Increasing Diversity Through Furthest Neighbor-Based Recommendation

Alan Said, Benjamin Kille, Brijnesh J. Jain, Sahin Albayrak
DAI-Lab
Technische Universität Berlin
{alan, kille, jain, sahin}@dai-lab.de

ABSTRACT

One of the current challenges concerning improving recommender systems consists of finding ways of increasing serendipity and diversity, without compromising the precision and recall of the system. One possible way to approach this problem is to complement a standard recommender by another recommender "orthogonal" to the standard one, i.e. one that recommends different items than the standard. In this paper we investigate to which extent an inverted nearest neighbor model, k-furthest neighbor, is suitable for complementing a traditional kNN recommender. We compare the recommendations obtained when recommending items disliked by people least similar to oneself to those obtained by recommending items liked by those most similar to oneself. Our experiments show that the proposed furthest neighbor method provides more diverse recommendations with a tolerable loss in precision in comparison to traditional nearest neighbor methods. The recommendations obtained by k-furthest neighbor-based approaches are almost completely orthogonal to those obtained by their k-nearest neighbors-based counterparts.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval - Information filtering, Retrieval models, Search process, Selection process; H.3.4 [**Information Technology and Systems Applications**]: Decision support

General Terms

Algorithms, Design, Experimentation, Measurement, Human factors

Keywords

collaborative filtering, recommender systems, nearest neighbor, diversity, movie recommendation, memory-based models

DDR'12 February 12, 2012 Seattle WA, USA.

1. INTRODUCTION & RELATED WORK

Collaborative filtering (CF) is widely used in information retrieval tasks, it calculates the relevance of an item for a user based on other user's rating information (or other interaction) on items co-rated by the group and the user. CF approaches are commonly categorized as either model-based or memory-based. In this work, we focus on the latter, creating item predictions for a user by finding users similar to that user (in terms of co-rated items), a so-called neighborhood. Using information from the neighborhood, we predict items of interest not previously seen by the user.

Neighborhood-based approaches commonly use measures such as Pearson's correlation coefficient or cosine similarity to find the neighborhoods. Memory-based methods are simple, flexible, and provide adequate results, but may suffer from the lack of unexpectedness, or serendipity. This is due to these methods trying to identify the most probable items by inferring them from the items of users most similar to oneself. Items the target user would like, yet unknown to the group of most similar users to her, would likely not be included in the recommendation list. The effect of this is that recommended items are often ones which are familiar to the target user, even though she has not rated them yet. However, items which are unknown to the target user could potentially increase the perceived usefulness of the system [4].

To overcome the drawback of recommending already known items, we propose the usage of an "orthogonal" recommender, i.e. one that recommends items which are not recommended by standard approaches. For this purpose we investigate the concept of furthest neighbors. Based on the proverb *The enemy of my enemy is my friend*, our approach recommends items which are disliked by those least similar to each user.

The furthest neighbor concept has been researched previously. Choi et al. [3] for instance used a dataset reduction model based on nearest and furthest neighborhood selection in order to minimize effects caused by sparsity.

Serendipity, or "unexpectedness" has been covered by Murakami et al. [7] where unexpectedness in search results is measured in relation to predictions of so-called primitive models (standard algorithms). The measure was further adapted by including the notion of *usefulness* by Adamopoulos and Tuzhilin [1] which used it to measure the performance of recommender systems. Furthermore, Castells et al. [2] identified a gap in the formalization of novelty and diversity metrics in recommender systems and derived metrics for taking the recommended item's position and relevance into consideration when evaluating the list of recommended items.

To our knowledge, no attempt at diversifying recommender systems based on furthest neighbors has been attempted previously. Therefore, in this paper, we propose *k-Furthest Neighbors* (kFN), which, when used for recommendation purposes, provides recommendations with higher serendipity and diversity when compared to more traditional approaches.

We investigate how neighborhood-based CF recommender systems perform when the neighborhood creation is inversed, i.e. by picking the furthest neighbors, and recommending their least favorite items. Our assumption is that kFN approaches generate more serendipitous recommendations.

