

Batch versus Sequential Active Learning for Recommender Systems

Toon De Pessemier

Sander Vanhove

Luc Martens

Toon.DePessemier@ugent.be

Sander.Vanhove@ugent.be

Luc1.Martens@ugent.be

Imec - WAVES - Ghent University

9052 Gent, Belgium

ABSTRACT

Recommender systems have been investigated for many years, with the aim of generating the most accurate recommendations possible. However, available data about new users is often insufficient, leading to inaccurate recommendations; an issue that is known as the cold-start problem. A solution can be active learning. Active learning strategies proactively select items and ask users to rate these. This way, detailed user preferences can be acquired and as a result, more accurate recommendations can be offered to the user. In this study, we compare five active learning algorithms, combined with three different predictor algorithms, which are used to estimate to what extent the user would like the item that is asked to rate. In addition, two modes are tested for selecting the items: batch mode (all items at once), and sequential mode (the items one by one). Evaluation of the recommender in terms of rating prediction, decision support, and the ranking of items, showed that sequential mode produces the most accurate recommendations for dense data sets. Differences between the active learning algorithms are small. For most active learners, the best predictor turned out to be FunkSVD in combination with sequential mode.

CCS CONCEPTS

• **Information systems** → **Information systems applications**; *Data analytics*; Data mining.

KEYWORDS

Recommender System, Active Learning, Cold Start Problem

Reference Format:

Toon De Pessemier, Sander Vanhove, and Luc Martens. 2021. Batch versus Sequential Active Learning for Recommender Systems. In *4th Workshop on Online Recommender Systems and User Modeling (ORSUM 2021), in conjunction with the 15th ACM Conference on Recommender Systems, October 2nd, 2021, Amsterdam, The Netherlands (Remote)*.

1 INTRODUCTION

Recommender systems are tools and techniques to assist users in the content selection process thereby coping with the problem of information overload. But the algorithms under the hood suffer from

the cold start problem, i.e., the impossibility to provide accurate recommendations to new users or users with a small number of ratings. Moreover it has become essential to design recommender systems that are able to handle user feedback in real-time [1]. Active learning addresses these problems by proactively selecting items to be presented to the user in order to acquire her ratings and hence improve the output of the recommender. Active learner is a special case of machine learning that classifies data samples. It can acquire additional training samples by interactively querying a user for the correct labels of new data samples. The algorithm determines what data samples it is least certain about, and then asks the user to label these samples.

Active learning can also be used for recommender systems [12]. The assumption is that some items can be rated by the user, because the user is familiar with the item and can assess whether it matches her preferences or the user has experienced the item in the past. In fact, these items might already be consumed by the user outside the content platform. For a movie recommender for example, the user might have seen a movie through another service, at a friend's house, in movie theater, etc., with the result that no rating has yet been provided in the video recommender. Although these items are not rated yet, the user is able to specify a rating, which is very useful data for the recommender. Therefore, an active learner will select some items whose ratings are very informative for the recommender. Next, the user is asked to rate these items, if possible, and the ratings are used for generating recommendations. But not all ratings have the same information value. The ratings of universally-liked items, i.e. items rated high by most people, are typically less informative. An additional rating for such a universally-liked item is probably again a high rating, but is not very informative for the user's preferences.

As a result, it is important to make a good selection of items that are offered to the user for rating. This is the task of the active learner. All those items can be selected at once, this is called the *batch mode*. This selection of all these items is based on the ratings available before the active learning process starts. But those items can also be selected one by one, we refer to this as the *sequential mode*. The selection of the first item is based on the ratings available before the active learning. The selection of the second item is based on these ratings, as well as the first rating obtained by the active learning. This continues in the same way for subsequent item selections.

So, for the selection of the n -th item, the active learner has $n-1$ additional ratings to make a decision.

