2011 IEEE International Conference on Privacy, Security, Risk, and Trust, and IEEE International Conference on Social Computing

# Network-Centric Recommendation: Personalization with and in Social Networks

Amit Sharma
Department of Computer Science
Cornell University
asharma@cs.cornell.edu

Dan Cosley
Information Science Department
Cornell University
danco@cs.cornell.edu

***Abstract*—People often rely on the collective intelligence of their social network for making choices, which in turn influences their preferences and decisions. However, traditional recommender systems largely ignore social context, and even network-aware recommenders don't explicitly support social goals and concerns such as shared consumption and identity management. We present relevant theories and research questions for a more *network-centric* approach to recommendations and introduce PopCore, a platform for studying them in Facebook. An initial 50-user study with PopCore gives insights into tradeoffs around the popularity, likeability, and rateability of recommendations made by a set of network-centric algorithms and to people's thoughts about the idea of network-centric recommendation.**

## I. Introduction

Collaborative filtering (CF) recommender systems, deployed in both research contexts such as GroupLens [1] and commerical sites such as Netflix [2] and Amazon [3], have helped people cope with the increasing scope of choices available in movies, books, music, and other domains for almost two decades. These systems have strong appeal, as they can deliver useful recommendations by leveraging only preference information, with no other knowledge of either people or recommended items required. A weakness of pure CF systems, however, is that they require lots of preference information to be effective. Recognizing this, hybrid systems [4] use information such as demographics or product attributes to help make recommendations when data is sparse.

The rise of social networking and information sharing systems such as Facebook and Twitter adds new resources that recommender systems might leverage. People mention things they like and share links in status updates, explicitly reveal demographic and preference information in their profiles, and make connections to friends and acquaintances. *Network-aware* recommender algorithms and systems that leverage this information (e.g., [5], [6]) are a promising route to improving the utility of CF-based information filtering.

However, most of this work poses recommendation in social media as still doing collaborative filtering for individuals, just with network data added to improve prediction accuracy. We believe there is value in taking a more *network-centric* approach to information filtering and recommendation. By this we mean thinking about the role and design of recommender systems that act in a social networking context, rather than the dominant use cases in e-commerce sites. This invites two main questions. First, the properties, practices, values, and opportunities that appear when recommendations are situated inside a social network may change the algorithms or interfaces we build. Second, considering how algorithms, interfaces, people, and networks affect each other around recommendations leads to a number of questions at the intersection of computer science and social science.

In this paper, we make two main contributions. The first is to explain why a network-centric approach matters. Theories from the area of social network analysis, including homophily, influence, and diffusion, suggest new approaches for making recommendations and modeling users. Social-psychological factors such as the role of shared experiences in affirming friendships and the expression of opinions to manage identity present new goals for and constraints on network-centric recommender systems. Finally, observations of practice around making recommendations in social networks both online and off inform what is likely to be valued in these systems.

Our second contribution is to introduce PopCore, a Facebook application designed as a platform for studying network-centric recommendations. We present the system as it stands and demonstrate its research potential by deploying a number of network-aware algorithms to investigate the effects of popularity, personalization, similarity, and tie strength on how strongly, how often, and how well people react to recommendations for movies, television, and books. We also collect initial data about people's thoughts about self-presentation and recommendations in social networks. These data complement earlier studies that compare various network-aware recommendation algorithms [7], [8] and give insight into factors that will be important to network-centric recommendations. Our hope is that PopCore becomes both a well-used social recommendation platform and a community resource for conducting experiments in social recommendation with real users.

## II. Why network-centric recommender systems?

The dominant use case for recommender systems so far is as a tool for individual information filtering, particularly in e-commerce sites such as Netflix and Amazon. Collaborative filtering systems have traditionally framed the problem of making recommendations as choosing the best possible items for an individual user out of the space of all known items, given a set of users, items, and users' ratings of items. A natural

978-0-7695-4578-3/11 $26.00 © 2011 IEEE
DOI

approach, then, is to make predictions about unrated items and return the items with the highest predicted ratings as a ranked set of recommendations. This frames recommendation as a classification problem [9] that many researchers have attacked using various machine learning algorithms [10].

Framing the problem in terms of effective algorithms and individual transactions leads to a number of assumptions about algorithms, metrics, and interfaces for recommender systems that may be less appropriate or important in network-centric recommenders. Below, we discuss a set of related theories and considerations that, though not complete, provides a number of useful resources for design and research in this area.

### A. *Useful theories and forces*

One aspect of recommendation that is different in social networks is that the ties provide information. Homophily suggests that people with similar characteristics have a higher chance of being connected to each other [11]. People tend to spend time with and live nearer people of the same race, gender, educational background, occupation, class, and other demographic and socio-economic variables. Thus, network ties might capture useful similarity information [7], [8], [12], especially when combined with indicators of tie strength such as interaction frequency [13], ratings of trust [14], or implied computed trust [15], [16].

Social connections convey more information than just helping to identify potentially similar users. Our social networks impact our preferences [17]: ideas, preferences, fads, and opinions all diffuse through the network [18], changing minds (or at least informing receptive ones) as they go. But except in news recommendation, where the idea of an information need being satiated is explicitly addressed, most CF work assumes that preferences are more or less stable. However, it is known that people are inconsistent re-raters and that ratings are influenced by showing predicted ratings [19] and others' reviews [20]. This suggests that explicitly considering the influence of both interfaces and other people and modeling how preferences change over time may be valuable.

