Whom Should I Trust? The Impact of Key Figures on Cold Start Recommendations

Patricia Victor,
Chris Cornelis
Appl Math & CS
UGent, Krijgslaan 281 (S9)
9000 Gent, Belgium
Patricia.Victor@UGent.be
Chris.Cornelis@UGent.be

Ankur M. Teredesai
Institute of Technology
UW Tacoma, 1900 Pacific Ave
Tacoma, WA, USA
ankurt@u.washington.edu

Martine De Cock
Appl Math & CS
UGent, Krijgslaan 281 (S9)
9000 Gent, Belgium
Martine.DeCock@UGent.be

ABSTRACT

Generating adequate recommendations for newcomers is a hard problem for a recommender system (RS) due to lack of detailed user profiles and social preference data. Empirical evidence suggests that the incorporation of a trust network among the users of the RS can leverage such 'cold start' (CS) recommendations. Hence, new users should be encouraged to connect to the network as soon as possible. But whom should new users connect to? Given the impact this choice has on the delivered recommendations, it is critical to guide newcomers through this early stage connection process. In this paper, we identify key figures in the trust network (in particular mavens, connectors and frequent raters) and investigate their influence on the coverage and accuracy of a collaborative filtering RS. Using a dataset from Epinions.com, we demonstrate that the generated recommendations for new user are more beneficial if they connect to an identified key figure compared to a random user.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence; H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval

General Terms

Algorithms, Human Factors

Keywords

trust network, recommender system, cold start problem

1. INTRODUCTION

Systems that guide users through the vast amounts of online information are gaining tremendous importance. Among such applications are recommender systems [10], which, given some information about their users' profiles and relationships, suggest items that might be of interest to them. One of the most widely used recommendation techniques is collaborative filtering (CF) [9], which typically works by identifying users whose tastes are similar to those of the particular user and recommending items that they have liked. However, CF RSs still pose important challenges, one of their main weaknesses being the *cold start problem*: new users have not rated a significant number of items, and cannot properly be linked with similar users[1]; hence, accurate and adequately personalized recommendations are difficult to generate.

It is only recently that the CS problem has received attention from the RS community (see e.g. [6, 8]). One of the promising directions suggests that the incorporation of a *trust network* (in which the agents are connected by trust scores indicating how much they trust and/or distrust another agent) can significantly help alleviate the CS problem, primarily because the information included in trust statements about a RS's user can be propagated and aggregated, and hence more people and products can be matched [6, 13].

Since the trust information in such a *trust-enhanced RS* has a significant direct influence on the delivered recommendations (both amount and quality), users must be encouraged to connect to other users in the trust network as soon as possible (see e.g. [2, 6]). But, this is not so obvious for a CS user as he often does not know which users will have the best impact on the generated recommendations. Therefore, we investigate techniques to identify *key figures* in the trust-enhanced RS network, and evaluate the benefits for a new user to connect to one or more key figures. More specifically, we will analyze their influence on the number and accuracy of the generated recommendations. To evaluate the techniques we propose in this paper, we use a large dataset from Epinions[2], a prominent e-commerce site that gives users the opportunity to include other users (based on their quality as reviewers of all kinds of consumer goods) in their own 'web of trust' (WOT). The results can be generalised to other trust-based RSs.

In Sect. 2, we describe classical CF RSs and explain how trust-enhanced RSs can help alleviate the CS problem. To benefit from these trust algorithms, a new user needs to

[1]Note that the phrase *cold start* has also been used to describe the situation where recommendations are required for items that have never been rated (see e.g. [11]).

[2]`www.epinions.com`

SAC'08 March 16-20, 2008, Fortaleza, Ceará, Brazil

know which users are best to connect to. In this paper, we identify different user classes in the RS network as mavens (knowledgeable users who write a lot of reviews), connectors (with a lot of connections in the trust network) and frequent raters (who rate a lot of reviews). The characteristics of the new users and these key figures in the dataset are analyzed in Sect. 3. In Sect. 4, we show that it is more beneficial for new users to connect to such key figures rather than making random connections. We conclude the paper with a discussion of future research directions.

