text{if } i = n_w \end{cases}
\end{aligned}$$

where,

$$i = \left\lfloor \frac{c(u, v) \cdot (n_w + 1)}{100} \right\rfloor$$

and,

$$f = \frac{c(u, v) \cdot (n_w + 1)}{100} - i$$

i is an integer and f is a fraction greater than or equal to 0 and less than 1.

We may think of $t(u, v)_w$ as the value $t(u, v)$ transformed such that instead of being in reference to the disposition to trust of agent u , it is now in reference to the disposition to trust of agent w .

Instead of reporting $t(u, v)$, an agent u calculates $c(u, v)$ and communicates this percentile value to agent w . Given $c(u, v)$, agent w determines $t(u, v)_w$ and considers that as the recommended value.

Continuing the example from above, consider $d_{Bob} = \langle 0.2, 0.3, 0.3, 0.3, 0.5, 0.5, 0.5, 0.6, 0.8 \rangle$. Then:

$$\begin{aligned}
t(\text{Alice}, \text{Carol})_{Bob} &= d_{Bob}[i] + f \cdot (d_{Bob}[i + 1] - d_{Bob}[i]) \\
&= d_{Bob}[4] + 0.17 \cdot (d_{Bob}[5] - d_{Bob}[4]) \\
&= 0.3 + 0.17 \cdot (0.5 - 0.3) = 0.33
\end{aligned}$$

where,

$$\begin{aligned}
i &= \left\lfloor \frac{c(\text{Alice}, \text{Carol}) \cdot (n_{Bob} + 1)}{100} \right\rfloor \\
&= \left\lfloor \frac{41.67 \cdot (9 + 1)}{100} \right\rfloor = 4
\end{aligned}$$

and,

$$\begin{aligned}
f &= \frac{c(\text{Alice}, \text{Carol}) \cdot (n_{Bob} + 1)}{100} - i \\
&= \frac{41.67 \cdot (9 + 1)}{100} - 4 = 0.17
\end{aligned}$$

The implementation of the functions *percentile* and *trustvalue* is based on the method for estimation of percentiles given by NIST [29].

5 Experiments

5.1 Experiment Design

Our objective is to test if the trust values recommended through the subjectivity-eliminated trust recommendation method are of higher quality than those given by the unmodified trust recommendation method in which trust values are conveyed without any alteration. The quality of a recommended trust value is stated as its closeness to the trust value that the source agent would assign to the target agent if it had direct experience with it.

Given a web of trust, we find paths of length 2 such that there also exists a direct edge from the source agent to the target agent. For such an instance, not only can we calculate the subjectivity-eliminated recommended trust value but we also know what value the source agent has assigned to the target agent based on direct experience. We therefore have a reference value with which we can compare the values given by the subjectivity-eliminated trust recommendation method and the unmodified trust recommendation method.

If the value given by the subjectivity-eliminated trust recommendation method is closer to the reference value than the one given by the unmodified trust recommendation method, we consider the experiment run as a success (hit) for our method. If the opposite is true, we consider it a failure (miss). If both values are the same or are within a small range (0.05) of each other, we count neither a hit nor a miss.

To facilitate the discussion we establish the following terminology:

- α – recommended trust value given by the unmodified trust recommendation method which does not take subjectivity into account
- β – recommended trust value derived from the subjectivity-eliminated trust recommendation method
- γ – trust value depicting the source agent’s trust in the target agent based on direct experience

Given G , a web of trust, and z , the minimum number of outgoing edges for source and recommender agents, the experiment is algorithmically described in Figure 1.

As discussed in Section 4.1, an agent must have made a significant number of trust value assignments in the past for a close approximation of their disposition to trust. z represents this number. We experiment with different values in Section 5.3.

Given a large and diverse web of trust we can assume that there will be both hits and misses. If the number of hits is significantly larger than the number of misses, we have an indication that the method is effective. On the contrary if the number of misses is considerably greater than the number of hits or if there is no significant pattern then we may infer that the method is ineffective.

The experiment has been implemented using the Java Graph library (JGraphT). When determining an alternate path, the first path returned by Dijkstra’s algorithm that meets the given criteria is used. In the following sections we describe a web of trust and proceed with experiment runs.

