A Research Platform for Recommendation within Social Networks

Amit Sharma
Dept. of Computer Science
Cornell University
Ithaca NY 14853
asharma@cs.cornell.edu

ABSTRACT

Recommendations within a network do affect, and get affected by, the information flow and the social connections within the network. Thus, designing a network-centric recommender system requires understanding people's preferences, their social connections, as well as the characteristics of the network they inhabit. This creates a major challenge in research on network-centric recommendations—exploring questions around networks and recommendation is hard because they invariably depend on the interaction between more than one user. We describe a research platform that we have built that helps us answer network-centric research questions. We present an overview of the system and demonstrate its usefulness through an example study involving directed suggestions between pairs of participants. As a useful side-effect, it is also helping us collect data about people's preferences in social networks.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Information Filtering*; H.1.2 [**Models and Principles**]: User/machine systems—*Human Factors*

General Terms

Experimentation, Human Factors

1. INTRODUCTION

Recent research on recommendation using social networks has taken two main approaches—augmenting collaborative filtering with social data [7,10], or using only the first-degree connections for recommendation [3, 6]. However, the gains reported by using social signals to recommend items are only slight [1], except for network-specific tasks such as friend recommendation.

One of the reasons for mixed results with *network-aware* approaches may be that people's explicit social connections have little to do with their interests in a particular domain. Thus, in domains where social connections overlap with user interest, social recommendation may be useful, in others, not so much [13]. This claim is supported by evidence that recommendations based on implicit networks constructed from domain-specific user activity give good results on predicting users' preferences [5, 15].

However, we argue that there is more to explicit social networks that offline measures of recommender performance may not capture. For instance, studies with Facebook and Google users have found that showing the name of a particular friend with a recommendation (music or news) can alter a user's perception of it. [9, 12]. Recommendations based on social connections may also help users navigate their social network and help them become more aware of the interests and preference of people in it. As often happens in the offline world, friends of a user may also like to recommend items directly to him [2]. In such cases, and others such as group recommendation, these recommendations can help support shared experiences as well as influence the interpersonal relationships between people.

In addition, a user's preferences are not static; they are continuously being influenced by their network. The structure and properties of a network, as a whole, are also affected by the connections between people and their activities within the network [4]. These factors suggest that there is value in considering recommendations as embedded within a social network, rather than being served in isolation. We call this approach to recommendation *network-centric* [11], in contrast to the network-aware approaches described earlier.

The principles of network-centric recommendation are based on the observation that recommendations do affect, and get affected by, the information flow and social connections within a network. Thus, designing a network-centric recommender system requires understanding people's preferences, their social connections, and the characteristics of the network they inhabit.

Framed this way, a challenge of network-centric recommendation is that it makes designing and evaluating systems hard, because they invariably depend on interactions between more than one user. This paper describes how we can use PopCore[1], a research platform we have built, to explore questions around networks and recommendation. As a network-centric system built on top of Facebook's social

