To appear in Proc. Sixth Workshop on eBusiness (WeB2007)

Social Network Collaborative Filtering: Preliminary Results

Rong Zheng
New York University
44 West 4th Street
New York, USA
+1 2129980827
rzheng@stern.nyu.edu

Foster Provost
New York University
44 West 4th Street
New York, USA
+1 2129980806
fprovost@stern.nyu.edu

Anindya Ghose
New York University
44 West 4th Street, 8-181
New York, USA
+1 2129980807
aghose@stern.nyu.edu

Abstract

This paper reports on a preliminary empirical study comparing methods for collaborative filtering (CF) using explicit consumers' social networks. As user-generated social networks become increasingly important and visible in technology-mediated consumer interactions, we can begin to ask how the rich associated information can be used to improve inference. Theories from social psychology have long discussed that social relationships are likely to connect similar people. If the social similarity is in line with the recommendation task, the social network may provide a small, dense set of "recommenders" for CF. To our knowledge this is the first study to show effects of social-network information for estimating purchase behavior with CF. We examine a data set of consumers that contains a social network of consumer-selected friends, as well as their purchases from a large online retailer. We examine two ways to incorporate social-network information into CF: using proximity in the social network to modify the traditional CF, and using the social network to restrict the set of recommenders selected. The results show that social network proximity does not seem to improve recommendations. On the other hand, CF with social-network members selected as recommenders predicts purchases far better than CF with the recommenders not socially connected.

Keywords: Social Network, Collaborative Filtering, Recommender System, Data Mining

1. Introduction

Recommender systems generate product recommendations in order to reduce consumers' search costs in light of the increasing product variety on the Internet (Resnick and Varian 1997). Using past information about consumers and products, these systems identify promising future interactions between consumers and products, and present to users information about items they are most likely to be interested in. The goal of a commercial recommender system includes both increasing sales and increasing user satisfaction. In the context of explicit social network information, this paper examines *collaborative filtering* (CF), perhaps the most well-known recommendation technique, which has been used widely in e-commerce applications and especially in academic research (Adomavicius and Tuzhilin 2005). The underlying principle behind "user-based" CF is to find customers who purchased the same or similar items. Importantly, CF techniques are very expensive computationally, having to compare large numbers of users to each other—which is one reason why Amazon.com instead uses alternative recommendation techniques (Linden et al., 2003).[1] Another significant problem from which

[1] Consider that Amazon.com has 50 million or more active customers; a single "long-tail" category like books can have over a million items.

current CF methods suffer is data sparsity. Similarity is computed based on a few data points which lead to poor quality recommendations.

Explicitly represented, user-generated social networks are becoming increasingly important in technology-mediated consumer interaction. Consumers have flocked to social-networking sites, revealing various information about themselves, and connecting themselves to others. One important potential ramification of the increased availability and accessibility of information about consumers is that it may be useful to help users to find items that they would like to purchase. This has obvious implications for e-commerce firms, such as Amazon.com, who can increase sales via product recommendations. Additionally, users themselves may be able to find products better by examining the preferences implicit or explicit in the information revealed by prior purchasers of a product (e.g., via product reviews) or by their social-network neighbors.

The link between networks of individuals and recommender systems has been made previously. Weng et al. (2006) explicitly address the problems of sparsity and scalability, employing a notion of trust based on correlations between ratings. Similarly, Massa et al. (2004) quantified *trust* using data from Epinions.com—where specifically, trust is based on feedback on other people's ratings. This additional information was shown to be effective in addressing cold start problems of CF systems and reducing computational cost.

However, it was not until recently that some researchers started to evaluate systematically the potential of explicit social networks[2] for making product recommendations. Off-line social networks have been the subject of intense academic research for many decades, but studies have been limited by the difficulty of gathering either real social links or purchase information or both (Hill et al., 2006). Hill et al. (2006) examined a large data set including both a communication-based social network and specific information on a targeted marketing campaign. They showed that the social network could be used to great advantage to determine the consumers likely to purchase, even in the presence of sophisticated modeling using consumer-specific information. However, their setting did not include a variety of different products, and so was not amenable to traditional recommendation techniques. *FilmTrust,* a movie recommender system, (Golbeck, 2006) provided a platform for people to rate movies and at the same time make friends; as with the other studies, the results show that friend information can be added into CF system and make better recommendations.

