Towards Diversity in Recommendations using Social Networks

Swapneel Sheth, Jonathan Bell, Nipun Arora, Gail Kaiser
Department of Computer Science, Columbia University, New York, NY 10027
{swapneel, jbell, nipun, kaiser}@cs.columbia.edu

ABSTRACT

While there has been a lot of research towards improving the accuracy of recommender systems, the resulting systems have tended to become increasingly narrow in suggestion variety. An emerging trend in recommendation systems is to actively seek out diversity in recommendations, where the aim is to provide unexpected, varied, and serendipitous recommendations to the user. Our main contribution in this paper is a new approach to diversity in recommendations called "Social Diversity," a technique that uses social network information to diversify recommendation results. Social Diversity utilizes social networks in recommender systems to leverage the diverse underlying preferences of different user communities to introduce diversity into recommendations. This form of diversification ensures that users in different social networks (who may not collaborate in real life, since they are in a different network) share information, helping to prevent siloization of knowledge and recommendations. We describe our approach and show its feasibility in providing diverse recommendations for the MovieLens dataset.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Information Filtering*

General Terms

Algorithms, Human Factors

Keywords

Recommender Systems, Collaborative Filtering, Diversity, Serendipity

1. INTRODUCTION

Recommender Systems have become increasingly commonplace. Systems like Amazon [2], Facebook [8], Last.fm [15], and Netflix [20] provide recommendations in a variety of domains and have millions of users. As interest in recommender systems has grown in industry, there has been an increasing research effort surrounding recommender systems in academia. Much research has focused on recommendation techniques for groups of users at a time [5, 11, 12], applications of recommender systems to novel domains [4, 13, 19], and improving the recommendation accuracy [10, 21, 25]. However, as accuracy has become an increasingly important criteria, recommender systems have become narrow in suggestion variety, falling victim to the "portfolio effect" [1] - preferring "safe" recommendations that are more similar to users' past activities over "serendipitous" recommendations, yielding many suggestions for largely similar items.

User studies have also shown that recommendation accuracy alone does not always result in high user satisfaction: users are most interested in how meaningful a recommendation is, regardless of its "accuracy" [16, 29]. A new emerging research trend in recommender systems is to seek diversity, which introduces new and relevant recommendations that differ from the rest of the recommendation pool [17, 29].

In this paper, we propose a new approach to diversity in recommendations using social networks that we call "Social Diversity". By utilizing the underlying diversity of preferences between different social networks, we are able to bring a new form of diversity to recommender systems. Social Diversity uses the inherently diverse interests of different groups of people as a kernel to introduce diverse recommendations.

Our work on Social Diversity is motivated by our three previous recommender systems: genSpace [19], which is targeted towards the domain of bioinformatics; Retina [18], which is targeted towards CS1 (introduction to programming) courses; and COMPASS [22], which is targeted towards providing multi-core optimizations to programmers. Each of these systems is community-driven and uses collaborative filtering (CF) algorithms, which predict user interests based upon their history, and the history of other users. While extracting a reference architecture and discovering best practices from these systems [23], we realized that an important addition to any such system would be a domain-agnostic approach to add recommendation diversity. We aim to provide recommendations that can cross barriers between social groups and allow users to explicitly request diverse recommendations that they would normally not get, such as "Show me what genomics analysis tools are being used by researchers in Europe," "Show me what kinds of errors students from other sections had," or "Show me the parallelization techniques that people at some other organization are using."

We found that most existing systems do not directly address questions of this nature (e.g., regarding different social networks) and this has motivated our approach to diversity, which we discuss in Section 2. We describe our empirical results in Section 3 and finally, we conclude the paper with an overview of related work in Section 4.

2. APPROACH

Recommender systems typically utilize collaborative filtering (CF) or content-based algorithms for generating recommendations. In CF systems, recommendations are generated by either finding "people like you," or items similar to those that you like. Content-based systems generate recommendations using item metadata such as genre or language of a movie. Without a method for adding diversity, these systems typically become narrow in suggestion variety but recent research has begun to explore solutions to this increasingly important area.

