Exploriometer: Leveraging Personality Traits for Coverage and Diversity Aware Recommendations

Chatzicharalampous Evangelos
Informatics Dept
Aristotle University
54124 Thessaloniki, Greece
evang.chatz@gmail.com

Zigkolis Christos
Informatics Dept
Aristotle University
54124 Thessaloniki, Greece
chzigkol@csd.auth.gr

Vakali Athena
Informatics Dept
Aristotle University
54124 Thessaloniki, Greece
avakali@csd.auth.gr

ABSTRACT

Since the first introduced Collaborative Filtering Recommenders (CFR), there have been many attempts to improve their performance by enhancing the prediction accuracy. Even though rating prediction is the prevailing paradigm in CFR, there are other issues which have gained significant attention with respect to the content and its variety. **Coverage**, which constitutes the degree to which recommendations cover the set of available items, is an important factor along with **diversity** of the items proposed to an individual, often measured by an average dissimilarity between all pairs of recommended items.

In this paper, we argue that coverage and diversity cannot be effectively addressed by conventional CFR with pure similarity-based neighborhood creation processes, especially in sparse datasets. Motivated by the need for including wider content characteristics, we propose a novel neighbor selection technique which emphasizes on variety in preferences (to cover polyphony in selection). Our approach consists of a new metric, named "**Exploriometer**", which acts as a personality trait for users based on their rating behavior. We favor users who are explorers in order to increase polyphony, and subsequently coverage and diversity; but we still select similar users when we create neighborhoods as a solid basis in order to keep accuracy levels high. The proposed approach has been experimented by two real-world datasets (MovieLens[1] and Yahoo! Music[2]) with coverage, diversity and accuracy aware recommendations extracted by both traditional CFR and CFR enhanced with our neighborhood creation process. We also introduce a new metric, inspired by the Pearson Correlation Coefficient, to estimate the diversity of recommended items. The derived results demonstrate that our neighbor selection technique can enhance coverage and diversity of the recommendations, especially on sparse datasets.

Categories and Subject Descriptors

H.2.8 [Information Systems]: Data Mining

[1] Movielens 10 M—http://www.grouplens.org/node/73

[2] Yahoo! Movies R4—http://webscope.sandbox.yahoo.com/

WWW 2015 Companion, May 18–22, 2015, Florence, Italy.
ACM 978-1-4503-3473-0/15/05.
http://dx.doi.org/10.1145/2740908.2742140

General Terms

Algorithms, Experimentation.

Keywords

Collaborative Filtering Systems, Diversity, Coverage, Neighbor Selection.

1. INTRODUCTION

Collaborative Filtering Recommenders (CFR) rely on the assumption that *similar* users may exhibit *similar* preferences and, given that, they produce personalized suggestions with information originating from like-minded users [1]. In such cases, identifying the ideal peers for a target user would result in more accurate recommendations. Following this rationale, CFR have showed great power in predicting ratings over the years, thus achieving high levels of accuracy against other types of recommenders [2]. However, improving accuracy only has proven to be insufficient in the attempt to increase the quality of the produced recommendations [3] [4].

Moving beyond accuracy metrics, researchers have focused on other important concepts such as **Coverage** and **Diversity**: the former defines the degree to which final recommendations cover the entire set of available items [5], while the latter defines the variety of recommended items [6] [7] [8]. A recommender with high coverage delivers to users a more detailed and careful recommendation based on wider investigation of the item space. Diversity, on the other hand, is supposed to let the recommender to act more dynamically and lively by providing non-trivial recommendations [6] [7]. This work is motivated by the fact that both concepts can contribute on significant improvements in recommendations' quality indicators.

Traditional collaborative recommenders have proven to be inadequate to meet such needs and the main reason for that lies in their way of selecting the so-called like-minded peers. A great majority of CFR reach to suggestions through user neighborhoods characterized by *"a consensus of tastes"* where the *k*-closest neighbors (i.e. users with similarity on tastes) are assembled together. Such a process reduces the variety of items which may be suggested, because like-minded neighbors tend to rate the same items, i.e. resulting in less coverage [3]. In addition, CFR lack on having neighborhoods with the so-called *"polyphony of tastes"*, which is a trait enabling the consideration of more diverse suggestions. These phenomena worsen when data are sparse, a common issue in most real-world rating datasets. The inherent inability of CFR to address sparse data magnifies the aforementioned drawbacks with respect to their neighborhood generation, since neighbors who show little knowledge of an item space would inevitably provide a very narrow perspective to the CFR of what users may like [9] [10].

