

Diversity and Novelty in Social-Based Collaborative Filtering

Dimitris Sacharidis
E-Commerce Research Division
TU Wien
Austria
dimitris@ec.tuwien.ac.at

ABSTRACT

Social-based recommenders seek to exploit the mechanisms of homophily and influence observed in social networks in order to provide more accurate recommendations. The way they achieve this is by enforcing similar preferences among users that are socially connected. It is thus reasonable to question whether such approaches lead to the formation of echo chambers, i.e., social groups with a narrow set of preferences and which receive recommendations with low diversity and novelty. This work studies this research question and quantifies the diversity and novelty of existing methods. An important finding is that it is possible to increase accuracy without sacrificing diversity and novelty.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

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

ACM Reference Format:

Dimitris Sacharidis. 2019. Diversity and Novelty in Social-Based Collaborative Filtering. In *27th Conference on User Modeling, Adaptation and Personalization (UMAP '19)*, June 9–12, 2019, Larnaca, Cyprus. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3320435.3320479>

1 INTRODUCTION

Social connections play a crucial role in guiding our choices and decisions. The mechanisms of *homophily* and *social influence* observed in social networks [21], suggest that our preferences and tastes tend to become similar to those of whom we interact with in our everyday life [6, 12, 25]. Based on this premise, several *social-based recommender systems* [8, 13–19, 28, 29] seek to exploit social connections in order to improve the recommendation accuracy, but also increase coverage, and address the cold-start user problem.

The vast majority of social-based recommenders apply collaborative filtering (CF) under the notion that a user’s model should be similar not only to that of her neighborhood but also to that of her social circle. The prevalent way to implement this in model-based

CF is to constrain each user’s model (e.g., the latent factors in matrix factorization) to be similar to those of her friends, a technique called *social regularization* [13, 17, 29].

While social-based recommenders may increase the accuracy of recommendations, one reasonable concern is about their role in creating echo chambers [4]. When the system explicitly forces one’s preferences to be alike her friends’, we can expect even more pronounced manifestations of filter bubbles compared to those observed in conventional recommenders [22].

This work studies the effect of social-based recommendations in terms of the novelty and diversity experienced by users, and makes the following contributions.

- We present simple definitions of novelty and diversity that better capture how users perceive recommendations with respect to the social groups they belong to.
- We present a novel social-based recommender that considers the similarity of users in terms of the structure of the social network.
- We present an experimental study that evaluates the accuracy, diversity, and novelty of social-based recommenders. We find that our proposed recommender improves on accuracy, while exhibiting similar diversity and novelty with the state of the art.

2 RELATED WORK

Social-Based Recommenders. 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 [18, 19], 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 [16] and in [15].

The SocialMF model introduced in [8] 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., $R \approx U^T V$. 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '19, June 9–12, 2019, Larnaca, Cyprus

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6021-0/19/06...\$15.00

<https://doi.org/10.1145/3320435.3320479>

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 [17] 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 [8], 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 [8], 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 *pairwise 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 [5, 13, 23, 29]. The current state of the art method extends the local low-rank matrix approximation (LLORMA) ensemble method [11] 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.

Diversity and Novelty. Measuring the accuracy in capturing the user's historical preferences conveys only one aspect of the quality of a recommender system [20]. Many researchers call for aiming towards other aspects such as unexpectedness [1], serendipity [7], diversity [30], and novelty [27]. Such aspects are closely related to each other and can be classified as conveying either a notion of *diversity* or of *novelty*. As nicely posed in [3], the former relates to differences within parts of the experience (e.g., a recommendation list), while the later to differences between present and past experiences (e.g., between a recommendation list and the historically rated items).

Crucial to the concepts of diversity and novelty is a measure of *similarity* between items. This can be computed based on item content [30], or in a pure CF setting, based on the users that have interacted with these items [24], or even based on the latent factors of items computed by a model-based method such as MF.

As both concepts relate to lists of items, whereas many recommenders are trained in a pointwise or a pairwise manner (e.g., seeing one or two historical interactions at a time), including diversity and novelty objectives in the training process is not possible.