The goal of this study is to investigate the difference between batch and sequential active learning and the impact on the accuracy of the final recommender, which is fed with rating data of the active learner. Five active learning algorithms are evaluated. To estimate how informative an item is, some of these learners need a recommender algorithm that predicts how the user will evaluate the item; we refer to this as a “predictor algorithm”. The active learning algorithms are combined with three different predictor algorithms, and the impact of this predictor algorithm is evaluated. Active learning strategies are typically evaluated offline [2]. In this study, two publicly available dataset are used for the evaluation.

Related work has shown that active learning approaches significantly increase the number of ratings acquired from the user and the recommendation accuracy [3]. Many different active learning strategies have been proposed in literature [5]. Some of these early strategies are very simple, such as selecting the most popular items, the items with the highest variance, selecting random items, or a selection based on another aggregated statistic [10]. Later, some more advanced active learning strategies have been proposed [3][7]. These advanced strategies analyze the available ratings of each individual user, and calculate which item rating would be the most informative to obtain as the next one.

In addition, an active learning strategy has been proposed to combine the natural rating acquisition with an active learning algorithm [4]. Elahi et al. compared the recommender accuracy obtained with various of these active learners, but they did not distinguish between selecting items in batch or sequentially [5]. In online tests, the effect of the active learning technique on the user experience has been investigated [10]. However, according to our knowledge, the difference between active learning in batch and sequential mode has not yet been investigated for different algorithms.

The remainder of this paper is structured as follows. Section 2 defines the active learning algorithms that were used in this study. The used recommendation algorithms are described in Section 3. Section 4 discusses the experiment, with details about the data set, the data flow in the evaluation process, and the evaluation metrics. The results are reported in Section 5. Finally, Section 6 draws conclusions.

2 ACTIVE LEARNER ALGORITHMS

The goal of an active learner algorithm is to identify data set items, whose ratings are the most informative for a recommender system. Universally-liked items are typically less informative than controversial items in view of deriving user preferences and profiling. Two types of active learners can be distinguished: non-personalized and personalized active learners.

Non-personalized active learners identify items that cover different user preferences, based on certain item metrics, such as rating variance, rating entropy, or item popularity, calculated with aggregated rating data of the community. Items with a high variance or entropy are informative because different users rate these items in a different way. Popular items are informative because they are useful for calculating user-user similarities, which are based on

items rated by both users. Non-personalized active learners provide for each user the same set of “most informative items”, regardless of the individual user’s historical ratings.

Personalized active learners, also called prediction-based active learners [5], identify items for which a rating would be informative, given the individual user’s historical ratings. These personalized active learners provide for each user a personalized set of “most informative items”, calculated based on the user’s available ratings. For example, a rating for a specific item is less informative if a rating for a very similar item is already available. This personalized set of informative items is often calculated using an alternative recommendation algorithm, referred to as “the predictor algorithm”. In this study, the focus is on personalized active learners, which are the most advanced ones.

2.1 Binary Prediction Active Learner

The Binary Prediction Active Learner is using an external predictor to rank items. The goal is to identify items that have a high probability that the user can rate them. Asking users to rate unknown items is not meaningful, since users will skip the ratings process or provide unreliable ratings. A recommendation algorithm can act as a predictor of user’s item selections. Indeed, in an offline evaluation with datasets, a recommendation algorithm is often assessed on its ability to predict which items will be consumed by the user (e.g., with the precision and recall metric) [11].

Binary Prediction first transforms the rating matrix to a binary matrix with the same number of rows and columns, by mapping null entries (missing ratings) to 0, and not-null entries (available ratings) to 1 [3]. So, this binary matrix specifies which items have been rated by which users. Next, a recommendation algorithm is applied to this binary matrix, resulting in a prediction score for each item. This prediction score is used to rank the items. Items with a higher prediction score have a higher probability to be consumed by the user.

The idea of the Binary Prediction Active Learner is that the items with a higher score also have a higher probability that the user is familiar with these items, and as a result, can provide a rating. In this study, we utilized the Binary Prediction in combination with three different predictors described in Section 3: FunkSVD, Item-Item Collaborative Filtering, and User-User Collaborative Filtering.

