Recommending Based on Rating Frequencies

Fatih Gedikli
Department of Computer Science
44227 Dortmund
Germany
fatih.gedikli@tu-dortmund.de

Dietmar Jannach
Department of Computer Science
44227 Dortmund
Germany
dietmar.jannach@tu-dortmund.de

ABSTRACT

Since the development of the comparably simple neighborhood-based methods in the 1990s, a plethora of techniques has been developed to improve various aspects of collaborative filtering recommender systems like predictive accuracy, scalability to large problem instances or the capability to deal with sparse data sets. Many of the recent algorithms rely on sophisticated methods which are based, for instance, on matrix factorization techniques or advanced probabilistic models and/or require a computationally intensive model-building phase. In this work, we evaluate the accuracy of a new and extremely simple prediction method (RF-Rec) that uses the user's and the item's most frequent rating value to make a rating prediction. The evaluation on three standard test data sets shows that the accuracy of the algorithm is on a par with standard collaborative filtering algorithms on dense data sets and outperforms them on sparse rating databases. At the same time, the algorithm's implementation is trivial, has a high prediction coverage, requires no complex offline pre-processing or model-building phase and can generate predictions in constant time.

Categories and Subject Descriptors

H.1.2 [**Models and Principals**]: User/Machine Systems—*human factors*; H.3.3 [**Information Search and Retrieval**]: Information filtering—*Selection process.*

General Terms

Algorithms, Performance.

Keywords

Collaborative filtering, Computational efficiency, Accuracy, Sparse data sets, Evaluation.

1. INTRODUCTION

Pure collaborative filtering recommender systems (RS) only rely on a given user-item rating matrix to make rating

predictions for items that the *active user* has not seen yet. Early neighborhood-based recommendation schemes simply used the k nearest neighbors (kNN) as predictors for unseen items. Later on, a broad range of more advanced and sophisticated methods have been applied to better exploit the given rating information and to improve the recommendation process in one or the other dimension. Examples for such methods include matrix factorization, various probabilistic models, clustering techniques, graph-based approaches as well as machine learning techniques, based on, e.g., association rule mining, see also [1].

Typically, the more elaborate approaches outperform the commonly-used kNN baseline method in terms of accuracy in particular for sparse data sets or in terms of scalability as they rely on offline pre-processing or model-building phases. In [4], Lemire and Maclachlan formulate additional desirable features of a recommendation scheme such as that they are easy to implement, can be updated on the fly, are efficient at query time and are "reasonably" accurate. Their evaluation shows that the proposed *Slope One* family of item-based recommender algorithms, which is based on the computation of "popularity differentials between items for users", leads despite its simplicity to relatively accurate predictions (measured in terms of Mean Absolute Error). Due to its simplicity, different implementations of the algorithm in various programming languages and frameworks are available today.

In this paper, we propose an even simpler recommendation scheme, RF-REC, which is only based on the absolute frequencies of the different rating values per user and per item. The method is therefore trivial to implement, can generate predictions in constant time, does not require a computationally intensive offline model-building phase, and at the same time leads to competitive prediction coverage and accuracy results in particular for sparse data sets.

In the rest of the paper, we will first describe the RF-REC recommendation scheme in more detail and present results of an experimental evaluation on three commonly-used data sets.

2. RECOMMENDING BASED ON RATING FREQUENCIES

Let us illustrate the RF-REC recommendation scheme with a simplified and relatively sparse rating database shown in Figure 1. The goal in our example is to predict Alice's rating for item $I3$.

When adopting a user-based kNN scheme, probably no prediction can be made because only one relatively similar user $U1$ exists which could be taken as a predictor for Alice.

	I1	I2	I3	I4	I5	Average
Alice	1	1	**?**	5	4	2.75
U1	2		5	5	5	4.25
U2			1	1		1.00
U3		5	1	1	2	2.25
Average	1.50	3.00	2.33	3.00	3.67	

Figure 1: Example user-item rating matrix.

If we allow also such small neighborhood sizes, the prediction for Alice will usually consist of taking the neighbor's rating for $I3$ and using it for the prediction by making a weighted addition to Alice's average rating. Similarly, in an item-based kNN approach, Alice's rating value for item $I4$, whose rating vector is similar to the one of $I3$ will be taken as a predictor. In both cases, the prediction for Alice for item $I3$ will be rather high.