Our approach is to create Social Diversity in recommendations by utilizing users' memberships in social networks. These social networks can both be "real world" groups: networks created and explicitly joined by users, or "virtual" groups: meta groups automatically created by the system from usage and/or demographic information. Therefore, we do not require the presence of self-defined social networks to implement Social Diversity, and our approach is general enough to apply to most domains. Any CF recommender system that has additional user information such as demographics (age, gender, location) can benefit from our approach. The definition of these networks would, of course, vary with the domain. We would use both of these types of social networks to filter recommendations and provide diversity. Users could choose the source of their recommendations: globally from all users in the system, from members of virtual groups, from users in real world groups, or from any combination of sources.

The proportion of diverse results provided would be controlled by the system and/or by the individual users. Users could choose to increase or decrease the amount of diversity as needed. Diversity could be introduced automatically by the system, or directly in response to a request from a user. That is, our approach would allow users to specifically request to have their recommendations influenced by those for a specific social group. Groups could be automatically selected in many ways. For example, we could select groups such that we ensure a full spectrum of preferences are represented (combatting siloization), or to specifically show that preferences from similar (but not the same) groups are presented. Note that concerns and research challenges arising due to the privacy ramifications of Social Diversity are outside of the scope of this paper and we leave them to future work.

Consider a hypothetical collaborative filtering-based movie recommender system X similar to Netflix in which users are offered movie recommendations based on past movie rating history. Imagine that we have a user, "Francesco" who likes watching a certain kind of movie - Romantic Comedies - and he has watched (and liked) movies such as Notting Hill, You've Got Mail, and There's Something About Mary. Typical CF recommender systems will provide other similar Romantic Comedies as recommendations (an example of the Portfolio Effect [1]). This may make it very hard for Francesco to "broaden his horizons" with other kinds of movies.

Our approach using social networks would provide Francesco with diverse movie recommendations. A Socially Diverse recommender system will ensure that throughout a users' session, in addition to "normal" recommendations, he also receives suggestions from different, randomly selected social groups. This approach works to counteract information siloization - the clustering of preferences within specific social groups. In addition to standard recommendations, Francesco could ask to see, for example, movies that other students in the age group of 18-24 like. Francesco would not even need to be part of a particular social network to receive recommendations from it, if he wants - but he could still use them to broaden his horizons and intentionally diversify his personal recommendations. If he wants to watch movies that people in Minsk, Belarus or Abidjan, Ivory Coast watch or movies watched by people who like Foreign Language Films, he is free to do so (assuming, of course, that there is sufficient user data in these areas). Again, these diversified recommendations could also be pushed automatically to Fred without his implicit request so as to suggest new avenues of interest to him.

In Section 4, we compare Social Diversity to other existing approaches towards diversity in recommender systems. In the empirical results that follow, we describe our implementation details and show the feasibility of our approach to provide movie recommendations with the commonly-used MovieLens dataset.

3. EMPIRICAL RESULTS

We conducted evaluations of our approach offline, using the one hundred thousand ratings MovieLens dataset from the GroupLens Research Project at the University of Minnesota [9]. This dataset is commonly used in recommender system evaluations [14, 28] and includes 100,000 ratings from 943 users across 1682 movies. Each user is tagged with their age, gender, occupation, and zip code.

The goal of our evaluation was to validate that Social Diversity introduces real diversity into recommendation lists. We leave the evaluation of user satisfaction to a future user study, described briefly in Section 5. In our experiment, we created a simple user-based recommender system, which provided movie recommendations based upon what similar users watched. We then introduced diversity by simulating the user's desire to see results from another social group. For each user, we selected a random social group (e.g., 40-50 year old lawyers from New York), and made sure that some percentage of the recommendations were directed towards that social group. We varied the fraction of diverse recommendations from 0 to 100 percent and calculated metrics as described below. Each metric captures how successfully we are at perturbing the resulting recommendation list, and are commonly used in other work as described in Section 4.