Those shortcomings have been usually addressed by introducing randomness in the recommendation procedure [11], filtering out items which are too similar to items the user has already rated [12], or increasing the diversity of recommendations [8]. Also, artificial reduction of the sparsity has been proposed as an approach to increase the overall quality of the recommendations [10]. All these approaches attempt to improve the performance of the CFR by consolidating additional techniques to the recommendation generation progress. However, they can be seen as countermeasures to the foretold shortcomings, because none focuses on reconsidering the algorithm's rationale in terms of the user neighborhoods creation process.

In this paper, we argue that coverage and diversity cannot be effectively addressed by CFR with pure similarity-based neighborhood creation processes. Motivated by the need for polyphony, we propose an advanced approach of selecting neighbors with variety in tastes. Our approach introduces of a new metric, entitled "**Exploriometer**", which acts as a personality trait for users and is based on their rating behavior. A user's exploration trait is calculated as the mean rarity of his/her rated items when the level of rarity for each item is estimated by the number of users who have rated it. Users with low levels in this trait could be considered as "*Mainstream users*" who rate mostly popular items, while users with high values could be characterized as "*Domain Experts*" with a tendency to explore beyond the popular items. This work considers the latter as valuable members of a neighborhood, since they can widen CFR options when making recommendations.

Although accuracy has been criticized for its role in evaluation purposes, it still remains a significant metric for estimating recommender's quality. Therefore, in our proposed neighborhood creation process we select as neighbors, Domain Experts that also display high similarity with the user. The main contribution of this work is summarized in the following two major issues:

1. **Create neighborhoods with polyphony of tastes**: In this work, we search for explorers, i.e. users who have a broader view of the domain rather than preferring only popular/trivial items. We make use of such users to alter the synthesis of neighborhoods created by CFR in an attempt to provide recommendations. Our new enhanced neighbor selection process has the purpose to increase the levels of coverage and diversity without negatively affecting the accuracy of the recommendations.

2. **Validate the importance of diversity**: Aiming to provide a proof-of-concept, we proprose a novel metric, named *Recommendation Diversity*, to prove that the proposed recommendations are impacted and also characterized by variety of tastes.

We have carried out experiments with two real-world datasets (MovieLens and Yahoo! Music) and we have analyzed coverage, diversity and accuracy of the provided recommendations extracted by both traditional CFR and CFR enhanced with our neighborhood creation process. The evaluation of the proposed approach has been carried out with the use of existing metrics. In particular, we make use of prediction and catalogue coverage metrics [5], while for diversity we adopt the Mean Item Rarity [13] together with the proposed Recommendation Diversity metric. For measuring the accuracy, we make use of the Root Mean Square Error metric [14].

The rest of the paper is structured as follows. Section 2 briefly discusses several related studies. In Section 3, we formulate our proposal for optimizing the coverage and diversity of the Collaborative Systems. Also, we introduce the 'Recommendation Diversity' metric and describe how this metric estimates the diversity of items that form a recommendation. Section 4 describes the datasets used as well as our experimentation and showcases the results of the proposed methodology. Finally, in Section 6 conclusions and future work are highlighted.

2. RELATED WORK

Since the first collaborative filtering systems in the mid-90's [15] [16], there have been many attempts to improve their performance focusing mainly on rating prediction accuracy [17] [18]. Neighborhood selection techniques (e.g. similarity weighting of neighbors [19], top-N filtering and negative filtering [20] [17]) have played a vital role in such efforts. Although accuracy is one of the prevailing performance metrics in recommender systems, it is more than evident that higher predictive accuracy does not always correspond to higher levels of user satisfaction [21]. Therefore, there has been an increasing attention to other metrics such as Coverage and Diversity.

