Mining User Behavior in Social Recommender Systems

Carine Mukamakuza
TU Wien
carine@ec.tuwien.ac.at

Dimitris Sacharidis
TU Wien
dimitris@ec.tuwien.ac.at

Hannes Werthner
TU Wien
werthner@ec.tuwien.ac.at

ABSTRACT

Social recommender systems make use of the available information about social connections between users to improve the quality of the recommendations. The assumption is that if two users are connected, they are likely to have similar preferences, and thus the system should make similar recommendations. Recently many approaches have been proposed based around similar assumptions, whose validity however has not been systematically studied. In our work we make the first step towards examining whether there exist observable relationships between social connections and rating behavior in social recommenders. In particular, we examine publicly available datasets containing traces of rating behavior along with a social graph. Using techniques from social network analysis and statistics, we investigate whether heavy rates, having provided feedback on many items, are also popular, i.e., central in the social network, and vice versa. Our results indicate important connections between heaviness and popularity. Specifically, we find that heaviness implies popularity, and that the association is stronger among very heavy raters.

CCS CONCEPTS

• **Information systems** → **Social recommendation**; **Recommender systems**;

1 INTRODUCTION

The information overload is a characteristic of our society, making it complicated to make choices about what to consume, e.g., movie to watch, a restaurant to visit. Recommender systems attempt to help people take decisions by exploiting their stated preferences and the past behavior of them and other similar people. Conventionally, before the advance of such services, people used to resort to their social connections to seek for advice and expert opinions.

The key idea of *Social Recommender Systems* is to enhance recommendations by also drawing information from the social context of the user. The underlying assumption is that for a particular item, the decision process of a user depends on her individual preferences, but also on interpersonal influence from her social connections. Influential people may strongly affect the decisions of a person. Thus the structure of social network is important in trying to understand the social effect and the extent of its impact.

WIMS '18, June 25–27, 2018, Novi Sad, Serbia

ACM ISBN 978-1-4503-5489-9/18/06…$15.00
https://doi.org/10.1145/3227609.3227651

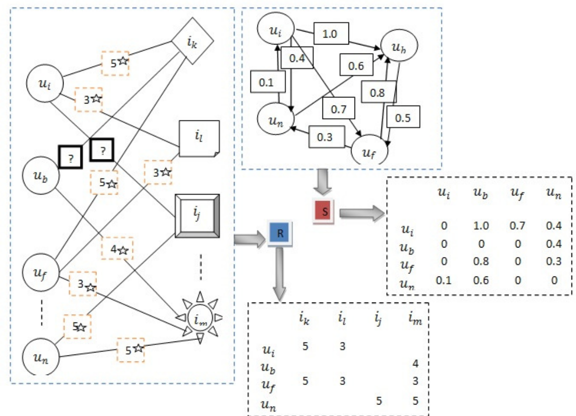

Figure 1: An example of user-item rating behavior and user-user social connections

In this work, we study *collaborative filtering* (CF) based social recommenders that draw information from two components. The first is the *rating behavior* represented by the ratings matrix, where each existing entry corresponds to the rating given to an item by user. The second component is the *social connections* conveyed by the social adjacency matrix, where entries portray the friendship strength between users. Social recommenders predict ratings using these two matrices, under the assumption that a user's behavior is influenced by her social connections.

An example of a social recommender is shown in Figure 1, which depicts the rating behavior of users, denoted as u_i, on items, denoted as i_j, on the left, and the social connections among users on the right. The former is captured by the rating matrix R, where a non empty entry R_{ij} corresponds to the rating given by user u_i on item i_j. The latter is conveyed by the social adjacency matrix S, where entries portray the friendship strength between users. Social recommenders draw upon information from both matrices to predict ratings, under the assumption that a user's behavior is influenced by her social connections.

Social recommendations is an active research area in the past few years, but relatively young. Existing approaches suffer from certain weaknesses, and often make explicit assumptions about the impact of social ties that they never validate[17]. They also fail to take into account the structure (local and global) of the social network, and how much an impact it has on the rating behavior. Our main ambition is the formulation and statistical analysis of the impact that social connections have in rating behavior at different levels. Can we predict how users rate items, and to what extent, purely by observing their position in the social network, and vice versa? An additional contribution is the theoretical evaluation of the assumptions made by state of the art social recommenders, and whether they hold in various domains. Ultimately, we would have a better understanding of what aspects of social connections exactly affect rating behavior. This will bring us initial ideas towards a more realistic model for social recommendations, based on observed and quantifiable types of social influence.