This paper reports on an empirical study comparing CF techniques that incorporate information on consumers' social networks. We compare social-proximity-based CF with traditional CF. A major distinction of our work is that we examine data on purchases, rather than data on ratings. In many practical situations purchases are available, but not ratings. Also the data span many product categories, adding a significant challenge for CF. We examine how helpful the social network information is specifically for estimating a consumer's propensity to purchase a product. In addition, we propose a new way for CF to incorporate the consumers' social network: select the users from which to make recommendations to be the social-network neighbors, under the presumption that these individuals are more likely to provide useful information, and thereby avoid (tremendous) computational expense and data sparsity.

In this study, we investigate two possible components of social-network-based collaborative filtering (SNCF): using the social network to model the proximity between consumers, which is

[2] In addition to the forementioned work, other prior work (e.g., Perugini et al. 2004; Huang et al. 2004) has noted that some recommendation systems induce an indirect interaction network over consumers. For the sake of clarity and focus, in this paper we distinguish *social* networks based on actual friend or acquaintance links, from induced consumer relationship networks based on co-purchase or other non-social relationships.

further incorporated into CF, and restricting the universe from which recommenders are selected. The results show that social network proximity does not seem to improve recommendations. On the other hand, CF with social-network members selected as recommenders predicts purchases far better than CF with the recommenders not socially connected.

2. Social-Network Collaborative Filtering (SNCF)

2.1 Proximity-based SNCF

In the typical setting, collaborative filtering (CF) exploits the interaction data between consumers and products and predicts products a consumer will purchase. The input of the problem is an $M \times N$ *interaction matrix* $T = (t_{ij})$ associated with M consumers $C = \{c_1, c_2, \ldots, c_M\}$ and N products $P = \{p_1, p_2, \ldots, p_N\}$. We focus on transactional data (rather than rating data). That is, a_{ij} can take the value of either 0 or 1, with 1 representing an observed transaction between c_i and p_j and 0 the absence of transactions. *User-based* CF algorithms first construct a consumer similarity matrix $W = (w_{st})$, s, t =1, 2, …, M. The similarity score w_{st} is calculated based on the row vectors of A using a vector similarity function. A high similarity score w_{st} indicates that consumers s and t may have similar preferences since they have previously purchased many common products. $W \cdot A$ gives potential scores of the products for each consumer.

An explicit consumer social network can be represented as a graph with nodes being the consumers and links being the social relationships among them. For this study, we introduce a straightforward modification. Specifically, instead of using similarity between past purchasing behavior to find consumers with similar purchasing preferences, *proximity-based SNCF* uses the distance between consumers in the social network. We adopt the standard graph-theoretic definition of distance of nodes: the minimal number of edges that link the nodes. Therefore for *proximity-based* SNCF the first step of the similarity computation is to find the minimum number of edges between two nodes. The input is a graph which is represented by the *adjacency matrix* $G = (g_{st})$, s, t =1, 2, …, M; g_{st} can be 1 or 0 depending on whether there is an edge between consumer s and consumer t. The output is a distance matrix $D = (d_{st})$ s, t =1, 2, …, M. These distances can be computed by *Dijkstra's* algorithm, or when the social-network links are unweighted, simply via a *breadth-first search* (which is the case for this study). Then, under the assumption that social influence will decay exponentially as the social-network distance increases, the distance matrix is transformed to the *influence* matrix $I = (i_{st})$ s, t =1, 2, …, M via: $i_{st} = \exp(-d_{st})$. In direct analogy to CF, the scores for the potential recommendations are calculated by $I \cdot A$. Below we will create different "versions" of the algorithm by limiting the span of influence to $d \leq k$, for particular values of k. So, for example, we can look at the influence only of direct neighbors by setting k=1.

2.2 Selection-based SNCF

Social theory tells us that social relationships are likely to connect similar individuals (McPherson et al., 2001). Therefore, restricting the whole consumer base to a smaller subset may be beneficial to CF methods. Intuitively, problems of data sparsity and computation cost will be reduced if the subset of data generates a denser matrix G. The procedure for *selection-based SNCF* is very simple. Restrict the set of recommenders to those who are in the target consumer's social network, and apply standard (or one's favorite) CF.

3. Data and Experimental Setup

We generate recommendations for a subset of 1206 customers from Amazon who have chosen to reveal their purchases on Amazon's site. The total set of purchases includes all revealed purchases made by these consumers over the three-month period between May and July 2007. In sum, a total of 11,773 distinct items were bought by these 1206 consumers. About 50% of the purchased items are books; 40% are CDs and DVDs, and the remaining 10% include products from other categories such as electronics, apparel, etc. Importantly for this study, one-half of these 1206 customers are interconnected by the "Amazon Friends" social network.[3] The set of 603 who have at least one "friend" who also has revealed her purchases, will be the social network chosen as recommenders and as targets for SNCF in our study.