These influence processes tend to lead to clusters of local preferences. "cyber-balkanization" in which people and groups tend to see only material they like [21]. This may be bad in domains such as civic participation where knowledge of and respect for other points of view are virtues, but probably fine for entertainment domains such as movies and books. This locality has important implications. For instance, coverage—the ability to make predictions about any item—is often an important goal of CF systems [22]. But it might be fine to focus on recommending items that the local community has experienced, occasionally importing "new" items through external processes or bridges to other subcommunities [23] that cause a member to consume a new item and thus make it available. If coverage is not a primary goal, the need for large amounts of preference data and the consequent sparsity and cold-start problems [24] fade a bit, while algorithms that reason carefully about small amounts of preference data embedded in a network become more interesting.

The criteria for a "good" recommendation may also change in network-centric recommendations. Often, enjoyment of a movie depends not just on personal preferences but on social aspects such as going to the theater together [25]. These experiences can also later support reminiscing and relationship maintenance. However, there is little research about the consumption of recommendations and even less about social consumption. Group recommenders typically look at functions for combining individuals' predictions to make group predictions [26], [27]. However, linear combinations of predicted ratings capture only a small amount of the social value of recommendations. Thus, considering how recommendations support not just individual consumption decisions but how they affect relationships, groups and communities is another fertile area of research. How can they cement friendships, support identity management, and enact (or thwart [28]) the influence processes at play?

Many of these influence processes are enacted through what Bernstein et al. [29] call directed suggestions. "You should eat/read/watch this food/book/movie!" is a staple statement through which knowledge and preferences have been communicated forever ("Avoid the red berries!"). More recently, passing information through social networks has become a frequent, though minority, use of social media [30]. For the most part, though, neither social media systems nor CF algorithms provide explicit support for making directed recommendations or collecting them for later use and consumption. This may be a mistake, since people often have nuanced knowledge of others' preferences and contexts.

### B. *Interesting, mostly open questions*

These considerations point to a number of interesting questions around how to use network data in recommending, what it means to make effective recommendations in a network context, and what the effects of these recommendations might be on both the people and the network.

**What is the algorithmic design space?** Traditional recommender systems use a variety of approaches: content versus collaborative, model versus memory-based, machine learning and data mining algorithms of all flavors. Adding social data and social forces makes the design space much larger. Would algorithms that consider tie strength, diffusion processes, local popularity, or group consumption be feasible and effective?

**Do network-based recommendations work?** Network-derived recommendations were perceived as more useful and relevant than similarity-based recommendation in a corporate context [8]. But it is not clear what contexts are well-suited to these recommendations, and what the tradeoffs are between coverage, accuracy, data, computation time, satisfaction, and other commonly used metrics.

**What does it mean to "work"?** There is a more general question about how to measure the effects of recommendations in networks. Embedded recommendation systems are different than recommendation algorithms, and the recommendations might affect tie formation and strength, the diffusion of preferences, and so on.

**How to explain social recommendations?** People value explanations in recommender systems [31], and social explanations may be especially powerful given the influence people have on each other and the value of shared experience. How can network information help people make sense of recommendations?

**How to support shared consumption and directed recommendation?** As noted earlier, social factors aren't limited to influence: shared consumption and directed recommendations are both fundamental parts of social life. How to integrate them with automated or computer-supported recommendations is an open question.

**How do preferences relate to identity?** Recommending in a social context raises issues of self-presentation that do not come up in anonymous CF systems. Do you want your friends to know that you like Barry Manilow? People carefully groom their public images, and expressed preferences are an important part of this in social situations [32].

**How to manage privacy?** Even so, in network-centric recommendations, people will learn more about the community and probably about individual members in the community. Thinking about algorithms and interfaces that balance utility, sharing, and privacy is another important question, especially given the potential leaking of preference information through recommendations and conversations about items [33].

## III. The PopCore platform

PopCore is a Facebook application we are building to both conduct experiments in network-centric recommendations and to be a useful tool in its own right. Our goal is to develop it into a flexible platform for both ourselves and other researchers to help address the questions above. We chose Facebook because it provides access to network and preference data as well as to people who can both use it and participate in experiments. Here we present the current instantiation of PopCore, which we used to conduct the studies described below.

### A. Data available from Facebook

Every Facebook user has a profile that contains information he or she provides, including friendship ties, demographic data, and information about liked movies, books, TV shows and other items. People express these preferences by pressing the "Like" button associated with a Facebook page for the given item; pressing Like causes the page for that item to be associated with the user's profile (Figure 1). Friends can see each other's profiles and can interact through a number of mechanisms including wall posts, messages, tagging each other in pictures, and commenting on these posts and messages.

When someone uses PopCore, we first ask for consent to participate in an experiment and to access their profile information. PopCore stores whatever data is visible for the gender, age, and Likes of the given user and their friends for movies, books, and TV shows. We focus on these categories because they are relatively popular. As a first study on the PopCore platform, we do not include music because of the wide range of granularities (song, album, artist, band), but certainly it would be an interesting domain addition in the future. PopCore also records the amount of recent interaction between a user and their friends.

