

Temporal Diversity in Recommender Systems

Neal Lathia[§], Stephen Hailes[§], Licia Capra[§], Xavier Amatriain[†]

[§]Dept. of Computer Science, University College London, Gower Street WC1E 6BT, UK

[†]Telefonica Research, Via Augusta 177, Barcelona 08021, Spain
n.lathia, s.hailes, l.capra@cs.ucl.ac.uk, xar@tid.es

ABSTRACT

Collaborative Filtering (CF) algorithms, used to build web-based recommender systems, are often evaluated in terms of how *accurately* they predict user ratings. However, current evaluation techniques disregard the fact that users continue to rate items *over time*: the temporal characteristics of the system's top- N recommendations are not investigated. In particular, there is no means of measuring the extent that the *same items* are being recommended to users over and over again. In this work, we show that temporal diversity is an important facet of recommender systems, by showing how CF data changes over time and performing a user survey. We then evaluate three CF algorithms from the point of view of the *diversity* in the sequence of recommendation lists they produce over time. We examine how a number of characteristics of user rating patterns (including profile size and time between rating) affect diversity. We then propose and evaluate set methods that maximise temporal recommendation diversity without extensively penalising accuracy.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Information Filtering

General Terms

Algorithms, Performance, Theory

Keywords

Recommender Systems, Evaluation

1. INTRODUCTION

Recommender systems have become essential navigational tools for users to wade through the plethora of online content. Many of them use collaborative filtering (CF) algorithms, that compute recommendations using the ratings that each user has input. CF algorithms are often evaluated according to how *accurately* they predict these ratings,

or by measuring the *precision* and *recall* of ranked recommendations [1]. However, the set of evaluation measures available disregard the fact that the process of inputting ratings and receiving recommendations happens *iteratively* over time. Users update their profiles in an incremental manner, rating content as they consume it, and typical CF algorithms are retrained regularly (e.g., weekly [2]) to reflect this. As recommender systems grow dynamically, a problem arises: current evaluation techniques do not investigate the temporal characteristics of the produced recommendations. Researchers have no means of knowing whether, for example, the system recommends *the same items* to users over and over again, or whether the most *novel* content is finding its way into recommendations. The danger here is that, as results may begin to stagnate, users may lose interest in interacting with the recommender system.

In this work, we investigate one dimension of recommendations: the diversity of top- N lists over time. We first examine why temporal diversity may be important in recommender system research (Section 2) by considering temporal *rating patterns* and the results of a user survey. Based on these observations, we evaluate three CF algorithms' temporal diversity from 3 perspectives (Section 3): by comparing the intersection of sequential top- N lists, by examining how diversity is affected by the number of ratings that users input, and by weighing-in the trade-off between accuracy and diversity over time. We finally *design* and *evaluate* a hybrid mechanism to promote temporal diversity (Section 4), comparing its performance to a range of baseline techniques. We conclude in Section 5 by discussing future research directions.

2. WHY TEMPORAL DIVERSITY?

We explore the importance of temporal diversity from two perspectives: (a) changes that CF data undergoes over time and (b) how surveyed users respond to recommendations with varying levels of diversity.

(a) **Changes Over Time.** Recommender systems grow over time: new users join the system and new content is added as it is released. Pre-existing users can update their profiles by rating previously unrated content; the overall volume of *data* thus grows over time. As a consequence of the continuous influx of ratings, any summary *statistics* related to the recommender system's content may also change. We visualised these changes in the Netflix prize dataset¹, a large-scale collection of ratings for movies. The ratings in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGIR'10, July 19–23, 2010, Geneva, Switzerland.

Copyright 2010 ACM 978-1-60558-896-4/10/07 ...\$5.00.