Instead, techniques resort to post-processing, or *reranking*, of the items in order to achieve better diversity or novelty [2], often at the expense of degrading accuracy. We note that although the trade-off between accuracy and diversity/novelty in recommendations has been recently studied, e.g., [10, 26], to the best of our knowledge there is no related work in the context of social-based recommender systems.

3 SOCIAL-BASED DIVERSITY AND NOVELTY METRICS

In this section, starting from standard concepts [3], we introduce definitions of diversity and novelty that are suitable for social-based recommender systems. For the following, assume there exists a function $d(i, j)$ that measures the distance (dissimilarity) between items i and j – its exact definition is orthogonal to our aim.

We define *individual diversity* for a particular user u as the average pairwise distance among the recommended items P_u , also called intra-list diversity [30]:

$$IDIV_u = \frac{1}{|P_u|(|P_u| - 1)} \sum_{i \in P_u} \sum_{j \in P_u} d(i, j).$$

In a similar spirit, we define *individual novelty* for a user u as the average pairwise distance between a recommended item and an item the user has interacted with (e.g., rated, purchased) in the past:

$$INOV_u = \frac{1}{|P_u||R_u|} \sum_{i \in P_u} \sum_{j \in R_u} d(i, j),$$

where R_u is the set of items user u has interacted with.

The previous definitions consider the diversity and novelty of recommendations with respect to users individually. Naturally, we can report the individual novelty and diversity averaged across all users, to obtain an overall sense of how diverse and novel are the recommendations. To better capture the effect of social-based recommendations on diversity and novelty, however, we need to define these concepts for sets of users that form a connected component of the social network; we refer to such a set of users as a *social group*.

We can then define *group diversity* for a social group g as the average pairwise distance among the items recommended to any user in the group:

$$GDIV_g = \frac{1}{|P_g|(|P_g| - 1)} \sum_{i \in P_g} \sum_{j \in P_g} d(i, j),$$

where $P_g = \cup_{u \in g} P_u$ is the set of recommendations to the entire social group.

Similarly, we define *group novelty* for group g as the average pairwise distance between an recommended to some group member and an item a group member has interacted with in the past:

$$GNOV_g = \frac{1}{|P_g||R_g|} \sum_{i \in P_g} \sum_{j \in R_g} d(i, j),$$

where $R_g = \cup_{u \in g} R_u$ is the joint interaction history of all group members.

These metrics are useful in different manners. To generate an overall view of group diversity and novelty, we may select a set of groups of interest and then report the average values of these

metrics. To assess fairness of the social-based recommender, we may compare diversity and novelty across social groups.

4 EVALUATION

Section 4.1 describes our experimental setup, while Section 4.2 presents our findings.

4.1 Setup

Dataset. The dataset we use for our evaluation, called Douban, concerns a popular Chinese social networking service¹ that allows users to connect to each other and provide content and ratings to movies, books, music, and events. The dataset is a subset² of the data compiled and published by the authors of [17], and includes ratings of movies on a scale of 1 through 5. There are 1,048,575 ratings given by 8,890 users on 23,185 movies, and 7,908 bidirectional connections among the users.

Recommenders. In our evaluation, we compare the following recommenders.

- **MF:** This is the base matrix factorization model used by all social-based recommenders. It performs a decomposition of the ratings matrix, including bias terms, in 50 latent factors.
- **S:** This model, introduced in [8], extends MF by including average social regularization based on the adjacency matrix S of the social graph. Specifically, 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 .

- **Sp:** This is pairwise social regularization based on S . The regularization term is:

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

- **SQ:** This model, introduced in [17], includes average social regularization based on matrix S where each term is weighted by the PCC similarity between a pair of users. The regularization term is:

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

where Q_{uv} is the PCC between users u and v .

- **SQp:** This model, also introduced in [17], includes pairwise social regularization based on the PCC-weighted adjacency matrix $S \circ Q$. The regularization term is thus:

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

- **SX:** This model is our contribution and includes average social regularization based on matrix S where each term is weighted by the SimRank node similarity [9] between a pair of users. SimRank computes the similarity of a pair of nodes in a spirit analogous to how PageRank computes the importance of a node. Therefore, SimRank encapsulates a more global view of the social structure compared to the

local view conveyed by the social (adjacency) matrix S . The regularization term is:

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

where X_{uv} is the SimRank between users u and v .