As a first important step towards our vision, in this paper we study a specific research question. Is there a connection between *heavy raters*, who have made a large number of ratings, and *popular users*, who have acquired many social connections in the system? To answer this question we go both ways, looking whether heavy users are popular and vice versa. Specifically, we employ techniques from social network analysis to determine different interpretations of "*popularity*" based on network centrality. On the other hand, the "*heaviness*" of a user has a single interpretation, the number of her ratings. Then, we seek for correlations between popularity and number of ratings in multiple ways. Our evaluation positively answers this question, suggesting a bilateral connection between very heavy and very popular users. In addition, we find strong evidence suggesting that heaviness implies popularity. Among studied popularity interpretations, the simple degree centrality has shown stronger correlations to the number of ratings a user has.

The remainder of this paper is structured as follows. Section 2 establishes the necessary background and overviews existing work. Section 3 describes methodology, Section 4 presents experimental results, while Section 5 concludes.

2 RELATED WORK

Collaborative Filtering. Social recommender systems borrow ideas from *Collaborative Filtering* (CF), which is the most commonly used method for making recommendations. In CF approaches, users and items with similar rating patterns are taken into account [8] to produce a recommendation for the target user.

The basic entity in CF is the *user-item ratings matrix*, composed of a set of items $I = i_1, \ldots i_n$ and a set of users $U = u_1, \ldots u_m$. The ratings matrix $R \in \mathbb{R}^{n \times m}$ contains the ratings given by users to items, where n represents the number of items and m number of users. CF exclusively uses the ratings in R to make recommendations.

Memory-based methods for CF are divided into two categories. *User-user* techniques make the assumption that users had similar tastes in the past they are most likely to have the same tastes in the future, i.e., user preferences then to remain constant and stable over the time. Then to predict ratings of a target user, they utilize the ratings to the target item by a set of the users whose similarity level is closer to the target user, the neighborhood. On the other hand, *item-item* methods uses the target user's profile to compute the target item's similarity to other items rated by the target user.

Model-based methods make predictions by learning parameters describing how ratings are generated. The most famous is the *Matrix Factorization* (MF) technique [9, 10, 19]. In its simplest incarnation, MF computes a low-rank approximation of the sparse ratings matrix R by multiplication of two matrices.

Social Recommender Systems. Social recommenders (SRS) operate similar to collaborative filtering systems but differ in that they make recommendations taking into account the social connections between users. That is, SRS make use of the ratings matrix R and the social adjacency matrix S. In the following, we review the most important related work, differentiating between *memory-based* SRS and *model-based* SRS. For an overview of this research area and other associated topics, we refer the reader to [7].

Memory-based social recommenders apply ideas from memory-based collaborative filtering to combine information from the social graph and the past user behavior. In Trust-aware Recommender systems (TaRS) [16], the idea is to treat the social neighborhood of the target user in a manner similar to the rating neighborhood in user-based CF. Following TaRS, several works have recently appeared. An experimental evaluation of several memory-based social recommenders is provided in [2]. The authors also propose to fuse recommendations from friends with recommendations from implicit social relations, and show that such an approach improves accuracy and increases coverage.

Homophily in social networks refers to the notion that similar users tend to be socially connected and vice versa. In the context of social recommenders, the work in [3] studies homophily on two online social media networks, BlogCatalog and Last.fm.by extracting communities based on the network ties. Similarly, [1] investigates the presence of homophily in three systems that combine tagging social media with online social networks.

The other important category in SRS is *model-based social recommenders*, where model-based collaborative filtering, and predominantly matrix factorization, approaches are used. On of the first works in this direction is SoRec [14] that extends the basic MF model to incorporate the social network. The social adjacency matrix S is factorized into a user-specific matrix U and a factor-specific matrix F, where matrix U is also part of the factorization of the ratings matrix. The latent feature vectors of users are then learnt based on both the rating and social network matrices.

Social trust ensemble [13] builds on the hybrid idea of [16], and defines a linear combination of basic MF predictions with social network predictions. The basic idea in Social Regularization [15] is to use the basic MF formula for predicting ratings, but use regularization terms to force the learned user feature vectors to be similar between friends. SoCo [12] combines contextual information and social network information to improve quality of recommendations. In the community-based models of [11], the idea of social regularization is taken one step further. A target user can belong to different communities and they should be regularized differently.