2.2 Decision Tree Active Learner

Decision trees can be used to decide which item of the data set to ask the user to rate. For this study, we adopted the efficient tree learning algorithm published by Yahoo [7]. In this algorithm, each node of the decision tree evaluates user preference toward a certain item, and directs the user along a labeled edge to one of its subtrees based on the user’s rating. There are three options corresponding to three edges:

- (1) The user likes the item by scoring it higher than the average rating value of the data set.
- (2) The user dislikes the item by scoring it lower than the average rating value of the data set.
- (3) The user does not know the item and is unable to rate it.

Based on the user’s ratings, the user follows a path from the root to a leaf that characterizes her. Thus after each split, the set of users

is split into three partitions, one for each option. Subsequently, the algorithm continues with the partition of remaining users in the next node. Figure 1 shows a graphical representation of this decision tree.

The item used in each node of the tree is selected based on an error metric, which is defined as the squared deviation between the predicted ratings at the node and the true ratings of the users represented by the node. For more details regarding the construction of the tree, we refer to the original paper [7].




Figure 1: Graphical representation of the Decision Tree Active Learner.

2.3 Highest Prediction Active Learner

The Highest Prediction Active Learner applies a traditional predictor algorithm on the available data set and then asks the user to rate the items with the highest predicted score [5]. The idea is that items with a high prediction score have a high probability to be familiar to the user, and can receive a rating. In case the item with the high prediction score receives a low rating from the user, this rating is a good correction of the user preferences. As traditional predictor, we used the same three algorithms as for the other active learners: FunkSVD, Item-Item and User-User Collaborative Filtering.

This active learner approach is well suited if the recommendation algorithm used for the final recommendations is very prone for the cold start problem and the predictor algorithm used in the active learner is not.

2.4 Impact Analysis Active Learner

Impact Analysis Active Learner searches for items that enable the recommender to make a prediction for other items that could not be predicted before [3]. A user-based collaborative filtering algorithm cannot generate a prediction for an item, if no similar user has rated it. Impact Analysis can identify user-item pairs for which the rating reveals additional user similarities. These user similarities can then be used to calculate rating predictions for other items.

Figure 2 shows the user-item pairs for which ratings are available on the left. On the right, the most informative rating according to Impact Analysis is indicated, it is the rating for the pair (User1,Item4). Users are represented as nodes (circles) and items are nodes as well (squares). If a user has rated an item, an edge between the two nodes is drawn. Collaborative filtering algorithms need four-node paths such as (User1-Item1-User2-Item2). Impact Analysis searches

for edges that create the most additional four-node paths. In Figure 2, a four-node path can be found from User1 to Item1 to User2 to eventually Item2. Because User1 and User2 both rated Item1, a collaborative filter is able to make a prediction about the preference of User1 for Item2, given that User2 is a neighbor of User1.

Suppose that the active learner wants to gain additional information about User1, so that a collaborative filter can make a more informed decision. Which Item is the best to ask the user to rate? Item1 is already rated by User1. Item2 does not create any new four-node paths. The active learner could ask User1 to rate Item3 or Item5, because they both create one new four-node path, through User3 and User4 respectively. However, as indicated on the right side of Figure 2, the most informative rating is for Item4. A rating of User1 for Item4 creates two new four-node paths for User1: (User1-Item4-User3-Item3) and (User1-Item4-User4-Item5). If User1 is able to rate Item4, the collaborative filter is able to make a prediction about Item3 and Item5.




Figure 2: Graphical representation of the ratings. According to Impact Analysis Active Learner, the rating (User1,Item4) is the most informative.

2.5 Lowest Prediction Active Learner

The Lowest Prediction Active Learner is a variation on the Highest Prediction Active Learner [5]. As indicated by the name, this learner asks the user to rate the items with the lowest prediction score. The probability that the user can rate the item might be lower, since the user dislikes the item according to the predictor. In the movie domain for example, it is likely that the user has not yet seen a movie if it does not fit into his sphere of interest. As a result, the user might be unable to rate this movie.