3.1 Metrics

The most common evaluation metrics used are precision, recall, and fall-out. Precision refers to the fraction of items recommended that are relevant to the user, and is calculated as:

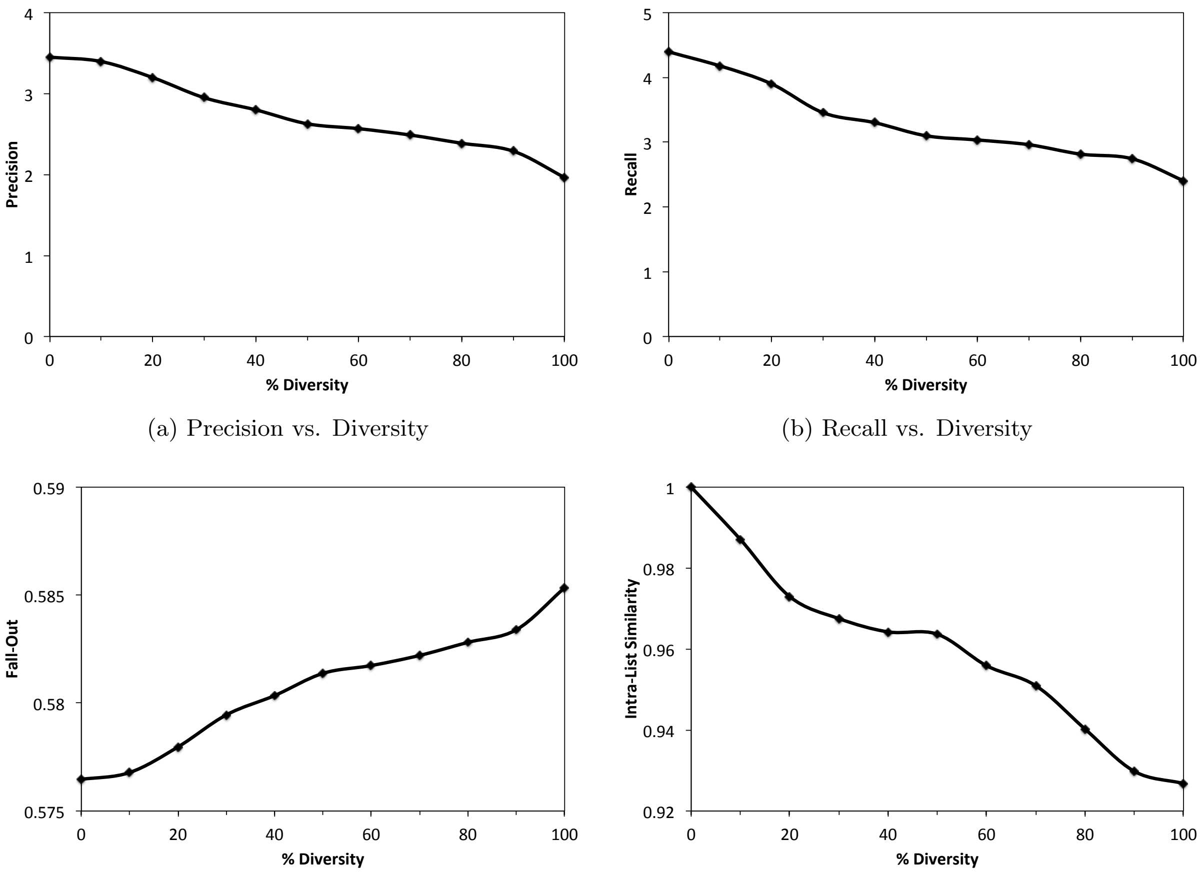

(a) Precision vs. Diversity

(b) Recall vs. Diversity

(c) Fall-Out vs. Diversity

(d) ILS vs. Diversity

Figure 1: MovieLens Empirical Results

$$\text{Precision} = 100 \times \frac{|\{\text{relevant items}\} \cap \{\text{all suggestions}\}|}{|\{\text{all suggestions}\}|} \quad (1)$$

Recall captures the fraction of relevant items recommended (out of all possible relevant items), and is calculated as:

$$\text{Recall} = 100 \times \frac{|\{\text{relevant items}\} \cap \{\text{all suggestions}\}|}{|\{\text{all relevant items}\}|} \quad (2)$$

Diverse recommender systems are expected to have poor precision and recall as by definition, a diverse recommender system provides results that would not be classified as relevant by any traditional means. Therefore, the expectation is that as we introduce more diverse recommendations, precision and recall will decrease, similar to results observed by Ziegler et al. [29].