The common assumption in all of the related work is that if two people are socially connected then they must have similar preferences. So the proposed methods enforce this assumption in their recommendation strategy. However, it is more reasonable to expect that the aforementioned assumption tends to hold but not in every case. Our goal is to investigate the validity of this assumption, and understand when it is likely to hold and when not. For this purpose we need to first identify when social connections, particularly *popularity*, imply specific rating behavior, particularly *heaviness* in terms of number of ratings, and vice versa. The findings of our work can then be applied to design more realistic and effective social recommenders — a task we defer to the future.

3 DATA AND METHODOLOGY

Section 3.1 presents the datasets used in our study, and defines the terms popularity and heaviness. Then Section 3.2 overviews the methodology we follow to answer our research question.

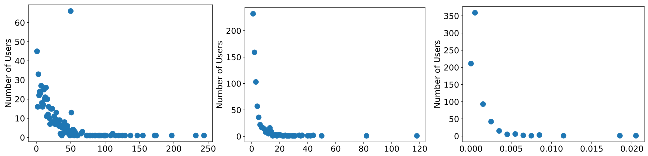

(a) NumRatings **(b) Degree** **(c) PageRank**

Figure 2: FilmTrust: Probability distribution of a user having specific values of NumRatings, Degree, and PageRank

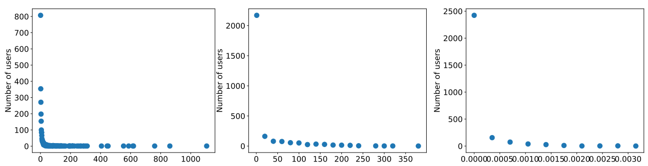

(a) NumRatings **(b) Degree** **(c) PageRank**

Figure 3: CiaoDVD: Probability distribution of a user having specific values of NumRatings, Degree, and PageRank

3.1 Datasets

In our study, we use two publicly available datasets collected from traces of user interaction in social recommenders. These data are commonly used in the literature and contain rating activity, i.e., a ratings matrix R, as well as information about the social connections among users, i.e., an adjacency matrix S.

The first dataset, FilmTrust[6], comes from a social networking site in which users can rate and review movies.[1] FilmTrust essentially contains two sub-datasets, a social network in addition to the user-item ratings. The social connections are bidirectional and capture the trust between users (trustee, trustor). Users can specify a level of trust, but due to sharing policy, we only know whether a connection exists.

FilmTrust contains 1,508 users, 2,071 items, 35,497 ratings, and 1,853 social connections. As there exist 635 users with no social connections, and 133 with no ratings history, we exclude them from our analysis. That is we only consider the 740 users that have rated at least one item and trust, or are trusted by, at least another person. The mean number of ratings per user is 23.5 with the minimum and the maximum being 1 and 244. The ratings scale is from 0.5 to 4 with a step 0.5, and the mean rating score over all ratings is 3.0.

The second dataset CiaoDVD [5], is collected on the Ciao website.[2] CiaoDVD contains the social connections among its users. Compared to FilmTrust, CiaoDVD is about an order of magnitude larger, having 17,588 users, 16,121 items, 72,665 ratings, and 40,133 social connections. However, there exist 12,930 users with no social connections, and 1,918 with no ratings history. We exclude them from our analysis, and study the 2,620 users with both pieces of information. The mean number of ratings per user is 12.57 with the minimum and the maximum being 1 and 1,106 respectively. The ratings scale is from 1 to 5, and the mean rating score over all ratings is 4.07.

This work investigates whether there exists a relationship between heavy raters and popular users. We define "*heaviness*" of a user in terms of the number of ratings (NumRatings) s/he has provided. Moreover, we define "*popularity*" of a user as the *centrality* of the node representing the user in the social graph. Every person has some degree of influence or importance within the social domain under consideration, and one expects such importance to surface in the structure of the social network; centrality is a quantitative measure that aims at revealing the importance of a node [4]. Here we consider two definitions of centrality.