- **SQp:** This model is our contribution and includes pairwise social regularization based on the SimRank-weighted adjacency matrix $S \circ X$. The regularization term is:

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

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.

Methodology. We perform 5-fold cross validation, splitting the dataset into train and test subsets with a ratio of 4:1. We train the recommender on the train subset and ask it (1) to predict the ratings for the user-item pairs in the test subset, and (2) to rank for each user her unrated (i.e., not in the train dataset) items. Then, based on (1) we measure the prediction accuracy in terms of root mean square error, denoted as RMSE; based on (2) we measure the mean NDCG per user at various ranking prefixes, when only the top- k recommendations are considered, which we denote as NDCG@ k .

Moreover, based on (2), we compute the diversity and novelty metrics introduced at Section 3 for certain groups of users in the social network. Specifically, we identify the most socially active users that have more than 10 connections; there are 48 such users. Then for each such user, we identify her ego network (i.e., that includes herself and her friends) which forms a social group. Overall, the 48 social groups cover 1,394 users. We request the top-10 recommendations for each of these users, and compute their average individual diversity and novelty. In addition, we measure the average group diversity and novelty over these 48 social groups.

4.2 Results

Table 1 presents the evaluation results on all tested recommenders. We report prediction accuracy (RMSE; lower values are better), ranking accuracy (NDCG@ k at five cut-off levels; higher values are better), and individual and group diversity and novelty (higher values are better). For each metric, we report the average value over 35 executions: 5 test datasets, and 7 randomly initialized training sessions per test dataset. We also present the standard deviation of the observed values as error terms. For each metric, i.e., column, the best value, including those that are statistically indistinguishable from the best, are shown in bold.

Regarding accuracy metrics, we observe that prediction and ranking accuracy metrics agree. The base MF method has the worst accuracy, while among social-based recommenders, SX is the best with more than 12% improvement in RMSE and up to 5% improvement in NDCG. In general, average appears to work better than pairwise social regularization (S/SQ/SX vs. Sp/SQp/SXp). The use of PCC weights in the social matrix does not seem to bring a significant benefit (S vs. SQ). On the other hand, weighing the social matrix by SimRank (SX) results in much more accurate recommendations.

¹<http://www.douban.com>

²The original dataset is no longer publicly available.

Table 1: Accuracy, Diversity, and Novelty metrics

	RMSE	NDCG@1	NDCG@2	NDCG@5	NDCG@10	NDCG@20	I-Diversity	G-Diversity	I-Novely	G-Novely
MF	0.952±0.011	0.802±0.002	0.798±0.002	0.794±0.001	0.792±0.001	0.801±0.002	0.370±0.015	0.380±0.014	0.403±0.006	0.413±0.005
S	0.867±0.001	0.829±0.001	0.822±0.001	0.819±0.001	0.817±0.001	0.826±0.000	0.412±0.003	0.404±0.004	0.425±0.001	0.426±0.001
Sp	0.890±0.003	0.810±0.002	0.807±0.001	0.805±0.001	0.804±0.001	0.813±0.001	0.317±0.015	0.288±0.014	0.387±0.005	0.396±0.004
SQ	0.868±0.002	0.827±0.002	0.822±0.001	0.818±0.001	0.816±0.001	0.825±0.001	0.415±0.003	0.407±0.003	0.425±0.001	0.426±0.001
SQp	0.884±0.004	0.812±0.002	0.808±0.001	0.805±0.001	0.805±0.001	0.814±0.000	0.358±0.016	0.328±0.015	0.402±0.007	0.409±0.006
SX	0.835±0.002	0.838±0.001	0.831±0.001	0.826±0.000	0.825±0.000	0.833±0.000	0.410±0.003	0.403±0.002	0.422±0.001	0.425±0.001
SXp	0.934±0.005	0.807±0.004	0.801±0.003	0.797±0.003	0.795±0.002	0.803±0.002	0.297±0.013	0.286±0.013	0.381±0.003	0.390±0.002

