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

Joshua Gorner

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

1 Problem

Zhang [1] has recently proposed a novel trust-based framework for systems including electronic commerce. This system relies on a model of the trustworthiness of advisors (other buyers offering advice to the current buyer) which incorporates estimates of each 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.

Specifically, Zhang's model incorporates a "personalized" approach for determining the trustworthiness of advisors from the perspective of a given buyer. The trustworthiness of each advisor is calculated using, in part, a private reputation value; that is, a comparison of the buyer's own ratings of sellers to that advisor's ratings among sellers that both users have had experience with. This is then combined with a public reputation value which remains consistent for all users, reflecting how consistent each advisor's ratings are with the community as a whole.

The system then determines personalized values for the trustworthiness of each seller, using both a private reputation value based on the user's own dealing with the seller, and a public reputation value calculated by weighting each advisor's rating for that seller by their advisor trustworthiness.

The model is intended to be applicable to any e-commerce marketplace where sellers can adjust the quality and price of their products. However, we primarily consider a procurement (reverse) auction system where buyers indicate what they wish to purchase, and then choose among several potential sellers, those offering goods in response to that request, based on price, quality, or other criteria.

In this research, we look one of the open questions stemming from Zhang's proposal, namely the determination of the optimal size of a user's advisor network — that is, the number of advisors on which the user relies.

The original model assumes that all buyers in the system will be considered. However, we believe that an overly large network of advisors could result in reduced inaccuracy in predicting the trustworthiness of a given seller, due to the inclusion of outlier data, even if these are weighted less. Moreover, if such a system was to be scaled up to potentially millions of users, it would not be computationally efficient to calculate trustworthiness values for all of them. That said, we also do not wish to restrict the size of the network so much that the remaining advisors have insufficient experience to accurately assess the sellers. [2]

We believe this work would help to extend the usefulness of Zhang's original model; specifically, our extension would help to improve the performance and accuracy of such a model in any future deployments. This in turn would make such a model more valuable to adopt, allowing users to have higher confidence in e-commerce overall. However, our work should generally provide a contribution to any researcher concerned with using social networks; we contend that the question of how large network should be is an important consideration.

Moreover, we are mindful of the impact this work could have with trust modelling more generally, as well as the broader area of multiagent systems. For example, this model could be adapted to other applications, including information sharing among peers for home healthcare tasks, where the ability to trust others, and indeed the opinions of others with regards to one's health and safety, is paramount.

2 Progress to Date

To date we have identified three potential methods which may allow us to optimize the advisor network size. The first two methods are fairly similar in that they reduce the network to some proportion of the top advisors for that buyer; the third may be used in combination with one of the aforementioned methods to optimize the amount of information that can be derived from a restricted-size network.

2.1 Trustworthiness Thresholding

In our first method, we first note that for each advisor a, a trustworthiness value $Tr(a)$ is calculated by each buyer b, with the said value falling in the range $(0, 1)$. We then define some threshold L $(0 \leq L \leq 1)$ representing the minimum value of $Tr(a)$ at which we will allow an agent to be included in b's advisor network. A buyer b will then only make use of those advisors where $Tr(a) \geq L$ in determining the public trustworthiness of any sellers of interest.

2.2 Maximum Number of Advisors

In the second method, we set a maximum number of advisors $max\_nbors \leq n$, where n is the total number of advisors in the system, for the advisor network of each buyer. More specifically we sort n advisors according to their trustworthiness value $Tr(a)$, in order from greatest to least, then truncate this set to the first max_nbors advisors. Similar to the previous method, the buyer b will then only make use of these max_nbors advisors in his or her public trustworthiness calculations for sellers.

2.3 Advisor Referrals

Our final method is based on the advisor referral scheme described in [3], and may be used in combination with either of the other methods described above. We posit that by allowing a buyer to indirectly access other advisors with pertinent information outside its advisor network, we can further optimize the network size. However, in doing so we must continue limit the number of advisors accessed through such referrals, in order to ensure some degree of computational efficiency.

For each advisor a_j in the advisor network of a buyer b, that is, the set $A_b = \{a_1, a_2, \ldots, a_k\}$, b checks whether 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 that may be calculated using formulae provided in Zhang's model.

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 may 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. The recursion may be repeated up to $\lceil log_k(|B|) \rceil$ network "levels", where B is the set of all buyers (advisors) in the system. Additional work is needed in order to determine whether setting a lower value for this maximum number of levels may be necessary for computational efficiency.

This would of course be repeated for each advisor until a full set of k advisors that have each had at least N_{min} interactions with s is found, or a smaller set if the recursion limit has been exceeded on one or more occasions.

3 Plan for Remaining Work

We have outlined above three possible modifications that may allow us to optimize the advisor network size, which may be tested in four possible combinations (thresholding with or without advisor referrals, and likewise for maximum network size). We plan to research some potential extensions or alternatives to these schemes in the next few months.

In the spring of 2010, we will then proceed to begin our test work. Initially, we anticipate testing the first two methods, by themselves, using a large set of data in order to determine empirically any difference in performance and accuracy compared with Zhang's original model. As in [1], this would likely take the form of a simulated marketplace incorporating a variety of buyers with varying levels of activity and honesty; we would also need to use different values of the threshold L and the maximum number of advisors n in order to determine an optimal value for our system and hence the maximum increase in accuracy, in any, over Zhang's model.

Finally, if at least one of these methods provides increased performance, we may then look at the advisor referral mechanism and determine whether this may provide even greater improvements. Again, this would necessitate testing the system against various combinations of parameters in order to determine an optimal algorithm.

At each stage of this evaluation we would need to consider both the method's accuracy in distinguishing between honest and dishonest buyers using a measure such as the Matthews correlation coefficient [4], as well as a measure of the computational performance of that method.

Our analysis will continue during the fall of 2010, possibly leading to changes and extensions to our model. We anticipate that our work will be completed either by the end of the year or in early 2011.

Certainly we have encountered a number of other ideas which may be able to help to improve Zhang's model, including the application of this model to time-sensitive tasks which may require a buyer to make a very quick decision; in such a scenario, the buyer will only have time to consult a very limited number of advisors before this decision must be made. Such concepts certainly merit further examination, although it is not clear at this point whether we will be able to examine these possibilities as part of our current research.

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. 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
4. 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