# Fairness and Diversity in Social-Based Recommender Systems


Dimitris Sacharidis
TU Wien
Vienna, Austria
dimitris@ec.tuwien.ac.at

Carine Pierrette Mukamakuza
TU Wien
Vienna, Austria
carine@ec.tuwien.ac.at

Hannes Werthner
TU Wien
Vienna, Austria
werthner@ec.tuwien.ac.at



## ABSTRACT

In social networks, the phenomena of homophily and influence explain the fact that friends tend to be similar. Social-based recommenders exploit this observation by incorporating the social structure in collaborative filtering techniques. In practice, these recommenders tend to make friends appear more similar compared to non-socially aware techniques. Various proposals have demonstrated the benefit of incorporating social connections. But at what cost? In this work, we show that there exist users that are mistreated in social recommenders. Specifically, their individual preferences are suppressed more compared to other users in their social circle. We seek to identify who they are and develop techniques that protect them, without severely affecting the effectiveness of the recommender.



## CCS CONCEPTS

• **Information systems → Recommender systems**.

## KEYWORDS

Social-Based Recommender Systems; Social Regularization; Fairness; Diversity; Novelty

**ACM Reference Format:**
Dimitris Sacharidis, Carine Pierrette Mukamakuza, and Hannes Werthner. 2020. Fairness and Diversity in Social-Based Recommender Systems. In *Adjunct Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct), July 14–17, 2020, Genoa, Italy*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3386392.3397603


## 1 INTRODUCTION

The mechanisms of *homophily* and *social influence* observed in social networks [14], suggest that our preferences and tastes are quite similar to those of whom we interact with in our everyday life [3, 6, 17]. Based on this premise, several *social-based recommender systems* [4, 7–13, 18, 19] seek to exploit social connections in order to improve the recommendation accuracy, but also increase coverage, and address the cold-start user problem.

Existing approaches in social-based recommenders extend collaborative filtering by enforcing similarity constraints between friends. The most common way to enforce this is to constrain the latent representations of users, extracted by model-based approaches such as matrix factorization, to be similar to those of her friends, a general approach called *social regularization* [7, 11, 19].

It has been shown that social regularization indeed helps in recommendation effectiveness. In this work, we ask at what cost. We suspect there are specific users which are mistreated, in that they are forced to become similar to users they would not otherwise be. Moreover, we suspect that there exist users who become more isolated that others, in that they tend to be more influenced by their social circle, creating thus social echo chambers [1].

In this work, we seek to design social-based recommenders that eliminate such phenomena. We start by identifying the types of users that are most likely to be treated unfairly. These are the so-call cold-start users, who have not provided as much feedback to the system. When we look at pairs of cold-cold users we see that they tend to be affected stronger than other pairs of users. Thus we propose a social regularization method that explicitly protects them.

Then, we turn our attention at what may cause social echo chambers. To some extent similarity among friends is desirable. What might cause echo chambers is the lack of diversity among them. Thus we propose another regularization approach that allows friends to be similar, but overall within a community it forces members to be more diverse.

## 2 BACKGROUND AND RELATED WORK

Social-based recommender systems make use of information from two sources, the user-item *rating matrix* $R \in \mathbb{R}^{m \times n}$, and the *social matrix* $S \in \mathbb{R}^{m \times m}$ corresponding to the adjacency matrix of the social network. Early work on social-based recommenders assumed that social connections conveyed trust between users of the system. In [12, 13], the authors propose a memory-based CF technique to integrate trust into recommendations, which is called Trust-aware Recommender System (TaRS). Matrix factorization (MF) techniques first appear in [10] and in [9].

Starting with a graph induced by explicit trust statements, one can define *local* and *global* metrics to quantify the trust between any two users. The former compute a subjective measure of trust, while the latter an objective measure of global reputation. In [13] the same authors present detailed evaluation results of their technique, which implements a simple local trust metric, called MoleTrust. The proposed algorithm predicts the rating based on a user-based CF technique, where instead of the user similarity, the user trust is used to determine the neighborhood and weigh the ratings. In all experiments, this technique resulted in higher accuracy (in terms of maximum absolute error) and coverage (in terms of number of predictable ratings) than standard user-based CF. Also, they find that hybrid techniques based on trust and similarity, and global trust metrics (such as PageRank) performed worse than pure local trust-based ones.