The intense focusing solely on rating prediction accuracy has revealed one of the most important problems in recommenders, i.e. the narrow rating prediction focus (over-specialization) [5]. Regarding this issue, empirical studies [22] indicated that consumers tend to choose diversity against popularity in recommendations. However, mainstream algorithms have a tendency to focus on certain parts of available item space and favor already popular items while completely dismissing long-tail items in most cases [7] [21]. Such phenomena of over-specialization result from low coverage and diversity and are often addressed by introducing randomness in the recommendation procedure [11], filtering out items which are too similar to items the user has already rated [12], or increasing the diversity of recommendations [8]. Interestingly, there have been studies [23] [24] where an inverted neighborhood model was presented. Such models are based on *k*-furthest neighbors to identify less ordinary neighborhoods in order to create more diverse recommendations by recommending items disliked by the least similar users.

In this paper, we argue that the over-specialization problem stems on the fact that the neighborhood selection process is typically similarity-based. For that reason, instead of introducing additional techniques to the recommendation creation process, as the ones mentioned above, we deal with over-specialization in recommenders by proposing an enhancement of the neighbor selection technique of CFR.

3. ENHANCED USER NEIGHBORHOODS

Traditional CFR search for strong similarity of tastes among users when creating neighborhoods. Although this offers a good basis to rely on in order to highlight accurate suggestions, the quality of a recommender cannot be ensured by accuracy itself. We argue that CFR should search for other neighborhood characteristics as well, thus, we propose the so-called "polyphony of tastes" as an additional prerequisite. This polyphony can be reached by incorporating 'explorers' (i.e. Domain Experts) into similarity-based neighborhoods. Therefore, we are interested in forming groups which exhibit strong similarity, but also contain users who have strong explorer's activity. To achieve that, our neighbor selection technique identifies the users who are 'explorers' and then selects as neighbors the 'explorers' who are most similar to the user. To approach this objective, we introduce the "Exploriometer", a metric that quantifies a level of a user's exploration activity (i.e. rating) with respect to rare items.

In order to provide a proof-of-concept for the Exploriometer and to quantify the "polyphony of tastes" among recommended items, we propose a metric which depends on users' ratings. We call this metric "Recommendation Diversity" and it is estimated by the sum of the reversed values of similarity among all possible pairs of items, divided by the number of all pairs. Intuitively, we estimate the "polyphony of tastes" recommendations by examining how dissimilar from each other are the items on that recommendation. Our proposed metric is based on Pearson Correlation Coefficient [26].

3.1 Exploriometer

The "Exploriometer" metric quantifies the level of a user's 'exploratory' actions (rates) on rare items. We consider users who have rated rare items useful as neighbors, since it is argued that one of the causes of low diversity is the long-tail items distribution [25] [6] [7]. Such statistics indicate that a small number of items are rated by a lot of users whereas the majority of items receive a small number of ratings. This phenomenon has a negative effect on the diversity of the collaborative systems due to the fact that those systems can only recommend items to a user that are rated by their neighbors and, thus rare items are less likely to be recommended. By selecting neighbors who are 'explorers', we enable the system to recommend rare items. For the estimation of the "Exploriometer" we initially define the rarity of an item.

Lemma 1: Given that the rarity of the *i*-th item is expressed by the next formula:

$$rarity(i) = 1 - \frac{|r(i)| - arg\min_j |r(j)|}{arg\max_j |r(j)| - arg\min_j |r(j)|} \qquad (1)$$

where $|r(i)|$ is the number of ratings of the *i*-th item.

Proof: Suppose that an item x has the most ratings. It applies that $|r(x)| = arg\max_j |r(j)|$ and therefore

$$\begin{aligned} rarity(x) &= 1 - \frac{|r(x)| - arg\min_j |r(j)|}{arg\max_j |r(j)| - arg\min_j |r(j)|} \\ &= 1 - \frac{arg\max_j |r(j)| - arg\min_j |r(j)|}{arg\max_j |r(j)| - arg\min_j |r(j)|} = 1 - 1 \\ &= 0 \end{aligned}$$