Fall-out is closely related to recall and captures the proportion of non-relevant items recommended (out of all possible non-relevant items) and is calculated as:

$$\text{Fall-out} = 100 \times \frac{|\{\text{non-relevant items}\} \cap \{\text{all suggestions}\}|}{|\{\text{non-relevant suggestions}\}|} \quad (3)$$

Our expectation is that as diversity increases, fall-out will increase as well, as it is measuring the proportion of recommendations returned that would not be typically classified as relevant.

We also use a variant of the Intra-List Similarity (ILS) metric as defined by Ziegler et al. [29]. We propose calculating ILS as an average value, rather than as the absolute sum suggested by Ziegler, so as to allow for easier comparison between users. This metric captures the average similarity between each item that a user has indicated a positive preference for and each suggested item. We calculate the similarity between two items as the euclidean distance between their ratee's vectors.

$$\text{ILS} = 100 \times \frac{\Sigma_{u \in P} \Sigma_{m \in R_u} \Sigma_{i \in u_p} D(m,i)}{\Sigma_{u \in P} |R_u|} \quad (4)$$

3.2 MovieLens Results

We created and executed our empirical results in the Apache Mahout toolkit [3], on a computer with Java 1.6.0_24, a 2.8 GHz Intel Core i7 and 8GB of RAM. We calculated precision, recall, and fall-off at 10 - that is, we withheld the top 10 preferences from each user for evaluation purposes, and then found the percentage of those 10 items that appeared in the top 10 recommendations for that user. We ran the experiment 100 times and averaged the results so as to hide

	Diversity Approach	Recommendation Class	Evaluation Data	Evaluation Metrics
Bradley and Smyth [6]	Bounded Random Selection, Greedy Selection, Bounded Greedy Selection	Content Based	JobFinder.ie	Similarity, Diversity, Relative Benefits
Ziegler et al. [29]	Topic Diversity	Content Based, CF	Book Crossing	Precision, Recall, ILS, User Study
Yu et al. [27]	Explanation based Diversity	Content Based, CF	del.icio.us and Yahoo Movies dataset	Pair-wise Result Set Distance
Kelly and Bridge [14]	Bounded-Greedy Algorithm	Conversational CF	MovieLens	Average number of recommendations to reach success
Zhang and Hurley [28]	Statistical Model to provide choices to developers	top-N Recommender Systems	MovieLens	Concentration Curve of Hit Distribution and Diversity, Similarity and Precision Distributions
Our approach	Social Diversity	CF	MovieLens	Precision, Recall, Fall-Out, ILS

Table 1: Comparison of Diversity in Recommendations

the randomness from selecting each users' top 10 movie selections (in the case of ties in the top 10). Note that when diversity is at 0%, we are evaluating a purely non-diverse CF system, and at 100%, we are providing only diverse results.

3.2.1 Precision and Recall

We analyzed the effect of our diversification method on precision and recall by varying the percentage of recommendations that we diversified. Figure 1a shows the effect of diversification on precision and Figure 1b shows the effect of diversification on recall. In both figures, the x-axis displays the percentage of diversity and the y-axis displays precision and recall, respectively. We observe that as the amount of diversity increases, both precision and recall decrease. As we go from 0% diversity to 100% diversity, precision reduces by 43.11% and recall decreases by 45.23%.

Thus, in general, our empirical results show that the accuracy of our recommender system decreases as we add diversity. These results are as expected and similar to the results shown by Ziegler et al., suggesting that we are achieving diversity. By being able to replicate their results from a Topic Diversity approach with our approach, we show the general feasibility of using Social Diversity in recommender systems.

3.2.2 Fall-Out and ILS

We also analyzed how effective our algorithm is by using two other metrics - fallout and intra-list similarity. These metrics capture how our diversification algorithm perturbs the results of a standard recommender system.

