Towards Context-Enriched Trust Prediction: A Proposal

Marcin Sydow

Polish-Japanese Institute of Information Technology,
Web Mining Lab,
Koszykowa 86, 02-008, Warszawa, Poland**
msyd@pjwstk.edu.pl

Abstract. This short paper describes an early stage of research aiming at improving trust prediction in social networks.
It builds on recent findings on observed correlation between trust and user similarity. A machine-learning approach to the trust prediction problem with use of contextual information is proposed and experimental work envisaged.
An overview of the existing results is presented and it indicates that the proposed approach constitutes a novel combination of ideas and, as such, has a potential to contribute to the previous research in the area.

Key words: trust metrics, context, experimentation, social networks

1 Introduction

There is an explosion of interest in on-line communities where users' personal preferences can be explicitly expressed. Trust management plays an important role in such systems since on-line users base their decisions on information expressed by other users. However, the amount of information contained in social networks to be analysed by the users grows exponentially, so that there is a strong need for developing and improving automatic techniques that support users in making their on-line activities.

An example of such automatic support is a recommender system like `epinions`[1] or `FilmTrust`[2]. The system, given the preferences of the users *and* trust values expressed among them, is capable of automatically recommending some items to particular users. In this way, the information about trust among the users significantly improves the decision-support process.

The mechanism mentioned above can work quite well given that the trust values are expressed by the users. However, since on-line communities grow constantly, significant part of the users are newcomers to the system so that they did not specify their trust to other users, yet.

** The work is supported by the Polish Ministry of Science grant N N516 4307 33.

[1] `http://www.epninions.com`

[2] `http://trust.mindswap.org/FilmTrust/`

Due to this problem, we study the issue which is somehow reverse to the recommendation problem mentioned earlier. Namely, we aim at studying the problem of how to better predict an unknown trust between a pair of users given some additional information available in the system which we call the *contextual* information.

The term *contextual information* in social networks can be interpreted in many ways. In the most general meaning, the context of trust in a social network can be understood as all the information available in the social network except the trust information itself.

In our research, however, we plan to initially narrow the definition of the contextual information to the recommendations (ratings) given by the users. Thus, our notion of the context-aware trust might differ from those considered elsewhere, e.g. in [12] or [10]. This limited understanding of the trust context definition is partially due to the fact that quite often only the recommendation data, except the trust data itself, is available in publicly accessible real datasets concerning trust-related research. In particular, we plan to build *user profiles* based on their recommendations, and subsequently treat such profiles as a special kind of contextual information in the task of trust prediction.

An example of a dataset, with both: trust and recommendation data available, is the epinion dataset [1]. We plan to start the experimentation with this particular dataset.

In future, if the results obtained for such a limited definition of the context of trust are encouraging and we have access to datasets with richer information, we plan to treat the context of trust in a broader meaning.

2 Motivation and Related Work

A recent study by Ziegler et al. [15] experimentally proves, that there exists a significant correlation between the trust expressed by the users and their similarity based on the recommendations they made in the system.

Our proposal of research builds on that finding and considers going a step further. Namely, using the recommendations made by the users as the contextual information we aim at *predicting* potentially unknown trust between users. The main difference is that in the cited paper merely the *correlation* between the trust and the recommendation-based context is measured, while in contrast, we plan to measure how much the context helps in a full machine-learning approach to predicting the unknown trust value.

More precisely, our research aims at experimentally comparing the predictive value of the three groups of attributes in the machine-learning task of the trust prediction in social networks. The three groups of attributes are:

1. the attributes based solely on the pure topology of the trust network
2. the attributes based on the contextual information. As was explained, we plan to begin with taking the recommendation-based information to build such attributes
3. the combined trust-based and contextual-based attributes

One of the main goals of the proposed research is to experimentally check our hypothesis which is as follows: *a trust-predictor build on the third, contextually-enriched group of attributes, significantly outperforms the predictors build on pure-trust and pure-contextual attributes, separately, which seems to be suggested by the results reported in [15]*

A good general introduction to the trust management issues is given in [11]. Trust and distrust propagation models are proposed in [4] and further enhanced in [13] in the special context of Web spam combating.