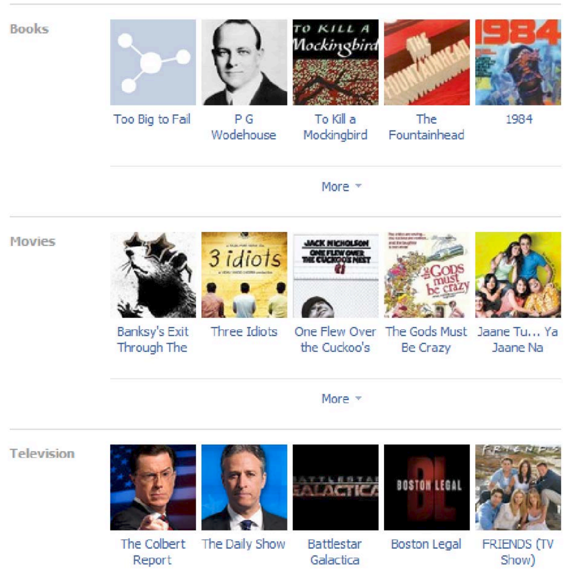


Fig. 1. Sample 'Arts and Entertainment' profile. This profile is visible to all friends. Depending on user privacy settings, it may even be public.

These data are noisy. Demographic information is optional; privacy settings may affect the communication we can see; people don't Like everything they like; pages are occasionally miscategorized or duplicated. We consider this a feature, not a bug. Studying recommendations in real use, with real users and real data, requires handling these issues in a way that is often obscured by dataset-based recommender research.

### B. Interface and data collection

The current version of PopCore has a simple interface tuned to presenting recommendations made by a number of different algorithms and collecting users' reactions to those recommendations (Figure 2). Recommendations are organized into tabs, each of which contains up to 20 recommended items. Users can rate a recommendation on a scale of 0.5 to 5 stars, to indicate that it is bad ("dislike") by clicking the red "no symbol", or to Like it. Though the Like button and a high rating both give positive evaluations, they convey different signals. To Like an item is to publicly identify oneself with it, while a rating serves as a private affirmation of interest.

Other than that, the interface has few features. In addition to the item name, contextual information from DBpedia[1] was also provided for items whenever the user hovered over an them. The "Instructions" and "Feedback" tabs are specific to our first study, which is described next.

[1] http://dbpedia.org

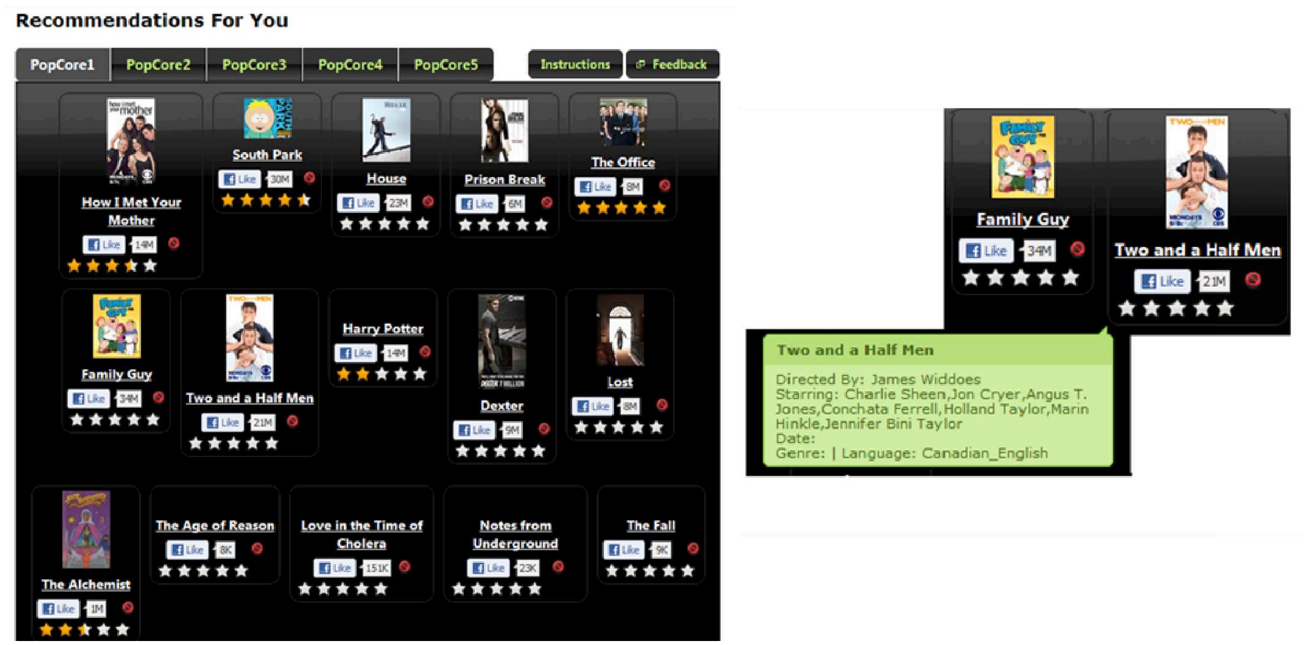


Fig. 2. The PopCore interface. There are five tabs, each containing recommendations from a different algorithm. At right, additional information is shown when a user hovers over the item. Users can give feedback by rating, Liking, or disliking an item.