We divide the purchases by timestamp. The 20% most recent purchases for each consumer are held out for prediction; the 80% older purchases will be used to make recommendations. Each recommender system calculates a score for each potential user/product pair, resulting in a ranked list of recommendations. For this study, we use standard precision/recall analysis to evaluate the quality of this ranked list of recommendations, examining the possible tradeoffs between the accuracy of recommendations (precision) and coverage of actual purchases (recall). We adopted standard evaluation measurement developed in (Breese, 98):

Precision: $p = \frac{\text{Number of } Hits}{\text{Number of recommendations}}$ Recall: $R = \frac{\text{Number of } Hits}{\text{Number of purchases made in the testing set}}$

Hits is the number of recommendations that correspond to actual purchased products in the testing set. *Precision* and *recall* measure how relevant the recommendations are. We should keep in mind that, in contrast to evaluations where users rate the desirability of all products, for this "actual purchase" prediction we do not expect very high precision or recall; predicting actual purchases is a very hard problem (Huang et al., forthcoming).

4. Results

Our first experiment compares traditional CF with *proximity-based* SNCF. For our 1206 consumers, **Figure 1** plots precision/recall curves for *proximity-based*[4] SNCF (as described above, using the social network to calculate the similarity between users) and regular CF for making recommendations. Clearly, the *proximity-based* adjustments result in inferior recommendations. Perhaps this should not be surprising: traditional CF is designed specifically for making recommendations from consumers with similar tastes in products. The friend relationship is likely to comprise different factors, which could add noise to mislead CF.

We should note that, in an absolute sense, the recommendation accuracy here is remarkable. Recall that this estimation corresponds to whether the CF can "predict"[5] whether or not a user will purchase the product, and further whether that purchase will show up in these data. Even if the CF were to predict correctly that a consumer would like a product, there are many reasons why that would not lead to a recorded purchase: the consumer may already have it, the consumer may purchase it elsewhere, the consumer may simply not have purchased it yet, etc. The precision for the top-100 recommendations is around 20%, and for the top-1000 recommendations is still around 10%—with a recall of 5%. Previous studies of the accuracy of recommendations to predict actually purchases do not come close to these precision/recall tradeoffs (e.g., Huang et al., forthcoming).[6]

[3] The 1206 customers are the intersection of the purchase-revealers and the Amazon Friends network; one-half have social-network neighbors who are also purchase revealers.

[4] The proximity-based SNCF in Figure 1 uses k=2; the results are similar for larger k. Using k=1 does not produce enough recommendations for competitive recall (early-curve precision is slightly better).

[5] See discussion of limitations below.

[6] The numbers are not completely comparable, but the implication is clear.

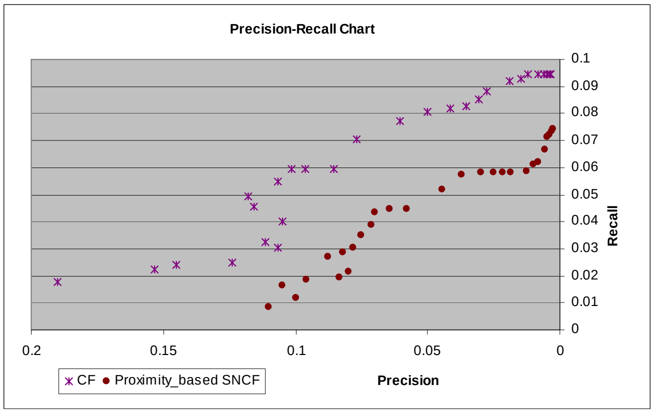

Figure 1. Precision/recall comparison between traditional CF and proximity-based SNCF.

Our second experiment compares traditional CF with *selection-based* SNCF. Specifically the same CF technique is applied respectively on two different data sets. One is the whole set of the consumers (the same 1206 as in the first experiment) and the other is the subset of 603 consumers who are interconnected by the social network. **Figure 2** shows the precision/recall comparison. When CF is applied to the whole data set, the results are substantially worse than with *selection-based* SNCF. The techniques have comparable precision, but for a given level of precision the recall is cut in half. What's worse, the running time increases super-linearly in the size of the user-base (Linden et al. 2003), and generally one would run CF on much more than just 1206 consumers.