¹<http://www.netflixprize.com>

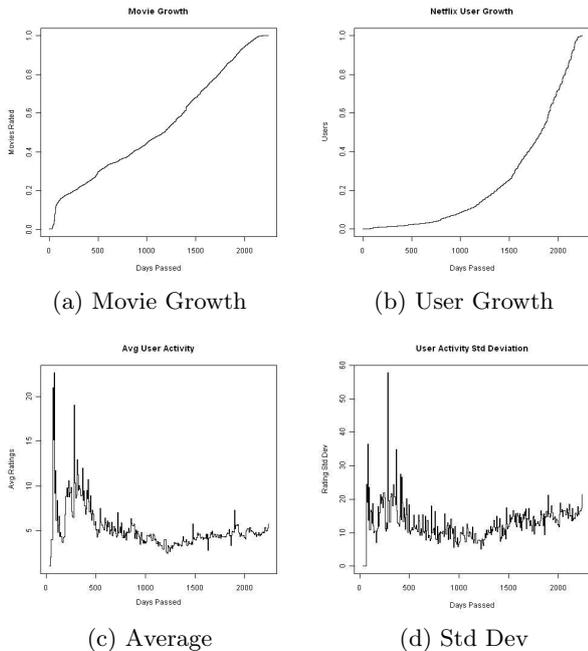


Figure 1: Growth of Netflix Dataset Over Time & Average/Standard Deviation of Returning Users' Ratings Input Per Week

this set span a total of 2,243 days. We plot the movie and user growth in Figures 1(a) and 1(b) respectively; from these plots we see that there is a continuous arrival of both new users and movies. While new movies seem to be added at a relatively linear pace, the number of users grows exponentially over time. We now turn to the rating frequency: in Figures 1(c) and 1(d) we plot the average and standard deviation of ratings input by returning users (i.e., users who have already previously rated at least once) per week. The plots show the high variability in how users interact with the recommender system. After the initial high fluctuation in average user ratings per week, the mean value flattens out at approximately five ratings per user per week. However, the standard deviation of this mean falls within [5, 20]: rating behaviour also varies over time, and viewing rating sets from a non-temporal viewpoint does not account for these changes. These changes affect the ratings' summary statistics: in [3], Koren shows how *global* summary statistics vary over time. Similarly, our previous work looks at how changes are further reflected in the global rating median and mode [4]. All the summary values fluctuate over time, reflecting how the overall distribution of ratings shifts as more users interact with the system.

What do we learn from observing these changes? The datasets do not only remain incredibly sparse, but they also do not stabilise; recommender systems continuously have to make decisions based on *incomplete* and *changing* data, and the range of the changes we observe in the Netflix data have a strong impact on the predictability of ratings. Furthermore, the continuous rating of content means that the data that an algorithm will be trained with at any particular time is likely to be different than data it trained with previously.

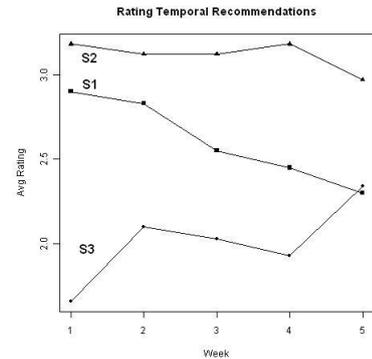


Figure 2: Survey Results for (S1) Popular Movies With No Diversity (S2) Popular Movies With Diversity and (S3) Randomly Selected Movies

The question we explore in this paper is: do these changes translate into different recommendations over time?

(b) User Survey. In order to determine whether temporal diversity is important for recommender system *users*, we designed three surveys that *simulate* systems that produce *popular movie* recommendations over the course of five “weeks.” We opted to recommend popular movies in order to avoid a variety of confounds that would emerge had we selected a personalised CF algorithm (e.g., the quality of the algorithm itself and the cumbersome process of asking users to rate films). Survey 1 (S1) and Survey 2 (S2) both recommended popular movies drawn from a list of the 100 all time most profitable box office movies². S1 had *no diversity*: it consistently recommended the top-10 box office hits. S2’s recommendations, instead, did change over time. Each week, approximately seven of the previous week’s ten recommendations would be replaced by other movies in the top-100 box office list. Lastly, Survey 3 (S3) recommended movies that were randomly selected from the Netflix dataset: the recommendation process included full diversity, but was very unlikely to recommend popular movies.