## IV. Exploring network-centric recommendations

A key question for network-centric recommenders is how to make recommendations in a world where, instead of seeing every possible user, the candidate set of similar users is pruned only to the friends of a user. Having hundreds of users instead of millions immediately changes the nature of the problem. The unary nature of ratings also renders many classic algorithms and similarity metrics less suitable.

In this study we explore algorithms that work in this data-parsimonious environment by exploiting network features. We focus on three main questions. One is whether network ties are useful for prediction compared to similarity between item preferences, which has been explored in Wikipedia [7] and an IBM social networking platform [8]. A second is to explore tradeoffs between plausible network-aware algorithms and the properties of the recommendations they generate. The third is to explore users' reactions to these recommendations and to the idea of a network-centric recommender.

We deployed six algorithms. The first three, *Network-Random*, *Network-Popular*, and *Overall-Popular*, serve as non-personalized baseline algorithms. *Timeline* leverages the streaming nature of social media content by presenting items recently liked by friends. The last two, *Network-Similar* and *Interaction-Strength*, address our first question. All of the algorithms are reasonably fast and suitable for on the fly computation with live data.

- *Network-Random*. This algorithm selects a random list of items from the Likes of a user's friends.
- *Network-Popular*. This algorithm selects the items with the most Likes in a user's friend network.
- *Overall-Popular*. This algorithm utilizes the global network of users, as opposed to a user's local network. Items are ranked according to the number of Likes received across the whole database of stored user information.
- *Timeline*. This algorithm mimics the chronological flow of information in a user's network. Items Liked by a user's friends within the last week are selected in reverse chronological order.
- *Network-Similar*. This algorithm rates friends of a user on the basis of similarity in taste (Likes) with the given user. Each friend is ranked based on the number of Liked items in common with the user (matches in the user-item matrix), normalized by the total number of Liked items by him/her (size of the friend's item row)[2]. An item receives a score equal to the sum of the similarity scores for all friends who have liked that item. Items are ranked in descending score order.
- *Interaction-Strength*. This algorithm rates friends based on frequency of interaction, a useful proxy for tie strength [13]. Friends are given an activity score equal to the number of comments they make on the last 50 items posted on a user's Facebook wall. An item receives a score equal to the sum of the activity scores for all friends who have Liked that item. As in *Network-Similar*, the items are ranked in descending score order.

Although all of the algorithms have some dependence on time, *Interaction-Strength* and *Timeline* are the most time-

[2]This is the Tversky index [34] with $\alpha = 1$ and $\beta = 0$. It is asymmetric, thus not a formal similarity measure, but effectively makes the the user who is receiving recommendations the "gold standard" to compare against.

sensitive. *Interaction-Strength* effectively refreshes its score continuously to focus on friends whom one is currently talking to. *Timeline* is susceptible to changes made every second, and hence is the most responsive to updates in live data. These algorithms are more reflective of the network's current state and are first steps towards buzz-based algorithms as described in [35] in an e-commerce setting. *Overall-Popular* also evolves slowly as more users join PopCore.

### A. Procedure

Users were recruited through a tool for recruiting student participants for experiments and via online requests for participation through mailing lists, posting links in facebook statuses, and so on. A total of 50 users participated (27 female, 23 male); Table I presents an overview of data we collected. Participants' age (visible in 51% of cases) varied from 18 to 35, and the number of their friends ranged from 82 to 2217.

Users first saw a consent screen and instructions. We then displayed recommendations from each of the algorithms in one of the five PopCore tabs in Figure 2. The *Network-Popular*, *Network-Similar*, *Interaction-Strength*, and *Overall-Popular* algorithms ran throughout the experiment. *Timeline* was replaced halfway through the experiment by *Network-Random* because *Timeline* often displayed very few recommendations due to a lack of Like activity by friends. The assignment of algorithms to tabs was randomized to minimize order bias. Evaluating items was not mandatory; users could choose to take an action on a recommendation or skip it. Participants were instructed to complete a questionnaire about their experience by pressing the "Feedback" button after they finished looking through the algorithms.

| Category | Mean | (Min, Max) |
|---|---|---|
| Age | 22.4 | $(18, 35)$ |
| Friends | 620 | $(82, 2217)$ |
| Movie Likes | 9.9 | $(0, 193)$ |
| Book Likes | 4.8 | $(0, 62)$ |
| TV Likes | 3.1 | $(0, 59)$ |

TABLE I
OVERVIEW INFORMATION OF DATA COLLECTED FROM FACEBOOK. AMONG THE THREE DOMAINS, MOVIES HAD THE MOST LIKED ITEMS.

### B. Measures

We define "usefulness" in three different ways. The first was to recommend items people would prefer, which we measured by looking at the average rating for each algorithm. The second was to make recommendations people could engage with, which we measured by looking at the total number of *actions* (ratings + Likes + dislikes) people took. The third was to make recommendations that people had strong reactions to, which we measured as (Likes + dislikes). The second and third metrics are inspired by the goal of many social media sites to generate user-created content and the role of expressing preferences as a tool for managing one's identity in the site.

