Optimizing Advisor Network Size in a Personalized Trust-Modelling Framework for Multi-Agent Systems

Joshua Gorner and Robin Cohen

David R. Cheriton School of Computer Science, University of Waterloo
Waterloo, Ontario, N2L 3G1
{jgorner, rcohen}@uwaterloo.ca

Abstract. This paper explores potential improvements to Zhang's personalized trust approach for e-commerce, in particular examining means of optimizing the number of advisors that each buyer maintains in their social network. We propose three such improvements, two directly relating to the size of the network (through either the use of a threshold or by setting a maximum network size), and a third which may indirectly reduce the necessary size by ensuring that relevant trust information from users outside this network remains available. We provide examples to illustrate these approaches, and propose future work to optimize certain aspects of these methods and evaluate their overall effectiveness.

Keywords: user modelling, multi-agent systems, trust modelling, social networks, electronic commerce

1 Introduction

Zhang [1] has recently proposed a novel trust-based framework primarily developed for use in agent-oriented electronic commerce. This system relies on a model of the trustworthiness of a particular buyer's advisors — that is, other buyers within the system that report on sellers — which incorporates estimates of the advisor's private and public reputations. Users create a social network of trusted advisors, and sellers will offer better rewards in order to satisfy trusted advisors and thus build their own reputations.

In this paper, we look at one of the open questions regarding the optimal size of a user's advisor network. Retaining a large network could be inefficient, and may result in reduced accuracy in determining trust in a seller: many potential advisors will not have similar tastes to the current user, so a large network could make predictions less accurate. However, with a smaller network, there is a greater risk that the advisors will have insufficient experience [2]. We suggest three modifications to Zhang's model which may help in this regard, and demonstrate how the revised approach would operate in each case.

2 Related Work in Selecting Advisors

2.1 Mechanism

The mechanism we examine here [1] uses a multi-stage "personalized" approach for representing reputation in an e-commerce system, summarized below.

We note at the outset that the model only considers two possible ratings, positive (1) or negative (0). A buyer, denoted by b, first constructs a measure of private reputation of advisors, based on the advisors' ratings for sellers that b has previously dealt with, and representing an estimation of the probability that an advisor a will give fair ratings to b. Each pair of ratings considered is weighted based on the amount of time that separates the submission of the two ratings using a "forgetting factor", λ ($0 \leq \lambda \leq 1$), such that a pair of ratings will have a greater weight if they are made within close time proximity. This measure is known as "private" reputation, since this evaluation makes use of the buyer's own experiences, and is calculated as shown in equation 1:

$$\alpha = N_p + 1, \qquad \beta = N_{all} - N_p + 1, \qquad R_{pri}(a) = E(Pr(a)) = \frac{\alpha}{\alpha + \beta} \tag{1}$$

In Equation 1, N_p represents the sum of the weights of all positive rating pairs (that is, pairs of identical ratings) for all sellers commonly rated by b and a, and N_{all} is the total sum of weights of all rating pairs involving b and a. If $\lambda = 0$, then N_p and N_{all} will be simply the counts of the applicable types of rating pairs.

Next, the public reputation of an advisor, or the probability that an advisor will provide "consistent" ratings, is calculated similarly, using equation 2:

$$\alpha' = N_c + 1, \qquad \beta' = N'_{all} - N_c + 1, \qquad R_{pub}(a) = \frac{\alpha'}{\alpha' + \beta'} \tag{2}$$

Here, N_c represents the number of ratings, provided by an advisor a, that are consistent with the majority of ratings provided for that seller up to the moment that this additional rating is submitted, while N'_{all} is the total number of ratings provided by a.

At this point, given some maximum acceptable level of error $\epsilon \in (0, 1)$ and level of confidence $\gamma \in (0, 1)$, we derive w, the reliability of the private reputation value, which we then use in our calculation of the overall trustworthiness of a. As can be seen, a more reliable private reputation will have a greater effect on the overall result:

$$N_{min} = -\frac{1}{2\epsilon^2} ln \frac{1 - \gamma}{2} \tag{3}$$

$$w = \begin{cases} \frac{N_{all}}{N_{min}} & \text{if } N_{all} < N_{min} \\ 1 & \text{otherwise} \end{cases} \tag{4}$$

$$Tr(a) = wR_{pri}(a) + (1 - w)R_{pub}(a) \tag{5}$$

Once this value has been calculated for each advisor, a similar approach can be taken for the trustworthiness of a given seller s. First the buyer b calculates her private reputation of s, or the probability that s will provide good service, based on b's past experiences with s. This makes use of the number of positive ratings, $N^b_{pos,i}$, and negative ratings, $N^b_{neg,i}$, she provided for s in each time window T_i, as well as the forgetting factor λ:

$$R_{pri}(s) = \frac{\sum_{i=1}^{n} N_{pos,i}^{b} \lambda^{i-1} + 1}{\sum_{i=1}^{n} (N_{pos,i}^{b} + N_{neg,i}^{b}) \lambda^{i-1} + 2} \tag{6}$$