Each survey was structured as follows. The users were first queried for demographic data. They were then offered the first week’s recommendations, represented as a list of ten movie titles (and the relative DVD covers and IMDB links) and asked to rate these top-10 recommendations on a 1-5 star scale. After submitting their rating, they were presented with a buffer screen containing thirty DVD covers, and had to click to continue to the subsequent week; this aimed at diverting users’ attention before presenting them with the next week’s recommendations. After rating all five week’s worth of recommendations, they were asked to comment on the recommendations themselves and answer a number of questions relating to diversity over time. Users were invited to participate in one or more of the surveys: S1 was completed 41 times, S2 had 34 responses, and S3 was completed 29 times. Due to the surveys’ anonymity, we do not know how many users completed more than one survey. We therefore treat each completed survey individually. Of the 104 total responses, 74% of the users were male, 10% were 18-21 years old, 66% were 22-30 years old, and 24% were between 31 and 50 years of age. On average, the users

²<http://www.imdb.com/boxoffice/alltimegross>

claimed to watch 6.01 ± 6.12 movies per month, and while 61% of them said they were familiar with recommender systems, over half of them claimed they used them less than once a month. On the other hand, 29% use recommender systems weekly or daily: our respondents therefore include a wide variety of movie enthusiasts and both people who do and do not use recommender systems.

We averaged the users’ ratings for each week’s recommendations and plot the results in Figure 2. The S2 results (popular movies with diversity) achieve the highest scores: on average, these five weeks of recommendations were rated 3.11 ± 0.08 stars. The low temporal standard deviation reflects the fact that the rating trend remains relatively flat; the average for each week is about 3 stars. S3’s results (randomly selected movies), were consistently disliked: the average rating peaks at 2.34 stars for week 5. In fact, some users commented on the fact that recommendations “appeared to be very random,” “varied wildly” and the system “avoid[ed] box office hits.” The main result of our surveys is reflected in S1’s results (popular movies with no diversity): as the same recommendations are offered week after week, the average ratings *monotonically* decrease. The average for week 1 was 2.9, which falls within the range of values measured in S2, while by week 5 the average score is 2.3, which is lower than the average score for the random movies. Not all commented on the lack of recommendations diversity; however, most modified their ratings for the recommendations as the lack of diversity persisted. This shows that when users rate they are not only expressing their tastes or preferences; they are also responding to the impression they have of the recommender system. In the case of S1, users commented on the fact that the algorithm was “too naive” or “not working,” and the lack of diversity “decreased [the respondent’s] interest.”

In order to test the statistical significance of the three different survey’s results, we performed an analysis of variance (ANOVA): in this case, the null hypothesis is that the ratings for each survey are the same. We can reject the null hypothesis with 99% confidence with p-values less than 0.01: the p-value we measured for the three methods is $9.72e^{-14}$. A pairwise t-test between each survey further shows that the ratings input for each survey cannot be attributed to the sampling of the study; we omit the boxplots due to lack of space.

The final part of the surveys asked users about qualities they sought in recommendations. Overall, 74% said it was important for recommender systems to provide results that *accurately* matched their taste (23% selected the ‘neutral’ option to this question). 86% said it is important for recommendations to *change over time*; in fact, 95% stated it is important that they are offered new recommendations. It thus quickly becomes apparent that temporal diversity is a highly important facet of recommender systems, both in terms of the direct responses and rating behaviour of the surveyed users. In the following sections, we evaluate the temporal diversity of three state of the art CF algorithms.

3. EVALUATING FOR DIVERSITY

Given the above, we now aim to examine how diverse CF algorithms are over time. We focus on three algorithms: a **baseline**, where a prediction for an item is that item’s mean rating, the **item-based k-Nearest Neighbour (kNN)** algorithm, and a **matrix factorisation** approach based on

Singular Value Decomposition (SVD). Due to limited space, we do not detail each algorithm; Adomavicius and Tuzhilin [5] thoroughly review these and many more approaches. We chose these algorithms since they not only reflect state-of-the-art CF, but also each manipulate the rating data in a different way and may thus produce varying recommendations.