Degree is the most intuitive interpretation of popularity, as it counts the number of (incoming or outgoing) connections a user has. In terms of the adjacency matrix S, the Degree of user u_i is

$$d_i = \sum_{k=1}^{m} (S_{ki} + S_{ik}).$$

PageRank [18] depends on the number of incoming connections of a user as well as their quality, with higher centrality users giving more importance to their outgoing connections; in some sense, the higher its PageRank is the more respected a user is. In terms of the adjacency matrix S, PageRank satisfies the equation

$$x_i = \alpha \sum_{k=1}^{m} \frac{S_{ki}}{\max\{d_k^{out}, 1\}} x_k + \frac{1-\alpha}{m},$$

where $d_i^{out} = \sum_{k=1}^{m} S_{ki}$ is the out-Degree centrality of user i, and α is the damping factor, typically set to 0.85.

Figure 2 shows three probability distributions in Filmtrust. First, Figure 2a depicts the probability (in raw numbers) of a user being heavy, i.e., giving a specific number of ratings, which we hereafter refer to as NumRatings. Then, Figure 2b shows the probability of a user being popular in terms of Degree; the mean Degree is 4.7, with min and max values of 1 and 118. Figure 2c draws the probability of a user being popular in terms of PageRank; the mean PageRank is 0.0012, with the min and max values of 0 and 0.21. These right-skewed distributions show that the majority of users give few ratings and have low centralities, and that at the same time there exist several users that are very heavy and very popular. Figure 3 presents the same distributions for CiaoDVD. The mean Degree is 21.75, with min and max values of 1 and 349, while the mean PageRank is 0.000241, with min and max values 0, 0.003440.

3.2 Methodology

To establish whether a relationship between heaviness and popularity exists, we consider two approaches. The first, termed *partitioning*, divides users into three groups A, B, C according to one of the attributes, either heaviness (NumRatings) or popularity (Degree and PageRank); the selected one is called *partitioning attribute*. These partitions contain roughly the same number of users (i.e., about one third), with partition A having users with low values in the partitioning attribute, while partition C consists of users with high values in the partitioning attribute. In each partition, we compute the mean of the other (non-partitioning) attribute; e.g., for partitions on heaviness, we compute the average popularity (Degree and PageRank). Then, we apply ANOVA to investigate whether the mean is significantly different across partitions. If that is the case, we further investigate whether the mean increases from partitions A through C. For this purpose, we perform the Tukey

[1] http://trust.mindswap.org/FilmTrust
[2] http://dvd.ciao.co.uk

Table 1: FilmTrust: Description of Partitions

	A	div	B	div	C
NumRatings	242	11	267	30	231
Degree	232	1	262	4	246
PageRank	241	5.4×10^{-4}	268	1.1×10^{-3}	231

Table 2: CiaoDVD: Description of Partitions

	A	div	B	div	C
NumRatings	808	2	978	6	957
Degree	899	2	930	6	914
PageRank	913	1.3×10^{-4}	913	1.6×10^{-4}	914

HSD test to check every pair of partitions and see if the difference of their means is significant.

The second approach, termed *ranking*, ranks users according to one of the attributes, either heaviness (NumRatings) or popularity (Degree and PageRank); the selected one is called *ranking attribute*. A simple way to determine connections between heaviness and popularity, is to count the number of common users appearing in the top positions according to each ranking attribute; the higher the number of common users, the stronger the relationship. Another way is to simply compute the (Pearson or Spearman) correlation coefficient of the heaviness and popularity attributes of users. However, since correlations across the entire set of users are most likely to be very low, it makes more sense to restrict the set of users considered. Therefore, we select either the very heavy (i.e., top-100 users by NumRatings), or the very popular (i.e., top-100 users by Degree or PageRank), and compute the correlation on this subset.

4 RESULTS

We present the results from applying the aforementioned methods.

4.1 Partitioning

To assess the relationship between NumRatings and centralities, we consider three distinct divisions, one per each attribute, NumRatings, Degree, PageRank. A division splits users into three partitions, A, B, C, in increasing value of the partitioning attribute. We first determine the lower and upper terciles (3-quantiles) of the partitioning attribute and divide accordingly. Partitions are thus balanced, with each containing roughly 1/3 of all users. Descriptions of the partitions are shown in Tables 1 and 2 for FilmTrust and CiaoDVD, respectively.