2. COLLABORATIVE FILTERING

Commonly, memory-based collaborative filtering utilizes the k-nearest neighbor approach to identify candidate items. Due to the nature of the algorithm, for users with average taste (i.e. liking mainstream items), using kNN often results in recommending items which are already known by the user. These recommendations can be of little benefit to the users, as the probability of being familiar with popular items is high. The recommendations might contain items which are of interest for the users, however no recommendation is actually necessary as they are already known. In order to improve the usefulness of a recommendation, the items which are recommended should not be previously known by the users. For this purpose, we propose an alternative version of kNN, kFN - k-Furthest Neighbors, based on the same premises as its traditional counterpart, it provides more diverse recommendations without increasing the complexity of the algorithm.

2.1 k-Furthest Neighbors

The k-furthest neighbor algorithm is very similar to its nearest neighbor cousin, but is in essence its antipole in terms of similarity. Fig. 1 shows a scenario where the kFN approach itentifies suitable recommendations based on dislikes. In order to identify the furthest neighbors of users, an augmented inverted similarity measure is used. The dissimilarity of two users is calculated by finding all users who have co-rated at least five items. One additional restriction on the ratings of the co-rated items is that they should not be in the middle of the rating scale, i.e. items rated with a 3 on a 1 to 5 rating scale should not be counted as co-rated. Those ratings do not contribute much to identify users with dissimilar tastes since no clear like/dislike of an item is expressed. Users who disagree strongly on a significant amount of items could still be very similar based on non-controversial items (i.e. those given average ratings). Next, standard similarity metrics (Pearson and cosine) are calculated between users, those least similar to each other form neighborhoods. In the recommendation process, the least liked items of the neighborhoods are recommended.

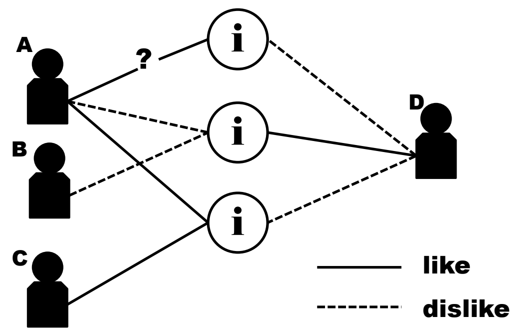

Figure 1: Users A and B, and A and C have similar taste. User C however does not share hers taste with others, here a mismatch in taste can be used to find interesting items.

3. EXPERIMENTS

In our experiments, we investigate the orthogonality and accuracy in terms of precision and recall of two kFN approaches towards their kNN counterparts.

3.1 Dataset

In our experiments we used a subset of the Movielens 1M100K [6] dataset. The dataset was a random sample of 1 million ratings from the dataset, excluding the 100 most popular movies and users with fewer than 40 ratings, resulting in a dataset with 44, 214 users and 9, 432 movies. This filtering was performed in order to remove the bias created by popular items. In the Movielens dataset, approximately 25% of all users have rated the most 3 popular movies with the highest rating [8], no matter how similar they are regarding less common movies. In a nearest neighbor scenario, this would mean that even very dissimilar users have high similarity when it comes to popular movies. This bias is due to the fact that users have higher probabilities to have rated items they like than items they do not like [5], simply put - people are more inclined to watch (and subsequently rate) items they think they will like.

The dataset was divided into a training and a validation set. Items selected into the test set had to have been rated above the corresponding user's average rating value, the rationale for this being that an item which has received a low rating is not a good recommendation.