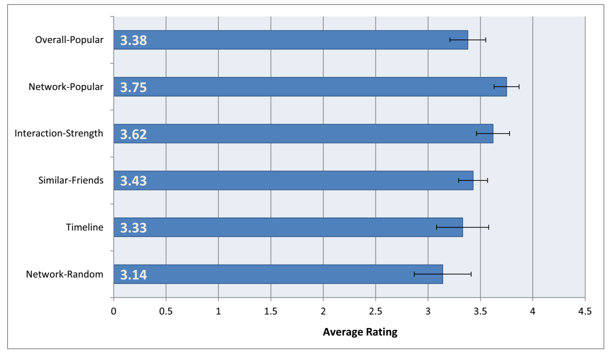


Fig. 3. Average rating for each algorithm (on a scale from 0.5-5). Error bars shown with 95% confidence interval. *Network-Popular* performs significantly better than the other algorithms.

The questionnaire had three main sections. The first focused on PopCore's usability. The second asked about the quality of recommendations, as well as how people thought they would like suggestions from each of our algorithms (we did not explain them during the experiment). The third asked more broadly about network-centric recommendation, including reactions to public versus private profiles and their overall desire for network-centric recommendations. Of the 50 participants, 33 completed the questionnaire.

### C. Results

Recommendations received a total of 1369 actions: 1136 ratings, 113 Likes, and 120 dislikes[3]. The overall average rating of rated items was 3.56. Most participants described the recommendations as generally useful, calling them *"pretty accurate"* or *"fairly good"*. However, some participants also noted that the recommendations mostly reflected *"popular and known classics"* and did not work as well in the case of books.

We first look at which algorithms tended to recommend highly-rated items. Figure 3 shows the average rating for each algorithm with 95% confidence intervals. A one-way within-subjects ANOVA showed that ratings differed significantly across the algorithms ($F(5, 108) = 5.16, p = 0.00027$). Post hoc Tukey HSD tests showed that *Network-Popular* performed better than other algorithms, $p < .05$. *Interaction-Strength* performs better than *Network-Similar* (mean 3.62 versus 3.43), but the difference is not statistically significant. Among the popularity-based algorithms, *Network-Popular* performed significantly better than *Overall-Popular*.

We now examine engagement, or how often users took action on each algorithm's recommendations, shown in Table II. A recommendation is considered as *shown* if the user viewed the tab containing that recommendation. Popularity-based algorithms (*Network-Popular* and *Overall-Popular*) have the highest action percentage, personalized algorithms (*Network-*

[3]Sometimes, multiple algorithms suggested the same item. Actions on those items were credited to all algorithms that recommended the item. Users made 1245 distinct actions (1025 ratings, 111 Likes, and 109 dislikes).

| Algorithm | Shown | Actions | | | |
|---|---|---|---|---|---|
| | | Total | Ratings | Likes | Dislikes |
| *Overall-Popular* | 373 | 255 (68.3%) | 209 | 20 | 26 |
| *Network-Popular* | 663 | 410 (61.8%) | 356 | 30 | 24 |
| *Interaction-Strength* | 561 | 260 (46.3%) | 215 | 20 | 25 |
| *Network-Similar* | 648 | 297 (45.8%) | 235 | 30 | 32 |
| *Timeline* | 173 | 67 (38.3%) | 60 | 2 | 5 |
| *Network-Random* | 350 | 80 (22.9%) | 61 | 11 | 8 |
| Overall | 2768 | 1369 (49.5%) | 1136 | 113 | 120 |

TABLE II
AGGREGATE STATISTICS FOR USER ACTIONS. POPULARITY-BASED ALGORITHMS ENABLE THE MOST ACTIONS; NON-PERSONALIZED ALGORITHMS HAVE THE FEWEST.

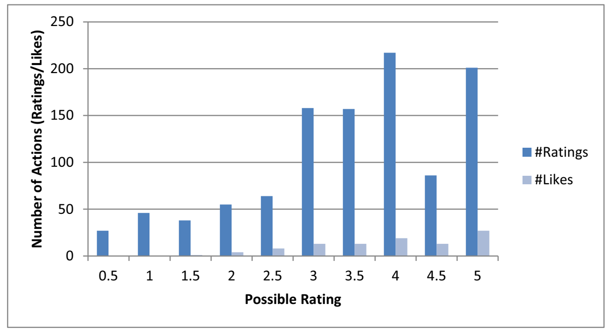


Fig. 4. Frequency of ratings and Likes spread over all possible rating values (0.5-5). Likes are comparatively rarer, and users tend to Like those items which they rate highly.

*Similar* and *Interaction-Strength*) are in the middle, and non-personalized algorithms (*Network-Random* and *Timeline*) have the lowest response. Among the network-centric algorithms, *Network-Popular* has the highest action percentage.

The relationship between popularity and action is also seen in the number of Likesan item has across all of Facebook, with *Network-Popular* fetching the most popular items (average number of Likes over $4M$) among the network-centric algorithms, while *Interaction-Strength* ($2.52M$) and *Network-Similar* ($2.57M$) having intermediate values. *Overall-Popular* reported the highest popularity overall at $14.1M$, and *Timeline* and *Network-Random* had averages of $1.8M$ and $0.9M$ respectively.