Does the mean Degree differ across NumRatings partitions? In the first experiment, we partition users according to their NumRatings, and compute the mean Degree in each partition. Then, we apply ANOVA to investigate whether the mean Degree is significantly different across partitions. The results for FilmTrust is shown in the top part Table 3, where an F value of 24.4 provides significant evidence against the hypothesis that the means are equal (p-value in the order of 10^{-11}).

Following this result, we investigate whether the mean Degree increases from partitions A through C. We apply the Tukey HSD test to check every pair of partitions and see if the difference of their mean Degree is significant. The difference of means and its corresponding 95% confidence interval for each pair are shown in the bottom part of Table 3. As suspected partitions A and B,

Table 3: FilmTrust: ANOVA and Tukey Test on Mean Degree among NumRatings Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	2942	1470.9	24.4	$\mathbf{5.18 \times 10^{-11}}$
Residuals	798	48102	60.3		

Pair	Diff. of Means	95% CI	p-value
B - A	0.861	$[0.716, 2.439]$	0.40
C - B	3.565	$[1.987, 5.143]$	$\mathbf{4 \times 10^{-6}}$
C - A	4.427	$[2.849, 6.004]$	$\approx \mathbf{0}$

Table 4: CiaoDVD: ANOVA and Tukey Test on Mean Degree among NumRatings Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	1969	984.6	340.3	$\mathbf{2 \times 10^{-16}}$
Residuals	2928	8472	2.9		

Pair	Diff. of Means	95% CI	p-value
B - A	1.08	$[-3.75, 5.92]$	0.86
C - B	19.80	$[14.96, 24.65]$	$\approx \mathbf{0}$
C - A	20.89	$[16.05, 25.73]$	$\approx \mathbf{0}$

Table 5: FilmTrust: ANOVA and Tukey Test on Mean NumRatings among Degree Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	16018	8009	11.49	$\mathbf{1.21 \times 10^{-5}}$
Residuals	783	545748	697		

Pair	Diff. of Means	95% CI	p-value
B - A	-0.313	$[-5.729, 5.103]$	0.99
C - B	9.729	$[4.317, 15.145]$	$\mathbf{8.1 \times 10^{-5}}$
C - A	9.416	$[3.91, 14.832]$	$\mathbf{1.4 \times 10^{-4}}$

containing non-heavy users, have mostly similar mean Degrees and no significant difference is observed. However, there is a significant difference when we compare either A or B with partition C of heavy raters.

Results of ANOVA and Tukey HSD test for CiaoDVD are shown in Table 4, where similar conclusions can be drawn. In general, heavy raters tend to be more popular (in terms of Degree) compared to others.

Does the mean NumRatings differ across Degree partitions? We also study the reciprocal association. The ANOVA analysis based on the mean NumRatings among partitions based on Degree is shown in Table 5, where an F value of 11.49 provides significant evidence against the hypothesis that the means are equal (p-value in the order of 10^{-5}. The Tukey HSD test shows that partitions A and B of non-popular users have mostly similar mean NumRatings and no significant difference is observed. However, there is a significant difference when we compare B with C, and of course A with C, implying that popular (in terms of Degree) users tend to be heavier raters. Results on CiaoDVD, Table 6, suggest an identical relationship.

Does the mean PageRank differ across NumRatings partitions? We repeat the previous setup, this time measuring popularity by means of PageRank. Tables 7 and 8 present the results on

Table 6: CiaoDVD: ANOVA and Tukey Test on Mean NumRatings among Degree Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	248535	124268	55.21	$\mathbf{2 \times 10^{-16}}$
Residuals	2784	6265827	2251		

Pair	Diff. of Means	95% CI	p-value
A - B	2.54	$[-2.43, 7.52]$	0.45
C - B	17.24	$[17.24, 12.26]$	$\approx \mathbf{0}$
C - A	19.79	$[14.81, 24.77]$	$\approx \mathbf{0}$

Table 7: FilmTrust: ANOVA and Tukey Test on Mean PageRank among NumRatings Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	0.0000774	3.869×10^{-5}	16.38	$\mathbf{1.06 \times 10^{-7}}$
Residuals	798	0.0018845	2.360×10^{-6}		

Pair	Diff. of Means	95% CI	p-value
B - A	0.000226	$[-8.6 \times 10^{-5}, 0.00053]$	0.20
C - B	0.000517	$[2.04 \times 10^{-4}, 0.00082]$	$\mathbf{3.2 \times 10^{-4}}$
C - A	0.000742	$[4.3 \times 10^{-4}, 0.00105]$	$\mathbf{1.0 \times 10^{-5}}$