All of the algorithms share a common theme: they produce predicted ratings that can then be used to recommend content. The idea is to use the predictions in order to generate a personalised ranking of the system’s content for each user (note that, in our scenario, items that have been rated cannot be recommended). However, it may be the case that items share the same predicted rating. For example, a number of items may all have 5-star predictions. In this case, the predicted rating alone is not conducive to a meaningful ranking. We solve this problem by introducing a *scoring function* to rank items, regardless of the model used to generate predicted ratings. The scoring function uses two pieces of information: the predicted rating, and the *confidence* in the prediction (i.e., number of data points used to derive it) as used in [6]. Assuming a 5-star rating scale, we first subtract the scale mid-point (3 stars) from the prediction and then multiply by the confidence:

$$s_{\hat{r},i} = (\hat{r}_{u,i,t} - 3.0) \times \text{confidence}(\hat{r}_{u,i,t}) \quad (1)$$

This scoring function ensures that items with high prediction and confidence are promoted, and low prediction with high confidence are demoted. For example, an item with a predicted 5 star rating, derived from 2 ratings, will be ranked lower than another item with a 4 star prediction based on 50 ratings. If two items had the same score, then we differentiated them based on their respective average rating date: the item that had been rated more recently is ranked higher. The greatest advantage of this method, as detailed in [6], is the heightened *explainability* of recommendations.

Methodology. In order to examine the sequence of recommendations produced by a system, we explore CF algorithms that iteratively re-train on a *growing dataset*. Given a dataset at time t and a window size μ (how often the system will be updated), we train the algorithm with any data input prior to t and then *predict* and *rank* all of the unrated items for each user. The t variable is then incremented by μ , and the entire process is repeated, except that now the latest ratings become incorporated into the training data. In other words, at time t we generate a set of top- N lists—corresponding to the top- N recommendations each user would receive—in order to then examine how the sequence of ranked items that we produce will vary as the system is updated. This method includes a number of advantages: we test the algorithms as data grows (and view more than a single iteration of this process), making predictions based only on ratings that are currently available. We simulate the iterative update of deployed systems, and stay true to the order users input ratings.

Since users do not necessarily log-in consistently to the system, we cannot be certain that each top- N list would have been viewed by each user. We therefore only generate a top- N list for the users who will rate at least one item in time $(t + \mu)$; we assume that if the user is rating an item then they have logged into the system and are likely to have seen their recommendations. The benefit of this is that we compare the current recommendations to those

that users are likely to have seen before. It remains possible that users viewed their recommendations without rating any content; however, given this uncertainty in the data, we only consider the scenario where there is evidence that the users have interacted with the system.

We use 5 subsamples of the Netflix dataset for cross-validation purposes. Each subsample includes 50,000 users and any rating input (by any user) prior to a pre-defined “edge” time $\epsilon = 500$ days after the first date in the Netflix dataset. Our final subsets have approximately 60,000 users; setting the ϵ value as we did is roughly equivalent to bootstrapping a new recommender system with 10,000 users rating content for over a year (we thus hope to avoid settings where our results will be skewed by system-wide cold start problems). We set the window size $\mu = 7$ days; the system will be updated weekly [2]. Since the Netflix data spans a total of 2,243 days, selecting $\mu = 7$ days allows us to explore 249 system updates.

3.1 Measuring Diversity Over Time

We define a way to measure the diversity between two ranked lists as follows. Assume that, at time t , a user is offered a set of 10 recommendations. The next time the user interacts with the system only 1 of the 10 recommendations is different. Therefore, the diversity between the two lists is $\frac{1}{10} = 0.1$. More formally, given two sets L_1 and L_2 , the set theoretic difference (or relative complement) of the sets denotes the members of L_2 that are not in L_1 :

$$L_2 \setminus L_1 = \{x \in L_2 | x \notin L_1\} \quad (2)$$

In our example above, only 1 of the 10 recommendations was not the same: the set theoretic difference of the two recommendation lists has size 1. We thus define the diversity between two lists (at depth N) as the size of their set theoretic difference over N :

$$\text{diversity}(L_1, L_2, N) = \frac{|L_2 \setminus L_1|}{N} \quad (3)$$

If L_1 and L_2 are exactly the same, there is no diversity: $\text{diversity}(L_1, L_2, N) = 0$. If the lists are completely different, then $\text{diversity}(L_1, L_2, N) = 1$. This measure disregards the actual ordering of the items: if a pair of lists are re-shuffled copies of each other, then there continues to be no diversity. However, we can measure the extent that recommendations change as a result of the same content being promoted or demoted by measuring diversity at varying depths (N).