On the other hand, suppose that the item y has the least ratings. It applies that $|r(y)| = arg\min_i |r(i)|$ and therefore

$$rarity(y) = 1 - \frac{|r(y)| - arg\min_j |r(j)|}{arg\max_j |r(j)| - arg\min_j |r(j)|} = 1 - \frac{arg\min_j |r(j)| - arg\min_j |r(j)|}{arg\max_j |r(j)| - arg\min_j |r(j)|} = 1 - 0 = 1.$$

In summary, the rarity metric is inversely proportional to the multitude of ratings and its range is [0, 1] where 0 correspond to the item with the most ratings (most popular) and 1 correspond to the item with the least ratings (most rare). The item rarity metric is used to calculate the "Exploriometer" as follows:

$$exploriometer(u) = \frac{\sum_{i \in I_u} rarity(i)}{|I_u|} \qquad (2)$$

where I_u is the set of items which are rated by user u and $|I_u|$ defines the number of items of this set. In principle, this metric represents the mean rarity of the items the user rated. Regarding performance, the complexity of calculating the above metric for all the users is $O(I \times U)$, where U is the total number of users and I is the total number of items.

3.2 Recommendation Diversity

As we mentioned earlier, we aim to increase recommendation coverage and diversity by enhancing the variety of tastes within the neighborhoods. In order to quantify how this enhancement affects the final recommendations, we propose a metric which estimates the diversity among the recommended items. We name this metric "Recommendation Diversity" and it incorporates Pearson Correlation Coefficient [26] as follows:

$$sim(i,j) = \frac{\sum_{u \in U_{i,j}} \left(ut(u,i) - \overline{ut}(i)\right)\left(ut(u,j) - \overline{ut}(j)\right)}{\sqrt{\sum_{u \in U_{i,i}} \left(ut(u,i) - \overline{ut}(i)\right)^2 \sum_{i \in I_{u,\dot{u}}} \left(ut(u,j) - \overline{ut}(j)\right)^2}}$$

where i, j denote the items that we are interested in their similarity, $U_{i,j}$ is the set of users who have rated both items i, j and $\overline{ut}(i)$ is the mean rating of item *i*. Intuitively, this metric considers the items to be similar if they received similar ratings from the users. We have considered this coefficient, because it is a metric widely adopted for the calculation of the similarity among items since it is used in item-based Collaborative Systems [2]. Therefore, based on this coefficient, we define "Recommendation Diversity" as follows:

$$RecDiv(rec(u)) = \frac{\sum_{i,j \in rec(u), i=1\ldots(|rec(u)|-1), j=i+1\ldots|rec(u)|} \frac{-sim(i,j)+1}{2}}{\sum_{i,j \in rec(u), i=1\ldots(|rec(u)|-1), j=i+1\ldots|rec(u)|} 1} \quad (3)$$

where rec(u) is the set of items recommended to user (recommendation) that we want to calculate the diversity and |rec(u)| is the number of those items. In essence, this formula is the sum of the reversed values of similarity among all the possible pairs of items of rec(u), divided by the number of all possible pairs. This metric is the mean dissimilarity of the recommended items with its estimation complexity being assessed to $O(I^2 \times U)$.

4. EXPERIMENTATION

In this section, we make use of two real-world rating datasets, one with movies and one with songs ratings, in order to evaluate our enhanced neighbor selection methodology. Initially, we provide details on these datasets and stating the reasons for choosing them. Then we present the experimentation procedure and conclude by describing and analyzing the derived results.

4.1 Datasets

Two datasets widely used in bibliography have been chosen. The one contains the user ratings of the movie recommendation system MovieLens and the other contains the user ratings of the music community Yahoo! Music. By using two datasets from different domains, we are proceeding to an experimentation which is not bound by a specific dataset or a specific domain.

The datasets properties are summarized in Table 1. Both our datasets follow the long tail distribution (see Figure 1), which is very common on ratings [25] [6]. Another characteristic of the input data that affects the performance of all recommendation systems is the density [9].