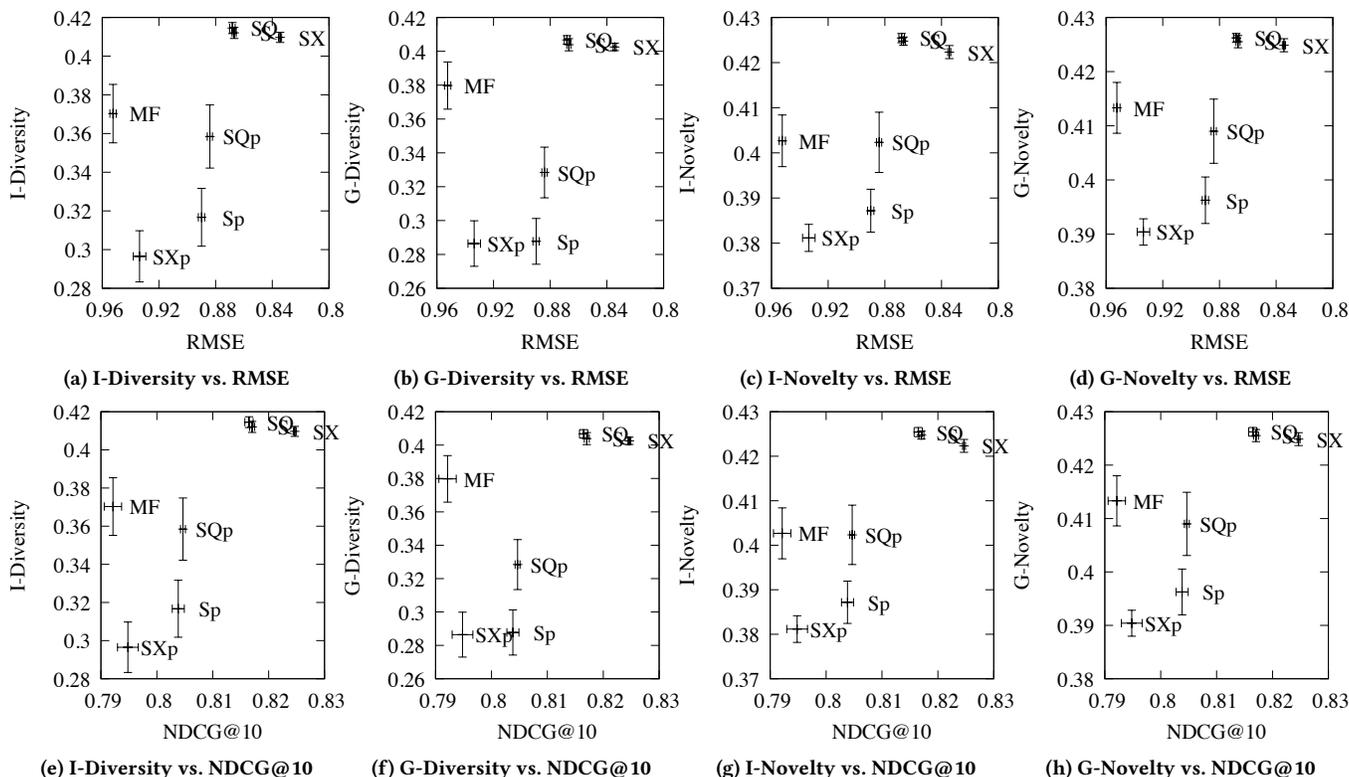


Figure 1: Trade-off between Accuracy and Diversity/Novelty

Regarding the diversity and novelty metrics, we also observe an agreement across metrics. The important observations here are the following. First, average social regularization, contrary to what one might expect, results in increased diversity and novelty, compared to base MF. It is only pairwise social regularization that suffers in this regard. Given that pairwise social regularization does not greatly improve accuracy, and at the same time degrades diversity/novelty, we conclude that average social regularization is preferable. Second, among the average social regularization recommenders (S, SQ, SX), there is no clear winner as all exhibit statistically indistinguishable values of diversity and novelty.