In our approach, in contrast, the predictions are based on absolute rating frequencies. The prediction function for a given user u and an item i in the RF-REC recommendation scheme is defined as follows:

$$pred(u, i) = \arg\max_{r \in possibleRatings}$$
$$\Big(\big(freqUser(u, r) + 1 + \mathbb{1}_{avg-user}(u, r) \big) *$$
$$\big(freqItem(i, r) + 1 + \mathbb{1}_{avg-item}(i, r) \big) \Big)$$

where $freqUser(u, r)$ is the frequency of ratings with value r of the target user u and $freqItem(i, r)$ is the frequency of ratings with value r of the target item i. $\mathbb{1}_{avg-user}(u, r)$ and $\mathbb{1}_{avg-item}(i, r)$ are indicator functions which return 1 if the given rating corresponds to the rounded average rating of the target user or target item accordingly and 0 otherwise.

The rationale of the prediction scheme is as follows. First, instead of taking rating *averages* into account for the calculations (as done in kNN-based approaches and also in Slope One) we rely on rating *frequencies*. Intuitively, this can be advantageous in case of extreme ratings, i.e., since Alice only gave very low and very high ratings and at the same time item $I3$ also only received extreme ratings. Incorporating user or item averages would move the predictions away from these extremes. Note that in [2] Herlocker et al. have also observed that high variance in the rating data can lead to decreased recommendation accuracy.

The "1" in the middle of our formula is used to avoid that in situations, in which a user has never given a rating (or an item never received a particular rating), the whole term is multiplied with zero. The indicator function in our scheme shall help in situations, in which several ratings have the same frequency counts. If this is the case and in addition one of these ratings corresponds to the average rating, we add some small extra weight to it, thus very slightly preferring the average rating.

In the example where the frequency of ratings of Alice are [1:2, 2:0, 3:0, 4:1, 5:1] and the rating frequencies of item $I3$ are [1:2, 2:0, 3:0, 4:0, 5:1] we would calculate:

Rating value 1: (2+1+0)*(2+1+0) = 9
Rating value 2: (0+1+0)*(0+1+1) = 2
...
Rating value 5: (1+1+0)*(1+1+0) = 4

Since the formula result for rating 1 is the highest, we would predict that Alice would give a "1" to item $I3$, which is strongly different from the ratings that we would predict with the other methods. When using the Slope One scheme, for example, 2.38 would be the predicted rating for Alice, which is slightly below her average.

Regarding *prediction coverage*, i.e., the question for what percentage of items a recommender can make predictions, note that in contrast to kNN approaches that often use similarity and neighborhood size thresholds, our recommendation scheme can make predictions if at least one rating for the target item or one rating by the user is available.

In order to measure the predictive accuracy of the method, we evaluated our approach on three popular data sets using a common experimental procedure and accuracy metric. The results of this evaluation are described in the following section.

3. EXPERIMENTAL EVALUATION

3.1 Algorithms, data sets and metrics

For evaluation purposes we have used two data sets from the movie domain and one data set from the book domain: we used the 100k-MovieLens rating database (100,000 ratings by 943 users on 1,682 movies, 0.9369 sparsity), a snapshot of the Yahoo!Movies data set (211,231 ratings by 7,642 users on 11,915 movies, 0.9976 sparsity) and a subset of the Book-Crossing data set [7] (100,000 ratings by 30,041 users on 37,429 books, 0.9999 sparsity; only explicit item ratings were taken into account)[1]. Due to memory limitations only a randomly chosen subset of the full Book-Crossing data set was considered[2]. The MovieLens data set only contains users who have rated at least 20 items; the minimum number of rated items per user in the Yahoo! data set is 10. The Book-Crossing data set only contains users with at least one rating.

The density level of the data sets were varied by using subsamples of different sizes, starting with very sparse settings up to the 90% data set, see also [5]. The smallest subsample contained 10% of the original data. In this subsample, the average number of ratings per user was around 10 for the MovieLens data set, 3 for the Yahoo!Movies data set and 2 for the Book-Crossing data. Further measurements were taken in steps of 10% up to the 90% data set, which corresponds to the usual 90% train/test ratio for *Mean Absolute Error* (MAE) measurements. The experiments were appropriately repeated in order to factor out randomness effects.