2. RELATED WORK

A CF algorithm [9] is used to arrive at a rating for an item i that is new to a user a. This new rating is based on a combination of the ratings of the nearest neighbours (similar users) already familiar with item i. The classical CF-formula is given by (1). The unknown rating $p_{a,i}$ for an item i for a user a is predicted based on the mean $\overline{r}_a$ of ratings by a for other items, as well as on the ratings $r_{u,i}$ by other users u for i. The formula takes into account the similarity $w_{a,u}$ between users a and u, usually calculated as Pearson's Correlation Coefficient (PCC) [4].

$$p_{a,i} = \overline{r}_a + \frac{\sum_{u=1}^{k} w_{a,u}(r_{u,i} - \overline{r}_u)}{\sum_{u=1}^{k} w_{a,u}} \qquad (1)$$

The effectiveness of CF based RSs is significantly affected by the number of ratings available for each user: the more ratings are available, the better the quality of the recommendations. Moreover, generating recommendations is only possible for users who have rated at least two items because the PCC requires at least two ratings per user. Since CS users, being new users, have rarely rated a significant number of items, and they usually constitute a sizeable portion of the RS's user community (see e.g. [7]), it is very important to address this problem.

Trust-enhanced RSs can alleviate the CS problem by using additional information coming from a trust network in which the users are connected by trust scores indicating how much they trust and/or distrust each other. A simple version of a trust-based RS formula is a special case of (1) in which the weights $w_{a,u}$ are replaced by trust information $t_{a,u}$ [6]. In this approach, trust is interpreted as a numerical value which ranges between 0 and 1 (absence/full presence of trust).

The main strength of trust-enhanced RSs is their use of *trust propagation operators*; mechanisms to estimate the trust transitively by computing how much trust an agent a has in another agent c, given the value of trust for a trusted third party (TTP) b by a and c by b. By combining this information with the available ratings, more users and (consequently) more items get covered by the RS, even if only few trust statements per user are available [6].

For a user a, we define the *rating coverage*, or coverage for short, as the ratio of the amount of items for which $p_{a,i}$ as in (1) can be calculated versus the total amount of items available in the RS. A prediction $p_{a,i}$ can be calculated when a trusts at least one user u to the degree $t_{a,u} \neq 0$ and u already rated i. [6] demonstrates that the coverage for CS users increases significantly when they connect to the trust network.

For any RS algorithm, coverage and accuracy results must be evaluated together; an increase in coverage is only beneficial when the accuracy does not drop significantly, while an accuracy increase is not useful when there are too few ratings that can be predicted. As shown in e.g. [2, 6], this trade-off turns out to be advantageous for trust-enhanced RSs, and especially for CS users. Golbeck [2] and Massa et al. [6] report that using only the information coming from trusted acquaintances, and from users who are trusted by trusted people in turn, makes the recommendations significantly more accurate and also more personalized. Hence, it is beneficial for a new user to connect to the trust network as soon as possible. But, as will be demonstrated in the following section, it is often the case that CS users in the classical sense (people who provided only a few product ratings) are also CS users in the trust sense, meaning that they issued only a few, or no trust statements at all. Therefore, we propose to guide the new users during the connection process by suggesting to connect to key figures who have a positive impact on the coverage while maintaining sufficient accuracy.

	CS1	CS2	CS3	CS4
% of review raters	36.52	12.32	6.85	4.47
% in LC	18.43	30.85	38.34	44.88
mean # trust rel	0.27	0.51	0.72	0.99
mean # distrust rel	0.03	0.05	0.06	0.09

Table 1: CS users in the dataset

3. CS USERS AND KEY FIGURES

Epinions.com is a popular e-commerce site where users can write reviews about products and assign a rating to them. The dataset from Guha et al. [3] contains 1 560 144 reviews (written by 326 983 users) that received 25 346 057 ratings by 163 634 different users. Reviews are evaluated by assigning a helpfulness rating which ranges from 'not helpful' (1/5) to 'most helpful' (5/5). Note that we do not have information about consumer products and product ratings, but work with reviews and review ratings instead; in other words, we evaluate a 'review recommender system'.

