

PopCore: A system for Network-Centric Recommendations

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

Meethu Malu
Information Science
Cornell University
Ithaca, NY 14850
mm956@cornell.edu

Dan Cosley
Information Science
Cornell University
Ithaca, NY 14850
danco@cs.cornell.edu

ABSTRACT

In this paper we explore the idea of *network-centric* recommendations. In contrast to individually-oriented recommendations enabled by social network data, a network-centric approach to recommendations introduces new goals such as effective information exchange, enabling shared experiences, and supporting user-initiated suggestions in addition to conventional goals like recommendation accuracy. We are building a Facebook application, PopCore, to study how to support these goals in a real network, using recommendations in the entertainment domain. We describe the design and implementation of the system and initial experiments. We end with a discussion on a set of possible research questions and short-term goals for the system.

Keywords

recommender systems, social recommendation, network-centric

1. INTRODUCTION

Users are increasingly disclosing information about themselves and their relationships on social websites such as Facebook, Twitter, and Google+. These data provide signals that have been used to augment traditional collaborative filtering techniques by making *network-aware* recommendations [8, 9]. Such recommenders use social data to support prediction, provide social context for the recommendations, and help alleviate the cold-start problem typically found in recommender systems. Much of their power comes from social forces, such as homophily, trust and influence, and thus these recommenders do not just provide better recommendations, they can also support the study of these forces. For example, in [4], the authors divide a user's social contacts into familiarity and similarity networks (proxies for trust and homophily, respectively), and study their relative impact on the quality of recommendation.

But we can take this a step farther. Just as a user's network can influence the recommendations he/she receives,

the recommendations, in turn, also can influence the network and alter the underlying social processes. For instance, new recommendations can alter the diversity of the set of items within a network while group recommendations can strengthen social ties. Thinking of recommendations as being embedded in a network, rather than informed by it, provides a new context for analyzing and designing recommender systems—and an important one, given people's increasing interaction and consumption in online social networks.

In this paper, we lay out our approach to exploring this *network-centric* approach to recommendation system design. We start by discussing new concerns such systems foreground, focusing on design goals that come from thinking about the social aspects of recommendations that are embedded in a network, compared to more individually-focused systems. Second, we introduce PopCore, the network-centric recommender system we are building in Facebook to support these goals. We have already deployed an initial proof-of-concept version to conduct initial experiments around network-aware algorithms [13]; here, we discuss how we are evolving the system to support the social design goals. We close by laying out issues that doing network-centric recommendations raise, most notably around the tension between social sharing, privacy, and identity management, and outlining the initial questions we hope to address as we design and build both the system and the community.

2. NETWORK-CENTRIC DESIGN GOALS

Thinking about recommendations as embedded in a social network raises a number of questions, ranging from using network data to improve individual recommendations to using people's behavior to study large-scale patterns of diffusion and other social science forces at work. Given our goal to design a useful network-centric recommender system, here we focus on design goals that capture social elements that are more salient than they would be in a typical e-commerce recommender application.

Directed Recommendations. An integral part of social experience is sharing information with others person-to-person. Such user-generated *directed suggestions* have been studied for link-sharing [2] and are ripe for study in other domains, integration with automated applications, and application to a social network context. Allowing directed suggestions might encourage people to be more active participants in the system and allow them ways to express their identity. These suggestions may also be more accurate than collaborative filtering for certain tasks [6], and in aggregate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Workshop on Recommender Systems and the Social Web '11 Chicago, IL
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.



Figure 1: A mockup of the PopCore interface. By default, recommendations are shown from all three domains, Movies, Books and TV. The controls at the top help a user decide the composition of the list of items, while the lower section provides contextual visualization for the recommendations.

support data mining and automated recommendation from this user-generated ‘buzz’ [10].