Table 8: CiaoDVD: ANOVA and Tukey Test on Mean PageRank among NumRatings Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	384331	192165	98.39	$\mathbf{2 \times 10^{-16}}$
Residuals	2928	5718398	1953		

Pair	Diff. of Means	95% CI	p-value
B - A	-1.4×10^{-6}	$[-3.3 \times 10^{-5}, 3.0 \times 10^{-4}]$	0.99
C - B	1.2×10^{-4}	$[9.1 \times 10^{-5}, 1.5 \times 10^{-4}]$	$\approx \mathbf{0}$
C - A	1.2×10^{-4}	$[8.9 \times 10^{-5}, 1.5 \times 10^{-4}]$	$\approx \mathbf{0}$

Table 9: FilmTrust: ANOVA and Tukey Test on Mean NumRatings among PageRank Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	17266	8633	12.73	$\mathbf{3.6 \times 10^{-6}}$
Residuals	801	543138	678		

Pair	Diff. of Means	95% CI	p-value
B - A	-0.239	$[-5.521, 5.043]$	0.99
C - B	9.948	$[4.666, 15.230]$	$\mathbf{3.3 \times 10^{-5}}$
C - A	9.709	$[4.427, 14.991]$	$\mathbf{5.3 \times 10^{-5}}$

FilmTrust and CiaoDVD, respectively. The findings are similar, except with slightly lower significance: heaviness implies popularity.

Does the mean NumRatings differ across PageRank partitions? Finally, we consider PageRank partitions and study whether they contain users with significantly different NumRatings. Results are presented in Tables 9 and 10. As in the case of Degree partitions, popularity implies heaviness.

4.2 Ranking

Based on the previous findings, we seek for further connections, this time among *very heavy raters* (top-100 users according to NumRatings) or *very popular users* (top-100 users according to Degree or PageRank). For FilmTrust that corresponds to about 13% of the users, while for CiaoDVD to about 4%.

Table 10: CiaoDVD: ANOVA and Tukey Test on Mean NumRatings among PageRank Partitions

	DF	Sum. Sq.	Mean Sq.	F value	Pr (>F)
partition	2	217992	108996	51.27	$\mathbf{2 \times 10^{-16}}$
Residuals	2739	5822399	2126		

Pair	Diff. of Means	95% CI	p-value
B - A	2.038	$[-3.02, 7.09]$	0.61
C - B	17.82	$[12.76, 22.88]$	$\approx \mathbf{0}$
C - A	19.86	$[14.80, 24.91]$	$\approx \mathbf{0}$

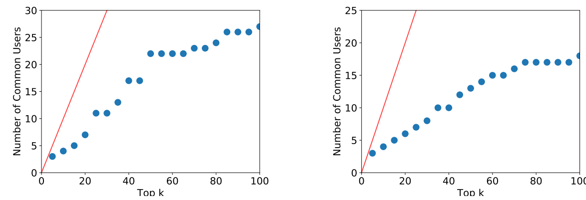

(a) NumRatings and Degree **(b) NumRatings and PageRank**

Figure 4: FilmTrust: Number of common users among the Top-K heaviest and most popular (Degree, PageRank) users

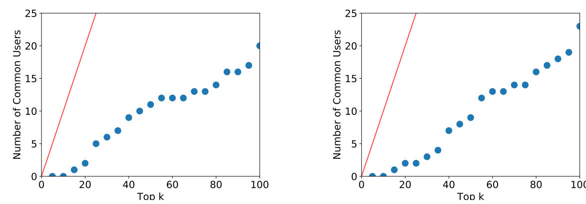

(a) NumRatings and Degree **(b) NumRatings and PageRank**

Figure 5: CiaoDVD: Number of common users among the Top-K heaviest and most popular (Degree, PageRank) users

How many common users exist among the top-100 heavy and the top-100 popular? First we consider the number of common users across these rankings, with the results shown in Figures 4 and 5 for the two datasets. We see that the number of common users increases with a much lower rate that the maximum possible (drawn as the red line). Hence there exist more common users among the really heavy and the really popular.