We focus on users who have evaluated at least one review. In this group, 59 767 users rated only one review, 20 159 only two, 11 216 exactly three and 7 322 exactly four. These cold start users constitute a large part of the Epinions community, namely about 60% of all review raters. The relative numbers of users are given in Table 1 where these CS users are denoted by CS1 (exactly one review), CS2 (two reviews), CS3 and CS4.

Besides evaluating reviews, users can also evaluate other users based on their quality as a reviewer. This can be done by including them in their WOT (i.e. a list of reviewers whose reviews and ratings were consistently found to be valuable[3]) or by putting them in their block list (i.e. a list of authors whose reviews were consistently found to be offensive, inaccurate or low quality[3], thus indicating distrust). These evaluations make up the Epinions WOT graph consisting of 131 829 users and 840 799 non self-referring trust or distrust relations (see also [3]). About 85% of the statements are labelled as trust, which is reflected in the average number of users in a WOT (5.44) and in a block list (0.94). Due to the large portion of trust statements, we will focus on trust information only in the remainder of the paper.

[3]See www.epinions.com/help/faq/, accessed on Sep 6 07

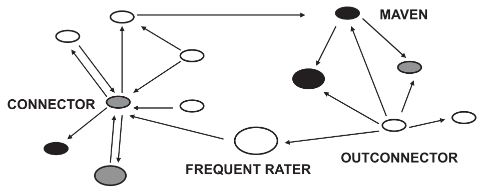

Figure 1: Key figures: example

The trust graph consists of 5 866 connected components (i.e. maximal undirected connected subgraphs). The largest component (LC) contains 100 751 users, while the size of the second largest component is only 31. Hence, in order to receive more trust-enhanced recommendations, users should connect to the largest component. But as shown in Table 1, this cluster does not even contain half of the cold start users. This, combined with the fact that cold start users evaluate only a few users (as shown in the third and fourth row of Table 1), illustrates that CS users in the classical sense are very often CS users in the trust sense as well. Better results can be expected when newcomers connect to a large component of the trust graph, but they may encounter difficulties in finding the most suitable people to connect to. Therefore, we define three user classes and locate them in the network.

The first class of key figures are *mavens*, people who write a lot of reviews. This term is borrowed from Gladwell's book [1] in which they are defined as knowledgeable people who want to share their wisdom with others. Out of the three user classes mavens are the most visible, and hence the ones which are the easiest to evaluate: the more reviews someone writes, the better a new user can form an opinion on him and put him in his personal WOT or not.

Unlike mavens, *frequent raters* are not always so visible. They do not necessarily write a lot of reviews but evaluate a lot of them, and hence are an important supplier for the recommender system: it is not possible to generate predictions without ratings. By including a frequent rater in a trust network, more reviews can be reached, which has a direct influence on the coverage of the system.

Mavens and frequent raters are not necessarily bound to the trust network, but *connectors* are: they connect a lot of users and occupy central positions in the trust network. Such users issue a lot of trust statements (many outlinks) and are often at the receiving end as well (many inlinks). Their strength lies not in their rating capacity or visibility, but in their ability to reach a large group of users through trust propagation. When a trust-based algorithm has to find a path from one user to another, a connector will be part of the propagation chain more often than a random user, and propagation chains containing connectors will on average be shorter than other chains. Shorter chains have a positive influence on the accuracy of the trust estimations and recommendations [2].