*UMAP '20 Adjunct, July 14–17, 2020, Genoa, Italy*



ACM ISBN 978-1-4503-7950-2/20/07…$15.00
https://doi.org/10.1145/3386392.3397603

In [10] the authors introduce a probabilistic matrix factorization technique for social-based recommendations, termed Social Recommender (SoRec). As in standard matrix factorization (MF), they decompose the ratings matrix into two $d$-dimensional latent feature matrices $U \in \mathbb{R}^{d\times n}$ for the users and $V \in \mathbb{R}^{d\times m}$ for the items, so that the predicted rating matrix is computed as $R \approx U^\intercal V$. In addition, they assume that the social graph can be decomposed with the same user matrix $U$ and a factor feature matrix $Z$, such that the social matrix is given by $S \approx U^\intercal Z$. Then, they derive the matrices $U, V, Z$ by minimizing an objective function similar to standard MF, where they seek to minimize the error between $S$ and $U^\intercal Z$, and $R$ and $U^\intercal V$, after standard regularization. In this model, the matrix $Z$ is uninterpretable.

In [9] a different probabilistic MF approach is taken, termed Recommendation with Social Trust Ensemble (RSTE). In the simplest model, they consider only recommendations from social trust relationships. The assumption is that the rating of an item $j$ by user $i$ is generated by the weighted (according to trust levels) sum of ratings given by/predicted for $i$'s friends. In other words, $R \approx SU^\intercal V$, where $U$, $V$ are similar to standard MF and are learned by minimizing the error between $R$ and $SU^\intercal V$, after standard regularization. In the ensemble model, the actual rating is assumed to be a smoothed sum of the standard non-social predicted rating $U^\intercal V$ and the trust-based predicted rating $SU^\intercal V$, i.e., $R \approx \alpha U^\intercal V + (1-\alpha)SU^\intercal V$. Note that equivalently, one could interpret this ensemble model as the social model where the adjacency matrix $S$ is modified to include $\alpha$ in its diagonal (instead of zero), indicating the level of trust to oneself, and all other entries are scaled by $1-\alpha$.

The SocialMF model introduced in [4] attempts to account for the effects of *selection* and *homophily* observed in social networks. The former indicates that users tend to connect to like-minded people, while the latter says that two friends develop similar interests over time. The key idea in SocialMF is that the user feature vectors of two friends in a MF model should be similar reflecting exactly selection and homophily. The authors call this effect trust propagation, although there is actually no propagation of trust values in the social graph. The predicted rating is as in standard MF, i.e., $R \approx U^\intercal V$. However, the $U_u$ feature vectors should additionally encode the social relationships of each user $u$. The assumption is that the estimate of the latent feature vector of user $u$ is the weighted average of those of his direct neighbors, i.e., $\widehat{U_u} \approx US_u$; here vector $S_u$ contains the $[0,1]$ trust values of users $u$. Therefore, the objective function should minimize the error between predicted and actual rating, but also the discrepancy between the user feature matrix $U$ and the aggregate matrix composed of the features of neighbors expressed in matrix form as $US$.

In [11] the authors emphasize the difference between trust relationships and friendships, making the argument that trust-based approaches are not suitable for social recommendations. Their input is matrix $S$ which is the binary (un-weighted) adjacency matrix of a given social network. However, in their models they weigh each edge by the *Pearson Correlation Coefficient* (PCC) similarity of the common ratings between the adjacent users. Therefore, one can construct a new matrix $S' = S \circ Q$ that contains the similarities of friends, instead of 0/1 values, where $Q_{uv}$ is the PCC between users $u$, $v$, and $\circ$ denotes the elementwise (Hadamard) product for matrices. Similar to the idea in [4], the goal is to constraint the feature vector of each user to be similar to those of its friends. The first model, which we call *average social regularization*, is essentially identical to [4], and makes the assumption that the users' feature matrix is similar to the *average* feature matrix inferred from friends given the modified matrix $S'$, i.e., $\widehat{U} \approx US'$. This is somewhat restrictive as it forces each user's features to be similar to the average features of her friends. The second model, which we call *individual social regularization*, relaxes this and assumes that the feature vector of a user is similar to the feature vector of her friend to the degree indicated by their rating similarity. Hence, for each pair of friends $u, v$ there is a regularization term constraining $\|U_u - U_v\|$ with a strength equal to the rating similarity $S'_{uv}$ between them.