---

[1]http://www.grouplens.org/node/73,
http://webscope.sandbox.yahoo.com,
http://www.informatik.uni-freiburg.de/~cziegler/BX
[2]All books with a lexicographically higher ISBN than 375811370 were filtered out from the data set.
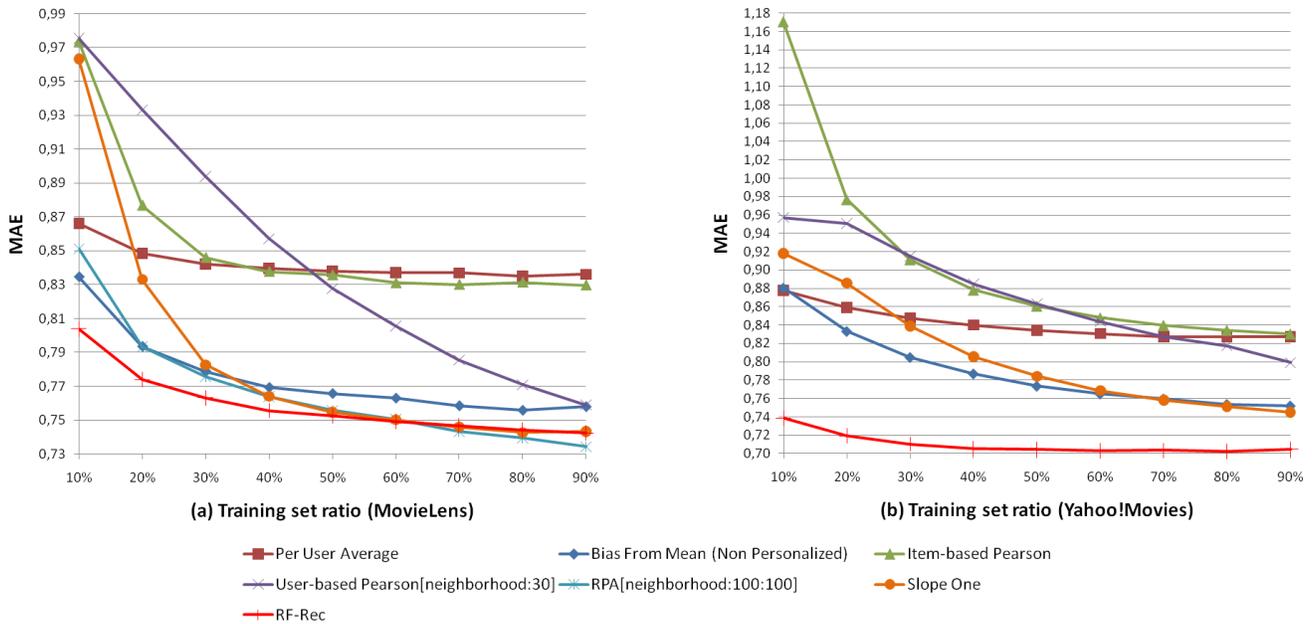
Figure 2: MAE values for different training set sizes: MovieLens (a), Yahoo!Movies (b).

We compared the following algorithms: user-based kNN (using default voting, Pearson similarity and the neighborhood size of 30 as suggested as optimal value in literature), item-based kNN (Mahout's item-based kNN-method implementation with Pearson similarity)[3], Slope One [4], Bias from Mean (Non Personalized), Per User Average, the recent recursive prediction algorithm (RPA) [6] (using larger, empirically determined neighborhood sizes of 100) and RF-REC.

3.2 Results

Figure 2 (a) shows the MAE values for different training set sizes for the MovieLens data set. Up to the 50% level, RF-REC has consistently better accuracy than all other techniques. Above that level, the accuracy of RF-REC (0.742) is comparable to Slope One (0.743) and RPA (0.734). Note that RF-REC due to its nature leads to 100% prediction coverage also for very sparse data sets. In contrast, the coverage of the user-based kNN method, for example, slowly increases from 44% to 95% when increasing training set ratio from 10% to 90%, see Table 1. In the experiments, we further varied the neighborhood size ns of the user-based kNN method. Increasing ns to 300 lead to the observation that the accuracy improved and was comparable to the one of RPA for training set sizes higher than 50%.

Figure 2 (b) shows the results for the Yahoo! data set. We can observe similar accuracy values also for this data set. In particular, the improvement of our RF-REC algorithm is even stronger on that data set. A possible explanation for this observation could be the different sparsity levels of the two data sets, i.e., assuming that RF-REC works particularly well for sparse settings, it is intuitive that even better results can be achieved on the sparser Yahoo! data set than on the MovieLens data set.

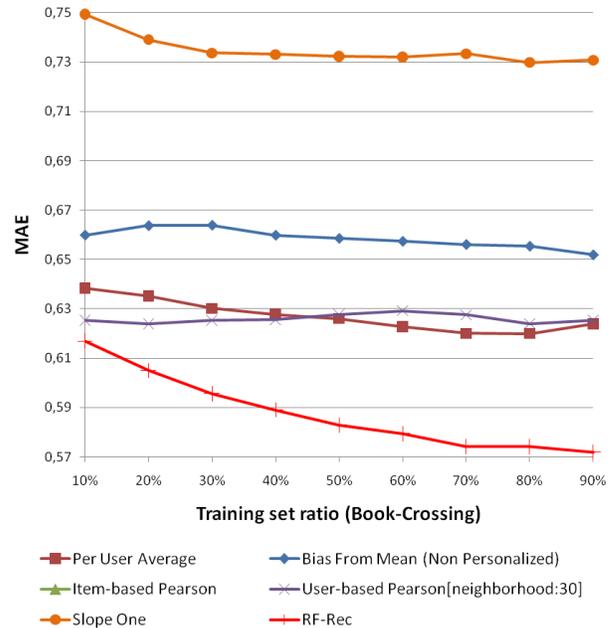Finally, in Figure 3, the results for the Book-Crossing data