Fig. 1 shows a diagram with examples of each type: the darker the node, the more reviews the user wrote (maven). The larger the node, the more reviews the user evaluated (frequent rater). The trust network is denoted by the arrows representing trust relations; connectors are characterized by many incoming and outgoing arrows. In the Epinions dataset, we define a maven as someone who has written at least 100 reviews (M-100+), a frequent rater as someone who evaluated at least 2500 reviews (F-2500+), and a connector as someone who has an in+out degree of at least 175 (C-175+). With these definitions, the community contains 1925 mavens, 1891 frequent raters and 1813 connectors. The tresholds are chosen such that the different sets have similar sizes; in Sect. 4 we experiment with other tresholds as well.

The sets of connectors and mavens share a large number of users, which is not surprising because mavens are visible through the reviews they write, making it more likely for others to connect to them with trust statements. The conditional probability P(M-100+|C-175+)≈0.52. More surprising is the relation between connectors and frequent raters, namely P(F-2500+|C-175+)≈0.64. The intersection of the maven set and the frequent rater set also contains many users (933), so there clearly is a strong overlap between the different groups of key figures. This indicates that users who are active on one front are often active on other fronts as well. These findings may be influenced by Epinions' 'Income Share program' and the benefits of being selected as a category lead, top reviewer or advisor[3]. Some of these classes are related to the key figures we defined, though our approach for identifying key figures only relies on objective data, while the selection in the Income Share program is partially subjective. Note that Epinions' interface also has an impact on the visibility and relatedness of the user classes.

Although the latter characteristics may be influenced by the specific situation, the three user classes can be detected in all kinds of trust-based RSs, and hence the results in the following section can easily be generalised.

4. INFLUENCE OF KEY FIGURES

In this section, we investigate the influence of key figures on the coverage and accuracy of CS recommendations.

4.1 Potential of key figures

First we only consider CS users who have exactly one key figure in their WOT, e.g. the set of CS2 users who are connected with exactly one maven of type M-1000+. Let us denote such a set as U. We represent a user of U by a and the corresponding key figure by k_a. $Rat_1(k_a)$ denotes the set of reviews rated by k_a, and $Acc_1(a, X)$ the set of accessible reviews that are rated by the users from a's WOT-list that do not belong to X, i.e. $Acc_1(a, X) = \bigcup \{Rat_1(u) | u \in WOT(a) \setminus X\}$. Note that $|Acc_1(a, \emptyset)|$ is the number of reviews for which a rating can be predicted with the trust based variant of (1). The percentage of these ratings accessible only through key figure k_a is

$$KCP_1(a) = 100 \cdot \frac{|Rat_1(k_a) \setminus Acc_1(a, \{k_a\})|}{|Acc_1(a, \emptyset)|}\% \qquad (2)$$

We call this the *key figure coverage portion*. This measure cannot be used in isolation from the size of the WOT-list. Assume e.g. that $WOT(c) = \{k_c, u_1, u_2, u_3\}$, $WOT(b) = \{k_b, u\}$, and that $KCP_1(b) = KCP_1(c) = 50\%$. Despite the equal KCP values, k_c is clearly the strongest key figure because k_c achieved 50% in the presence of 3 competitors in the WOT-list, while k_b has only one. Therefore, we compare $KCP_1(a)$ with the *expected coverage portion* $ECP_1(a)$ for a user a. The formula for $ECP_1(a)$ is given by

$$ECP_1(a) = \frac{100}{|WOT(a)|}\% \qquad (3)$$

The *extra coverage benefit* $ECB_1(U)$ for set U, provided by a specific key figure type, then becomes

$$ECB_1(U) = \frac{1}{|U|}\sum_{a\in U}(KCP_1(a) - ECP_1(a)) \qquad (4)$$

Note that we do not take into account users who have only one WOT member (i.e. the key figure), because in that case $ECP_1(a)$ is meaningless. In the remainder of the paper we write ECB_1 and KCP_1 instead of $ECB_1(U)$ and $KCP_1(a)$.