Several variations on the basic idea of social regularization have been proposed since then [2, 7, 15, 19]. The current state of the art method extends the local low-rank matrix approximation (LLORMA) ensemble method [5] in two ways: (1) the users and items comprising a local model are determined by the social network structure, instead of user-user and item-item rating similarities, and (2) pairwise social regularization is employed.

## 3 APPROACH

A good social-based recommender system should satisfy two desiderata. First, it should treat users *fairly*, without introducing biases for specific types of users. Second, it should seek to *avoid social echo chambers* where users' preferences become isolated from others. Before describing how to achieve these goals, we note that in this work, we care about how a recommender sees or treats users. We assume the recommender builds internal representations of users in terms of *latent factors*, or embeddings, as in matrix factorization techniques, for example. To measure the similarity of two users $u, v$ as seen from the perspective of the recommender, we define the *latent factor similarity*, termed LF-sim, using the cosine similarity of their latent representation normalized to $[0,1]$:

$$\text{LF-sim}(u,v) = \frac{1}{2}\left(1 + \frac{U_u^\intercal U_v}{\|U_u\|\|U_v\|}\right),$$

where $U_u, U_v$ are the vectors of the latent factors for users $u, v$ respectively. LF-sim essentially controls the output of the recommender: if two users have high latent factor similarity, they would get highly similar recommendations.

### 3.1 Treating Users Fairly

Social-based recommenders that apply social regularization operate on the assumption that socially connected users are like-minded. Generally speaking they enforce specific constraints on the learning process that discovers the latent factors of users. More precisely, the latent factors of a user are required to be similar to those of their friends. Through this process, the recommender sees friends as being more similar compared to when no social structure is considered.

To illustrate this effect, we compare the LF-sim between friends with and without social regularization. In Figure 1, for each pair of friends, we plot a point where the x coordinate is the LF-sim without, and the y coordinate is the LF-sim with social regularization. We

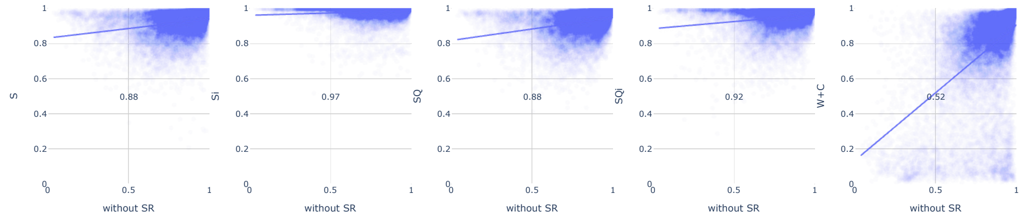


**(a) S** **(b) Si** **(c) SQ** **(d) SQi** **(e) W+C**

**Figure 1: LF-sim with and without social regularization for five social-based recommenders; Douban.**

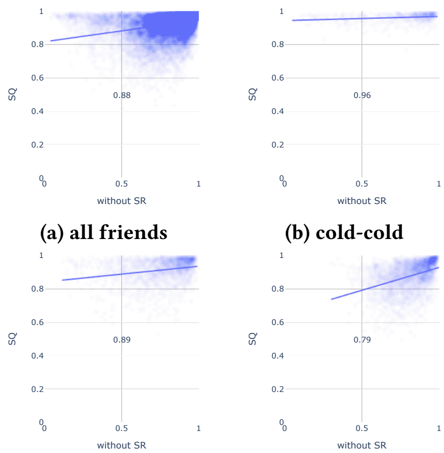


**(a) all friends** **(b) cold-cold**

**(c) cold-warm** **(d) warm-warm**

**Figure 2: LF-sim with and without social regularization for different friendship types; Douban.**

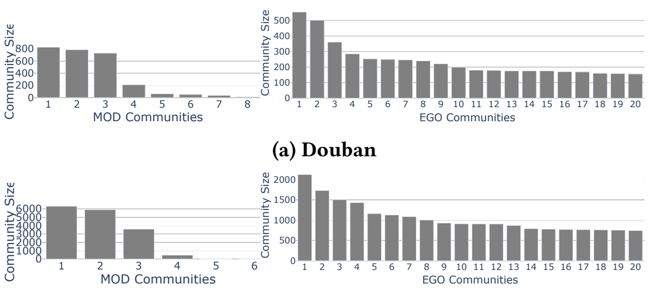


**(a) Douban**

**(b) Epinions**

**Figure 3: Community Sizes**

show the same plot for five different recommenders and for the Douban dataset used in the experiments.