The advantage of SNCF is much more striking if we examine these results carefully. Recall that the SNCF recommendations are a subset of the overall recommendations. It turns out that in the comparison the SNCF recommendations dominate the comparison; the recall is cut in half because only the SNCF recommendations play any considerable role in the accuracy. For comparison, **Figure 3** shows the performance of CF on the subset of 603 consumers who are not part of the social network (notice the scales of the axes). With the exception of their membership in the social network, these two sets of 603 customers are similar in terms of their visible characteristics (e.g., the number and variety of purchases). However, the precision and recall for traditional CF on the non-networked subset are an order of magnitude worse than on the networked subset—demonstrating just how much advantage is conferred by the selection of the social network recommenders. At least for data sets of this size, CF without social-network selection just can't compare with SNCF.

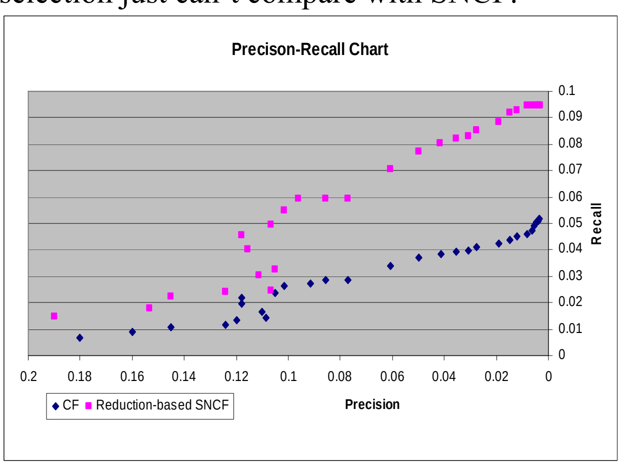

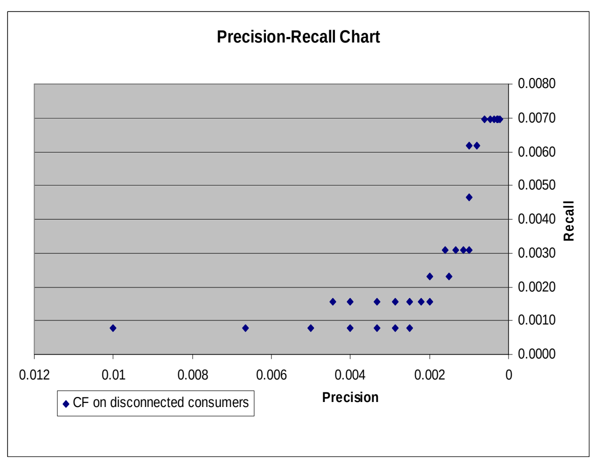

Figure 2. Precision/recall comparison between traditional CF and selection-based SNCF.

Figure 3. Precision/recall of traditional CF on disconnected consumers.

5. Other Related Literature

Large real-world networks such as the WWW, internet topology, social networks, biological networks, and linguistic networks have been extensively studied from a structural point of view. Typically, these studies address properties of the graph including its size, density, degree distributions, average distance, small-world phenomenon, clustering coefficient, connected components, community structures, etc. (Nowell et al. 2005). Online friendship and email graphs have been studied in the context of explaining and analyzing friendships (Kumar et al. 2004) and demonstrating the small-world and navigability properties of these graphs (Dodds et al 2003, Nowell et al. 2005, Adamic and Adar 2005). Co-author and co-occurrence information in online documents were also used to construct social network from which collaborators can be queried (Kautz et al, 1997). However, with the exception of Hill et al. (2006), none of this work has examined the impact of social-network-based relationships on members' affinity to purchase products, and in particular, none with the objective of designing recommender systems.

Prior work in recommender systems has postulated that "recommendations, however, are not delivered within a vacuum, but rather cast within an informal community of users and social context" (Perugini et al. 2004). Recent research (Huang et al. forthcoming, Mirza et al. 2003) improved the quality of recommendations by extending the direct co-purchase relationship to an indirect co-purchase network. A limitation of this stream of work is that only co-purchase behavior is counted for deriving consumer's purchase preferences. Our work contributes to this stream of research by demonstrating the impact of users' social information on the prediction of product purchases. In this regard, the implications of our work are related to the emerging stream of work on online word-of-mouth that captures how the inherent trust embedded in user-generated opinions and social information disclosures in online communities affect product sales (Chevalier and Mayzlin 2006, Forman et al. 2007). Studying a close-knit book-lover community, Ziegler et al. (2005) showed how topic diversification in the recommended set increases recommendation quality as perceived by users.

6. Discussion

Clearly these results are based on a single study, on a relatively small data set—and should be taken as a preliminary study. Nevertheless, they show convincingly that recommendations made by social-network-based collaborative filtering (CF) can be far superior to recommendations made by CF on a similar-size data set that does not represent a social network.