Figure 1c shows the effect of diversification on fall-out and Figure 1d shows the effect of diversification on ILS. The x-axis displays the percentage of diversity and the y-axis displays fall-out and ILS, respectively. We observe that the fall-out rate increases by 1.54% as diversity increases from 0% to 100% as we are showing more movie suggestions that are less traditionally relevant to the user. We also observe that the intra-list similarity decreases by 7.3%, confirming that Social Diversification introduces diversity to recommendations.

3.3 Discussion

In our studies, we observed a significantly higher percent difference in recall and precision (45.23% and 43.11%, respectively) than in fall-out and ILS. We attribute the greater change in accuracy and precision (as compared to fall-out) to the metrics themselves: fall-out is taken as a ratio to the number of non-relevant suggestions, while accuracy and precision are ratios to the number of relevant suggestions. In our experimental setup (and in general) there are usually far fewer relevant suggestions than non-relevant ones and hence, the percentage difference should be smaller for fall-out. We attribute the relatively small change in ILS to the inherently limited difference in movie preferences between the different demographic groups in the data set (e.g., there are "blockbusters" that ALL social groups enjoyed).

One potential limitation of our system is that it leverages the inherent diversity of preferences that exist in social networks. If this diversity of preferences did not exist, our approach towards diverse recommendations would be limited. We, however, believe that such diversity is inherently present and the feasibility of our approach is validated by our empirical results, which replicate previous work by Zeigler et al. [29].

4. RELATED WORK

This section describes the previous work done in the recommender systems community addressing diversification of recommendations. We summarize the related work in Table 1, which shows various diversity mechanisms and compares them to our approach.

Previous work in recommendation diversity has stressed the importance of diversification in recommendations [24]. User studies have shown that accuracy alone does not guarantee high user satisfaction in recommender systems [16, 29], resulting in several approaches that aim to introduce diversity.

Bradley and Smyth [6] focus only on content-based filtering systems and introduce three algorithms for diversity. Their overall strategy focuses on delivering recommendations ranked on a "quality" basis, optimizing the tradeoff between similarity and diversity. By avoiding outside information (e.g., metadata about items or users), this approach benefits from very low overhead, but potentially suffers from not fully leveraging all available information.

Zeigler et al. [29] focus on CF recommender systems and introduce the notion of "Topic Diversification". They ensure diversity by balancing suggestions across topics (categories) using a diversity metric called "Intra-List Similarity" to assess the degree of diversity in their recommendation lists. This approach has the potential to provide more relevant diversity (by clustering items into topics), but may have a higher overhead than a standard CF system from performing this categorization.

Kelly and Bridge [14] explore diversification algorithms for conversational (feedback based) CF systems. They propose a greedy algorithm for diversification that calculates the pairwise difference between each item in the result set [24]. This approach increases result quality from user-feedback while maintaining a notion of diversity, but is limited only to conversational collaborative filtering recommender systems.

Yu, Lakshmanan, and Amer-Yahia [26, 27] introduce Explanation-based Diversification, which focuses on providing diversity to recommender systems lacking any real-world attribute information. They use heuristic algorithms based on distance measures between "explanations", to increase diversity in recommendations. An inherent limitation to this approach is that that explanations may generate false positives as items having different explanations need not be diverse.

Zhang and Hurley [28] propose a statistical model for diversity in recommendation systems using a concentration index, which measures the ability of an algorithm to recommend novel items. They analyze various algorithms using the concentration index to determine which algorithms are more suited towards diversity. As Zhang and Hurley focus on top-N recommender systems, their approach may not work, in general, for all the different kinds of recommendations.

Netflix [20] includes a "Local Favorites" feature, which gives recommendations from the user's current geographical region such as New York or Seattle. Users can also choose to get recommendations from other regions by specifying the state, city, or zip code. Using this feature, users can choose to receive a list of recommendations that is based on the preferences of users in a specific geographic area. This approach can be viewed as diversity of recommendations and it is similar to our approach. Our approach, however, is not limited to only geographic regions: it is generic and can use a variety of social information such as age, gender, and past and current affiliations for diversity. Moreover, Netflix's "Local Favorites" feature operates separately from the rest of the recommendations: it is impossible for a Netflix user to have these results mixed in with his or her primary recommendation list.