The advantage of this active learner is that the recommender gets some negative feedback. In general, users are more likely to rate the items they like, than items they dislike. Since users typically chose for items they think they will like, the majority of the ratings in data sets is often positive. This problem is also referred to as the self-selection bias. As a result, the recommender has difficulties in learning the difference between what a user dislikes, and what the user has not yet tried.

3 RECOMMENDATION ALGORITHMS

The following three recommendations algorithms are used as predictor algorithm in the active learners.

3.1 User-User Collaborative Filtering

The User-User Collaborative Filtering algorithm calculates recommendations for the target user by identifying users that are similar in terms of rating behavior, also called neighbors [11]. The target user's preference for unrated items is estimated based on the ratings of these similar users for these unrated items. A standard collaborative filtering algorithm, as described in [11] has been implemented with the following parameters. Rating predictions are based on a minimum of 10 neighbors (to ensure accurate predictions) and maximum 30 neighbors (since more neighbors might not make the predictions more accurate, but increase the computational load and add more noise). For each of these neighbors, a minimum similarity value of 0.15, calculated with the Pearson correlation, is required.

3.2 Item-Item Collaborative Filtering

Item-Item Collaborative Filtering tries to identify similar items in terms of their ratings. Recommendations are calculated based on the similarities between items that received a high rating from the target user and items not yet rated by this target user. The standard algorithm is implemented with the following parameters. An item should have at least 1 neighbor, and a maximum of 20 neighboring items is used for each item. The similarity value of a neighboring item should be higher than 0.0. This threshold is rather low in order to find enough similar items for each item. In general, item-item similarities are more stable than user-user similarities, therefore less neighbors are used for item-item collaborative filtering.

3.3 Funk SVD

Funk SVD is a recommendation algorithm based on matrix factorization. Instead of an exact mathematical approach, it uses an iterative optimization based on gradient descent [6]. It is an effective algorithm that was successfully practiced by many others [9]. In our implementation, the amount of features per item is set to 25, the number of iterations is set to 100, and the learn rate is 0.001.

4 EXPERIMENTAL SETUP

The active learners are evaluated through an extensive experiment.

4.1 Data set

The evaluation is performed based on the Jester and the BookCrossing data set. The Jester data set is a rating set comprised of ratings about jokes, published by Goldberg et al. from the University of California [8]. Three different versions of the data set have been published. In this study, the data set is used that contains ratings for 100 jokes from 73,421 users summing to a total of 4.1 million ratings collected between 1999 and 2003. These ratings are ranging from -10.0 to +10.0 using a continuous scale. Each user in this data set has rated at least 15 jokes and the majority rated more than 35 jokes. So, compared to other data sets, this data set is rather dense. This allows the recommender to accurately learn user preferences, which is the reason for choosing this dataset.

The BookCrossing dataset [13] contains 278,858 users (anonymized but with demographic information) providing 1,149,780 ratings (explicit / implicit) about 271,379 books. Ratings are expressed on a scale from 1 to 10. This data set is very sparse, which makes it more difficult to learn user preferences.

4.2 Data Flow

The active learners of Section 2 have been evaluated, and the ones that need a predictor are tested in combination with the algorithms of Section 3. As final recommender algorithm, used to generate the recommendations for the users, the same algorithms of Section 3 were chosen. For each active learner, we evaluate the batch as well as the sequential mode. The evaluation is performed in the following subsequent phases:

- (1) *Select a user.* This user's ratings are removed from the training set and stored in a control data set. Except one rating of this user, randomly chosen, remains in the training set, and is used as input for the active learner. This simulates a new user, who has rated one item.
- (2) *Use the active learner.*
 - (a) *Train the active learner.* The current training set is used.
 - (b) *Select an item.* The active learner is used to calculate the next item that is the most informative for the user.
 - (c) *Add the rating to the training set.* If the rating for the selected user-item pair exists in the control set, it is marked as available and added to the training set. Otherwise, this rating is marked as missing.
 - (d) *Repeat item selection.* Items are selected until 6 ratings of this user could be added to the training set. (6 ratings + 1 initial rating showed to be a good number for generating recommendations for new users.)
 - In Batch mode \Rightarrow go to 2b.
 - In Sequential mode, and the rating is marked as available in 2c \Rightarrow go to 2a.
 - In Sequential mode, and the rating is marked as missing in 2c \Rightarrow go to 2b.
- (3) *Run the recommender.* Use the training set that includes the learned ratings from 2c to generate recommendations for the user.
- (4) *Repeat user selection.* This is repeated until 100 users are evaluated \Rightarrow go to 1.

4.3 Evaluation Metrics

The accuracy of the final recommendations is evaluated using three commonly-used metrics, Root Mean Square Error (RMSE), Precision, and Normalized Discounted Cumulative Gain (nDCG). RMSE is a rating prediction metric that measures the accuracy of an algorithm based on the error between the predicted (\hat{r}_{ui}) and real (r_{ui}) rating value of a user u , for an item i . It squares the individual errors, calculates the mean value over all ratings (n = number of ratings), and takes the square root of the mean to bring this to the scale of the ratings. The lower the value of the RMSE the better the rating prediction, where a value of zero stands for the perfect prediction.

$$RMSE = \sqrt{\frac{1}{n} \sum_{u \in U} \sum_{i \in I} (\hat{r}_{ui} - r_{ui})^2} \quad (1)$$

Precision is a decision support metric that measures the percentage of recommended items that are relevant. The higher the precision, the better the recommendations. A perfect selection of recommendations will obtain a precision value of 1.

$$\text{Precision} = \frac{\# \text{items relevant and recommended}}{\# \text{items recommended}} \quad (2)$$

The nDCG is a ranking metric that measures whether the algorithm ranks the items in the same order as the user. To calculate the nDCG, first the DCG of a list of recommendations is calculated. This is the discounted cumulative gain. It accumulates the gain of a list of items and multiplies the gain with a discount factor based on that item's position in the list.

$$\text{DCG}(L) = \sum_{i \in L} \text{gain}(i) * \text{disc}(i) \quad (3)$$

Where L is the ordered list of items offered to the user as recommendations. As common practice for data sets with ratings, the gain is defined as the rating the user u gave to item i :

$$\text{gain}(i) = r_{ui} \quad (4)$$

The discount factor for each item is proportional to its position within the list and defined as follows:

$$\text{disc}(i) = \frac{1}{\log_2(\text{rank}(i) + 1)} \quad (5)$$

Where $\text{rank}(i)$ is the position of the item within the list. Because of this discount factor, items that are presented first are more important than items at the back of the list.

This mimics the behavior of real users in the way that users' attention will be more drawn to and focused on the first items of the list. Gradually, their attention typically weakens over time towards the end of the list. After calculating the DCG of the proposed recommendation list, it can be normalized as follows:

$$\text{nDCG}(L) = \frac{\text{DCG}(L)}{\text{DCG}(L_{\text{ideal}})} \quad (6)$$

Where L_{ideal} is the ideal list to present to the user. This ideal list is a list of items, ordered according to the user's ratings from highest to lowest rating. These ratings can be available in the test set (or the control data set in our experiment) for evaluation. Of course, this ideal list is limited to the ratings that are available for the user and will almost never be the true ideal list for that user. The normalized discounted cumulative gain has a value between 0 and 1, the higher the better.

5 RESULTS

Figure 3 shows the RMSE of the User-User Collaborative Filtering algorithm that is fed with rating data as determined by the active learner. The results are based on the Jester dataset. So for the main recommender, User-User Collaborative Filtering is always used in this graph. For the predictor algorithm that is part of the active learner, different options are tested (FunkSVD, UserUser, and ItemItem). The lower the RMSE, the better the algorithm can predict the ratings, and typically, the better the recommendations. The box plots show the minimum, first quantile (25%), median, third quantile (75%) and maximum value of the RMSE for all users in the dataset.