3.2 Experimental Setup

We evaluated the potential of kFN being an orthogonal counterpart of kNN in order to be a good candidate for adding diversity in the context of recommendation. To this end, we evaluated the accuracy of both approaches using precision and recall as accuracy metrics [4]. To evaluate the

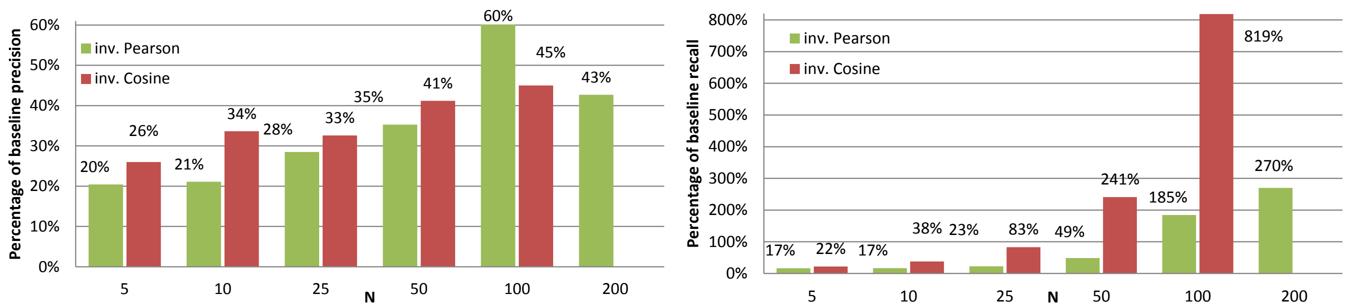

(a) Precision (b) Recall

Figure 2: The Precision@N (Fig. 2(a)) and Recall@N (Fig. 2(b)) compared to the respective kNN baselines (kNN Pearson compared to kFN Pearson; kNN cosine compared to kFN cosine) for $N = \{5, 10, 25, 50, 100, 200\}$ for the kFN recommenders.

(a) Precision

N	5	10	25	50	100	200
Pearson Similarity	0.007	0.011	0.017	0.028	0.041	0.090
Cosine Similarity	0.005	0.007	0.016	0.027	0.057	0

(b) Recall

N	5	10	25	50	100	200
Pearson Similarity	0.008	0.013	0.021	0.230	0.014	0.010
Cosine Similarity	0.002	0.006	0.007	0.006	0.005	0.004

Table 1: The precision (Table 1(a)) and recall (Table 1(b)) values for the baselines for which the percentages in Fig. 2 are calculated. Precision@N and Recall@N is calculated only for those users who have at least $2N$ ratings.

orthogonality of the approach, we used "overlap", which is the percental intersection of the lists of recommended items by both approaches. Experiments were conducted on kFN and kNN with Pearson's correlation coefficient as well as cosine similarity, precision (Precision@N) and recall (Recall@N) where calculated at N for all users who had at least 2N ratings.

3.3 Results and Discussion

Fig. 2(a) shows the percentages of precision obtained by both Pearson and cosine similarity-based kFN recommenders in comparison to their kNN counteparts, i.e. at $N = 5$ the Pearson-based kFN recommender reached a precision level corresponding to 20% of the regular Pearson-based one. The standard kNN recommenders outperform the kFN approaches, independent of N. Nevertheless, the kFN approaches perform at considerably high (20% to 60%) levels of the kNN precision values. Similarily, Fig. 2(b) shows the percental recall values of the kFN recommenders compared to their kNN baselines. Both kFN recommenders outperform their standard counterparts for $N >= 50$ for cosine and $N >= 100$ for Pearson. The baseline values for which percentages are calculated are shown in Table 1.

The orthogonality, expressed in overlap is shown in Fig. 3, where the percentage of overlapping recommended items between the kFN- and KNN-based recommenders is shown, for the kFN-kNN counterpars as well as between both kNN and kFN approaches. The highest number of overlapping items is found between the two inverted recommenders. The number of overlapping items between the kNN and kFN approaches is close to 0, meaning that the recommenders recommend practically disjoint lists of items. The implication of this being that, even though the precision values of the inverted recommenders are lower, they find items which the standard recommenders do not, and vice-versa.