SUBJECTIVITY-EXPERIMENT(G, z)

```

1  hits  $\leftarrow$  0
2  misses  $\leftarrow$  0
3  equals  $\leftarrow$  0
4  for all edges in  $G$ , whose source vertex (given as  $a_s$ ) and target vertex
   (given as  $a_t$ ) are not the same
5      do  $\gamma \leftarrow t(a_s, a_t)$ 
6          remove the edge  $(a_s, a_t)$ 
7          find an alternate path,  $p(a_s, a_t)$  from  $a_s$  to  $a_t$ , such that the length of
            $p(a_s, a_t)$  is equal to 2, that is,  $p(a_s, a_t) = \langle a_s, a_r, a_t \rangle$  where
            $a_r$  is a recommender vertex, and  $a_s$  and  $a_r$  have a minimum
           of  $z$  outgoing edges
8          if  $p(a_s, a_t)$  exists
9              then  $\alpha \leftarrow t(a_r, a_t)$ 
10                  $\beta \leftarrow \text{trustvalue}(\text{percentile}(t(a_r, a_t), d_{a_r}), d_{a_s})$ 
11                 if  $\alpha = \beta$  or  $|\alpha - \beta| < 0.05$ 
12                     then equals ++
13                 elseif  $|\beta - \gamma| < |\alpha - \gamma|$ 
14                     then hits ++
15                 elseif  $|\alpha - \gamma| < |\beta - \gamma|$ 
16                     then misses ++
17                 restore the edge  $(a_s, a_t)$ 
18  print hits, misses, equals

```

Fig. 1. Experiment design

5.2 Data Set

We generate a simulated web of trust based on a random graph [30] as described in Figure 2. n is the number of vertices in the graph, k is the number of outgoing edges of each vertex, and G is the generated graph.

As we discussed in Section 2, different source agents may assign different trust values to a target agent. This occurs due to their different dispositions to trust even though their individual experiences with the target agent are the same.

These ideas are reflected in the generation of this web of trust. The trustworthiness value q_{u_i} represents the experience that other agents would have with agent u_i . Since q_{u_i} remains constant for agent u_i , any agent that interacts with it has the same experience. Although this would not always be true in a real web of trust, placing this condition sets up a suitable controlled environment for our experiments. If there is an instance where the subjectivity-eliminated trust recommendation method is ineffective, we know that it is not because multiple agents may have assigned u_i different trust values due to different experiences, in which case subjectivity is irrelevant. The failure is in fact on part of the method.

GENERATE-WEB-OF-TRUST(n, k)

```

1  create an empty weighted directed graph,  $G(V, E)$ , where  $V$  is the set of
    vertices and  $E$  is the set of edges
2  populate  $V$  with  $n$  vertices, labeled  $u_i$ , where  $i = 0, 1, \dots, n - 1$ 
3  with each vertex  $u_i$ , associate a random trustworthiness value  $q_{u_i}$ 
    from the interval  $[0,1]$ 
4  with each vertex  $u_i$ , associate a random skew factor  $s_{u_i}$  from the interval  $[0,2]$ 
5  for each vertex  $u_i$ 
6      do select  $k$  random distinct vertices from  $V$ , refer to them as  $v_j$ , where
         $j = 0, 1, \dots, k - 1, u_i \neq v_j$ 
7      for each vertex  $v_j$ 
8          do create the edge  $(u_i, v_j)$  in  $E$ 
9              assign the weight  $power(q_{v_j}, s_{u_i})$  to  $(u_i, v_j)$ 
10 return  $G$ 

```

Fig. 2. Pseudo code for generating the web of trust

The skew factor represents the individual disposition to trust of each agent. Although different agents have the same experience with a given agent u_i , they each assign it a different trust value based on their own disposition to trust. If the skew factor s_{u_i} is less than 1, q_{v_j} would be skewed upwards. Otherwise if the skew factor s_{u_i} is greater than 1, q_{v_j} would be skewed downwards.

Weights or trust values are drawn from the set of real numbers between 0 and 1 therefore the resolution for expressing trust is high.

The resulting data set is a web of trust where we know that subjectivity in fact does exist.

The web of trust consists of n vertices and $n \cdot k$ edges. If the number of vertices is 1000 and $k = 100$, the total number of edges is $n \cdot k = 100,000$. A new web of trust is generated for each run according to the values of n and k under consideration. The number of outgoing edges for all vertices is exactly k , therefore $z = k$.

5.3 Experiment Runs and Observations

The results of two sets of experiment runs are given in Table 1 and Table 2. We note that with $n = 1000$, and $z = k = 180$, 95% of the time, the subjectivity-eliminated trust recommendation method gives better results than those given by the unmodified trust recommendation method (not considering instances when both methods give equal results).

We also note that increasing z improves the effectiveness of the method. However, increasing n while keeping z constant (that is, decreasing the connectivity of the graph) does not seem to deteriorate the effectiveness of the method.