The first four recommenders implement social regularization approaches from the literature. Specifically, S refers to average social regularization based on the adjacency matrix $S$ of the social graph, and was introduced in [4]. The regularization term is:

$$\sum_u \|U_u - \frac{1}{|\{v \in S_u\}|} \sum_{v \in S_u} U_v\|^2,$$

where $S_u$ denotes the friends of $u$. Si is individual social regularization based on $S$, and the regularization term is:

$$\sum_u \sum_{v \in S_u} \|U_u - U_v\|^2.$$

SQ denotes average social regularization based on matrix $S$ where each term is weighted by the PCC similarity $Q_{uv}$ between a pair of friends $u, v$, introduced in [11]. The regularization term is:

$$\sum_u \|U_u - \frac{1}{\sum_{v \in S_u} Q_{uv}} \sum_{v \in S_u} Q_{uv} U_v\|^2.$$

SQi is individual social regularization based on the PCC-weighted adjacency matrix $S \circ Q$, and is used in [5, 11]. The regularization term is:

$$\sum_u \sum_{v \in S_u} Q_{uv} \|U_u - U_v\|^2.$$

The last depicted recommender is our contribution, which we will describe later.

Ideally, we would like social regularization to have a small and uniform effect for all pair of friends. That translates into the points being distributed along the diagonal. In all plots, we also depict the linear regression line. The ideal state is achieved when this line is the diagonal $y = x$, meaning that the social-based recommender produces user representations that are mostly similar to what it would produce had it no knowledge of the social structure. To assess how far from the ideal situation a recommender is, we measure the size of the area under the linear regression line. The area size is shown in the center of the plots. The closer to 0.5 this size is, the more uniform the effect of social regularization is. In contrast, when the area has size close to 1 the recommender tends to make friends more similar than they would be without regularization. As seen in Figure 1, the first four recommenders tend to concentrate LF-sim around 1, as evidenced by points appearing close to the upper right corner and by an area size around 1. In contrast, the last recommender appears to not suffer from this effect.

To understand how we can combat the concentration of LF-sim around 1, we look deeper into who is stronger affected by social regularization. We say that a user is *warm* if they have given more feedback to the system than 75% of the users. Conversely, a user is *cold* if they have given less feedback than 25% of the users. We can now distinguish among three types of friendships: warm-warm (WW), warm-cold (WC), and cold-cold (CC) friends.

Figure 2 draws the LF-sim plot considering only all pairs of friends, and pairs belonging to each type; the recommender depicted is SQ, but all other existing recommenders generate similar plots. What we observe is that the effect is stronger for cold-cold friends, and weaker for warm-warm friends. In other words, cold-cold friends will be treated as practically similar by the recommender, even though the system only knows a few things about their preferences. It is important to note that this happens despite the fact that the SQ recommender explicitly regulates the effect of regularization according to feedback similarity (PCC). That is, two friends that have low PCC would end up having low LF-sim without social regularization, but due to social regularization they are forced to appear more similar.

Motivated by these observations, we want to design a social regularization approach that maintains the benefit of existing approaches (in that they increase recommendation accuracy) but without treating pairs of cold-cold friends unfairly compared to others. We propose to apply social regularization only when there is sufficient evidence, i.e., between warm-warm friends. Thus we propose a type of average social regularization, denoted as W, that only considers pairs of warm friends, leading to the term:

$$\sum_u \|U_u - \frac{1}{\sum_{v \in S_u} W_{uv}} \sum_{v \in S_u} W_{uv} U_v\|^2,$$

where $W_{uv}$ is 1 when both $u, v$ are warm, and 0 otherwise.