One of the limitations of this metric is that it measures the diversity between two lists; it only highlights the extent that users are being *sequentially* offered the same recommendations. In order to see how recommendations change, in terms of *new* items appearing in the lists, we define a top- N list’s *novelty*. Rather than, as above, comparing the current list L_2 to the previous list L_1 , we compare it to the set of *all* items that have been recommended to date (A_t):

$$\text{novelty}(L_1, N) = \frac{|L_1 \setminus A_t|}{N} \quad (4)$$

In this context, we define novelty in terms of what items have been previously recommended to a user. A list’s novelty will be high if all of the items have *never* been recommended before, and low if all of the items have been recommended at some point in the past (not just in the last update). We further define the *average diversity* δ_t and *average novelty*

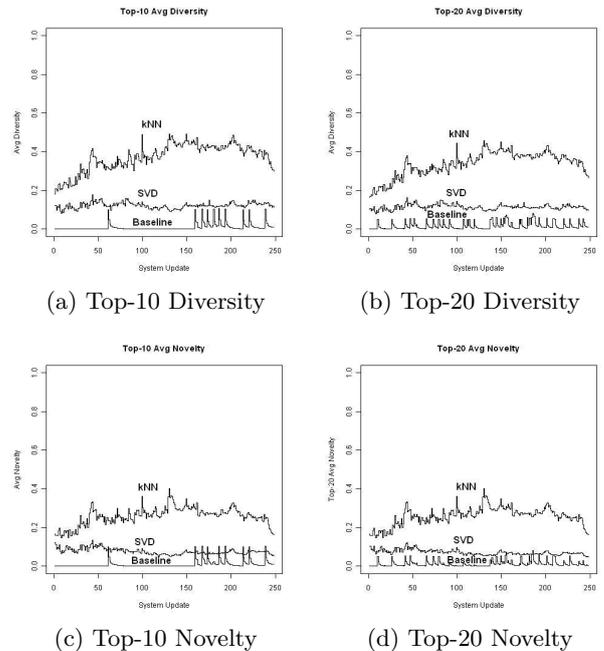


Figure 3: Top-10 and 20 Temporal Diversity and Novelty for Baseline, kNN and SVD CF

η_t that is generated by a given CF algorithm (at time t) as the average of the values computed between all the current top- N lists and the respective previous list for each user.

3.2 Results and Analysis

We computed δ_t and η_t for each of the 3 algorithms over all 249 simulated system updates outlined in Section 3, and report the results for the top-10 and top-20 recommendations in Figure 3. These results provide a number of insights into recommender system’s temporal diversity. As expected, the baseline algorithm produces little to no diversity. On average, users’ top-10 recommendations differ by (at most) one item compared to the previous recommendations. Both the factorisation and nearest neighbour approaches increment diversity; furthermore, the k NN algorithm is, on average, consistently more diverse than the sequence of recommendations produced by the SVD.

The novelty values (Figures 3(c) and 3(d)) are lower than the average diversity values. That means that when a different recommendation appears, it is more often a recommendation that has appeared at some point in the past, rather than something that has not appeared before. There are a variety of factors that may cause this; for example, new items may not be recommended because they lack sufficient ratings: the CF algorithm cannot *confidently* recommend them. However, this metric does not tell us whether the new recommendations are new items to the system, or simply content that has (to date) not been recommended. A full analysis of the novelty of recommendation warrants a closer inspection of when items join the system and when they are recommended. In order to focus our analysis, we thus separate the problems of recommending *new content* from that of *diversifying* sequential recommendations: in this work, we focus on the latter.