Table 1. Experiment runs with $n = 1000$

z, k	<i>hits</i>	<i>misses</i>	<i>equals</i>	$\frac{\textit{hits}}{\textit{hits}+\textit{misses}}$
10	0	0	0	-
20	0	0	0	-
30	16345	3568	6376	82%
40	39246	7371	15531	84%
50	80191	12439	29936	87%
60	141860	20251	50283	88%
70	223511	29094	85819	88%
80	332837	43046	130526	89%
90	488874	52617	180220	90%
100	674139	63542	253553	91%
110	903407	85568	331536	91%
120	1175525	97396	441145	92%
130	1520318	107460	554661	93%
140	1892642	137848	698261	93%
150	2383352	142981	830549	94%
160	2809821	181346	1084773	94%
170	3450976	195444	1242734	95%
180	4154572	203933	1448044	95%

Table 2. Experiment runs with $z = k = 100$

n	<i>hits</i>	<i>misses</i>	<i>equals</i>	$\frac{\textit{hits}}{\textit{hits}+\textit{misses}}$
1000	674139	63542	253553	91%
1200	673636	65947	251049	91%
1400	683320	64536	241659	91%
1600	680652	66192	246285	91%
1800	682642	64880	243262	91%

6 Discussion of Experiment Results / Future Work

The results of the experiment runs on the simulated web of trust provide a positive indication that the subjectivity-eliminated trust recommendation method is more effective than the unmodified method. Our method gives significantly better results when the number of outgoing edges of the agents is high. Even with relatively lower number of outgoing edges, the method still outperforms the one that does not account for subjectivity.

However, despite the strength of the results we can only consider them as a positive initial indication of the effectiveness of the method. Concrete conclusions are not feasible at this stage due to the reason that the simulated web of trust is a simplistic approximation of a real web of trust. Although the simulated web of trust takes into account real world issues such as rater bias, some other important aspects are simplified. For example, it is wired simply as a

random graph. Another simplification is the uniform distribution of the dispositions to trust and the trustworthiness of the agents. These simplifications may very well not have any impact on the effectiveness of the method however that is a hypothesis which needs to be tested.

An evident direction for future work is to test the method on a real web of trust or a closely approximated simulated web of trust. Some ideas for generating a more realistic web of trust include: 1) connectivity based on small-world [31] or scale-free networks [32], which are better representations of social networks, and 2) representing the dispositions to trust and the trustworthiness of the agents by a distribution such as normal or power-law. These ideas were not incorporated into this paper since their implementation is not straightforward. Most of the work on small-world and scale-free networks relates to undirected graphs. Methods for generating directed graphs are often for citation-like networks where older nodes do not have edges to newer nodes, which is not the case in a web of trust. So far we have also not come across any existing studies on the distributions of disposition to trust and trustworthiness.

7 Conclusion

This paper delved into the problem of subjectivity in trust recommendation, which we argued prevents the real meaning of a trust value from being conveyed by one agent to another. We presented a method for the elimination of subjectivity from trust recommendation that takes advantage of trust scores given as percentiles, which are equally meaningful among two agents. Experiments conducted on a simulated web of trust demonstrated that the method is highly effective for elimination of subjectivity from trust recommendation. The method is non-intrusive and does not require any change in how agents locally evaluate other agents. Furthermore, the method does not involve any third party mediation, thus making it suitable for decentralized networks. Validation of the experiment results on a real web of trust or a closely approximated simulated web of trust is proposed as future work. It is our hope that this paper will also serve as an introduction to the problem of subjectivity in trust recommendation and that it will inspire further research on this problem which has not received considerable attention.

References

1. Hoffman, K., Zage, D., Nita-Rotaru, C.: A survey of attack and defense techniques for reputation systems. Technical Report CSD TR 07-013, Department of Computer Science, Purdue University, IN, USA (2007)
2. Abdul-Rahman, A., Hailes, S.: Supporting trust in virtual communities. In: Proceedings of the 33rd Hawaii International Conference on System Sciences (January 2000)
3. Baras, J.S., Jiang, T.: Cooperation, Trust and Games in Wireless Networks. In: Advances in Control, Communication Networks, and Transportation Systems: In Honor of Pravin Varaiya. Birkhauser, Boston (2005)