Figure 3: MAE values for the Book-Crossing data set.

set are shown. The sparsity level of the Book-Crossing data set is even higher than the sparsity level of the Yahoo! data set. Again RF-REC has consistently better accuracy than all other techniques. The results of the item-based scheme are not shown in the figure because the best MAE result for this recommender is 1.181 (when using a train-ratio of 90%). Note that for this data set even the simple prediction scheme Per User Average achieves good result when compared with more complicated prediction schemes like Slope One and

---

[3]`http://lucene.apache.org/mahout`

Training set ratio	RF-REC	Slope One	User-based kNN	Item-based kNN
10%	100%	97%	44%	60%
20%	100%	98%	58%	94%
30%	100%	99%	70%	98%
..	..	..	..	..
90%	100%	99%	95%	99%

Table 1: Prediction coverage of recommendation approaches (MovieLens).

user-based kNN. The additional evaluation of our proposed RF-REC method on the book domain shows that the RF-REC method achieves good accuracy values even on data sets of different domains and sparsity levels.

Overall, these accuracy findings indicate that RF-REC has a constantly good performance which is quite independent of the training set ratio. It is despite its simplicity suitable to generate predictions with an accuracy which is comparable to existing approaches and is even better for sparse data sets, which can often be found in practice.

3.2.1 Computational complexity

In the RF-REC scheme, the "model-building" phase obviously consists of calculating the frequencies of the individual rating values per user and per item, which can be accomplished in a single scan of the matrix; the frequency statistics can be easily updated when new ratings are available. Given u users, i items and v possible rating values, the memory requirements for the model are constant: $(u*v) + (i*v)$. Also the calculation of predictions can be done with the formula from Section 2 in constant time. In absolute numbers, "model-building" requires less than 10 seconds even when hundreds of millions of ratings exist; predictions can be calculated in a few milliseconds on a standard desktop computer. We compared our method with Mahout's item-based kNN-method implementation on the 1 million MovieLens data set: model-building takes 500ms in our approach as opposed to 6 minutes with Mahout. Generating a prediction takes only 3ms in our framework (and 100ms with Mahout), which is a very important factor in high-traffic recommenders in which up to 1,000 parallel requests have to be served [3].

4. SUMMARY

In this work we proposed a new, frequency-based recommendation scheme that leads to good predictive accuracy and is at the same time highly scalable and very easy to implement. We have evaluated our proposed scheme on three different data sets and two diverse application domains. Our future work includes the evaluation of the approach on the Netflix data in order to compare it to the results of more recent methods; in addition, we will also compare the predictive accuracy of the different methods based on precision and recall. Besides this, we are experimenting with other variations of the scheme which for instance include weighting factors for the user and item rating frequencies.

Our evaluation demonstrated that comparably good results can be achieved with simple methods and that the additional accuracy improvements that can be achieved with the help of "classical" methods and even some of the more recent ones such as RPA are actually relatively low. Overall, the extra bit of predictive accuracy (measured in terms of

MAE) may therefore not justify the computational efforts required in particular by the recursive prediction scheme RPA.

In general, we hope that light-weight approaches like the herein proposed RF-REC method help to further promote the use of recommender systems in practice; by making the software used in our experiments publicly available[4], we also hope to contribute to the comparability of different algorithms since our study revealed that relevant algorithmic details and parameters are often not reported in sufficient detail in many research papers.

5. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.

[2] J. L. Herlocker, J. A. Konstan, and J. Riedl. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW'00)*, pages 241–250, New York, NY, USA, 2000.

[3] D. Jannach and K. Hegelich. A case study on the effectiveness of recommendations in the mobile internet. In *Proceedings of the 2009 ACM Conference on Recommender Systems (RecSys'09)*, pages 41–50, New York, NY, USA, 2009.

[4] D. Lemire and A. Maclachlan. Slope one predictors for online rating-based collaborative filtering. In *Proceedings of the 5th SIAM International Conference on Data Mining (SDM'05)*, pages 471–480, Newport Beach, CA, 2005.

[5] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web (WWW'01)*, pages 285–295, Hong Kong, 2001.

[6] J. Zhang and P. Pu. A recursive prediction algorithm for collaborative filtering recommender systems. In *Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys'07)*, pages 57–64, Minneapolis, MN, USA, 2007.

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

---

[4]http://ls13-www.cs.uni-dortmund.de/rec_suite.zip