	MovieLens	Yahoo! Music
Multitude of ratings	10 millions	717 millions
Multitude of items	10681	136000
Multitude of users	71567	1.8 million
Ratings domain (with step)	0,5-5 (0.5)	1-5 (1)

Table 1: Datasets Properties

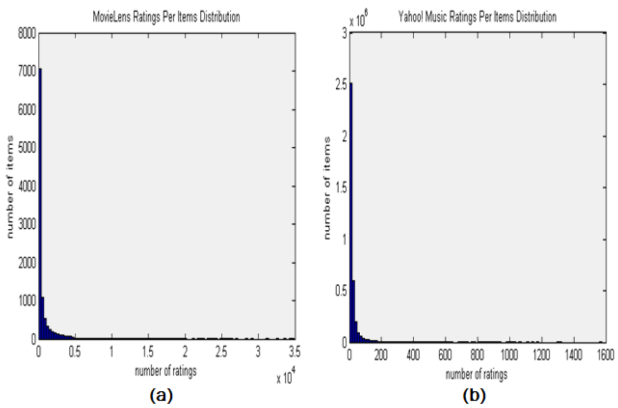

Figure 1: Ratings per item distribution for MovieLens (a) and Yahoo! Music (b) datasets

Density expresses the percentage of the total number of ratings ($I \times U$) which are known to the system (i.e. are included in the dataset) [27] and it is estimated by:

$$density = \frac{\#\ of\ nonzero\ entries}{Total\ \#\ of\ entries} \qquad (4)$$

where "# of nonzero entries" is the number of known ratings and "Total # of entries" is the size of $I \times U$. This characteristic affects the accuracy, the coverage and the diversity of recommendation systems, especially the CFR [9], because a sparse dataset does not contain many ratings per user, thus it is difficult to determine their likings. Also, a dateset with low density does not contain many ratings per item which makes it difficult to identify users with common ratings. MovieLens has *density* = 0.0130 and Yahoo! Movies has *density* = 0.0029. We can see that Yahoo! Movies is sparser than MovieLens making it interesting to observe their differences in performance.

4.2 Experimentation & Evaluation

In order to assess the effectiveness of the proposed methodology we split the datasets to training set (given as input to the system) and test set (used for evaluation). We analyze the difference of the produced recommendations between a conventional CFR and three CFR that use our modified neighbor selection technique. The evaluation is executed in both datasets and there are 3 different splits of the datasets to training-test set (25%-75%, 50%-50%, 75%-25%), the average results from those 3 splits are presented below.

For each CFR that uses our modified neighbor selection technique, a different exploriometer threshold is used in order to qualify-select users as "explorers". These thresholds are set as (i) 75%, (ii) 50% and (iii) 25% of the total users selected as explorers. By making the threshold gradually stricter, we better demonstrate the effects of the metric. The number of neighbors selected for a user is a very important parameter for the performance of a CFR, because the neighbors hold the information (ratings) with which the system will select the recommended items. For that reason, we want the sizes of neighborhoods to be the same between all four execution modes. In this way, the results will be dependent on the quality of the selected neighbors and not their quantity. To keep the size of the neighborhoods equal between the four recommenders, the similarity thresholds are set such as every neighborhood contains 10% of the total users.

For brevity reasons, when we describe the results for those methodologies, we call the conventional CFR that utilize only the similarity metric for neighbor selection (a) as Standard, and the systems that utilize our methodology with the three different thresholds as Weak, Medium and Strong Filtering respectively.

4.3 Results

In order to study the effects of our enhanced neighbor filtering, we begin by observing the neighbors that are selected with and without its use:

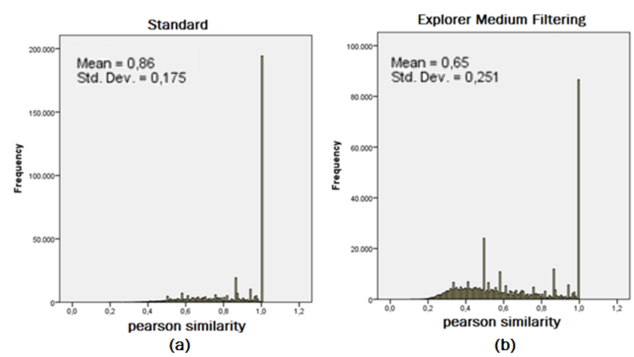

Figure 2: Neighbors Pearson Similarity Distribution for Standard (a) and Explorer Medium Filtering (b)

Figure 2 shows the distribution of neighbors according to their Pearson similarity with the user *u* for whom the system produces recommendations. We observe that when the explorer filtering is used (i.e. Figure 2(b)), the selected neighbors are less similar with *u* and thus their likings are not identical with u. This leads to an increase in Recommendation Diversity as observed in Figure 3.