Are NumRatings and Degree correlated? We investigate whether heaviness and popularity (in terms of Degree) are correlated among the 100 most popular users or the 100 heaviest raters. For FilmTrust, Figure 6a shows the values of Degree and NumRatings for each user among very popular users (according to Degree), while Figure 6b shows the corresponding scatter plot for the very heavy raters. In both figures we draw the linear regression line, and also measure Pearson and Spearman's correlation coefficients. The very popular users have weak Pearson and Spearman correlation values of 0.25 and 0.27 with low significance (p-values of 0.01 and 0.07). In contrast, the very heavy users have weak Pearson but strong Spearman correlation values of 0.3 and 0.67 with high significance (p-values of 0.002 and ≈ 0).

Similar results hold for the CiaoDVD dataset, shown in Figure 7. The very popular users exhibit non-significant weak correlation between heaviness and popularity, while the correlation in very heavy users is strong (Pearson and Spearman 0.31 and 0.44) and

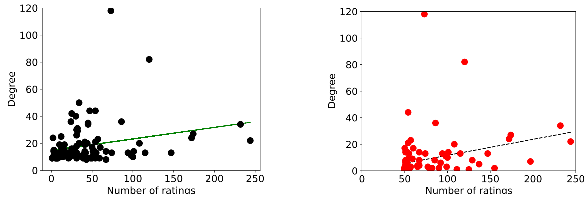

(a) Very Popular Users **(b) Very Heavy Raters**

Figure 6: FilmTrust: Scatter Plots (NumRatings, Degree)

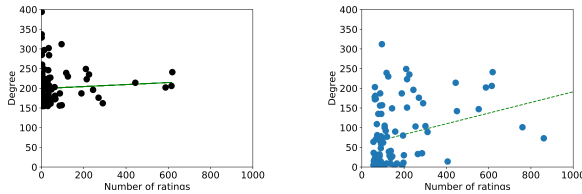

(a) Very Popular Users **(b) Very Heavy Raters**

Figure 7: CiaoDVD: Scatter Plots (NumRatings, Degree)

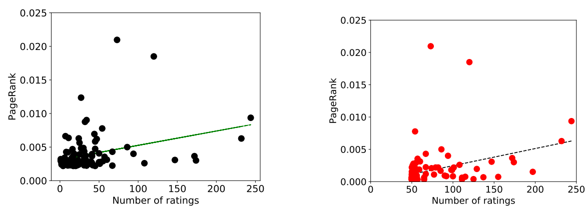

(a) VeryPopular Users **(b) Very Heavy Raters**

Figure 8: FilmTrust: Scatter Plots (NumRatings, PageRank)

significant. These results imply that (1) overall there is a weak association between NumRatings and Degree in the very popular and the very heavy raters, and (2) NumRatings and Degree are strongly correlated, in a non-linear sense, for the very heavy raters; the heavier the rater is, the more popular s/he becomes.

Are NumRatings and PageRank correlated? We repeat the previous setup but this time define popularity by PageRank. Figure 8 shows the results for FilmTrust, where the very popular users have an insignificant weak correlation among heaviness and popularity. On the other hand, the very heavy raters exhibit moderate to strong correlations (Pearson and Spearman 0.35 and 0.60) with high significance (p-values 0.004 and ≈ 0). Similar in CiaoDVD (scatter plots in Figure 9), heaviness and popularity among very heavy raters is moderately (Pearson and Spearman 0.37 and 0.45) correlated with high significance.

As a conclusion, we note that we have observed moderate to strong correlations among heavy users (top-100 by NumRatings) between their heaviness (NumRatings) and their popularity (Degree and PageRank). This correlation is not so much linear, as is rank-based (higher Spearman than Pearson correlation values).

5 CONCLUSIONS AND FUTURE WORK

This work makes the first step towards studying the effects of social connections in rating behavior in social recommenders. We have identified important strong connections between heaviness and popularity in social recommenders. In particular, the connection

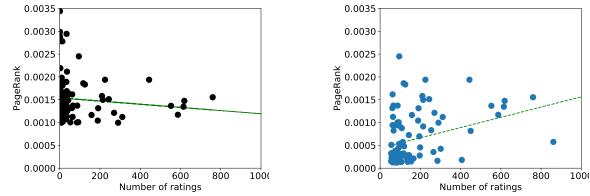

(a) VeryPopular Users **(b) Very Heavy Raters**

Figure 9: CiaoDVD: Scatter Plots (NumRatings, PageRank)

is stronger when we consider the very heavy raters, with strong evidence suggesting that heaviness implies popularity.