5. CONCLUSION

Recommender systems are typically evaluated through precision and recall statistics. While precision and recall can be used to demonstrate accuracy of a recommender system, they ignore the possibility for serendipitous suggestions. Other metrics such as ILS and fall-out can help to evaluate the diversity of recommendation lists, but only capture how much they may differ from traditional recommendations, without capturing the usefulness of the recommendations. An important challenge with diversity in recommender systems is evaluating the efficacy and benefits of the various diversification algorithms and unfortunately, there is no agreed upon metric for this evaluation yet (akin to accuracy for normal recommendations). The perfect metric for measuring recommender diversity is user satisfaction: how much users actually like the results being returned. We would ideally like to be able to empirically evaluate this metric without having to conduct user studies, which can be time consuming.

We are in the process of creating a live experiment with our system genSpace, studying the usefulness of our socially diverse recommendations in a bioinformatics application. genSpace is a plugin to an open-source Java-based platform for integrated genomics research called "geWorkbench" [7]. geWorkbench is used by researchers in computational biology and bioinformatics to run complex analyses on large data sets such as DNA, protein, and gene sequences. It contains over 70 different tools, making it potentially difficult for a new user to determine what tools to use and in what order. genSpace uses CF and assists users by providing recommendations such as the next analysis tool to use.

We have begun integrating Social Diversity into our existing recommender system for genSpace, and are planning to evaluate the benefits of our approach to diversity in this context. We are particularly interested in investigating the most effective networks (real or virtual) to serve as sources for diversity. We hope that our experiments and user studies will contribute to a better understanding of user needs and result in a metric to capture them.

In this paper, we introduced "Social Diversity" in recommender systems, which uses social and demographic information to provide serendipitous recommendation results. We validated the feasibility of our approach to diversity with an empirical study on the MovieLens dataset. Finally, we discussed the metrics available for evaluating diverse recommender systems and feel that understanding user needs is essential to building and evaluating diverse recommender systems. We believe that this area will provide useful avenues for further research.

6. ACKNOWLEDGEMENTS

The authors are members of the Programming Systems Laboratory, funded in part by NSF CNS-0717544, CNS-0627473 and CNS-0426623, and NIH 2 U54 CA121852-06.

7. REFERENCES

[1] K. Ali and W. van Stam. Tivo: making show recommendations using a distributed collaborative filtering architecture. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '04, pages 394–401, New York, NY, USA, 2004. ACM.

[2] Amazon.com. `http://www.amazon.com`.

[3] Apache. Mahout. `http://mahout.apache.org/`.

[4] A. Begel, K. Y. Phang, and T. Zimmermann. Codebook: discovering and exploiting relationships in software repositories. In *ICSE '10: Proceedings of the*

32nd ACM/IEEE International Conference on Software Engineering, pages 125–134, New York, NY, USA, 2010. ACM.

[5] S. Berkovsky and J. Freyne. Group-based recipe recommendations: analysis of data aggregation strategies. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, pages 111–118, New York, NY, USA, 2010. ACM.

[6] K. Bradley and B. Smyth. Improving recommendation diversity. In *Proceedings of the Twelfth Irish Conference on Artificial Intelligence and Cognitive Science, Maynooth, Ireland*, pages 85–94, 2001.

[7] A. Califano, A. Floratos, M. Kustagi, and J. Watkinson. geWorkbench: An Open-Source Platform for Integrated Genomics. `http://www.geworkbench.org`.

[8] Facebook. `http://www.facebook.com`.

[9] GroupLens Research Group. Movielens data sets. `http://grouplens.org/node/12`.

[10] A. Gunawardana and C. Meek. Tied boltzmann machines for cold start recommendations. In *RecSys '08: Proc. of the 2008 ACM conference on Recommender systems*, pages 19–26, 2008.