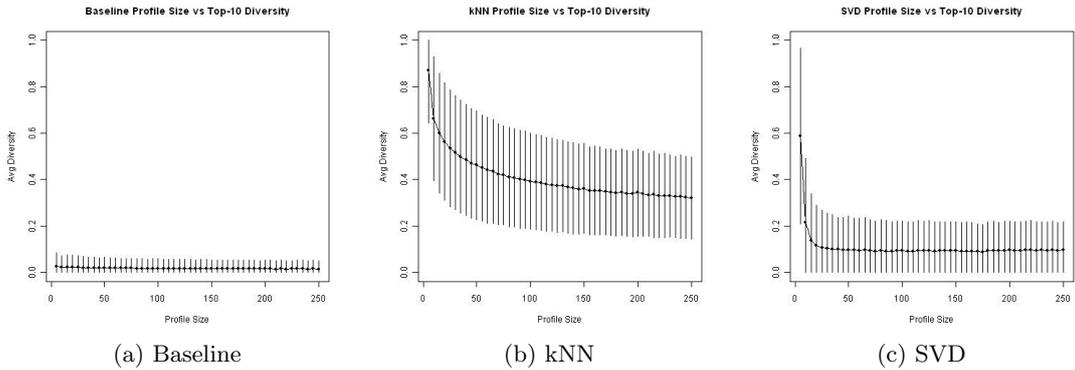


Figure 4: Profile Size vs. Top-10 Temporal Diversity for Baseline, kNN and SVD CF

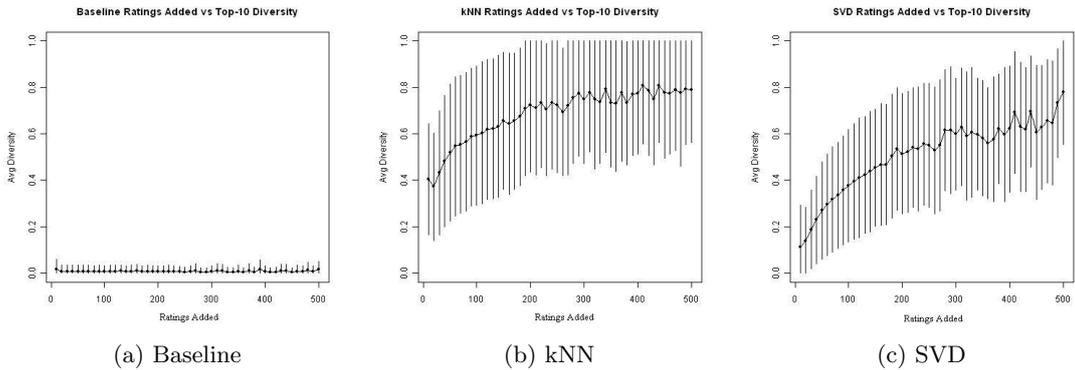


Figure 5: Ratings Added vs. Top-10 Temporal Diversity for Baseline, kNN and SVD CF

Both Figure 3(a) and 3(b) also look very similar: the diversity values for the top-10 and top-20 recommendations are nearly the same. In order for this to happen (i.e., for a comparison between two top-10 lists and two top-20 lists to produce the same value) there must be *more* diversity between the larger lists. For example, if only 1 item changes in the top-10, the diversity is $\frac{1}{10} = 0.1$, and the pair of top-20 lists will only produce this diversity value if 2 items have changed, $\frac{2}{20}$. What this means is that not all of the changes in the item rankings are occurring in the top-10: new items are also being ranked between the 11th and 20th positions.

At the broadest level, we thus observe that (a) both the baseline and SVD produce less temporal diversity than the k NN approach, and (b) across all CF algorithms, diversity is never higher than approximately 0.4. However, these are averaged results across many users, who may be each behaving in very different ways: we now perform a finer grained analysis of temporal diversity to explore the relation between users and the diversity they experience.

Diversity vs. Profile Size. The metric in Section 3.1 does not factor in the fact that the distribution of ratings per user is not uniform. Some users have rated a lot of items, while others have very sparse profiles. Users’ *profile size* (i.e., the number of ratings per user) may affect their recommendation diversity. We therefore binned the above temporal results according to users’ current profile size and then averaged the diversity of each group. We plot the re-

sults in Figure 4. The baseline (Figure 4(a)) continues to show next to no diversity, regardless of how many items users have rated. The rationale behind this is that the only profile information that the baseline factors in when it computes recommendations is whether the user has rated one of the popular items; results will only be diverse if the user rates all the popular content. The k NN (Figure 4(b)) and SVD (Figure 4(c)) results, instead, show a negative trend: diversity tends to reduce as users’ profile size increases.

Diversity vs. Ratings Input. Our temporal diversity metric is based on *pairwise* comparisons; we compare each sequential pair of top- N lists. One factor that may thus play an important role when determining how diverse a pair of lists will be from one another is *how much* the user rates in a given session. For example, one user may log in and rate two items while another may log in and rate fifty; the temporal diversity that each user subsequently experiences may be affected by these new ratings. We therefore binned users according to how many new ratings they input, and plot the results in Figure 5. As before, the baseline remains unaffected by how many new ratings each user inputs. The k NN (Figure 5(b)) and SVD (Figure 5(c)), instead, show a positive trend. These results can be interpreted as follows: the more you rate now, the more diverse your *next* recommendations will be in the next session.