REFERENCES

[1] Luca Maria Aiello, Alain Barrat, Rossano Schifanella, Ciro Cattuto, Benjamin Markines, and Filippo Menczer. 2012. Friendship prediction and homophily in social media. *TWEB* 6, 2 (2012), 9:1–9:33. https://doi.org/10.1145/2180861.2180866

[2] Shaikhah Alotaibi and Julita Vassileva. 2016. Personalized Recommendation of Research Papers by Fusing Recommendations from Explicit and Implicit Social Network. In *Proceedings of the 24th ACM Conference on User Modeling, Adaptation and Personalisation (UMAP 2016)*. http://ceur-ws.org/Vol-1618/IFUP_paper_2.pdf

[3] Halil Bisgin, Nitin Agarwal, and Xiaowei Xu. 2010. Investigating Homophily in Online Social Networks. In *2010 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2010, Toronto, Canada, August 31 - September 3, 2010, Main Conference Proceedings*. 533–536. https://doi.org/10.1109/WI-IAT.2010.61

[4] Paolo Boldi and Sebastiano Vigna. 2014. Axioms for Centrality. *Internet Mathematics* 10, 3-4 (2014), 222–262. https://doi.org/10.1080/15427951.2013.865686

[5] Guibing Guo, Jie Zhang, Daniel Thalmann, and Neil Yorke-Smith. 2014. ETAF: An extended trust antecedents framework for trust prediction. In *2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2014*. 540–547. https://doi.org/10.1109/ASONAM.2014.6921639

[6] Guibing Guo, Jie Zhang, and Neil Yorke-Smith. 2013. A Novel Bayesian Similarity Measure for Recommender Systems. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence IJCAI*. 2619–2625. http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6615

[7] Ido Guy. 2015. Social Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer, 511–543. https://doi.org/10.1007/978-1-4899-7637-6_15

[8] Rong Jin, Joyce Y. Chai, and Luo Si. 2004. An automatic weighting scheme for collaborative filtering. In *SIGIR*. 337–344. https://doi.org/10.1145/1008992.1009051

[9] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *SIGKDD*. 426–434. https://doi.org/10.1145/1401890.1401944

[10] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *IEEE Computer* 42, 8 (Aug. 2009), 30–37. https://doi.org/10.1109/MC.2009.263

[11] Hui Li, Dingming Wu, Wenbin Tang, and Nikos Mamoulis. 2015. Overlapping Community Regularization for Rating Prediction in Social Recommender Systems. In *RecSys*. 27–34. https://doi.org/10.1145/2792838.2800171

[12] Xin Liu and Karl Aberer. 2013. SoCo: a social network aided context-aware recommender system. In *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013*. 781–802. http://dl.acm.org/citation.cfm?id=2488457

[13] Hao Ma, Irwin King, and Michael R. Lyu. 2009. Learning to recommend with social trust ensemble. In *SIGIR*. 203–210. https://doi.org/10.1145/1571941.1571978

[14] Hao Ma, Haixuan Yang, Michael R. Lyu, and Irwin King. 2008. SoRec: social recommendation using probabilistic matrix factorization. In *CIKM*. 931–940. https://doi.org/10.1145/1458082.1458205

[15] Hao Ma, Dengyong Zhou, Chao Liu, Michael R. Lyu, and Irwin King. 2011. Recommender systems with social regularization. In *WSDM*. 287–296. https://doi.org/10.1145/1935826.1935877

[16] Paolo Massa and Paolo Avesani. 2007. Trust-aware recommender systems. In *RecSys*. 17–24. https://doi.org/10.1145/1297231.1297235

[17] Carine Pierrette Mukamakuza. 2017. Analyzing the Impact of Social Connections on Rating Behavior in Social Recommender Systems. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*. ACM, New York, NY, USA, 322–326. https://doi.org/10.1145/3079628.3079706

[18] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank Citation Ranking: Bringing Order to the Web.

[19] Ruslan Salakhutdinov and Andriy Mnih. 2007. Probabilistic Matrix Factorization. In *NIPS*. 1257–1264. http://papers.nips.cc/paper/3208-probabilistic-matrix-factorization