To get a more complete picture of the performance of the trust-enhanced RS, we also investigate the effect of key figures on the accuracy by calculating the *accuracy change* AC_1 for each CS group. AC_1 is the average $AC_1(a)$, which is obtained by subtracting the mean absolute error (MAE) obtained after excluding the ratings provided by the key figure, from the MAE when taking into account all available ratings. A positive AC_1 denotes higher prediction errors when taking into account the ratings provided by the key figure. To calculate the MAE's we use the leave one out method. Since reviews are rated on a scale from 1 to 5, the extreme values that AC_1 can reach are -4 and 4.

The results for the first experiment can be found in Table 2. A column/row corresponds to a specific user group/key figure, e.g., a M-100 is a maven who wrote at least 100 and maximum 499 reviews. Including a key figure clearly is very advantageous (an ECB_1 of 29% on average). Frequent raters appear to be the best asset for a CS user's WOT, with an average ECB_1 of 34% and over 40% for F-50000+. This is not surprising because they are the real suppliers of the RS. However, it may be difficult for newcomers to evaluate them because they are not always very visible. Since this is the main characteristic of mavens, it might be easier for a cold start user to issue a trust statement about such users. With an average ECB_1 of 26%, they score well too.

Similar conclusions apply to connectors (26%). We claimed that including a connector in a WOT yields shorter propagation chains because they connect more users and reach more reviews. Therefore, besides the above experiment (*level 1*, L1), we also measured the coverage by using information obtained by propagating trust one step (*level 2*, L2). This means that if a trusts b and b trusts u, $t_{a,u}$ in the trust-based variant of (1) equals 1. To only take into account reviews that cannot yet be reached on L1, we define $Acc_2(a, X) = \bigcup\{Rat_2(u)|u \in WOT(a) \setminus X\} \setminus Acc_1(a, X)$ with $Rat_2(u) = Acc_1(u, \emptyset) \setminus Rat_1(u)$. KCP and ECB then become

$$KCP_2(a) = 100 \cdot \frac{|Rat_2(a_k) \setminus Acc_2(a, \{a_k\})|}{|Acc_2(a, \emptyset)|}\% \qquad (5)$$

$$ECB_2(U) = \frac{1}{|U|}\sum_{a\in U}(KCP_2(a) - ECP_1(a)) \qquad (6)$$

The coverage results for L2 can be found in Table 3. Note that these results are lower than their L1-counterparts, which indicates that some of the reviews reached through the connector are also reached through other WOT members. However, one must realize that the amount of new reviews that is provided by a C-1000+ via one step propagation is more than 10 times the amount delivered by a C-1000+ on the first level. For instance, for CS1 users the L1 set contains 63 139 reviews, while 731 859 new reviews are reached on L2. Hence, it is clear that trust propagation and connectors have a large positive impact on the coverage of the RS.

Table 2: Evaluation for frequent raters (F), mavens (M) and connectors (C) on L1, Experiment 1

	ECB_1				AC_1		
TYPE (#)	CS1	CS2	CS3	CS4	CS2	CS3	CS4
F-100000 (2)	48	46	47	50	-0.23	0.04	0.08
F-50000 (36)	32	38	38	34	-0.04	-0.09	0.05
F-10000 (459)	31	33	36	36	0.16	-0.02	0.00
F-2500 (1891)	19	20	17	20	-0.06	0.03	-0.03
M-1000 (11)	39	36	38	38	0.05	-0.14	-0.12
M-500 (77)	18	23	22	25	0.02	-0.01	0.04
M-100 (1837)	19	18	20	16	0.16	0.08	0.04
C-1000 (47)	37	36	38	38	0.01	0.04	0.02
C-500 (253)	18	23	22	25	0.06	-0.05	0.05
C-175 (1813)	19	18	20	16	0.01	0.03	-0.04