Explicit liking and disliking actions, also shown in Table II, give ameasure of particularly strong reactions to a recommendation. Strong reactions were relatively rarer (16.9% of total actions), with Dislikes being about as common as Likes (8.7% versus 8.2%). Out of the 113 items Liked, 98 items were also rated by the same user. A histogram of the number of Likes and ratings binned with rating values (Fig. 4) shows that users tended to Like those items which they also rated highly, echoed by one user's choosing to Like *"the ones I like the most (as far as books and movies go)"*. Over 60% of Liked items are rated 4.0 or above. However, Fig. 4 also shows that only about 12% of the highly rated items are Liked. We posit that the relative rarity of Likes compared to high ratings may be because of the identity aspects of publicly liking an item, or an attempt at spam prevention (posts about a Liked item often show up in a user's newsfeed).

A chi-square test comparing the proportion of strong reactions to rating showed that some algorithms generate stronger reactions than the others ($\chi^2(5, 1369) = 12.25, p = 0.032$). *Network-Similar* and *Interaction-Strength* had more strong reactions than *Network-Popular*, ($\chi^2(1, 967) = 6.207, p = 0.013$). We do not interpret *Network-Random* or *Timeline* because they have less data[4].

[4]We also ran the test with just the four more active algorithms; the differences are still unlikely to be by chance ($\chi^2(3, 1222) = 7.653, p = 0.054$).

## V. DISCUSSION

Overall, network-centric algorithms can be a viable approach to recommendation. As with [7], [8], we find that network ties provide a useful way to choose potential content and neighbors for information filtering tasks, complementing (and in some ways, being better than) item-based similarity measures. The behavioral and questionnaire data shed some light on a number of the questions we posed earlier around the design space for algorithms, the need to understand the dynamics around identity management and privacy in social networks, the value of new metrics and behavioral signals, and the potential for supporting directed recommendations.

**Algorithmic tradeoffs.** As one might expect, the popularity-based algorithms, *Network-Popular* and *Overall-Popular*, tended to recommend items that people were able to act on more often. *Network-Popular* leads to recommendations that are significantly better-rated than *Overall-Popular*, suggesting that there is value in tapping into local knowledge. Its recommendations are also better-liked and acted on more often than the personalized network algorithms of *Network-Similar* and *Interaction-Strength*. On the other hand, people were more likely to have strong reactions to the personalized algorithms. Generating activity—particularly visible activity—is likely to be important for network-centric recommenders, as a driver for use and influence, and as a tool to help create social interactions around both the items and the system itself.

Fine-tuning the information considered by *Interaction-Strength* and *Network-Similar* might improve their ability to choose items people could act on while retaining their ability to create strong reactions. *Interaction-Strength* could take in richer factors, such as length of interaction, intensity of words used in communication, number of mutual friends etc. to better estimate the tie strength of friends [13], while *Network-Similar* could incorporate demographic information and other attributes of users to choose better neighbors [36].

We explicitly measured popularity, rateability, and likeability of the recommendations. Other factors also mattered, such as novelty. Because the algorithms tended to recommend more popular movies, they were not so helpful for discovering new content: *"I had seen a bunch of them and I liked the* *movies that I had seen, but there was nothing suggested that I hadn't seen that made me say 'Oh, I want to go see that now!'"* We thought the kind of "buzz-based" algorithm that *Timeline* represents would inject novelty and were surprised at how badly it performed. We first thought it might be *too* novel, presenting unfamiliar recommendations without the face validity that builds users' trust [37]. Instead, a combination of relatively infrequent Liking behavior and a tendency to Like older content caused *Timeline* to have little novel content. Questionnaire respondents were also skeptical of *Timeline*. We asked how often people would like recommendations of items "Most recently liked or rated by your friends" (*Timeline*), "Most popular among your friends with similar tastes" (*Network-Similar*), "Most popular among your closest friends" (*Interaction-Strength*), or "Suggested manually by your friends" (directed suggestions). On a five-point Likert scale where 1 is worst, the average rating for *Timeline* was 2.87, almost a full point below the other three; a chi-square test that binned responses into negative (1-2), neutral (3), and positive (4-5) bins showed that *Timeline* is significantly less interesting ($\chi^2(6, 30) = 13.16, p = 0.04$).

Since our main thrust was to evaluate network-based signals, we explicitly chose to avoid content information in our algorithms, although users wished we had provided more content information to explain recommendations and it would be interesting in future work to use DBPedia information or rating profiles in other services such as Netflix to make hybrid recommendations [4]. We also do not compare against classic collaborative filtering algorithms with large datasets, although that too is interesting future work[5].

**Privacy and identity.** Questionnaire data also shed some light into people's practices and beliefs around privacy- and identity-related aspects of network-centric recommendations. We might have predicted the problems with *Timeline* if we had had more insight into how people Like items. Liking is not usually planned; insteaed, it arises spontaneously through other activities: *"When a friend likes something similar and it shows up in someone's newsfeed"*, *"If I am on a webpage describing [an item] I really liked and the like button is easily accessible."* For a few, it is an occasional, batch process: *"I think every once and a while I go on a spree where I overhaul my likes, adding new ones and sometimes deleting them."* These Likes are often reserved for things that are of special value: *"I like only exceptional stuff."* None of this worked in *Timeline*'s favor.