Figure 1 illustrates the trade-off between accuracy and diversity/novelty. Note that the axis for RMSE is reversed with better (lower) values on the right. Therefore, in all figures, the best recommender is the one that is closer to the upper right corner of the

plane. It is evident that SX offers the best accuracy by a considerable margin, while having diversity/novelty on par with the best.

5 CONCLUSION

Social-based recommenders exploit the effects of homophily and social influence among users to improve the accuracy of standard collaborative filtering. This work has studied the effect of such systems in the diversity and novelty of the recommendations they make. Our results indicate that a certain approach (average social regularization) can actually increase the diversity and novelty of users when measured individually, and when examined with respect to the social groups to which users belong. Moreover, our proposed social-based recommender results in significantly more accurate recommendations while not sacrificing diversity and novelty.

REFERENCES

- [1] Panagiotis Adamopoulos and Alexander Tuzhilin. 2014. On Unexpectedness in Recommender Systems: Or How to Better Expect the Unexpected. *ACM TIST* 5, 4 (2014), 54:1–54:32.
- [2] Jaime G. Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *ACM SIGIR*. ACM, 335–336.
- [3] Pablo Castells, Neil J. Hurley, and Saul Vargas. 2015. Novelty and Diversity in Recommender Systems. In *Recommender Systems Handbook*. Springer, 881–918.
- [4] 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.
- [5] 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>
- [6] 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>
- [7] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *ACM RecSys*. ACM, 257–260.
- [8] 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>
- [9] Glen Jeh and Jennifer Widom. 2002. SimRank: a measure of structural-context similarity. In *ACM KDD*. ACM, 538–543.
- [10] Michael Jugovac, Dietmar Jannach, and Lukas Lerche. 2017. Efficient optimization of multiple recommendation quality factors according to individual user tendencies. *Expert Syst. Appl.* 81 (2017), 321–331.
- [11] 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.
- [12] 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.
- [13] Hui Li, Dingming Wu, Wenbin Tang, and Nikos Mamoulis. 2015. Overlapping Community Regularization for Rating Prediction in Social Recommender Systems. In *ACM RecSys*. 27–34. <https://doi.org/10.1145/2792838.2800171>
- [14] Hao Ma. 2013. An experimental study on implicit social recommendation. In *ACM SIGIR*. 73–82. <https://doi.org/10.1145/2484028.2484059>
- [15] 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>
- [16] 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>
- [17] 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>
- [18] 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
- [19] Paolo Massa and Paolo Avesani. 2007. Trust-aware recommender systems. In *ACM RecSys*. 17–24. <https://doi.org/10.1145/1297231.1297235>
- [20] Sean M. McNee, John Riedl, and Joseph A. Konstan. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI Extended Abstracts*. ACM, 1097–1101.
- [21] 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.
- [22] Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren G. Terveen, and Joseph A. Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *WWW*, Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel (Eds.). ACM, 677–686. <https://doi.org/10.1145/2566486.2568012>
- [23] Dimitrios Rafailidis and Fabio Crestani. 2017. Learning to Rank with Trust and Distrust in Recommender Systems. In *ACM RecSys*. ACM, 5–13.
- [24] Marco Túlio Ribeiro, Nivio Ziviani, Edleno Silva de Moura, Itamar Hata, Anísio Lacerda, and Adriano Veloso. 2014. Multiobjective Pareto-Efficient Approaches for Recommender Systems. *ACM TIST* 5, 4 (2014), 53:1–53:20.
- [25] 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>
- [26] Nava Tintarev, Emily Sullivan, Dror Guldin, Sihang Qiu, and Daan Odjik. 2018. Same, Same, but Different: Algorithmic Diversification of Viewpoints in News. In *UMAP (Adjunct Publication)*. ACM, 7–13.
- [27] Saul Vargas and Pablo Castells. 2011. Rank and relevance in novelty and diversity metrics for recommender systems. In *ACM RecSys*. ACM, 109–116.
- [28] 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>
- [29] Huan Zhao, Quanming Yao, James T. Kwok, and Dik Lun Lee. 2017. Collaborative Filtering with Social Local Models, 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>
- [30] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *WWW*. ACM, 22–32.