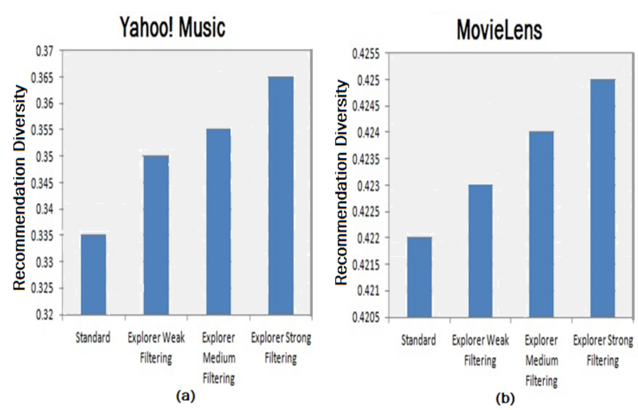

Figure 3: Recommendation Diversity for Yahoo! Music (a) and MovieLens datasets (b)

The first noticeable point of Figure 3, concerning the Standard CFR, is the vast difference in Recommendation Diversity between the two datasets. This is merely due to the difference in density; the higher density of the MovieLens dataset results in higher Recommendation Diversity. Regarding the use of Exporiometer, we observe that the higher the threshold used, the higher the Recommendation Diversity on both datasets, especially on Yahoo! Music dataset where Recommendation Diversity was very low initially. The same results can be extracted from the Mean Item Rarity metric (see Figure 4(b)), since it also quantifies the diversity of the recommended items.

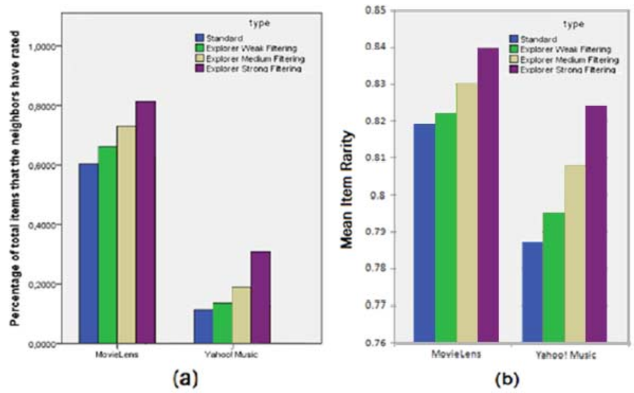

Figure 4: (a): Mean percentage of total items that are rated by neighbors of u, (b) Mean Item Rarity.

Also, due to the selection of neighbors who are less similar to the user, the system has a wider variety of items that is able to recommend. This is verified by our experiments. Figure 4(a) presents the mean percentage of total items that are rated by the neighbors of user *u* and, thus the system can predict the ratings of *u* and possibly recommend them. We observe that in the Yahoo! Music dataset the increase of the percentage of total items rated by neighbors of *u* when the explorer filtering is used is greater and it exceeds 100% because the original percentage was very low due to the high sparsity of the dataset.

The increase in the multitude of items that are rated by the neighbors also leads the system to predict a wider range of ratings from the test set, since the system can predict the user's ratings only for the items that were rated by their neighbors (see Figure 5(a)).