Next we derive the public reputation of the seller, the probability that the seller will provide good service given all advisors' past experiences with s, taking into account b's own model of trustworthiness of each advisor a_j. We first make use of the following discounting functions to determine b's trust of ratings provided by each a_j:

$$D_{pos_i}^{a_j} = \frac{2Tr(a_j) N_{pos,i}^{a_j}}{(1 - Tr(a_j))(N_{pos,i}^{a_j} + N_{neg,i}^{a_j}) + 2}, D_{neg_i}^{a_j} = \frac{2Tr(a_j) N_{neg,i}^{a_j}}{(1 - Tr(a_j))(N_{pos,i}^{a_j} + N_{neg,i}^{a_j}) + 2} \tag{7}$$

The public reputation of s is itself calculated as follows:

$$R_{pub}(s) = \frac{[\sum_{j=1}^{k} \sum_{i=1}^{n} D_{pos,i}^{a_j} \lambda^{i-1}] + 1}{[\sum_{j=1}^{k} \sum_{i=1}^{n} (D_{pos,i}^{a_j} + D_{neg,i}^{a_j}) \lambda^{i-1}] + 2} \tag{8}$$

Finally the overall trustworthiness of the seller s may be calculated:

$$w' = \begin{cases} \frac{N_{all}^{b}}{N_{min}} & \text{if } N_{all}^{b} < N_{min} \\ 1 & \text{otherwise} \end{cases}$$

$$Tr(s) = w' R_{pri}(s) + (1 - w') R_{pub}(s) \tag{9}$$

where N_{min}, the minimum number of ratings needed for the buyer b to be confident about the private reputation value it has of the seller s, is calculated according to equation 3, but is not necessarily the same value used in equation 4.

The model also includes an incentive mechanism, whereby honest advisors are rewarded by better offers from sellers, and in turn these sellers receive better reputations and ultimately more customers. While interesting in its own right, this mechanism does not directly affect our current work, and therefore we do not discuss this part of Zhang's model further.

2.2 Towards a Method for Determining Advisors

The Beta Reputation System (BRS) [3] uses beta population density functions in order to combine feedback from multiple sources and subsequently derive reputation ratings. The initial evaluation gives some broad estimates about the number of ratings required to obtain an accurate group rating; however, these are at best snapshots of what size network *might* be needed to obtain a stable rating under certain conditions. Ultimately, the evaluation provides little insight into what an optimal network size would be under the BRS.

TRAVOS [4] calculates trust using probability theory, taking into account any past interactions between agents, or alternatively reputation information provided by third parties. Here, each agent is expected to maintain a trust assessment for every other agent in the system that it has interacted with in the past. Under the circumstances, it seems that restricting the size of the advisor network is unrealistic. Indeed, to our knowledge, nothing of this nature have been explored with respect to the TRAVOS system.

Some potential methods for limiting the number of advisors may be derived from the work done in [2] to evaluate various design choices in a collaborative filtering algorithm. The first method, *correlation thresholding* [5], sets an minimum correlation weight that an advisor must have in order to be considered part of the user's "neighbourhood". However, if the threshold is set too high, then the neighbourhood may be very small, limiting the possibilities for predictions. In fact, for the data set examined in [2], correlation thresholding yielded declines in both coverage and accuracy compared to a non-thresholded algorithm.

The second method discussed, *best-n-neighbors*, as used in the GroupLens [6] system among others, picks a maximum number of neighbours to use, *max_nbors*. The neighbours chosen would be those with the highest correlation to the instant user. In [2], it was shown that a neighbourhood of 20 to 50 users (out of a population of 943) was found to provide an acceptable level of accuracy, providing an appropriate balance between sufficient coverage and eliminating inaccuracies.