Respondents were about evenly split about the value of having a public profile associated with PopCore. It was seen as useful *"to quickly search for relevant movies to watch from the profile of [people] with the same liking"* or for *"using Facebook for pre-dating research."* This last quote perhaps explains why many people *"think it's better to have it private."* Others interpreted our question to mean allowing their friends to see PopCore's recommendations for them, which was worrisome because they feared inaccurate portrayals: *"The recommendations that I got were ok, but not so point on that I think it would represent me, and I don't see a reason for my friends to see an only ok representation of what I like on Facebook."* The recommendations themselves also brought privacy concerns to mind for some users: *"Sounds very freaky how some of the things came on to my page! I thought I had never uttered them on Facebook!!!!!!! I shall be more careful henceforth."*

**Better metrics and signals.** We used the average rating in the same way that one might test an algorithm's ability to make accurate predictions against a static dataset. This misses some of the potential of online experiments with live users. Recommendations are not just predictions, they are decision support tools, and metrics that explore how they affect actual behavior might be a useful complement to offline experiments and metrics [38]. These behavioral signals could give insight into actual use practices and support the study of dynamic processes around influence and diffusion.

For instance, PopCore could have a Netflix queue-like tool for remembering items a user intends to see. Adding recommended items to the queue becomes a powerful signal that the recommendations actually affected users' behavior, helping solidify an intention to consume the items someday. By helping people manage the recommendations they receive, such a queue might encourage users to return more often to PopCore, providing both practical and experimental benefit. The Like button already allows user tastes to be diffused through the network; it would also be interesting to think about what the effects of a dislike button might be for supporting and studying identity management as well as how diffusion processes might work when both positive and negative sentiment are in play.

**Directed suggestions.** Participants' reactions suggested that the overall recommendations had a word-of-mouth feel: *"I would say the recommendations aligned fairly well with my tastes (and those of my friends). There were several movies and show suggestions friends have made in person, that also showed up in PopCore."*, *"I was surprised to see Hindi movies in recommendations. Although I haven't seen many of the movies, I have overheard my friends discussing them."*

Thus, allowing people to make user-initiated directed suggestions to friends seems to be a useful addition to network-centered recommenders and to PopCore. We will need to be careful to avoid the problem of spamming people's news feeds that some applications suffer (Lonely Cow, anyone?), but done right, these recommendations would provide strong behavioral signals, reveal latent knowledge of people's preferences, and perhaps recruit new users who would participate in experiments. Growing PopCore is crucial for making it an ongoing platform that we, and other researchers, can use in the future.

## VI. Conclusion

In this paper, we have discussed a number of useful ideas, theories, and questions to motivate doing network-centric recommendation research. We presented PopCore, our platform

[5]These datasets also are usually single-domain, while we are combining movie, book, and television likes and recommendations.

for doing this research in a real setting. Our initial study shed light on properties of network-centric algorithms and how they might support or hinder users' enagement, as well as revealing information about people's behaviors and attitudes toward recommendations in social networks.

Thinking about how recommendation changes when it is embedded in social networks raises questions and opportunities for research around algorithm design, social networks, social science more generally, and interface design. Network-centric recommendations provide a domain that invites looking at the bigger picture of how recommendations interact with the world. How does interface design affect influence processes? How might accounting for influence affect algorithm design? And how do the recommendations generated by those algorithms in turn affect the people and the networks that receive them? As is often the case, the most interesting questions arise at the intersections. We hope other researchers will join us in crossing them.

## ACKNOWLEDGEMENT

We would like to acknowledge Meethu Malu for her design work and NSF support via grants 0835451 and 0910664.

## REFERENCES

[1] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, "GroupLens: an open architecture for collaborative filtering of netnews," in *Proc. CSCW*, 1994, pp. 175–186.

[2] R. M. Bell and Y. Koren, "Lessons from the Netflix prize challenge," *SIGKDD Explor. Newsl.*, vol. 9, pp. 75–79, December 2007.

[3] G. Linden, B. Smith, and J. York, "Amazon.com recommendations: Item-to-item collaborative filtering," *IEEE Internet Computing*, vol. 7, pp. 76–80, 2003.

[4] R. Burke, in *The adaptive web*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds. Berlin, Heidelberg: Springer-Verlag, 2007, ch. Hybrid web recommender systems, pp. 377–408.

[5] H. Ma, H. Yang, M. R. Lyu, and I. King, "SoRec: social recommendation using probabilistic matrix factorization," in *Proc. CIKM*, 2008, pp. 931–940.

[6] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King, "Recommender systems with social regularization," in *Proc. WSDM*, 2011, pp. 287–296.

[7] 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, pp. 160–168.

[8] 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*, 2009, pp. 53–60.

[9] C. Basu, H. Hirsh, and W. Cohen, "Recommendation as classification: using social and content-based information in recommendation," in *Proc. AAAI*, 1998, pp. 714–720.