These results add support to the results of Hill et al. (2006) that social networks can enable technology-based methods to predict purchase behavior, and to our knowledge this is the first study to show the effect in a CF setting. Second, they demonstrate a very effective way to scale up CF, a technique previously thought to be inapplicable to large user bases (Linden et al., 2003). Specifically, CF can be scaled up by using social networks to scale down the user base used to make recommendations to a particular user. In our experiments, doubling the size of the user base did not improve the precision of the recommendations; recall was cut in half, and computational cost more than doubled (because the number of users and the number of products each double). Furthermore, the results on the larger user base only look good because the SNCF results are embedded—separated out, the non-social network results are an order of magnitude worse. Our on-going work examines different and larger data sets, e.g., from large social networking sites, to assess the robustness of these results.

These preliminary results have notable limitations. Due to selection bias, it may be that the set of consumers represented in this rather small social network is not representative of the general consumer (who participates in a social network). A more technical limitation is that our notion of "prediction" is based on selecting the most recent 20% of each consumer's purchases. It is possible that some of these purchases are actually later in time than some of the purchases used to make the recommendations. Finally, it may well be that the social network itself is being used as a recommendation system by the users, and the reason they are buying these products is that it was purchased by their social-network neighbors. This alone would be a remarkable finding, if it could be verified.

References

1. Adamic, L., and E. Adar. How to search a social network. *Social Networks*, 27(3):187–203, 2005.
2. Adomavicius, G., and A. Tuzhilin, Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions, *IEEE Trans. Knowl. Data Eng*. 17(6): 734-749, 2005.
3. Breese, J.S., D. Heckerman and C, Kadie. Empirical analysis of predictive algorithms for collaborative filtering. in *14th Conference on Uncertainty in Artificial Intelligence*, 1998, 43-52.
4. Chevalier, J. and D. Mayzlin, The Effect of Word of Mouth on Sales: Online Book Reviews, *Journal of Market Research*, 2006.
5. Dodds, P., R. Muhamad, and D. J. Watts. An experimental study of search in global social networks. *Science*, 301:827–829, 2003.
6. Forman, C., A, Ghose and B. Wiesenfeld. Examining the Relationship between Reviews and Sales: The Role of Reviewer Identity Disclosure in Electronic Markets, *NYU CeDER Working Paper*.
7. Golbeck, J. 2006. Generating Predictive Movie Recommendations from Trust in Social Networks. *Proceedings of the Fourth International Conference on Trust Management*.
8. Hill, S., F. Provost, and C. Volinsky. Network-based Marketing: Identifying likely adopters via consumer networks, *Statistical Science* 21 (2) 256–276, 2006.
9. Huang, Z., H. Chen, D. Zeng, Applying Associative Retrieval Techniques to Alleviate the Sparsity Problem in Collaborative Filtering. *ACM Trans. on Information Systems*, 22(1), 116–142 (2004).
10. Huang, Z., D. Zeng, H. Chen, A Comparative Study of Recommendation Algorithms in E-commerce Applications, *IEEE Intelligent Systems*, forthcoming.
11. Kumar, R., J. Novak, P. Raghavan, and A. Tomkins. Structure and evolution of blogspace. *Communications of the ACM,,* 47(12):35-39, 2004.
12. Kautz, H., Selman, B., and Shah, M. Referral Web: combining social networks and collaborative filtering. *Communications of the ACM,*40(3): 63-65, 1997.
13. Liben-Nowell, D., J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *PNAS*, 102(33):11623–11628, 2005.
14. Linden, G., B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. IEEE Internet Computing, page 7680, 2003.
15. Massa, P. and P. Avesani. Trust-aware collaborative filtering for recommender systems. Proc. of International Conference on Cooperative. 2004
16. McPherson, M., L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.
17. Mirza, B.J., B.J. Keller, and Ramakrishnan, N. 2003. Studying Recommendation Algorithms by Graph Analysis. *Journal of Intelligent Information Systems*, 20(2), 131–160, 2003.
18. Perugini, S., M. A. Goncalves, & E. A. Fox 2004. A connection centric survey of recommender system research. *Journal of Intelligent Information Systems* 23[1], 2004.
19. Resnick, P. and H. Varian. Recommender systems. *Communications of the ACM,* 40 (3). 56-58, 1997.
20. Weng, J., C. Miao, and A. Goh. Improving collaborative filtering with trust-based metrics. *Proc. of the 2006 ACM symposium on Applied computing*, pages 1860–1864, New York, NY, USA, 2006.
21. Ziegler, C.N., S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Procs of the 14th international conference on WWW*, pages 22–32, 2005.