[1]http://www.popcore.me

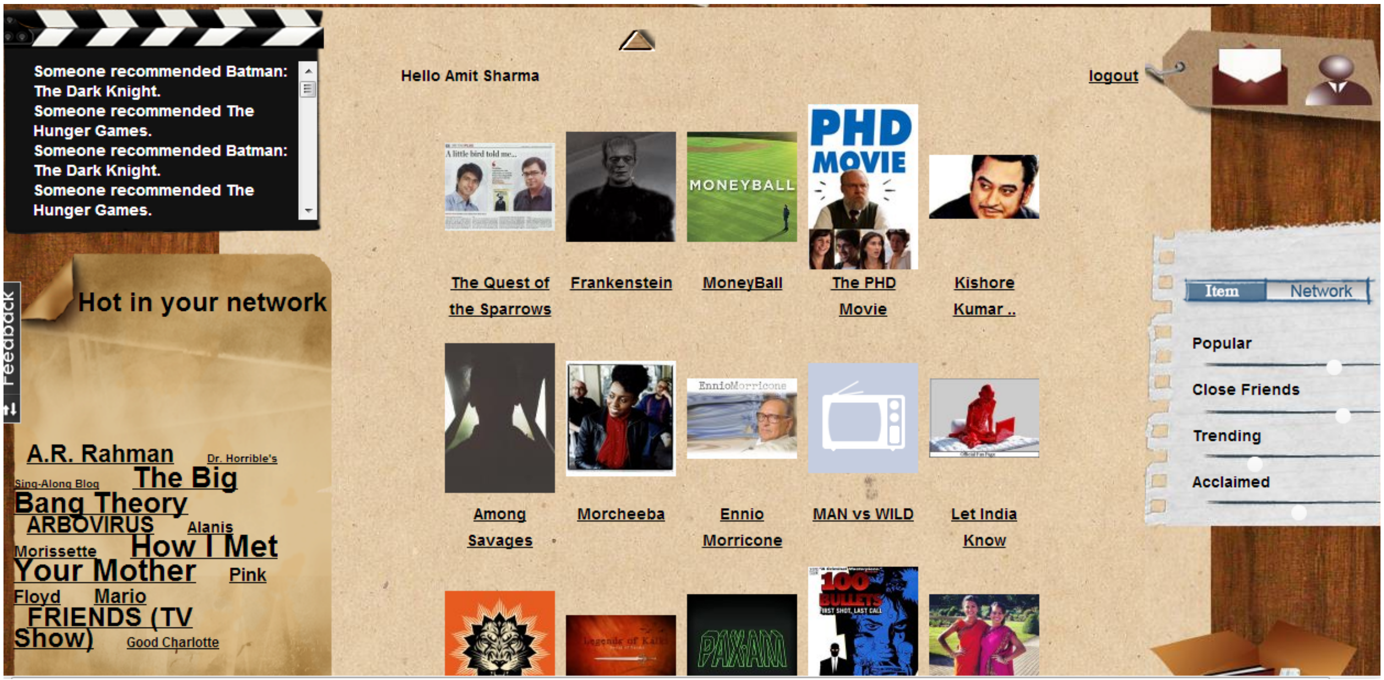

Figure 1: The PopCore interface. Recommended items, occupying the center of the screen, are a mix of algorithmic recommendations and items directly suggested to a user by her friends. Sliders on the right help a user control the sources for recommendation. Visualizations on the left help in network awareness, in this case a word cloud of popular items among a user's friends.

network, PopCore serves as a research tool to support live evaluations of network-centric recommendations, and collection of user reactions and feedback in a real network setting. It is also a functional recommender system in its own right, much like the MovieLens system[2].

2. OVERVIEW OF SYSTEM DESIGN

Figure 1 shows a screenshot of PopCore, which we first proposed two years ago in this workshop [14]. It uses Facebook as the underlying network and covers the entertainment domain, including movies, music, books and television shows.

The interface has three main parts: recommendations in the center, visualizations on the left and user controls on the right of the screen. Recommended items are a mix of recommendations computed from items Liked by a user's friends', and *directed suggestions* from his/her friends. These directed suggestions are one of the key features of the system—items can be recommended manually by users to their friends. These suggestions serve two purposes. First, they utilize friends' knowledge to bring interesting recommendations. Second, they allow people to express their preferences to their friends and support conversations and shared experiences on the recommended items.

For each recommendation, a user may Like it on Facebook, rate it, recommend it to some friends or add it to a personal queue (Figure 2). When a user clicks on the *Recommend* button for an item, he can choose one or more of his friends to recommend the item to. For convenience, we also offer a

[2] http://www.movielens.org

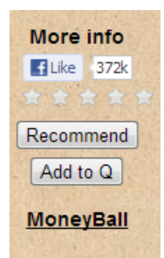

Figure 2: Possible actions on a recommended item, shown when a user clicks on it. A user may Like it on Facebook, rate it, recommend it to a friend or add it to her personal queue.

list of recommended recipients who PopCore predicts might be interested in the item (based on the similarity between the item and the user's friends).

Visualizations, such as the item cloud of most popular items Liked by a user's friends and updates of recent activity on the app (Figure 1), help create increased network awareness for the user. User controls on the right, which allow a user to tailor recommended items based on their popularity, social closeness or similarity of the people connected with those items, also help users find interesting content and navigate their friends' interests.

3. POPCORE AS A RESEARCH PLATFORM

We now describe how we are using the system as a test bed for research. In addition to the main website, we have a created a separate labs website[3]. It runs on the same core infrastructure, but has custom code that allows us to run experiments. Participants are asked to sign up and give permission for each experiment separately.

The core infrastructure of PopCore supports an ego-centric view of the network for each user, collecting Likes and network information about the user and his friends that can be used for a variety of studies. In our first experiment, we showed users a variety of different recommendation algorithms, some of which used past Likes of the users' friends. Users were asked to rate and Like items and react to the idea of network-centric recommendation [11]. We found that an algorithm suggesting the most popular items among a user's friends performed the best among those that used ego network information. It was also significantly better rated than an algorithm based on overall network popularity. Users' reactions to the recommendations confirmed that network ties can provide a useful way to choose potential content and neighbors for information filtering tasks.