A potential supplement to finding the optimal number of advisors is derived from [7], which discussed reputation management in a social network making use of an *advisor referral* mechanism. In this mechanism, a buyer's agent would consult its "neighbour" agents, each of which might either provide advice on the question itself, provide references to other appropriate advisors, or both, depending on the question. As a result, a buyer would be able to benefit from the information held by the pool of advisors without having a large number of neighbours [1]. It then stands to reason that this method could be used in combination with network-size optimization to provide a smaller advisor network size.

3 Analysis

3.1 Limiting the Network Size

The results in [2] would appear to point towards setting a maximum number of advisors as the preferable method of restricting the advisor network size, as opposed to correlation (or, in our case, trust) thresholding. However, we cannot overlook the distinction between correlation for collaborative filtering and reputation. While similarity with a buyer may indirectly impact on that buyer's private reputation of an advisor, the private reputation of a seller only relates to the buyer's ratings for that seller, ignoring similarity, while similarity is not a factor at all in public reputation. Hence we propose that both options, trustworthiness thresholding and maximum number of advisors, should be thoroughly examined.

That said, neither technique could be directly applied to Zhang's model for advisor reputation; the trustworthiness values must be calculated for all possible advisors before the buyer can proceed to calculate seller reputation.

Our application of these techniques in the seller reputation model are formalized as follows.

Trustworthiness Thresholding Choose some threshold L $(0 \leq L \leq 1)$ which represents the minimum advisor trustworthiness value $Tr(a)$ required to be included in the advisor network. We then define the set $A_{L,b} = \{a_1, a_2, \dots, a_k\}$ consisting of all advisors for which $Tr(a) \geq L$ for a particular buyer b. We then use the subset $A_{L,b,s}$, consisting of the advisors in $A_{L,b}$ that have provided ratings for the seller s, in place of the previously-defined set $\{a_1, \dots, a_k\}$, the set of all advisors that have provided ratings for s, in Zhang's Algorithm 2 (the seller reputation algorithm outlined in Equations 7 to 9 in this paper).

Maximum Number of Advisors For a particular buyer b, after having calculated the personalized trustworthiness of each advisor for b as per the first part of Zhang's model, we sort the list of all n advisors from greatest trustworthiness value to least, in the set $\{a_1, a_2, \dots, a_n\}$. We choose some maximum number of advisors for each buyer, $max\_nbors \leq n$, and then truncate this set to the set $A_b = \{a_1, a_2, \dots, a_{max\_nbors}\}$. We thus obtain the set of max_nbors advisors that have been calculated to be the most trustworthy for b. Again, the subset of A_b that has provided ratings for the seller s is used in place of the larger set $\{a_1, \dots, a_k\}$ in Zhang's Algorithm 2.

3.2 Advisor Referrals

We also wish to consider the possibility of combining one or both of the above methods with the advisor-referral technique suggested in [7] and discussed above. We diverge somewhat from the original mechanism insofar as Zhang's model does not require us to query each advisor for a recommendation. Rather, the buyer has access to each advisor's ratings for a given seller s via a central server, and uses this data to determine the public (or network) reputation for the seller.

We thus consider that advisors can "advise" by allowing buyers to make use of each advisor's own private reputation for a certain seller. In this case, an advisor "referral" system could be implemented using a variant of the measure used to weight private reputation in the original model. This would work as follows: For each advisor a_j in the advisor network of b, that is, the set $\{a_1, a_2, \dots, a_k\}$, b checks whether advisor a_j is an acceptable advisor for the seller s. This will be the case if $N_{all}^{a_j} \geq N_{min}$, where $N_{all}^{a_j}$ is the number of ratings provided by an advisor a_j for s, and N_{min} is some minimum number of ratings (which may be calculated using equation 3).

If a_j is not an acceptable advisor (that is, if $N_{all}^{a_j} < N_{min}$), the algorithm will query a_j's advisor network, sorted from most trustworthy to least trustworthy from the perspective of a_j, in order to determine, in a similar fashion, which (if any) of these advisors meet the criteria to be a suitable advisor for s. The first such advisor encountered that is itself not either (a) already in the set of acceptable advisors; or (b) in A_b — since this would imply that the recommended advisor would be added in any event at a later stage — will be accepted.

If none of the advisors of a_j meet the above criteria, this step would be repeated at each subsequent level of the network — that is, the advisors of each member of the set of advisors just considered — until an acceptable, unduplicated advisor was identified.