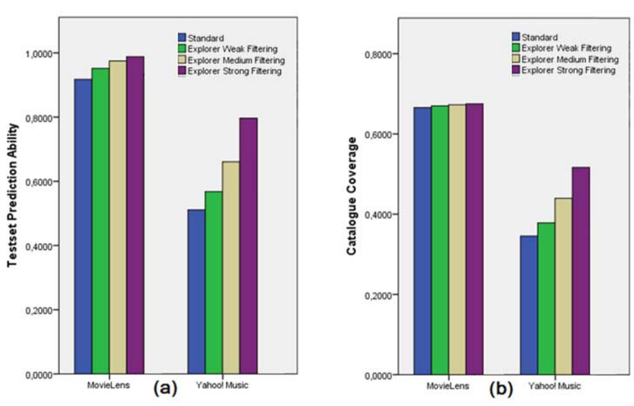

Figure 5: (a): Percentage of test set that the system can predict, (b): Catalogue Coverage (on 10-item recommendations)

This subsequently increases the possibility of recommending rare items leading to a greater percentage of total items which are recommended to the sum of users. The latter can be observed on Figure 5(b) that displays the catalogue coverage of the system, which is a metric that depicts the percentage of total items that the system recommends to its users. In detail, for Yahoo! Music only 34% of the total items were recommended initially (Standard) and that percentage increased up to 51% with explorer strong filtering, which means that there is a greater variety of items recommended to users.

Finally, there is not a statistical significant decrease of the accuracy, even if the neighbors are less similar to the user. The reason is that the variety of neighbors' opinions on how much the user *u* may be interested on an item compensate for the decreased similarity between the user and their neighbors. For example, if the user *u* has rated highly mostly action movies and in the evaluation set they have rated highly a documentary which is commonly highly perceived; when only similar users are used as neighbors, those users would appreciate action movies and they may disregard a documentary. If the neighbor list, on the other hand, includes users with a wider set of interests, it is more likely that a neighbor would have rated the documentary highly and the system would accurately predict the user's rating.

5. CONCLUSIONS

In this paper we have introduced an enhanced neighbor selection technique and we explored this technique in terms of its affecting on primarily the coverage but also the diversity of the generated recommendations. Our research of optimizing the coverage of CFR leads us to the conclusion that it is a complex problem which requires understanding of the total recommendation generation progress. We also observed that one of the main reasons for low coverage and diversity is the low density of the dataset.

The enhanced neighbor filtering managed to increase more efficiently the coverage when the density is low and this leads to the conclusion that the neighbor filtering is highly recommended when a CFR is recently initiated and has not collected many ratings from its users. In addition, the use of the proposed neighbor filtering is recommended when the set of items is far greater than the ability of the users to rate those items. For example, the Yahoo! Music dataset contained such a great number of songs that the users could only rate a small fraction with leads inevitably to a sparse dataset and low coverage and diversity. The identification of the density threshold that our neighbor filtering technique uses is a subject of future research. A subject of future research is also the optimal use of the "Exploriometer"; for example, weighting the "Exploriometer" and other similarity metrics to choose the neighbors may bring even better results. This weighting could also be user-driven thus providing more personalized recommendations.

6. ACKNOWLEDGMENTS

We would like to acknowledge Yahoo! Labs for providing this research endeavor with the dataset of Yahoo! Movies R2 via Webscope.

7. REFERENCES

[1] **Su, X., & Khoshgoftaar, T. M.** A survey of collaborative filtering techniques. . *Advances in Artificial Intelligence.* 2009.

[2] **Linden, G., B. Smith, and J. York.** Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing.* Jan.-Feb. 2003.

[3] **Sean M. McNee, J. R. and Konstan, J. A.** Accurate is not always good: How accuracy metrics have hurt recommender

systems. *In proceedings of the ACM WebKDD Workshop.* 2006.

[4] **Bobadilla, J., Hernando, A., Ortega, F., & Bernal, J.** A framework for collaborative filtering recommender systems. *Expert Systems with Applications, 38, 14609–14623. http://dx.doi.org/10.1016/j.eswa.2011.05.021.* 2011.

[5] **Ge, M., Delgado-Battenfeld, C., & Jannach, D.** Beyond accuracy: evaluating recommender systems by coverage and serendipity. *In Proceedings of the fourth ACM conference on Recommender systems (pp. 257-260). ACM.* 2010, September.

[6] **Brynjolfsson, E., Hu. Y J., and Simester, D.** Goodbye Pareto Principle, Hello Long Tail: The effect of Search Costs on the Concetration of Product Sales. *Net Institute Working Paper.* 2007.