In the study described above, participants did not know the social nature of the recommendations. Subsequently, we have used the platform to study how presenting a social explanation along with a recommendation can affect user ratings for music [12]. Here, we showed users different kinds of explanations involving their friends, such as "X and 3 other friends like it", along with a recommended musician. Our results show that social explanation has only a secondary effect on the ratings; the primary effect is that of a user's expectation of liking an item. Based on the findings, we provided a generative mixture model for a user's decision process on a recommendation.

Besides live experiments, data collected for the above studies helps us reason about the properties of networks and how they might affect recommendation. In one such study, we compared offline performance of algorithms using preferences of friends or the whole network, and found that algorithms using just the friends' information give comparable results to those using the full network [13]. Further, the data helped us investigate how the prevalence of *locality* of preference (concentration of Likes for an item in parts of the network) roughly correlates with the performance of friends-only algorithms on three different domains, suggesting that this locality is an important phenomenon and resource for recommender systems.

Till now, we have described how PopCore can be a useful platform for experiments involving a user and social data. However, the system is designed to support participation and interaction between multiple users. We now describe an ongoing experiment as an example of how PopCore can be a useful tool for conducting experiments with more than one participant.

3.1 Directed suggestions: A paired experiment

We consider the practise of sharing items between people, more commonly known as "word of mouth". Intuitively, suggesting an item to another person may involve two factors: having an opinion about the item, and having an understanding of the receiver's preferences. Understanding the processes behind *directed* suggestions, specifically around item and receiver selection, can help design recommendation systems that support such suggestions.

[3]http://labs.popcore.me

One of our first goals is to simply compare the recommendation quality of directed suggestions versus algorithmic recommendations. In the following paragraphs, we present how the PopCore platform can be useful for investigating this question.

Experiment Design: Previous research on helping people share news items in RSS feeds suggests that manual recommendations to friends can be useful [2]. Our approach is somewhat like Krishnan et al., who asked people who did not know a target user to make recommendations based on the target user's ratings [8]; however, instead of providing a list of ratings from a stranger, we ask people to make recommendations for their friends based on what they already know about them.

A simple way to compare the quality of recommendations would be to analyze logs of user ratings in PopCore and compare the performance of directed suggestions and algorithmic recommendations. However, this comparison will depend greatly on the particular recommender algorithm being used. One way to control for the recommendation algorithm is to allow only those directed suggestions that are also recommended by the algorithm we are comparing against. Then, among the algorithmically recommended items, we would be able to compare the ratings between items that were manually suggested and items that were not. Such a comparison would point to the quality of manual and algorithmic recommendations, as perceived by the user.

For ease of logistics, we design a movie recommendation experiment involving pairs of participants. The system generates its top 10 recommendations for each participant in a pair. It then randomly shuffles the 20 recommendations for a pair into a combined list, and shows the combined list to each participant. Similar to the main PopCore interface, participants are free to rate them or suggest them to their partner. We show this combined list to participants so that both of them rate and suggest from the same set of recommendations. This is done to ensure that both partners can choose to start the study at different times and still rate each other's potential suggestions in a single session. In addition, comparing ratings for the same item from both sender and receiver could tell us whether people tend to recommend items that they themselves rate highly, or those that may be relevant based on the other's preferences.

Experiment Flow: We run the study using the labs version of PopCore. In the first stage, a participant invites one of his friends as his partner. Once the other participant accepts the invitation, she enters the experiment and rates/suggests 20 items. The first participant then gets an email and he can go ahead and rate/suggest the same 20 items. During the experiment, we highlight the other participant's name whenever a user wishes to make a directed suggestion (Figure 3), in order to make the directed suggestions easy to send out. Thus, at the end of the experiment, we receive ratings for the items from both participants and

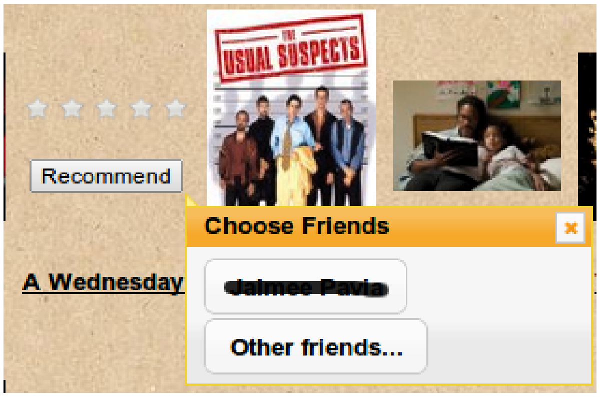

Figure 3: The interface for directly suggesting movies to a friend. In the live system, we would show a list of friends who the system predicts might be good to share with; for the experiment, we alter this behavior to show only the name of their experimenal partner.

a list of the items that they shared to one another. Participants are also presented with a questionnaire that asks about how close they are to the other participant, and the reasons they receive or send recommendations in general.