Once the full set of acceptable advisors has been determined, the "network" reputation would be calculated as in Zhang's model, using the advisor trustworthiness values held by the buyer b. This, of course, assumes that the seller s has had sufficient past interactions with the various advisors in the network such that there are at least k buyers that have each had at least N_{min} interactions with s, which is not guaranteed. If only a smaller number of acceptable advisors can be found, the system will simply use this reduced set to determine the network reputation.

To ensure broad coverage of the network while preventing infinite recursion, we limit the number of network "levels" calculated to at most $\lceil log_k(|B|)\rceil$, where B is the set of all buyers (advisors) in the system. However, we note that practically, in a large scale system, the number of levels may need to be smaller in order for this algorithm to be computationally efficient; we will leave such a decision for later work.

We summarize this mechanism in pseudo-code format as Algorithm 1.

4 Examples

4.1 Using Zhang's Model

As in [1], we consider the case where a buyer b wishes to assess the trustworthiness of a particular seller s_0 with whom the buyer has had little or no experience. For the purposes of this simplified example, we assume that there are three available advisors from which s_0 may seek advice, namely a_x, a_y, and a_z.

We assume initially that, among sellers that b has had past dealings with, each of these advisors has provided ratings only for the five sellers $(s_1, s_2, s_3, s_4, s_5)$, and has rated each of the sellers at most once in each time window in the sequence T, where T_1 is the most recent time window. The ratings may be either positive (1) or negative (0); a dash (-) indicates that no rating was provided during the indicated time window. The ratings provided by each advisor for these sellers are listed in Table 1. The buyer b has also provided some ratings for the sellers, as indicated in the same table; note here that b does not provide ratings for every seller each time window.

Algorithm 1 Selecting Advisors to Buyer b for Trustworthiness of Seller s Using Referrals

$A_b = \{a_1, a_2, \ldots, a_k\}$; {advisors in b's advisor network}
$A_s = \{\}$; {set of advisors that are suitable for providing advice regarding seller s}
N_{min} = minimum number of ratings for a to be a suitable advisor regarding s;
$maxnetlevel = \lceil log_k(|B|) \rceil$ {the maximum number of search iterations}
for $j = 1$ to k **do**
  $N_{all}^{a_j}$ = total number of ratings provided by a_j for s;
  if $N_{all}^{a_j} \geq N_{min}$ **then**
    append a_j to A_s;
  else
    $netlevel = 2$; {the number of connections between b and the advisors being searched}
    a_x = **null**; {the desired suitable advisor in place of a_k}
    A_c = the set of advisors for a_j sorted from most to least trustworthy (as per a_j);
    while a_x == **null** and $netlevel \leq maxnetlevel$ **do**
      $A_n = \{\}$; {the set of advisors to be considered in the next round, if necessary}
      for all a_c in A_c **do**
        $N_{all}^{a_c}$ = total number of ratings provided by a_c for s;
        if $N_{all}^{a_c} \geq N_{min}$ **and** $a_c \notin A_b$ **and** $a_c \notin A_s$ **then**
          $a_x = a_c$;
          break;
        else
          add the set of advisors for a_c to A_n;
        end if
      end for
      $netlevel + +$;
      $A_c = A_n$
    end while
    if $a_x \neq$ **null then**
      append a_x to A_s;
    end if
  end if
end for

Table 1. Ratings of Sellers Provided by Advisors and Buyer b

	a_w					a_x					a_y					a_z					b				
T	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5
s_1	0	0	0	0	1	1	1	1	1	1	1	1	0	1	0	0	0	0	0	0	1	1	1	1	1
s_2	0	0	1	1	1	1	1	1	1	1	0	1	1	0	1	0	0	0	0	0	1	1	1	1	-
s_3	0	1	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	0	0	0	1	1	1	-	-
s_4	0	0	0	1	1	1	1	1	1	1	0	1	0	1	0	0	0	0	0	0	1	1	-	-	-
s_5	0	1	1	1	1	1	1	1	1	1	1	0	0	0	1	0	0	0	0	0	1	-	-	-	-

We derive the trustworthiness values for each advisor using equations 1 through 5. For simplicity, in these calculations, we follow the method used in the examples provided in [1]. First, we will only consider pairs of ratings provided during the same time window, and thus assume that the forgetting factor as defined previously is $\lambda = 0$. For the determination of N_c, we assume for simplicity that any rating of 1 provided by the advisor is a "consistent" rating. Finally, in equation 3 we use $\gamma = 0.8$ and $\epsilon = 0.15$, leading to $N_{min} = 51$. The pertinent values are shown in Table 2.