Diversity and Time Between Sessions. The previous analysis was concerned with how diversity is influenced

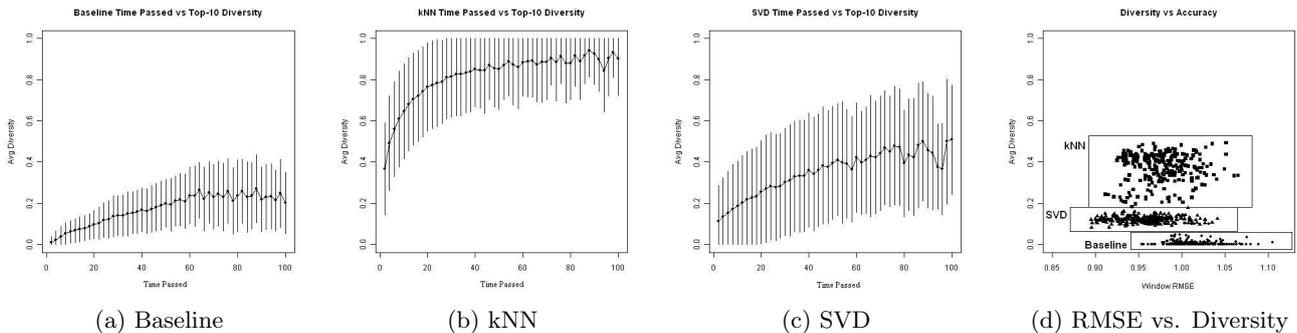


Figure 6: Time Passed vs. Top-10 Temporal Diversity for Baseline, kNN and SVD CF and Comparing Accuracy with Diversity

by a *single* user rating content. However, users do not rate alone: an entire community of users rate content over extended periods of time. We highlight this point with an example: some users may consistently log in and rate items every week; others may rate a few items now and not return for another month (and, in their absence, other users will have continued rating). In other words, diversity may be subject to the *time* that has passed from when one list and the next are served to the user. In order to verify this, we binned our diversity results according to the number of weeks that had passed between each pair of lists, and plot the results in Figure 6. In this case, all three of our algorithms show a positive trend: the longer the user does not return to the system, the more diversity increases. Even the baseline diversity increases: if a user does not enter the system for a protracted period of time, the popular content will have changed.

Lessons Learned. Overall, *average temporal diversity* is low. Ranking content based on popularity offers next to no diversity, while the *kNN* method produces the largest average temporal diversity. Larger *profile sizes* negatively affect diversity; it seems that users who have already rated extensively will see the least diverse recommendations over time. Pairwise diversity between sequential lists is largest when users rate many items before receiving their next recommendations; users should be encouraged to rate in order to change what they will be recommended next. Diversity will naturally improve as users extend the *time between* sessions when they interact with the system (even popular content eventually changes).

The final question we ask is how diversity relates to accuracy, the metric of choice is traditional CF research. To do so, we take the predictions we made at each update, and compute the Root Mean Squared Error (RMSE) between them and the ratings the visiting users will input. We then plot RMSE against average diversity in Figure 6(d). A plot of this kind has four distinct regions: low accuracy with low diversity (bottom right), high accuracy with low diversity (bottom left), low accuracy with high diversity (top right), and high accuracy with high diversity (top left). We find that the results for each algorithm cluster into different regions of the plot, corresponding to the different diversity results that they obtain. In terms of RMSE, different algorithms often overlap; for example, *kNN* CF is sometimes less accurate than the baseline. The baseline sits toward the

bottom right of the plot: it offers neither accuracy nor diversity. The SVD, on the other hand, tends to be more accurate than the baseline, although there is little diversity gain. The *kNN* results, finally, sit between the two others—in terms of accuracy—and above them when considering diversity.