Shared Experiences and Conversation. For many items, especially in the entertainment domain, enjoyment depends not just on personal preferences but also on social experiences such as enjoying the content with other people [3]. Given an item such as a movie, is it possible to predict the people who may join you for it? This is slightly different from group recommendations, which are typically aimed at a predefined group of people [1], and we expect that leveraging network information will make them more effective than earlier approaches that combined individual lists of recommendations [11]. Conversation is another social experience, and since people who disagree about movies have livelier conversation [7], algorithms might focus on recommending items that evoke strong reactions, or even “anti-recommendations”, along with the traditional goals of accurate recommendation. Systems aimed at individuals are unlikely to want to recommend hated items, but people often like to talk about them, and this propagation of negative opinion may also help others avoid bad experiences.

Network Awareness. Negative information is a specific kind of awareness, and people have a broad interest in awareness of what is happening in their social network [5]. From the point of view of information, taste, and fashion, it’s useful to know who the opinion leaders are, who are active and effective recommenders, what items are becoming hot or not, and who is knowledgeable about a given topic [12]. Thus, supporting social interaction not just between individuals but at the network level is likely to be valuable in a network-centric recommender system.

3. POPCORE: THE PLATFORM

We now discuss how we are starting to realize these goals in PopCore, a Facebook application we are developing for providing and studying network-centric recommendations. We chose Facebook because it provides us both network and preference data (though Likes), and also supports a diverse set of domains for items. PopCore works by fetching a user and her friends’ profile data on Facebook (subject to the user’s permission) and providing recommendations based on those signals. Currently, we restrict PopCore to the entertainment domain, including movies, books, TV shows and music. These categories have a fair amount of activity and broad popular appeal.

3.1 System Description/Design

We decided on a simple three-part interface, as shown in Fig. 1. The center section contains the main content to be shown (a list of items), while the top and bottom sections show content-filtering controls and contextual visualizations respectively. Each of the interface components, from the logo on down, is designed to support both the goals outlined above and the collection of interesting data to study.

PopMix. The top section is the control panel *PopMix*, which allows a user full control of the type of items shown in the content section. The controls are designed to be intuitive, inspired by a common interface metaphor, a music equalizer. Just like the music mixer allows a user to set sound output according to his tastes, the PopCore interface gives the user control over the domain, genre, popularity and other parameters he may choose. In order to account for temporal preferences, we also include a special *recency knob*

that allows users to select the proportion of recent versus older items shown. In addition, users may also view items expected to be available soon. For such current and ‘future’ items, users may notify and invite their friends (chosen manually or from a system-recommended list). This supports the goal of shared consumption.

Eventually, we plan to implement filters that allow people to control social network parameters as well. For instance, a user may also choose the relative importance/proportion of network signals for recommendation, such as link-distance of people from the user, interaction strength, age, location of people. People may also select a subset of people manually, or a named group of people, in which case recommendation morphs more into a stream of items from those sources.

Stackpiles. The middle section shows a number of views relevant to user tasks such as getting automated recommendations, directed suggestions, and remembering suggestions to follow up on. The top right corner of this section contains the tab-buttons to switch views as shown in Fig. 1.

The primary view presents automatically generated suggestions filtered on the user’s PopMix settings. The recommendation algorithm is a ranking algorithm that ranks a user’s friends based on their relevance to the user on a list of parameters, such as interaction strength and number of commonly Liked items. The most popular items according to this weighted user popularity are then chosen. In a given view, a user is shown a list of items arranged as cards in distributed stackpiles. Items are grouped into stackpiles based on their similarity, using k-means clustering over their attributes. The number of piles and their distribution is generated dynamically.

On flipping an item card, a user gets options to Like, Dislike, or rate the movie on a scale from 0.5-5. Giving a number of ways to interact with the movie supports rich data collection, and in the case of Dislike, the idea of anti-recommendations. Users may also directly suggest an item to one or more friends using the PopCore button. These suggestions are sent to the target users as a Wall Post or private message, based on the user’s preferences. PopCore members can also see these suggestions as a view in the content window by clicking the “N” button at its upper right. People can type in any friend; PopCore also suggests people that may be a good fit for enjoying the item with.

The other main view is a user’s personal library, which contains a user’s ‘For later’ list and the list of items for which the user has provided strong feedback. The ‘For later’ list may be thought of as a non-linear queue (and can be accessed through the “Q” button). The list benefits from the same stackpiling metaphor, thus allowing a user more visual and organized view of his/her library. Items are stackpiled based their similarity and recency in the list by default, however, users have full control to customize the groups.