Given that the kNN and kFN approaches seem to produce orthogonal recommendations, combining them could potentially improve recommendation quality considerably. The difficulty lies in how to combine the two disjoint lists. Standard ensemble methods, where the intersection of several algorithms is used, cannot work in this scenario. One method which could be used for the combination of these recommendations would be to alternate between the recommenders with a probability proportional to the precision values. In this case, if applied to $N = 10$ for the cosine recommenders in Fig. 2(a), the inverted recommendations could be recommended roughly one forth of the time. The result of this would be more diverse recommendations, with relatively small loss in precision.

	inv. Pearson vs. Pearson	inv. Cosine vs. Cosine	inv Pearson vs. inv Cosine	Pearson vs. Cosine
5		0,38%	11,92%	4,62%
10		0,17%	5,45%	1,40%
25	0,09%	0,01%	1,56%	0,40%
50	0,07%	0,03%	0,88%	0,13%
100	0,05%	0,03%	1,00%	0,02%
200			0,95%	

Figure 3: The overlaps between the lists of recommended items for the four recommenders at $N = \{5, 10, 25, 50, 100, 200\}$. Circle size reflects the percentage of items that overlap between the recommenders, the color reflects absolute numbers (ranging from 1 to 131). No number or circle indicates no overlap present between the compared approaches.

4. CONCLUSION & FUTURE WORK

In this work, we have shown that the phrase *the enemy of my enemy is my friend*, when applied to similarity calculations, does seem to hold. At least in terms of finding interesting, but perhaps unexpected items. We have shown that items which one's "enemies" dislike can very well be good recommendations.

The set of items recommended by the kFN approach is close to completely orthogonal to the set provided by standard approaches, with a tolerable loss in precision and recall. This orthogonality provides more diverse and unexpected items than those found through traditional approaches. This diversity ensures more robust recommendations minimizing recommendation of items simply due to their popularity.

Given the promising lists of items obtained by the kFN models, we plan on investigating how traditional nearest neighbor approaches could be combined with the presented furthest neighbor concept in order to minimize the loss of precision and recall, in combination with high levels of diversity and unexpectedness.

5. ACKNOWLEDGMENTS

The authors would like to express their gratitude to Benjamin Fields from musicmetric for a helpful discussion, feedback and support.

The work in this paper was conducted in the scope of the KMulE project which was sponsored by the German Federal Ministry of Economics and Technology (BMWi).

6. REFERENCES

[1] Panagiotis Adamopoulos and Alexander Tuzhilin, 'On unexpectedness in recommender systems: Or how to expect the unexpected', in *Proc of RecSys'11 Intl. Workshop on Novelty and Diversity in Recommender Systems*, (2011).

[2] P. Castells, S. Vargas, and J. Wang, 'Novelty and Diversity Metrics for Recommender Systems: Choice, Discovery and Relevance', in *International Workshop on Diversity in Document Retrieval (DDR 2011) at the 33rd European Conference on Information Retrieval (ECIR 2011)*, (April 2011).

[3] Sang Hyun Choi, Young-Seon Jeong, and M.K. Jeong, 'A hybrid recommendation method with reduced data for large-scale application', *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Trans. on*, (September 2010).

[4] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl, 'Evaluating collaborative filtering recommender systems', *ACM Trans. Inf. Syst.*, **22**, 5–53, (January 2004).

[5] Benjamin M. Marlin and Richard S. Zemel, 'Collaborative prediction and ranking with non-random missing data', in *Proceedings of the third ACM conference on Recommender systems*, RecSys '09, pp. 5–12, New York, NY, USA, (2009). ACM.

[6] MovieLens. MovieLens 10M/100k Data Set README. retrieved April 2011.

[7] Tomoko Murakami, Koichiro Mori, and Ryohei Orihara, 'Metrics for evaluating the serendipity of recommendation lists', in *Proceedings of the 2007 conference on New frontiers in artificial intelligence*, JSAI'07, pp. 40–46, Berlin, Heidelberg, (2008). Springer-Verlag.

[8] Alan Said, Brijnesh J. Jain, and Sahin Albayrak, 'Analyzing weighting schemes in collaborative filtering:cold start,post cold start and power users', in *27th ACM Symposium On Applied Computing (SAC '12)*, New York, NY, USA, (2012). ACM.