Table 3: Evaluation for connectors on L2

	ECB_2				AC_2		
TYPE (#)	CS1	CS2	CS3	CS4	CS2	CS3	CS4
C-1000 (47)	9	14	9	13	0.07	0.03	0.05
C-500 (253)	6	8	8	11	-0.01	0.00	0.02
C-175 (1813)	11	9	13	8	-0.01	0.00	-0.04

Let us now return to Table 2. The more active the key figure is, the more advantageous it is to have such a user in a WOT. Users who are connected with F-50000+ have a larger KCP_1 than users connected with F-2500, but this relationship also occurs with mavens and connectors. This confirms that users who are active on one front (being a maven or connector) are often active on other fronts as well (being a frequent rater boasting the number of accessible reviews).

An increase in coverage can only be beneficial when the accuracy does not drop significantly. AC_1 is shown in the right part of Table 2, which demonstrates that the absence or presence of a key figure in a WOT does not significantly change the accuracy. In other words, the key figures have a positive effect on the coverage, while maintaining sufficient accuracy. Note that no results are generated for the CS1 group: (1) uses the mean of a user's ratings, but the leave one out method already hides the sole rating of a CS1 user.

While AC_1 measures the accuracy change for reviews that are immediately accessible through users of a WOT-list (L1), AC_2 measures the accuracy change for reviews that become accessible through trust propagation (L2). AC_2 is obtained by subtracting the MAE of the predictions generated by information reached through TTPs other than the connector, from the MAE of the predictions based on all TTPs (including the connector). Table 3 shows that there are no significant changes in the accuracy. Combined with the fact that much more reviews can be reached on L2, this demonstrates the advantage of connectors for newcomers.

4.2 Benefit over random users

To take into account a larger group of users, we also conducted an experiment with the group of CS users who have no key figure in their WOT. For instance, 84.36% of the CS4 users have no F-2500 in their WOT, as opposed to 7.34% whose WOT contains exactly one. The goal of this experiment is to investigate the effect of adding a key figure to a CS user's WOT. To this aim, we calculate for each CS

Table 4: Evaluation, Experiment 2

TYPE	ECB_3 CS1	CS2	CS3	CS4	AC_3 CS2	CS3	CS4
F-100000	8	14	17	22	0.014	−0.008	−0.002
F-50000	8	14	17	21	0.009	0.010	−0.001
F-10000	8	14	17	21	0.006	0.002	0.001
F-2500	8	14	17	22	0.003	0.005	0.001
M-1000	8	13	16	20	0.005	0.010	0.007
M-500	7	12	15	18	0.004	0.006	0.006
M-100	6	10	12	15	0.001	0.005	0.01
C-1000	8	12	15	18	−0.009	0.000	−0.001
C-500	7	12	14	18	0.005	0.011	0.000
C-175	7	11	13	16	0.000	0.004	0.001

user the difference between the *extra user coverage portions* $EUC(a, B_1)$ and $EUC(a, B_2)$, in which B_1 represents a set of specific key figures and B_2 the set of all active users (users who evaluated at least one user or one review, hence this set contains key figures as well). $EUC(a, B)$ is defined as

$$EUC(a,B) = \frac{100 \cdot |Rat_1(b) \setminus Acc_1(a,\emptyset)|}{|Acc_1(a,\emptyset)| + |Rat_1(b) \setminus Acc_1(a,\emptyset)|}\% \quad (7)$$

in which b is a randomly selected user of B that is added to a's WOT-list. The extra coverage benefit, denoted by $ECB_3(U)$, is the average of $EUC(a, B_1) - EUC(a, B_2)$ over all a's.

AC_3 is calculated as the mean of $AC_3(a)$, which is the difference between the MAE obtained after adding the key figure and the MAE after taking into account the actual ratings only. Hence, positive accuracy changes denote higher prediction errors when taking into account the key figure.

The results for each CS group are shown in Table 4. The ECB_3 of all three types is about 14%; connecting to a key figure is clearly more beneficial than connecting to a random user. Note that the argument that more active key figures yield a higher ECB_3 applies here as well.