Note that participants have no way of detecting which items were suggested to them by their partner. Thus, the ratings can be assumed to pertain to the quality of the recommendation, rather than the social connection between the sender and the recipient.

Initial Results: We are currently in the process of collecting data. We are recruiting participants from a mix of college student pools and Amazon mechanical turk. From an analysis of the initial data on 27 pairs, we find that items which were manually suggested have a higher average rating than those that were algorithmically recommended. This supports the notion that manual recommendations are useful, and in some cases, better than an algorithmic recommendation. However, it may also be the case that people tend to recommend highly popular items which are more likely to get higher ratings. We hope to use offline data to corroborate these findings.

In addition, we find that the average rating for manual recommendations is higher for the senders than receivers. This result suggests that users may be sharing items that they like more frequently than those which they think can be relevant for the recipient.

4. CONCLUSION

We have presented PopCore, a research platform for conducting user experiments with recommendations in social networks. Our goal for this paper is to show how the system has helped us conduct experiments that require social data and/or active participation of users embedded in a network. We envision it as an open platform for research, for which we would like to discuss potential ways of collaboration at the workshop. For instance, one of the ways is to allow other researchers to setup experiments on PopCore. In addition, we also hope to share the data we are collecting, keeping in mind the potential challenges around user privacy in generating those datasets.

5. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grants IIS 0910664 and IIS 0845351.

6. REFERENCES

[1] O. Arazy, N. Kumar, and B. Shapira. A theory-driven design framework for social recommender systems. In *Journal of the Assoc. for Info. Sys.*, 2010.

[2] M. S. Bernstein, A. Marcus, D. R. Karger, and R. C. Miller. Enhancing directed content sharing on the web. In *Proc. CHI*, pages 971–980, 2010.

[3] J. Chen, R. Nairn, L. Nelson, M. Bernstein, and E. Chi. Short and tweet: Experiments on recommending content from information streams. In *Proc. CHI*, 2010.

[4] D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri. Feedback effects between similarity and social influence in online communities. In *Proc. KDD*, 2008.

[5] J. Golbeck and J. Hendler. Filmtrust: movie recommendations using trust in web-based social networks. In *Proc. IEEE Consumer Comm. and Networking Conf.*, 2006.

[6] I. Guy, M. Jacovi, A. Perer, I. Ronen, and E. Uziel. Same places, same things, same people? Mining user similarity on social media. In *Proc. CSCW*, CSCW '10, pages 41–50, New York, NY, USA, 2010. ACM.

[7] I. Konstas, V. Stathopoulos, and J. M. Jose. On social networks and collaborative recommendation. In *Proc. SIGIR*, 2009.

[8] V. Krishnan, P. K. Narayanashetty, M. Nathan, R. T. Davies, and J. A. Konstan. Who predicts better?: results from an online study comparing humans and an online recommender system. In *Proceedings of the 2008 ACM conference on Recommender systems*, 2008.

[9] C. Kulkarni and E. Chi. All the news that's fit to read: a study of social annotations for news reading. In *Proc. CHI*, 2013.

[10] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *Proc. WSDM*, 2011.

[11] A. Sharma and D. Cosley. Network-centric recommendation: Personalization with and in social networks. In *Proc. IEEE SocialCom*, 2011.

[12] A. Sharma and D. Cosley. Do social explanations work? Studying and modeling the effects of social explanations in recommender systems. In *Proc. WWW*, 2013.

[13] A. Sharma, M. Gemici, and D. Cosley. Friends, strangers, and the value of ego networks for recommendation. In *Proc. ICWSM*, 2013.

[14] A. Sharma, M. Malu, and D. Cosley. Popcore: A system for network-centric recommendations. In *Proc. RSWEB Workshop, RecSys*, 2011.

[15] T. Zhou, J. Ren, M. Medo, and Y.-C. Zhang. Bipartite network projection and personal recommendation. *Phys. Rev. E*, 76, Oct 2007.