[7] **Fleder, D. and Hosanagar, K.** Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science, 55(5), pp. 697-712.* 2009.

[8] **Ziegler, C., McNee, S. M., Konstan, J. A., and Lausen, G.** Improving recommendation lists through topic diversification. *In Proceedings of the 14th International Conference on World Wide Web. 22-32.* 2005.

[9] **Grčar, Miha, et al.** *Data sparsity issues in the collaborative filtering framework.* Springer Berlin Heidelberg. s.l. : Springer Berlin Heidelberg, 2006.

[10] **Wilson, D. C., Smyth, B., & Sullivan, D. O.** Sparsity reduction in collaborative recommendation: A case-based approach. *International journal of pattern recognition and artificial intelligence, 17(05), 863-884.* 2003.

[11] **Balabanović, M., & Shoham, Y.** Fab: content-based, collaborative recommendation. . *Communications of the ACM, 40(3), 66-72.* 1997.

[12] **Billsus, D., & Pazzani, M. J.** User modeling for adaptive news access. *User modeling and user-adapted interaction, 10(2-3), 147-180.* 2000.

[13] **Palit, G. P. and Taillie, C.** Diversity as a concept and its measurements. *J. Amer. Statist. Assoc. 77,379, 548-561.* 1982.

[14] **Hyndman, Rob J. Koehler, Anne B.** Another look at measures of forecast accuracy. *International Journal of Forecasting: 679–688. doi:10.1016/j.ijforecast.2006.03.001.* 2006.

[15] **D. Goldberg, D. Nichols, B. M. Oki, and D. Terry.** Using collaborative filtering to weave an information tapestry. *Communications of the ACM.* 1992, 35(12):61-70.

[16] **Rich, E.** User modeling via stereotypes. *Cognitive science.* 1979, Vols. 3(4):329-354.

[17] **Desrosiers, Christian, and George Karypis.** A comprehensive survey of neighborhood-based recommendation methods. *Recommender systems handbook.* Springer US, 2011, Vols. 107-144.

[18] **Koren, Y., & Bell, R.** Advances in collaborative filtering. *In Recommender Systems Handbook (pp. 145-186).* Springer US, 2011.

[19] **Said, A., Jain, B. J., & Albayrak, S.** Analyzing weighting schemes in collaborative filtering: cold start, post cold start and power users. *In Proceedings of the 27th Annual ACM Symposium on Applied Computing (pp. 2035-2040).* ACM, 2012.

[20] **Herlocker, J. L., Konstan, J. A., Borchers, A., & Riedl, J.** An algorithmic framework for performing collaborative filtering. *In Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval (pp. 230-237).* ACM, 1999.

[21] **Jannach, D., Lerche, L., Gedikli, F., & Bonnin, G.** What recommenders recommend–an analysis of accuracy, popularity, and sales diversity effects. *In User Modeling, Adaptation, and Personalization (pp. 25-37).* Springer Berlin Heidelberg, 2013.

[22] **A. Ghose, P. Ipeirotis, and B. Li.** Designing ranking systems for hotels on travel search engines by mining user-generated and crowd-sourced content. *Marketing Science.* 2012.

[23] **Said, A., Kille, B., Jain, B. J., & Albayrak, S.** Increasing diversity through furthest neighbor-based recommendation. . *Proceedings of the WSDM, 12.* 2012.

[24] **Said, A., Fields, B., Jain, B. J., & Albayrak, S.** User-centric evaluation of a k-furthest neighbor collaborative filtering recommender algorithm. . *In Proceedings of the 2013 conference on Computer supported cooperative work (pp. 1399-1408).* ACM, 2013.

[25] **Anderson, C.** The long tail. *New York: Hyperion.* 2006.

[26] **Resnick, P., N. Iakovou, M. Sushak, P. Bergstrom, and J. Riedl.** GroupLens: An open architecture for collaborative filtering of netnews. *In Proceedings of the 1994 Computer Supported Cooperative Work Conference.* 1994.

[27] **Sarwar, B., G. Karypis, J. Konstan, and J. Riedl.** Item-based Collaborative Filtering Recommendation Algorithms. *In Proc. of the 10th International WWW Conference,.* 2001.