Sequential active learning is clearly better than batch active learning. In sequential mode, informative items are selected one by one, and the ratings are moved from the control set to the training set. The selection of the most informative item is based on the available historical ratings (i.e. the original training set), as well as the ratings for the previously selected items that are informative for the recommender (i.e. ratings moved from control set to training set). So, sequential active learning can adjust the selection of informative items, one by one. In batch mode, all the informative items have to be selected at once, based only on the historical ratings. But ratings for these informative items cannot be taken into account. So, batch active learning has less information to make a decision, and as a result, it makes a suboptimal selection of informative items.

Comparing the different active learning algorithm shows that the Lowest Prediction Active Learner has the best results. But the difference with other active learners such as Highest Prediction and Binary Prediction are small.

When we investigate the influence of the predictor used for the active learning algorithm (FunkSVD, ItemItem, or UserUser), we do not see major differences in the accuracy of the final recommendations. But FunkSVD combined with sequential mode gives consistently good results for all active learners. All these predictors are able to distinguish between the items users like and items they dislike. Minor differences in the predictions of these algorithms have only a limited impact on the final recommendations.

Figure 4 shows the precision of the User-User Collaborative Filtering algorithm combined with different active learners for the Jester dataset. Again, a clear difference between sequential and batch mode can be recognized. Sequential mode obtains significantly better results in terms of precision. The differences between the active learning algorithms are small. There is no clear winner of the active learning algorithms.

Figure 5 shows the nDCG of the User-User Collaborative Filtering algorithm combined with different active learners for the Jester data set. The higher the nDCG, the better the algorithm ranks the items, and typically, the better the recommendations. From these box plots, similar conclusions can be derived. More accurate recommendations are achieved using sequential active learning than using batch active learning. The differences in rank accuracy for the various active learners are small; but the best results (in terms of nDCG) are obtained with the Impact Analysis active learner. FunkSVD as predictor combined with sequential mode yields very accurate recommendations; except for the Binary Prediction active learner, this was not the best combination.

Figures 6, 7, and 8 repeat this analysis for FunkSVD instead of User-User collaborative filtering as main recommendation algorithm. Figure 9 evaluates the nDCG for Item-Item Collaborative filtering on the Jester dataset. These graphs confirm the results for other recommendation algorithms. Active learning in sequential mode outperforms batch mode in terms of recommender accuracy. But the disadvantage of the sequential mode is that after each item selection, the active learner needs to re-run, which can be computationally expensive. We also conclude that the differences between the active learning algorithms in terms of accuracy are small.

The analysis was also performed for the BookCrossing data set. Figures 10, 11, and 12 show the nDCG for respectively User-User

collaborative filtering, FunkSVD, and Item-Item collaborative filtering as main algorithm combined with various active learners. An interesting conclusion can be drawn here: the results obtained with active learning in sequential mode are not significantly different than with the active learning in batch mode. This was also concluded in terms of the precision and RMSE for the BookCrossing data set (not shown in this paper). The sparsity of the BookCrossing data set can be the reason that sequential mode is not better than batch mode.

6 CONCLUSION AND FUTURE WORK

Active learning algorithms carefully select informative items and ask users to rate these in order to get a more detailed view on the users' preferences and being able to generate accurate recommendations. Two modes have been compared for active learning: the batch mode, that selects all items at once, and the sequential mode, that selects items one by one and calculates in between these selections which item would be the most informative one. Five different active learning algorithms have been evaluated, and three of them were combined with three different prediction algorithms. For a dense data set, the results showed that an active learner in sequential mode enables a recommender to learn more efficiently than an active learner in batch mode. So, if users are asked to rate some items, it is better to select these items one by one, than selecting all items at once. Based on the user's ratings for the first items, the sequential active learner can better decide which subsequent items are the most informative. For a sparse dataset, no significant differences between batch and sequentially mode were obtained. For active learners that use an underlying prediction algorithm, FunkSVD turned out to be a good choice, but the accuracy differences with other predictors are small. In future work, active learners will be tested in combination with other recommendation algorithms and additional active learning algorithms will be explored.