[10] X. Su and T. M. Khoshgoftaar, "A survey of collaborative filtering techniques," *Adv. in Artif. Intell.*, vol. 2009, pp. 4:2–4:2, January 2009.

[11] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual Review of Sociology*, vol. 27, no. 1, pp. 415–444, 2001.

[12] I. Konstas, V. Stathopoulos, and J. M. Jose, "On social networks and collaborative recommendation," in *Proc. SIGIR*, 2009, pp. 195–202.

[13] E. Gilbert and K. Karahalios, "Predicting tie strength with social media," in *Proc. CHI*, 2009, pp. 211–220.

[14] H. Ma, I. King, and M. R. Lyu, "Learning to recommend with social trust ensemble," in *Proc. SIGIR*, 2009, pp. 203–210.

[15] J. Golbeck and J. Hendler, "Inferring binary trust relationships in Web-based social networks," *ACM Trans. Internet Technol.*, vol. 6, pp. 497–529, November 2006.

[16] B. Liu and Z. Yuan, "Incorporating social networks and user opinions for collaborative recommendation: local trust network based method," in *Proc. the Workshop on Context-Aware Movie Recommendation*, ser. CAMRa '10, 2010, pp. 53–56.

[17] P. Singla and M. Richardson, "Yes, there is a correlation: - from social networks to personal behavior on the web," in *Proc. WWW*, 2008, pp. 655–664.

[18] J. Leskovec, A. Singh, and J. Kleinberg, "Patterns of influence in a recommendation network," in *Advances in Knowledge Discovery and Data Mining*, W.-K. Ng, M. Kitsuregawa, J. Li, and K. Chang, Eds. Springer Berlin / Heidelberg, 2006, vol. 3918, pp. 380–389.

[19] D. Cosley, S. K. Lam, I. Albert, J. A. Konstan, and J. Riedl, "Is seeing believing? How recommender system interfaces affect users' opinions," in *Proc. CHI*, 2003, pp. 585–592.

[20] C. Danescu-Niculescu-Mizil, G. Kossinets, J. Kleinberg, and L. Lee, "How opinions are received by online communities: a case study on amazon.com helpfulness votes," in *Proc. WWW*, 2009, pp. 141–150.

[21] M. V. Alstyne and E. Brynjolfsson, "Electronic communities: global village or cyberbalkanization?" in *Proc. ICIS*, 1996, pp. 80–98.

[22] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Trans. Inf. Syst.*, vol. 22, pp. 5–53, January 2004.

[23] R. S. Burt, *Structural Holes: The Social Structure of Competition*, 1992.

[24] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock, "Methods and metrics for cold-start recommendations," in *Proc. SIGIR*, 2002, pp. 253–260.

[25] P. Brandtzag, A. Folstad, and J. Heim, "Enjoyment: Lessons from Karasek," in *Funology*, ser. Human-Computer Interaction Series, M. Blythe, K. Overbeeke, A. Monk, and P. Wright, Eds. Springer Netherlands, 2005, vol. 3, pp. 55–65.

[26] J. F. McCarthy and T. D. Anagnost, "MusicFX: an arbiter of group preferences for computer supported collaborative workouts," in *Proc. CSCW*, 1998, pp. 363–372.

[27] M. O'Connor, D. Cosley, J. A. Konstan, and J. Riedl, "PolyLens: a recommender system for groups of users," in *Proc. ECSCW*, Norwell, MA, USA, 2001, pp. 199–218.

[28] C. Schwind, J. Buder, and F. W. Hesse, "I will do it, but I don't like it: user reactions to preference-inconsistent recommendations," in *Proc. CHI*, 2011, pp. 349–352.

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

[30] M. Naaman, J. Boase, and C.-H. Lai, "Is it really about me? Message content in social awareness streams," in *Proc. CSCW*, 2010, pp. 189–192.

[31] J. L. Herlocker, J. A. Konstan, and J. Riedl, "Explaining collaborative filtering recommendations," in *Proc. CSCW*, 2000, pp. 241–250.

[32] A. Voida, R. E. Grinter, N. Ducheneaut, W. K. Edwards, and M. W. Newman, "Listening in: practices surrounding iTunes music sharing," in *Proc. CHI*, 2005, pp. 191–200.

[33] D. Frankowski, D. Cosley, S. Sen, L. Terveen, and J. Riedl, "You are what you say: privacy risks of public mentions," in *Proc. SIGIR*, 2006, pp. 565–572.

[34] A. Tversky, "Features of similarity," *Psychological Review*, vol. 84, pp. 327–352, 1977.

[35] N. Parikh and N. Sundaresan, "Buzz-based recommender system," in *Proc. WWW*, 2009, pp. 1231–1232.

[36] 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*, ser. CSCW '10. New York, NY, USA: ACM, 2010, pp. 41–50.

[37] K. Swearingen and R. Sinha, "Beyond Algorithms: An HCI Perspective on Recommender Systems," in *ACM SIGIR Workshop on Recommender Systems*, New Orleans, USA, 2001.

[38] D. Cosley, S. Lawrence, and D. Pennock, "REFEREE: an open framework for practical testing of recommender systems using researchindex," in *Proc. VLDB*, Hong Kong, China, 2002, pp. 35–46.