The problem of pure link-based classification and appropriate link-based features are discussed by Liben-Nowell et al. [7] which proposes some measures for analysing proximity of nodes in a social *undirected* network representing co-authorship. Karamon et al. continues this research [5], however in this work the problem of *node classification* rather than *link classification* is discussed.

Correlation between trust and user profile similarity is discussed in [15] and further studied – by means of survey-based experiments – in [2]. An algorithm based on Bayesian network for trust inference is discussed in [6].

In the time of writing, the author discovered an unpublished draft paper concerning the problem of trust-based rating prediction and, inversely, similarity-based trust prediction by Matsuo et al. [9]. However, the current version of the draft does not seem to be completed (at the time of writing), for example, some important formulae seem to need to be syntactically clarified. However, since there seems to be a remarkable overlap with what is proposed in our research it seems necessary to address the issue how they relate to each other. In particular, in its current form, the mentioned draft:

- proposes general topology-based characteristics for trust computation, which, for example, never propagate further than through two links. In contrast, this paper proposes features that are especially designed to trust (and distrust) propagation measurement and are not constrained in such a way
- unfortunately does not address the problem of the extreme rating-data sparsity. In practice, in most of the available datasets the ratings data is too sparse to be directly processed. Our approach is aware of this issue and proposes how to try to overcome this problem
- does not seem to take into account the distrust concept

It would be very interesting, though, to compare the final results of the mentioned paper, and the research envisaged here, when they are ready.

The `epinions dataset` [1] was studied by Massa et al. [8] where also the interesting notion of "controversial users" was introduced in the context of the adversarial behaviour of some users.

The context of trust is extended by the user revision history in [14] and by the user's provenance in [3].

3 Proposed Experiments

As motivated at the beginning of the section 2, building on some recently published results, as explained above, we propose to experimentally test how much contextual information helps in trust prediction.

We propose the following experimental scheme:

1. Data: take a real dataset concerning trust network enhanced with some kind of contextual data such as ratings, user profiles, etc. At the beginning we plan to experiment with the publicly available `epinions dataset` [1], which contains both trust network and user recommendations, the latter one will be used as the contextual information
2. Sampling: from the dataset, take a random sample of user pairs for subsequent trust-prediction task. Due to the sizes of available real datasets, the epinions dataset in particular, sampling seems to be necessary, since repeating the experiment on *all* pairs of users in the dataset[3] seems to be computationally too expensive.
3. Training-testing split: as a common technique in supervised learning, to avoid evaluating the trust-prediction results on the same examples on which the predictor was trained (which is referred to as the *over-fitting* problem) split the sample into the *training* and *testing* subsets. We plan to tune the training/testing size ratio experimentally. The splitting should be *stratified* i.e. the proportion of the 'trusting' and 'distrusting' user pairs in sub-samples should be close to that observed in the whole sample.
4. Pure trust-topology prediction baseline: predict the value of trust (or distrust) between the nodes in the sample, based only on the trust-network information (*without* taking into account the actual known trust value on (u, v)). We propose to apply special trust-propagation techniques [4, 13].
5. Pure contextual-based prediction baseline: Compute contextual features, based on the additional information available in the dataset, such as user recommendations, similar to those described in [15, 2] and predict the trust using only these features
6. Contextually-enriched predictor: combine the trust-based and context-based features and finally train the combined trust predictor on the same sample as above.
7. Evaluation: compare the performance of the three prediction schemes above with use of some standard prediction-performance measures, such as prediction accuracy, the F-measure, etc.

The details about how the planned experiments relate to the previous research in the area and in which way the proposed experiments are novel are described in the following section.

[3] which is $O(n^2)$ for the number of users being n

4 Planned Contributions

The experimental work proposed in this paper is expected to contribute to the existing related research (see Section 2) in the following ways:

- supervised machine-learning approach, to the trust-prediction problem instead of a simple linear correlation measurement [15]. It seems that at least 2 different algorithms should be tested to have more reliable results, e.g. the Weka[4] implementation of the C4.5 decision tree and SVM.
- introduce some features based on especially designed trust-propagation models [4, 13]. Previously, only some general link-topology characteristics [7, 5] were used in this context.
- due to the extreme sparsity of the contextual data, when computing the ratings-based similarity between the users, we propose to first *cluster* the rated items with a method which is aware of the bi-partite structure of the rating graph. Independently, some dimension-reduction techniques such as SVD can be subsequently considered, if needed, to make the user profiles more overlapping. If a *hierarchical categorisation* of the rated items is available in the datasets, it should be used as well (as in [15]). Similar techniques concerning clustering the users could be considered. Due to this, the experiments can take into account also the users with little (or no) overlap in the rating-based profiles, in contrast to the previously reported approaches.
- in case such data is available, introduce a *distrust* prediction, which is a mathematically different task to trust prediction (e.g. distrust does not seem to allow for simple transitive propagation, in contrast to trust)
- when using rating data, take into account the true meaning of the values. Some previous approaches seem to not distinguish between 'positive' and 'negative' interpretation of the numerical values assigned to the ratings (and apply e.g. a dot-product similarity measure which is wrong in this context)
- on later stages of the research, if the textual ratings are available in the dataset, consider textual similarity of the textual reviews to compute more user similarity based features

To the best knowledge of the author, the combination of the propositions listed above is novel and can make the results of the proposed experimental research a potentially valuable contribution to the current research in the area.

5 Conclusions

Based on recently reported experimental results concerning observed significant correlation between user similarity and trust in social networks, a plan of further exploration of this direction is proposed here.

We propose a machine-learning approach to build a trust predictor and experimentally measure how much the contextual information can improve the

[4] `http://www.cs.waikato.ac.nz/ml/weka`

trust prediction accuracy compared to some trust prediction models which are based solely on pure topology of trust networks. We plan to start experimentation on the epinions dataset [1] and initially limit the contextual information to user recommendations, since such data is currently at our disposal.

References

1. `http://www.trustlet.org`.
2. Jennifer Golbeck. Trust and nuanced profile similarity in online social networks. Technical report.
3. Jennifer Golbeck. Combining provenance with trust in social networks for semantic web content filtering. In *IPAW*, pages 101–108, 2006.
4. R. Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412, New York, NY, USA, 2004. ACM.
5. J. Karamon, Y. Matsuo, H. Yamamoto, and M. Ishizuka. Generating social network features for link-based classification. In *PKDD*, pages 127–139, 2007.
6. U. Kuter and J. Golbeck. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *AAAI*, pages 1377–1382, 2007.
7. David Liben-Nowell and Jon Kleinberg. The link prediction problem for social networks. In *Proceedings of CIKM '03*, pages 556–559, NY, USA, 2003. ACM.
8. P. Massa and P. Avesani. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *AAAI*, pages 121–126, 2005.
9. Yutaka Matsuo and Hikaru Yamamoto. Measuring bidirectional effects on trust and rating on online social networks. Unpublished.
10. R. Neisse, M. Wegdam, P. Dockhorn Costa, and M. J. van Sinderen. Context-aware management domains. In B. Hulsebosch, G. Lenzini, and M. Wegdam, editors, *Proceedings of Context Awareness and Trust 2007 (CAT07), First International Workshop on Combining Context with Trust, Security and Privacy, Moncton, Canada*, volume 269 of *CEUR Workshop Proceedings*, pages 42–47, Nederlands, July 2007.
11. Sini Ruohomaa and Lea Kutvonen. Trust management survey. In *Proceedings of the iTrust 3rd International Conference on Trust Management, 23–26, May, 2005, Rocquencourt, France*, pages 77–92. Springer-Verlag, LNCS 3477/2005, May 2005.
12. Sini Ruohomaa and Lea Kutvonen. Making multi-dimensional trust decisions on inter-enterprise collaborations. In *Proceedings of the Third International Conference on Availability, Security and Reliability (ARES 2008)*, pages 873–880, Barcelona, Spain, March 2008. IEEE Computer Society.
13. Baoning Wu, Vinay Goel, and Brian D. Davison. Propagating trust and distrust to demote web spam. In *Workshop on Models of Trust for the Web*, Edinburgh, Scotland, May 2006.
14. Honglei Zeng, Maher Alhossaini, Li Ding, Richard Fikes, and Deborah L. McGuinness. Computing Trust from Revision History. In *Proceedings of the 2006 International Conference on Privacy, Security and Trust*, October 2006.
15. Cai-Nicolas Ziegler and Jennifer Golbeck. Investigating interactions of trust and interest similarity. *Decis. Support Syst.*, 43(2):460–475, 2007.