ACKNOWLEDGMENTS

This research received funding from the Flemish Government (AI Research Program).

REFERENCES

- [1] Marie Al-Ghossein, Pierre-Alexandre Murena, Talel Abdesslem, Anthony Barré, and Antoine Cornuéjols. 2018. Adaptive collaborative topic modeling for online recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 338–346.
- [2] Diego Carraro and Derek Bridge. 2020. Debaised offline evaluation of active learning in recommender systems. In *The Thirty-Third International Flairs Conference*.
- [3] Mehdi Elahi, Matthias Braunhofer, Francesco Ricci, and Marko Tkalcić. 2013. Personality-based active learning for collaborative filtering recommender systems. In *Congress of the Italian Association for Artificial Intelligence*. Springer, 360–371.
- [4] Mehdi Elahi, Francesco Ricci, and Neil Rubens. 2012. Adapting to natural rating acquisition with combined active learning strategies. In *International Symposium on Methodologies for Intelligent Systems*. Springer, 254–263.
- [5] Mehdi Elahi, Francesco Ricci, and Neil Rubens. 2014. Active learning strategies for rating elicitation in collaborative filtering: A system-wide perspective. *ACM Transactions on Intelligent Systems and Technology (TIIST)* 5, 1 (2014), 1–33.
- [6] Simon Funk. 2006. *Netflix update: Try this at home*. <https://sifter.org/~simon/journal/20061211.html>2020.3.21
- [7] Nadav Golbandi, Yehuda Koren, and Ronny Lempel. 2011. Adaptive bootstrapping of recommender systems using decision trees. In *Proceedings of the fourth ACM international conference on Web search and data mining*. 595–604.
- [8] Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. 2001. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval* 4, 2 (2001), 133–151.
- [9] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 426–434.
- [10] Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K Lam, Sean M McNee, Joseph A Konstan, and John Riedl. 2002. Getting to know you: learning new user preferences in recommender systems. In *Proceedings of the 7th international conference on Intelligent user interfaces*. 127–134.
- [11] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. Springer, 1–35.
- [12] Neil Rubens, Mehdi Elahi, Masashi Sugiyama, and Dain Kaplan. 2015. Active learning in recommender systems. In *Recommender systems handbook*. Springer, 809–846.
- [13] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*. 22–32.




Figure 3: RMSE value of user-user collaborative filtering combined with active learners with different algorithms on the Jester dataset. Lower is better.




Figure 4: The precision of user-user collaborative filtering combined with active learners with different algorithms on the Jester dataset. Higher is better.




Figure 5: nDCG value of user-user collaborative filtering combined with active learners with different algorithms on the Jester dataset. Higher is better.




Figure 6: RMSE value of the FunkSVD algorithm combined with active learners with different algorithms on the Jester dataset. Lower is better.




Figure 7: The precision of the FunkSVD algorithm combined with active learners with different algorithms on the Jester dataset. Higher is better.




Figure 8: nDCG value of the FunkSVD algorithm combined with active learners with different algorithms on the Jester dataset. Higher is better.




Figure 9: nDCG value of item-item collaborative filtering combined with active learners with different algorithms on the Jester dataset. Higher is better.




Figure 10: nDCG value of user-user collaborative filtering combined with active learners with different algorithms on the Bookcrossing dataset. Higher is better.




Figure 11: nDCG value of FunkSVD combined with active learners with different algorithms on the Bookcrossing dataset. Higher is better.




Figure 12: nDCG value of item-item collaborative filtering combined with active learners with different algorithms on the Bookcrossing dataset. Higher is better.