4. Beth, T., Borcherding, M., Klein, B.: Valuation of trust in open networks. In: Gollmann, D. (ed.) ESORICS 1994. LNCS, vol. 875. Springer, Heidelberg (1994)
5. Capra, L.: Engineering human trust in mobile system collaborations. In: Proceedings of the 12th ACM SIGSOFT International Symposium on Foundations of Software Engineering, Newport Beach, CA, USA (2004)
6. Griffiths, N.: Task delegation using experience based multidimensional trust. In: Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, The Netherlands (2005)
7. Meng, K., Zhang, X., chun Xiao, X., du Zhang, G.: A bi-rating based personalized trust management model for virtual communities. In: Proceedings of the 2006 IEEE International Conference on Networking, Sensing and Control (ICNSC 2006) (2006)
8. Marsh, S.P.: Formalising Trust as a Computational Concept. PhD thesis, Department of Mathematics and Computer Science, University of Stirling, Scotland, UK (1994)
9. Golbeck, J., Hendler, J.: Filmtrust: Movie recommendations using trust in web-based social networks. In: Proceedings of the 3rd IEEE Consumer Communications and Networking Conference (CCNC 2006), Las Vegas, NV, USA (2006)
10. Gambetta, D.: Can We Trust Trust? In: Trust: Making and Breaking Cooperative Relations, pp. 213–237. Department of Sociology, University of Oxford (2000)
11. Toivonen, S., Lenzi, G., Uusitalo, I.: Context-aware trust evaluation functions for dynamic reconfigurable systems. In: Proceedings of the WWW'06 Workshop on Models of Trust for the Web (MTW 2006) (May 2006)
12. Jonker, C.M., Treur, J.: Formal analysis of models for the dynamics of trust based on experiences. In: Garijo, F.J., Boman, M. (eds.) MAAMAW 1999. LNCS, vol. 1647. Springer, Heidelberg (1999)
13. Levien, R.: Attack resistant trust metrics. Manuscript, University of California - Berkeley (2002), <http://www.levien.com/thesis/compact.pdf>
14. Resnick, P., Zeckhauser, R.: Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *The Economics of the Internet and E-Commerce* 11 (2002)
15. McKnight, D.H., Choudhury, V., Kacmar, C.: Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research* 13(3), 334 (2002)
16. Rotter, J.B.: A new scale for the measurement of interpersonal trust. *Journal of Personality* 35(4), 651 (1967)
17. Rotter, J.B.: Generalized expectancies for interpersonal trust. *American Psychologist* 26(5), 443 (1971)
18. Boon, S.D., Holmes, J.G.: The Dynamics of Interpersonal Trust: Resolving Uncertainty in the Face of Risk. In: *Cooperation and Prosocial Behaviour*, p. 190. Cambridge University Press, Cambridge (1991)
19. McCord, M., Ratnasingam, P.: The impact of trust on the technology acceptance model in business to consumer e-commerce. In: Proc. of the Intl. Conf. of the Information Resources Mgmt. Association: Innovations Through IT, New Orleans, LA, USA (May 2004)
20. Kaluscha, E.A.: The Importance of Initial Consumer Trust in B2C Electronic Commerce - A Structural Equation Modeling Approach. PhD thesis, University of Klagenfurt, Austria (2004)
21. Ziegler, C.N., Lausen, G.: Spreading activation models for trust propagation. In: Proceedings of the 2004 IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE 2004) (2004)

22. Golbeck, J., Parsia, B., Hendler, J.: Trust networks on the semantic web. In: Klusch, M., Omicini, A., Ossowski, S., Laamanen, H. (eds.) CIA 2003. LNCS, vol. 2782, pp. 238–249. Springer, Heidelberg (2003)
23. Chen, R., Yeager, W.: Poblano: A distributed trust model for peer-to-peer networks. Technical report, Sun Microsystems (2000)
24. Abdul-Rahman, A., Hailes, S.: A distributed trust model. In: Proceedings of the 1997 Workshop on New Security Paradigms, Langdale, Cumbria, United Kingdom (1997)
25. Christie, R., Geis, F.L.: Studies in Machiavellianism. Academic Press, New York (1970)
26. Hirsh, H., Basu, C., Davison, B.D.: Learning to personalize. Communications of the ACM 43(8), 102 (2000)
27. Langley, P.: Adaptive user interfaces and personalization (2008), <http://www.isle.org/~langley/adapt.html> (retrieved January 5, 2009)
28. Fog, A.: Branch prediction in the pentium family (2008), <http://x86.org/articles/branch/branchprediction.htm> (retrieved January 5, 2009)
29. National Institute of Standards and Technology (NIST): Nist/sematech e-handbook of statistical methods - percentiles (2008), <http://www.itl.nist.gov/div898/handbook/prc/section2/prc252.htm> (retrieved January 5, 2009)
30. Bollobas, B.: Random Graphs. Cambridge University Press, Cambridge (2001)
31. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature 393 (June 1998)
32. Albert, R., Barabasi, A.L.: Statistical mechanics of complex networks. Reviews of Modern Physics 74 (January 2002)