[11] I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, and S. Ofek-Koifman. Personalized recommendation of social software items based on social relations. In *Proceedings of the third ACM conference on Recommender systems*, RecSys '09, pages 53–60, New York, NY, USA, 2009. ACM.

[12] J. Hannon, M. Bennett, and B. Smyth. Recommending twitter users to follow using content and collaborative filtering approaches. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, pages 199–206, New York, NY, USA, 2010. ACM.

[13] R. Holmes, T. Ratchford, M. P. Robillard, and R. J. Walker. Automatically recommending triage decisions for pragmatic reuse tasks. In *Proceedings of the 24th IEEE/ACM International Conference on Automated Software Engineering*, pages 397–408, 2009.

[14] J. P. Kelly and D. Bridge. Enhancing the diversity of conversational collaborative recommendations: a comparison. *Artif. Intell. Rev.*, 25:79–95, April 2006.

[15] Last.fm. `http://www.last.fm`.

[16] S. M. McNee, I. Albert, D. Cosley, P. Gopalkrishnan, S. K. Lam, A. M. Rashid, J. A. Konstan, and J. Riedl. On the recommending of citations for research papers. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, CSCW '02, pages 116–125, New York, NY, USA, 2002. ACM.

[17] S. M. McNee, J. Riedl, and J. A. Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI '06 extended abstracts on Human factors in computing systems*, CHI '06, pages 1097–1101, New York, NY, USA, 2006. ACM.

[18] C. Murphy, G. E. Kaiser, K. Loveland, and S. Hasan. Retina: Helping Students and Instructors Based on Observed Programming Activities. In *Proc. of the 40th ACM SIGCSE Techn. Symp. on CS Education*, pages 178–182, March 2009.

[19] C. Murphy, S. Sheth, G. Kaiser, and L. Wilcox. genSpace: Exploring Social Networking Metaphors for Knowledge Sharing and Scientific Collaborative Work. In *1st International Workshop on Social Software Engineering and Applications (SoSEA)*, pages 29–36, September 2008.

[20] Netflix. `http://www.netflix.com`.

[21] Y.-J. Park and A. Tuzhilin. The long tail of recommender systems and how to leverage it. In *RecSys '08: Proc. of the 2008 ACM Conf. on Recommender systems*, pages 11–18, 2008.

[22] S. Sethumadhavan, N. Arora, R. B. Ganapathi, D. John, and G. E. Kaiser. Compass: A community-driven parallelization advisor for sequential software. In *Proceedings of the 2009 ICSE Workshop on Multicore Software Engineering*, IWMSE '09, pages 41–48, Washington, DC, USA, 2009. IEEE Computer Society.

[23] S. Sheth, N. Arora, C. Murphy, and G. E. Kaiser. weHelp: A Reference Architecture for Social Recommender Systems. In *Software Engineering (Workshops)*, pages 461–472, 2010.

[24] B. Smyth and P. McClave. Similarity vs. diversity. In *Proceedings of the 4th International Conference on Case-Based Reasoning: Case-Based Reasoning Research and Development*, ICCBR '01, pages 347–361, London, UK, 2001. Springer-Verlag.

[25] A. Umyarov and A. Tuzhilin. Improving rating estimation in recommender systems using aggregation- and variance-based hierarchical models. In *Proceedings of the third ACM conference on Recommender systems*, RecSys '09, pages 37–44, New York, NY, USA, 2009. ACM.

[26] C. Yu, L. Lakshmanan, and S. Amer-Yahia. It takes variety to make a world: diversification in recommender systems. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT '09, pages 368–378, New York, NY, USA, 2009. ACM.

[27] C. Yu, L. V. S. Lakshmanan, and S. Amer-Yahia. Recommendation diversification using explanations. In *Proceedings of the 2009 IEEE International Conference on Data Engineering*, pages 1299–1302, Washington, DC, USA, 2009. IEEE Computer Society.

[28] M. Zhang and N. Hurley. Statistical modeling of diversity in top-n recommender systems. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 01*, WI-IAT '09, pages 490–497, Washington, DC, USA, 2009. IEEE Computer Society.

[29] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, WWW '05, pages 22–32, New York, NY, USA, 2005. ACM.