Table 2. Trustworthiness of Advisors a_w, a_x, a_y, and a_z for Buyer b

a_j	N_p	N_{all}	α	β	R_{pri}	N_c	N'_{all}	α'	β'	R_{pub}	w	$Tr(a)$
a_w	5	15	6	11	0.353	14	25	15	12	0.556	0.294	0.497
a_x	15	15	16	1	0.941	25	25	26	1	0.963	0.294	0.957
a_y	8	15	19	8	0.529	11	25	12	15	0.444	0.294	0.469
a_z	0	15	1	16	0.059	0	25	1	26	0.037	0.294	0.0434

We proceed to the calculation the trustworthiness of a seller s_0. As a preliminary matter, we remember that the buyer b has not provided any ratings in the past for s_0, and therefore $R_{pri}(s) = \frac{1}{2}$. Of our four advisors, only a_w, a_x and a_z have provided ratings for the seller s_0, as indicated in Table 3(a). The subsequent Table 3(b) indicates how these ratings translate into positive and negative amounts, while Table 3(c) shows how these ratings are discounted based on the advisor trustworthiness values calculated earlier.

Table 3. Ratings of s_0 Provided by a_w, a_x, a_z

(a) Ratings

T_i	T_1	T_2	T_3	T_4	T_5
a_w	1	0	1	0	1
a_x	0	0	0	1	1
a_z	1	1	1	1	1

(b) Amounts of Ratings

T_i	T_1	T_2	T_3	T_4	T_5
$N^{a_w}_{pos,i}$	1	0	1	0	1
$N^{a_w}_{neg,i}$	0	1	0	1	0
$N^{a_x}_{pos,i}$	0	0	0	1	1
$N^{a_x}_{neg,i}$	1	1	1	0	0
$N^{a_z}_{pos,i}$	1	1	1	1	1
$N^{a_z}_{neg,i}$	0	0	0	0	0

(c) Discounted Amounts of Ratings

T_i	T_1	T_2	T_3	T_4	T_5
$D^{a_w}_{pos,i}$	0.397	0	0.397	0	0.397
$D^{a_w}_{neg,i}$	0	0.397	0	0.397	0
$D^{a_x}_{pos,i}$	0	0	0	0.937	0.937
$D^{a_x}_{neg,i}$	0.937	0.937	0.937	0	0
$D^{a_z}_{pos,i}$	0.0294	0.0294	0.0294	0.0294	0.0294
$D^{a_z}_{neg,i}$	0	0	0	0	0

Using equation 8, we may then find the public reputation of s_0. In keeping with the examples provided in [1], we remove our previously-stated simplification that only compared ratings in the same time window, and thus set a forgetting factor of $\lambda = 0.9$:

$$R_{pub}(s_0) = \frac{\sum_{i=4}^{5} 0.937 * 0.9^{i-1} + 0.397 * (0.9^0 + 0.9^2 + 0.9^4) + \sum_{i=1}^{5} 0.0294 * 0.9^{i-1} + 1}{\sum_{i=1}^{5} 0.937 * 0.9^{i-1} + \sum_{i=1}^{5} 0.397 * 0.9^{i-1} \sum_{i=1}^{5} 0.0294 * 0.9^{i-1} + 2} = 0.4480$$

Finally, since the buyer has not dealt with s_0 before, the weight for the private reputation w' is zero, meaning we can immediately conclude that $Tr(s_0) = 0.4480$.

4.2 Reputation Thresholding

We now turn to exploring the effects of the modifications proposed in this paper by first examining how setting a minimum reputation threshold would affect the size of our network and our seller reputation model. We choose several potential values for the threshold L and indicate, based on the results in the previous section regarding advisor trustworthiness, how many advisors would be included in the buyer b's advisor network in this case. The results are shown in Table 4.

Table 4. Advisor Network Size with a Correlation Threshold

L	0	0.2	0.4	0.6	0.8	1
$A_{L,b}$	$\{a_w, a_x, a_y, a_z\}$	$\{a_w, a_x, a_y\}$	$\{a_w, a_x, a_y\}$	$\{a_x\}$	$\{a_x\}$	$\{\}$
$\lvert A_{L,b}\rvert$	4	3	3	1	1	0

Trivially, when $L = 0$, all advisors will be included in the network, and $Tr(s_0) = 0.4480$. For $L = 0.2$ and $L = 0.4$, the advisor network consists of a_w, a_x, and a_y, of which only a_w and a_x contribute ratings for s_0. We refer to the resulting trustworthiness value as $Tr_{w,x}(s_0)$, which we calculate using Equation 8 to be 0.439.