Visualizations. The bottom section contains visualizations of network activity around items that support the network awareness goals described earlier. The default view is a word cloud showing items weighted by the number of user’s friends who have Liked those items. Other visualizations include showing the friends who have contributed the most to the content shown to a user (either through directed recommendations, or algorithmically) along with the items that have been recommended (Fig. 2), or a timeline showing the entry and growth of recent items in a user’s network. The goal of these visualizations is to help the user navigate

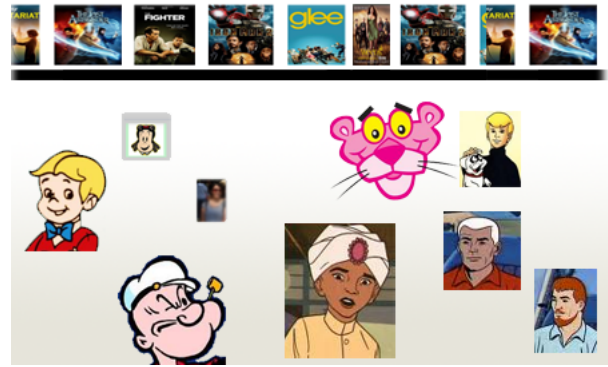


Figure 2: A visualization showing aggregated behavior among the user’s friends weighted by the amount of recommendations they make, with a detailed view of each item that has been recommended in the user’s network.

the multi-part social activity information in a clear, intuitive fashion.

4. ISSUES AND RESEARCH QUESTIONS

We conclude by discussing the major issues we expect around deploying a real network-centric recommender.

4.1 Trading off social and private elements.

A primary issue is that having access to more information enhances the social discovery and consumption experience, but there is a direct trade-off with privacy. For example, the visualization component is designed to show individual activity about either items or people, and aggregate information about the other. “Activity” in this case might represent making or receiving directed suggestions, rating items, getting recommendations, adding items to one’s queue, and so on. Consider showing items as the detail, people as the aggregate, and queuing as the activity. Users might want to know which items their friends are intending to consume, but it may often be the case that an individual using the system will queue a sequence of movies. Her picture will grow as the stream of movies changes. This will immediately convey her queuing behavior to others, and hence her privacy has been compromised.

Identity management also comes into play. Having Likes and Dislikes visible to all friends makes it easier for a user’s friends to follow his/her interests, but does it then affect the Liking behavior based on concerns about privacy and identity management? Similarly, a queue is a definite indication of interest and making it accessible to others will directly benefit shared experiences and co-operation, but it is unclear whether users would want to have a public queue. For now, we have decided to have everything except Likes and Dislikes private, but give the user an option to selectively enable items for sharing whenever an action is taken, with the hope of balancing identity, privacy, and discovery without imposing too much work.

4.2 Long-term goals and short-term questions.

The other major issue we see is that building out a network-centric recommender while building up its userbase promises to consume a fair amount of time. Thus, our short-term goal is to answer questions that need no or limited social interaction while the system and userbase develops.

Tradeoffs in doing network-centric recommendations. A network-centric approach affords fast algorithms, real-time capabilities, and modest user requirements compared to conventional collaborative filtering's use of large datasets, but it places a lot of emphasis on a person's immediate social network. This reduces the pool of available items, and may also lead to a possible loss of diversity among the items recommended. We plan to pit our algorithm against state-of-the-art collaborative filters and compare the performance of both in terms of the activity generated around recommendations and users' satisfaction with the automated recommendations they receive from each. Eventually we hope to develop recommendation strategies that use recommendations computed both on the full dataset and in a network-centric way in the user's local network.

Interpreting actions and developing metrics. PopCore provides a wide variety of actions that users can take with an item, including putting it in their queue, publicly Liking it or Disliking it, or suggesting it to friends. All these actions may convey signals that can be used to both improve the quality of recommendations and also evaluate them, although we need to learn to interpret them. What's the difference between a "Like" (which is public) and a 5-star rating (which is probably not)? Sharing an item provides an indication of "interestingness", but unlike ratings does not provide a definite scale of enjoyment, and in fact people may share disliked items.