### 3.2 Avoiding Social Echo Chambers

The aforementioned social regularization, as seen in the last plot of Figure 1 treats friends more fairly compared to the other approaches. However, there is another perspective that we should consider when designing a fair social-based recommender. A recommender should not amplify preferences and isolate users in social echo chambers. To make this requirement more concrete, we seek to define the *social influence* on an individual, capturing the degree to which an individual is affected by its social circle.

Consider a community $c$ and one of its members $u \in c$. As discussed before, in the viewpoint of the recommender, $u$ is similar to other users to the degree specified by the LF-sim metric. Therefore, for $u$ we can define their LF-sim based neighborhood of most similar users. If this neighborhood contains many other members from $c$, then we can claim the social influence of community $c$ on $u$ is high. Put in other words, if the most similar (in the eyes of the system) users to $u$ are from $u$'s community, then this community $c$ has relatively high influence on $u$.

We formalize this with the following definition. Fix a community $c$ and a user $u \in c$. Let $\mathcal{N}_k(u)$ be $u$'s LF-sim neighborhood, i.e., the set of the $k$ most similar, in terms of LF-sim, users to $u$. Then, the *community influence* (CI) of $c$ on $u$ is:

$$\text{CI}@k(u; c) = \frac{1}{k} |c \cap \mathcal{N}_k(u)|,$$

i.e., the proportion of $u$'s LF-sim neighbors that are in $c$. When CI is 1, then user $u$ has no similar user outside their community; put differently, community $c$ acts as a social echo chamber for user $u$.

Note that some degree of community influence is desirable, i.e., it is likely that a user is anyway similar to some of its friends by means of the homophily and social influence phenomena. To identify this desirable degree, we can look at the community influence observed when the recommender is agnostic of any social connections. As with the case of LF-sim, we want to compare CI with and without social regularization. Therefore, we compute the *community influence change* ($\Delta$CI) of $c$ to $u$ is:

$$\Delta\text{CI}@k(u; c) = \text{CI}@k(u; c) - \text{CI}_{w/o}@k(u; c)$$

where $\text{CI}_{w/o}@k(u; c)$ denotes CI without social regularization.

We now have a way ($\Delta$CI) to quantify the degree to which a social-based recommender is creating social echo chambers. Next, we want to ensure that a recommender seeks to keep the community influence change close to zero. To achieve this, we counter-balance the effect of social regularization. The idea is to allow friends to have similar representations, but additionally to require that their representations are *diverse*. Here diversity is with respect to the latent representation of the community, defined as the average representation of its members. We can thus achieve the benefit of social regularization, in that friends are treated similarly, and avoid echo chambers at the same time, in that members are dissimilar to the average community member.

To achieve the goal of diversity we introduce a regularization term, denoted as C for community diversity:

$$\sum_c \sum_{u \in c} \frac{1}{2} \left(1 + \frac{U_u^\top U_c}{\|U_u\| \|U_c\|}\right),$$

where $U_c = \frac{1}{|c|} \sum_{u \in c} U_u$ is the community latent representation. This C term approaches zero when each community member becomes distinct from the average community member, achieving the desired behavior.

## 4 EVALUATION

**Datasets.** The first dataset we use for our evaluation, called Douban, concerns a popular Chinese social networking service[1] that allows users to connect to each other and provide content and ratings to movies, books, music, and events. The dataset[2] contains 912,479 ratings (on a scale of 1 through 5) given by 2,874 users on 39,694 movies, and includes 48,552 bidirectional connections among the users. The second dataset, called Epinions, concerns a Web review site of products, where users indicate trust relationships among them. The dataset[3] contains 512,774 ratings (on a scale of 1 through 5) given by 16,564 users on 129,329 items, and includes 556,921 bidirectional connections among the users.

**Methods.** Our evaluation seeks to compare social regularization techniques. For the purpose of this comparison, the underlying recommender engine responsible for producing the latent user representations is orthogonal. In these experiments, we employ a base matrix factorization technique. The results of this baseline, non-socially aware, recommender will be denoted as MF. On top of this method, we employ the four existing social regularization methods, discussed in Section 3 and denoted as S, Si, SQ, SQi. Moreover, we study our own approach, denoted as W+C, which involves the W term for fair social regularization and the C term for avoiding echo chambers, as defined in Section 3. For all tested methods, we fix the set of hyperparameters (batch size, learning rate, regularization strength) to the values that optimize the performance (in terms of RMSE) of the base matrix factorization model.