For $L = 0.6$ and $L = 0.8$, the advisor network consists solely of a_x, and therefore the seller trustworthiness $Tr_x(s_0)$ (again by Equation 8) would be 0.697. Finally, for $L = 1$, the advisor network is the empty set and, trivially, $Tr_{empty}(s_0) = \frac{1}{2}$.

4.3 Maximum Number of Advisors

If we instead elect to use a maximum number of advisors, we would have the results shown in table 5, with the advisor network representing the max_nbors advisors most trusted by the buyer b. For comparison, we indicate the minimum trustworthiness value of the advisors in the network, to show the maximum threshold L that could be used to get the same result using thresholding.

Table 5. Trustworthiness of s_0 Using a Maximum Number of Advisors

max_nbors	0	1	2	3	≥ 4
A_b	$\{\}$	$\{a_x\}$	$\{a_w, a_x\}$	$\{a_w, a_x, a_y\}$	$\{a_w, a_x, a_y, a_z\}$
$min(Tr(a))$	undefined	0.957	0.497	0.469	0.0434
$Tr(s_0)$	0.5	0.697	0.439	0.439	0.4480

4.4 Advisor Referrals

We now examine the addition of our advisor referral mechanism to this system. To do so, we introduce a new advisor into the system, a_v, as well as another seller, s_6. To this point a_v has only provided ratings for s_6, while b has not provided any ratings for that seller; therefore there are no commonly-rated sellers for a_v and b, and thus $Tr(a_v) = 0.5$ from the perspective of b.

We also assume, as in the $max\_nbors = 3$ or $L = 0.4$ cases described above, that the advisor network for b consists of $\{a_w, a_x, a_y\}$ — a_v is too new to have been considered as a potential advisor in that case, although for purposes of demonstration we assume that a_v has somehow been included in the advisor networks of some of the other advisors. Finally we set N_{min}, the minimum number of ratings for an advisor to be considered acceptable for a given seller, as 3. The ratings that have been given by each advisor, and the resulting discounted amounts, are as shown in Table 6.

Given this information, the buyer b will examine its advisor network and find that a_w and a_x are indeed acceptable advisors for s_6, since both have achieved at least N_{min} interactions with s_6. However, a_y has only had one interaction with s_6, and would therefore not be considered an acceptable advisor. The buyer will then look to a_y's advisor network to identify an appropriate substitute.

Suppose then that a_y also has a three-agent advisor network consisting of a_v, a_x, and a_z, with trustworthiness values 0.5, 0.6, and 0.7 respectively. This information will be gathered by b as the ordered list $\{a_z, a_x, a_v\}$. The buyer will then iterate through the set, discarding

Table 6. Ratings of s_6 Provided by Advisors

(a) Ratings

T_i	T_1	T_2	T_3	T_4	T_5
a_v	1	1	0	1	1
a_w	0	1	1	0	-
a_x	1	0	1	-	-
a_y	0	-	-	-	-
a_z	1	1	-	-	-

(b) Amounts of Ratings

T_i	T_1	T_2	T_3	T_4	T_5
$N^{a_v}_{pos,i}$	1	1	0	1	1
$N^{a_v}_{neg,i}$	0	0	1	0	0
$N^{a_w}_{pos,i}$	0	1	1	0	-
$N^{a_w}_{neg,i}$	1	0	0	1	-
$N^{a_x}_{pos,i}$	1	0	1	-	-
$N^{a_x}_{neg,i}$	0	1	0	-	-
$N^{a_y}_{pos,i}$	0	-	-	-	-
$N^{a_y}_{neg,i}$	1	-	-	-	-
$N^{a_z}_{pos,i}$	1	1	-	-	-
$N^{a_z}_{neg,i}$	0	0	-	-	-