Exploring cross-domain recommendations. The network-centric approach relies heavily on people and their connections, and less on the items. This suggests that we may be able to cross-recommend items based on a user's network information and his/her preferences in a related domain, a task for which collaborative filters have not been so successful. Designing algorithms for cross-domain recommendation within a network is an interesting question in itself.

Social explanations. Right now PopCore uses data harvested from Wikipedia to present additional information about items to help people make decisions. However, that data does not explain why the recommendation was made, which is a commonly wanted feature in real world recommendation systems [14]. Using network information to help justify automated recommendations may be a powerful feature, given the way people rely on this information to make decisions already.

Once we have built the userbase we will be in a better position to ask questions about the explicit social elements we are designing for. Comparing directed to automatic recommendations, studying the value of awareness of network activity around items, exploring how recommendations and consumption propagate in the networks, and developing effective metrics for measuring social outcomes are all questions that we hope to address in the long term, and that we think are key for recommender systems as they move into social networks.

Acknowledgement.

We would like to acknowledge support from NSF grant

IIS 0910664.

5. REFERENCES

- [1] S. Amer-Yahia, S. B. Roy, A. Chawlat, G. Das, and C. Yu. Group recommendation: Semantics and efficiency. *Proc. VLDB Endow.*, 2:754–765, August 2009.
- [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] P. Brandtzaag, A. Folstad, and J. Heim. Enjoyment: Lessons from karasek. In M. Blythe, K. Overbeeke, A. Monk, and P. Wright, editors, *Funology*, volume 3 of *Human-Computer Interaction Series*, pages 55–65. Springer Netherlands, 2005.
- [4] I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, and S. Ofek-Koifman. Personalized recommendation of social software items based on social relations. In *Proc. RecSys*, pages 53–60, 2009.
- [5] A. N. Joinson. Looking at, looking up or keeping up with people?: Motives and use of Facebook. In *Proc. SIGCHI*, CHI '08, pages 1027–1036, New York, NY, USA, 2008. ACM.
- [6] V. Krishnan, P. Narayanashetty, M. Nathan, R. Davies, and J. Konstan. Who predicts better? Results from an online study comparing humans and an online recommender system. In *Proc. RecSys*, pages 211–218, Lausanne, Switzerland, 10/23/2008 2008. ACM.
- [7] P. J. Ludford, D. Cosley, D. Frankowski, and L. Terveen. Think different: Increasing online community participation using uniqueness and group dissimilarity. In *Proc. SIGCHI*, CHI '04, pages 631–638, New York, NY, USA, 2004. ACM.
- [8] H. Ma, I. King, and M. R. Lyu. Learning to recommend with social trust ensemble. In *Proc. SIGIR*, pages 203–210, 2009.
- [9] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *Proc. WSDM*, pages 287–296, 2011.
- [10] H. Nguyen, N. Parikh, and N. Sundaresan. A software system for buzz-based recommendations. In *Proc. SIGKDD*, KDD '08, pages 1093–1096, New York, NY, USA, 2008. ACM.
- [11] M. O'Connor, D. Cosley, J. A. Konstan, and J. Riedl. PolyLens: A recommender system for groups of users. In *Proc. ECSCW*, pages 199–218, Norwell, MA, USA, 2001.
- [12] N. S. Shami, Y. C. Yuan, D. Cosley, L. Xia, and G. Gay. That's what friends are for: Facilitating 'who knows what' across group boundaries. In *Proc. GROUP*, GROUP '07, pages 379–382, New York, NY, USA, 2007. ACM.
- [13] A. Sharma and D. Cosley. Network-centric recommendation: Personalization with and in social networks. In *IEEE International Conference on Social Computing*, Boston, MA, USA, 2011.
- [14] K. Swearingen and R. Sinha. Beyond Algorithms: An HCI Perspective on Recommender Systems. In *ACM SIGIR Workshop on Recommender Systems*, New Orleans, USA, 2001.