**Metrics.** For each social-based recommender, we measure recommendation accuracy in terms of: (1) the root mean square error (RMSE) over all predictions; and (2) the mean normalized discounted cumulative gain (nDCG) at various cut-off levels, where nDCG@k

[1] http://www.douban.com

[2] Accessed from http://smiles.xjtu.edu.cn/Download/Download_Douban.html.

[3] Accessed from http://www.trustlet.org/epinions.html.

denotes the metric's value when only the top-$k$ recommendations are considered. We evaluate the fairness and diversity of a recommender by measuring LF-sim within communities, and computing the community influence and its degree of change, as defined in Section 3.

We also measure the novelty and diversity of users individually, and of communities as a whole using the metrics introduced in [16]. Specifically, let $d(i, j)$ denote a function that measures the dissimilarity/distance between items $i$ and $j$; in this work, we choose a content-independent approach, where $d(i, j)$ is defined based on the cosine similarity of the latent factor vectors of $i$ and $j$. For a user $u$, let $P_u$ denote their set of recommended items, and $R_u$ the set of items they have interacted with (e.g., rated, purchased) in the past. Then, the *individual diversity* for $u$ is the average pairwise distance among the recommended items (a.k.a. intra-list diversity [20]): $IDIV_u = \frac{1}{|P_u|(|P_u|-1)} \sum_{i \in P_u} \sum_{j \in P_u} d(i, j)$. The *individual novelty* for $u$ is the average pairwise distance between a recommended item and an item from their history: $INOV_u = \frac{1}{|P_u||R_u|} \sum_{i \in P_u} \sum_{j \in R_u} d(i, j)$.

Now consider a community $c$. Its *community diversity* is the average pairwise distance among the items recommended to any community member: $CDIV_c = \frac{1}{|P_c|(|P_c|-1)} \sum_{i \in P_c} \sum_{j \in P_c} d(i, j)$, where $P_c = \cup_{u \in c} P_u$ is the set of recommendations to the entire community. Similarly, we define the *community novelty* for $c$ as the average pairwise distance between an recommended to some community member and an item a community member has interacted with in the past: $CNOV_c = \frac{1}{|P_c||R_c|} \sum_{i \in P_c} \sum_{j \in R_c} d(i, j)$, where $R_c = \cup_{u \in c} R_u$ is the community's interaction history.

**Communities.** From each dataset, we extract two types of communities. EGO communities are the ego networks of the top-20 users with the most social connections. MOD communities are extracted by a greedy approach for creating communities of maximum modularity. To draw meaningful conclusions we only keep MOD communities of at least 10 users. The sizes of the communities are shown in Figure 3.

**Results.** In the first round of experiments, we compare accuracy and user, community diversity and novelty. Results are shown in Tables 1 for Douban; similar finding hold for Epinions. The best value in each metric is highlighted bold; in case the second best value is statistically close it is also highlighted. Overall, we see that social based recommenders improve prediction accuracy (RMSE) over the baseline recommender. However, this is not the case with ranking accuracy (nDCG) where often the baseline has comparable or better effectiveness. In terms of individual and community diversity, Si and SQi perform best. For individual and community novelty, SQi and W+C are the best. In conclusion, we see that our recommender does not sacrifice recommendation effectiveness, while achieving best or second-best values in novelty and diversity.

Next we calculate the LF-sim of each pairs of friends for the communities. Results are shown in Figures 4 and 5 as boxplots summarizing the distribution of data points. What we observe is that all social-based recommenders, except our approach, tend to concentrate LF-sim near 1. In contrast, our approach mimics the distribution of LF-sim of the base recommender MF; the only case where this does not hold completely is in the last few MOD communities, which are very small. Overall, the median LF-sim,


**(a) Douban**

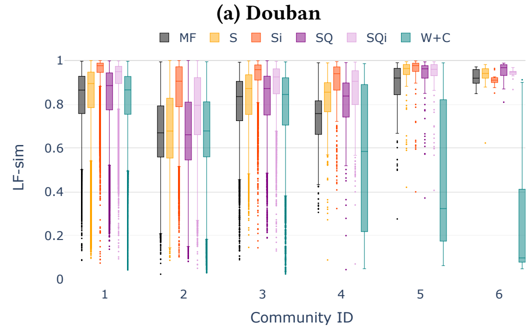


**(b) Epinions**

**Figure 4: LF-sim for MOD communities**

shown as the middle line in the box, of W+C is close to that of MF, and considerable lower than of the other recommenders.