Because we compute the MAE by predicting existing ratings and CS users rate very few reviews, there is only a small chance that an added key figure will provide a rating for a review which is rated by the CS user but not by other members of his WOT. The small accuracy changes may therefore indicate that the extra ratings provided by the key figure do not significantly affect the predictions (that can already be generated by the ratings of actual WOT members).

5. CONCLUSION & FUTURE WORK

Each key figure has its own characteristics; mavens are easy to evaluate, frequent raters provide a lot of ratings, and connectors help to reach more users and reviews. We have shown that they have a large positive impact on the recommendations for CS users: the coverage increases significantly while the accuracy of the recommendations does not change significantly. We have also demonstrated that it is indeed beneficial to guide newcomers through the connection process; including a randomly chosen user is less advantageous than adding an identified key figure to a CS user's WOT.

Our future work goes in several directions. We want to investigate the potential of other key figures, e.g. by using social network analysis measures. Another research path is the incorporation of distrust information into the recommendation process. Distrust could e.g. be exploited to alleviate the sparsity problem: through specific propagation operators that can handle trust as well as distrust (see e.g. [3, 5, 12]), more users and items could be reached. Distrust can also be used to debug a WOT: suppose that a trusts b completely, b fully trusts c and a completely distrusts c. The latter information ensures that the propagated trust result (viz. a trusts c) is invalid and that a will not use information coming from c in the future. As such, trust and distrust-enhanced algorithms could be used to filter out false positives generated by other techniques such as CF. However, much ground remains to be covered in this domain.

6. ACKNOWLEDGMENTS

Patricia Victor and Chris Cornelis would like to thank the Institute for the Promotion of Innovation through Science and Technology in Flanders and the Research Foundation-Flanders for funding their research. We thank Epinions.com for making the data available, in particular R. Guha, R. Kumar, P. Raghavan and A. Tomkins.

7. REFERENCES

[1] M. Gladwell. *The Tipping Point: How Little Things Can Make a Big Difference*. Little Brown, 2000.

[2] J. Golbeck. *Computing and applying trust in web-based social networks*. PhD thesis, 2005.

[3] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proc. of WWW2004*, pages 403–412, 2004.

[4] J. Herlocker, J. Konstan, L. Terveen, and J. Riedl. Evaluating collaborative filtering recommender systems. *ACM T Inform Syst*, 22(1):5–53, 1999.

[5] A. Jøsang and S. Knapskog. A metric for trusted systems. In *Proc. of NIST-NCSC1998*, pages 16–29, 1998.

[6] P. Massa and P. Avesani. Trust-aware collaborative filtering for recommender systems. *LNCS*, 3290:492–508, 2004.

[7] P. Massa and B. Bhattacharjee. Using trust in recommender systems: an experimental analysis. *LNCS*, 2995:221–235, 2004.

[8] S.-T. Park, D. Pennock, O. Madani, N. Good, and D. DeCoste. Nave filterbots for robust cold-start recommendations. In *Proc. of SIGKDD2006*, pages 699–705, 2006.

[9] P. Resnick, N. Iacovou, M. Suchak, P. Bergstorm, and J. Riedl. Grouplens: An open architecture for collaborative filtering of netnews. In *Proc. of CSCW1994*, pages 175–186, 1994.

[10] P. Resnick and H. Varian. Recommender systems. *Commun ACM*, 40(3):56–58, 1997.

[11] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock. Methods and metrics for cold-start recommendations. In *Proc. of SIGIR2002*, pages 253–260, 2002.

[12] P. Victor, C. Cornelis, and M. De Cock. Enhanced recommendations through propagation of trust and distrust. In *Proc. of WI-IAT2006 Workshops*, pages 263–266, 2006.

[13] J. Weng, C. Miao, and A. Gohl. Improving collaborative filtering with trust-based metrics. In *Proc. of SAC2006*, pages 1860–1864, 2006.