(c) Discounted Amounts of Ratings

T_i	T_1	T_2	T_3	T_4	T_5
$D^{a_v}_{pos,i}$	0.4	0.4	0	0.4	0.4
$D^{a_v}_{neg,i}$	0	0	0.4	0	0
$D^{a_w}_{pos,i}$	0	0.397	0.397	0	-
$D^{a_w}_{neg,i}$	0.397	0	0	0.397	-
$D^{a_x}_{pos,i}$	0.937	0	0.937	-	-
$D^{a_x}_{neg,i}$	0	0.937	0	-	-
$D^{a_y}_{pos,i}$	0	-	-	-	-
$D^{a_y}_{neg,i}$	0.375	-	-	-	-
$D^{a_z}_{pos,i}$	0.0294	0.0294	-	-	-
$D^{a_z}_{neg,i}$	0	0	-	-	-

a_z as an unacceptable advisor (having provided only two ratings for s_6), and also a_x as it is already in b's advisor network. Finally, b would then accept a_v as the third advisor, as it has an acceptable level of experience with s_6 but is not part of b's own advisor network.

As in the previous examples, b does not itself have enough experience with s_6 to generate a private reputation. Therefore, using the above information for the set of advisors $\{a_v, a_w, a_x\}$, the forgetting factor $\lambda = 0.9$, and Equation 8, we find that $Tr(s_6) = 0.6655$.

If b had not used advisor referrals but instead relied solely on its existing advisor network, namely $\{a_w, a_x, a_y\}$, it would have obtained a significantly different result — $Tr(s_6) = 0.5549$. However, the latter result makes much less use of the experience within the network for s_6 than did the one incorporating advisor referrals.

5 Discussion

In this paper, we have outlined three potential improvements to Zhang's personalized trust-modelling approach — trustworthiness thresholding, maximum number of advisors, and advisor referrals — all of which aim to reduce the computational complexity required to derive recommendations for trustworthy sellers from its advisors, and to improve the accuracy of these recommendations.

Again, much work remains to show that any or all of these approaches would be effective in improving on Zhang's approach. A more in-depth experiment using a larger set of data is required in order to verify which of these methods will provide the best performance - if, indeed, any of them is superior to the results in Zhang's original model. These will likely be in similar format to the comparative evaluations provided in [1], including marketplace simulations involving a number of buyer / seller agents with varying levels of honesty.

We will first seek to determine optimal parameters, or ranges of parameters, for both the threshold and *max_nbors* methods. Subsequently we will determine how these optimal versions compare to each other and to Zhang's original model. Finally, presuming that at least one of these modifications provides improved performance, we will attempt to implement advisor referrals on that system; again, we may need to try a number of parameter values in order to determine an optimal method.

For each stage of the evaluation, we should consider how well each method performs at correctly distinguishing between honest and dishonest agents. As in [1], we may compare performance using the Matthews correlation coefficient [8], which provides a single measure for the quality of binary classifications, such as for honest and dishonest sellers.

Other open questions remain: we might consider, for example, Zhang's suggestion to apply this model to time-sensitive tasks which may require a buyer to make a very quick decision; here, the buyer would only have time to consult a limited number of advisors.

Finally, we have found very limited work in the past on the effects of the size of the advisor network, or indeed other characteristics such as advisor referrals, on the usage of trust-based approaches of this nature. This research, into improvements to a particular system, may only have limited application to trust in general; it is likely that further work will be required to generate some more concrete principles in this area.

6 Acknowledgements

This research was supported in part by funding from the Natural Sciences and Engineering Research Council of Canada.

References

1. Zhang, J.: Promoting Honesty in E-Marketplaces: Combining Trust Modeling and Incentive Mechanism Design. PhD thesis, University of Waterloo (2009)
2. Herlocker, J., Konstan, J.A., Riedl, J.: An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms. Information Retrieval **5**(4) (October 2002) 1386–4564
3. Jøsang, A., Ismail, R.: The beta reputation system. In: 15th Bled Electronic Commerce Conference. (2002)
4. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In: Fourth International Joint Conference on Autonomous Agents and Multiagent Systems. (2005) 997–1004
5. Shardanand, U., Maes, P.: Social information filtering: Algorithms for automating "word of mouth". In: Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems. (1995) 210–217
6. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J.: Grouplens: An open architecture for collaborative filtering of netnews. In: Proceedings of the 1994 Conference on Computer Supported Collaborative Work. (1994)
7. Yu, B., Singh, M.P.: A social mechanism of reputation management in electronic communities. In: Proceedings of Fourth International Workshop on Cooperative Information Agents. (2000) 154–165
8. Matthews, B.W.: Comparison of the predicted and observed secondary structure of t4 phage lysozyme. Biochimica et Biophysica Acta **405**(2) (October 1975) 442–451