In the last experiment, we compute the community influence change ($\Delta$CI) at $k = 20$ for each member of a community. Results are shown in Tables 2 and 3. A desirable property of a social-based recommender is to have $\Delta$CI concentrated around zero. Among all recommenders, our approach W+C achieves this goal better.

# 5 CONCLUSION

This work studied issues of fairness and diversity in social-based recommenders. While existing social regularization techniques may polarize users, we present two novel ideas that achieve the benefits of increased recommendation effectiveness while treating users more uniformly and reducing the presence of social echo chambers.

# REFERENCES

[1] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, S1 (2016), 298–320.

[2] Rana Forsati, Iman Barjasteh, Farzan Masrour, Abdol-Hossein Esfahanian, and Hayder Radha. 2015. PushTrust: An Efficient Recommendation Algorithm by Leveraging Trust and Distrust Relations. In *ACM RecSys*. 51–58. https://doi.org/10.1145/2792838.2800198

[3] Mike Gartrell, Xinyu Xing, Qin Lv, Aaron Beach, Richard Han, Shivakant Mishra, and Karim Seada. 2010. Enhancing group recommendation by incorporating social relationship interactions. In *GROUP*, Wayne G. Lutters, Diane H. Sonnenwald, Tom Gross, and Madhu Reddy (Eds.). ACM, 97–106. https://doi.org/10.1145/1880071.1880087

[4] Mohsen Jamali and Martin Ester. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In *ACM RecSys*. 135–142. https://doi.org/10.1145/1864708.1864736

[5] Joonseok Lee, Seungyeon Kim, Guy Lebanon, Yoram Singer, and Samy Bengio. 2016. LLORMA: Local Low-Rank Matrix Approximation. *Journal of Machine Learning Research* 17 (2016), 15:1–15:24.

[6] Kevin Lewis, Marco Gonzalez, and Jason Kaufman. 2012. Social selection and peer influence in an online social network. *Proceedings of the National Academy of Sciences* 109, 1 (2012), 68–72.

[7] Hui Li, Dingming Wu, Wenbin Tang, and Nikos Mamoulis. 2015. Overlapping Community Regularization for Rating Prediction in Social Recommender Systems.

**Table 1: Accuracy, Diversity, and Novelty metrics; Douban**

| | RMSE | nDCG@1 | nDCG@2 | nDCG@5 | nDCG@10 | nDCG@20 | I-Div | C-Div | I-Nov | C-Nov |
|---|---|---|---|---|---|---|---|---|---|---|
| MF | 1.041±0.016 | 0.776±0.002 | 0.773±0.001 | 0.773±0.001 | 0.781±0.002 | 0.798±0.002 | 0.120±0.011 | 0.221±0.021 | 0.232±0.007 | 0.263±0.007 |
| S | **1.019**±0.008 | 0.770±0.003 | 0.766±0.002 | 0.770±0.003 | 0.780±0.003 | 0.799±0.003 | 0.116±0.009 | 0.191±0.015 | 0.216±0.002 | 0.247±0.003 |
| Si | 1.023±0.014 | 0.764±0.002 | 0.762±0.001 | 0.762±0.002 | 0.772±0.000 | 0.792±0.001 | **0.184**±0.060 | **0.251**±0.051 | 0.234±0.024 | 0.256±0.027 |
| SQ | 1.020±0.023 | 0.770±0.008 | 0.766±0.007 | 0.765±0.008 | 0.775±0.007 | 0.796±0.005 | 0.169±0.044 | **0.242**±0.033 | 0.235±0.012 | 0.256±0.012 |
| SQi | **1.013**±0.022 | 0.779±0.005 | 0.776±0.002 | 0.777±0.001 | 0.785±0.000 | 0.803±0.001 | **0.170**±0.039 | 0.227±0.043 | **0.247**±0.019 | **0.275**±0.019 |
| W+C | 1.022±0.012 | **0.787**±0.003 | **0.785**±0.003 | **0.786**±0.001 | **0.792**±0.001 | **0.809**±0.001 | 0.110±0.002 | 0.204±0.019 | **0.241**±0.000 | **0.278**±0.001 |


**(a) Douban**


**(b) Epinions**

**Figure 5: LF-sim for EGO communities**

**Table 2: Community Influence Change; Douban**

| | EGO ΔCI@10 | ΔCI@20 | ΔCI@50 | MOD ΔCI@10 | ΔCI@20 | ΔCI@50 |
|---|---|---|---|---|---|---|
| S | 0.020 | 0.021 | 0.021 | 0.042 | 0.041 | 0.034 |
| Si | 0.109 | 0.109 | 0.103 | 0.062 | 0.059 | 0.050 |
| SQ | 0.014 | 0.015 | 0.014 | 0.054 | 0.047 | 0.040 |
| SQi | 0.090 | 0.096 | 0.090 | 0.046 | 0.046 | 0.043 |
| W+C | **0.009** | **0.012** | **0.012** | **0.015** | **0.010** | **0.006** |

**Table 3: Community Influence Change; Epinions**

| | EGO ΔCI@10 | ΔCI@20 | ΔCI@50 | MOD ΔCI@10 | ΔCI@20 | ΔCI@50 |
|---|---|---|---|---|---|---|
| S | 0.007 | 0.009 | 0.007 | 0.017 | 0.012 | 0.010 |
| Si | 0.105 | 0.101 | 0.098 | 0.027 | 0.024 | 0.023 |
| SQ | 0.012 | 0.011 | 0.009 | 0.019 | 0.016 | 0.013 |
| SQi | 0.064 | 0.062 | 0.059 | 0.020 | 0.020 | 0.021 |
| W+C | **0.001** | **0.001** | **0.002** | **0.008** | **0.007** | 0.**006** |

In *ACM RecSys*. 27–34. https://doi.org/10.1145/2792838.2800171
[8] Hao Ma. 2013. An experimental study on implicit social recommendation. In *ACM SIGIR*. 73–82. https://doi.org/10.1145/2484028.2484059
[9] Hao Ma, Irwin King, and Michael R. Lyu. 2009. Learning to recommend with social trust ensemble. In *ACM SIGIR*. 203–210. https://doi.org/10.1145/1571941.1571978
[10] Hao Ma, Haixuan Yang, Michael R. Lyu, and Irwin King. 2008. SoRec: social recommendation using probabilistic matrix factorization. In *ACM CIKM*. 931–940. https://doi.org/10.1145/1458082.1458205
[11] 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
[12] Paolo Massa and Paolo Avesani. 2004. Trust-Aware Collaborative Filtering for Recommender Systems. In *CoopIS*. 492–508. https://doi.org/10.1007/978-3-540-30468-5_31
[13] Paolo Massa and Paolo Avesani. 2007. Trust-aware recommender systems. In *ACM RecSys*. 17–24. https://doi.org/10.1145/1297231.1297235
[14] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27, 1 (2001), 415–444.
[15] Dimitrios Rafailidis and Fabio Crestani. 2017. Learning to Rank with Trust and Distrust in Recommender Systems. In *ACM RecSys*. ACM, 5–13.
[16] Dimitris Sacharidis. 2019. Diversity and Novelty in Social-Based Collaborative Filtering. In *ACM UMAP*.
[17] Lara Quijano Sánchez, Juan A. Recio-García, Belén Díaz-Agudo, and Guillermo Jiménez-Díaz. 2013. Social factors in group recommender systems. *ACM TIST* 4, 1 (2013), 8:1–8:30. https://doi.org/10.1145/2414425.2414433
[18] Xiwang Yang, Harald Steck, and Yong Liu. 2012. Circle-based recommendation in online social networks. In *ACM KDD*. 1267–1275. https://doi.org/10.1145/2339530.2339728
[19] Huan Zhao, Quanming Yao, James T. Kwok, and Dik Lun Lee. 2017. Collaborative Filtering with Social Local Models. In *ICDM*, Vijay Raghavan, Srinivas Aluru, George Karypis, Lucio Miele, and Xindong Wu (Eds.). IEEE Computer Society, 645–654. https://doi.org/10.1109/ICDM.2017.74
[20] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *WWW*. ACM, 22–32.