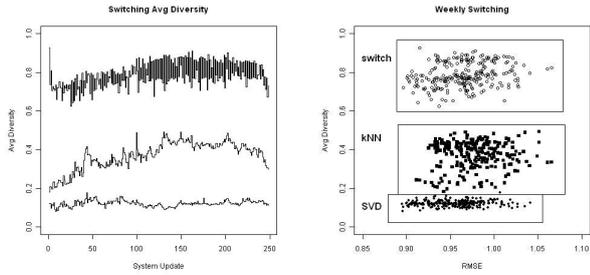
Based on what we have observed, it seems that improving the temporal diversity of a recommender system is an important task for system developers. In the following section, we describe and evaluate a number of techniques that may be implemented to do so. We then discuss the potential implications that modifying top-*N* lists to promote diversity may have.

4. PROMOTING TEMPORAL DIVERSITY

The easiest way of ensuring that recommendations will be diverse is to do away with predicted ratings and simply rank items randomly. However, *diversity* then comes at the cost of *accuracy*: recommendations are no longer personalised to users’ tastes. The random survey (Section 2) showed that this is not a viable option, since the recommendations were rated very low. We can thus anticipate that, when promoting diversity, we must continue to take into account users’ preferences. We do so with two methods: temporal hybrid switching, from a system (a) and user (b) perspective, and re-ranking individual users’ recommendations (c).

(a) Temporal Switching. Many state of the art approaches to CF combine a variety of algorithms in order to bolster prediction accuracy [5]. However, as described by Burke [7], another approach to building hybrid CF algorithms is to *switch* between them. Instead of combining prediction output, a mechanism is defined to select one of them. The rationale behind this approach is as follows: given a set of CF algorithms, that each operate in a different way, are likely to produce *different* recommendations for the same user; the top-*N* produced by a *kNN* may not be the same as that produced by an SVD. We thus switch between the two algorithms: we cycle between giving users *kNN*-based recommendations one week, and SVD-based recommendations the following week.

We plot the top-10 diversity over time for this switching method in Figure 7(a). Diversity has now been incremented to approximately 0.8: on average, 8 of the top-10 recommendations ranked for each user is something that was not recommended the week before. How does this affect accuracy? Intuitively, the overall accuracy that the system will



(a) Switching Temporal Diversity (b) Switching RMSE vs Diversity

Figure 7: Diversity (a) and Accuracy (b) of Temporal Switching Method

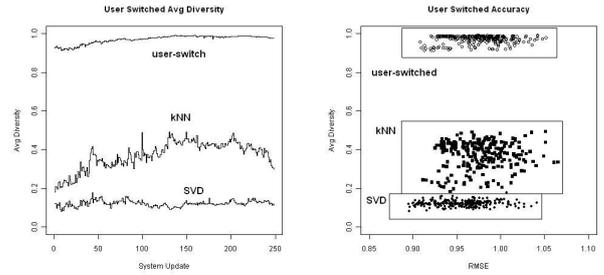
achieve will be somewhere between the accuracy of each individual algorithm. We compare the accuracy and diversity of our switching technique in Figure 7(b). The results for the switching method now cluster into two groups; each group lies *above* the candidate algorithms we selected. In other words, accuracy fluctuates between the values we reported for k NN and SVD CF, but the fact that we are switching between these two techniques ensures that diversity has been greatly increased.

(b) Temporal User-Based Switching. The method described in the previous section is very straightforward: the system changes the CF algorithm that is used from one week to the next in order to favour diversity. However, this method does not take into account how users behave; in particular, we previously noted that not all users have *regular* sessions with the recommender system. In fact, if their sessions were every other week, then the switching technique described in the previous section would be of no use at all. We therefore also tested a user-based switching algorithm. It works as follows: the system keeps track of *when* a user last appeared, and *what* algorithm was used to recommend content to that user during the last session. When the user reappears, the system simply picks a different algorithm from what it previously used. As before, we switched between using an item-based k NN and an SVD-based approach in our experiments. The results are shown in Figure 8.

The temporal diversity (Figure 8(a)) is now near 1: on average, users are being offered different recommendations to what they were shown the last time they interacted with the system. On the other hand, accuracy (Figure 8(b)) now falls between the k NN and SVD results. In other words, we sacrifice the low-RMSE of the SVD, but still do better than simply using the k NN approach: in return, the average diversity has been greatly amplified.

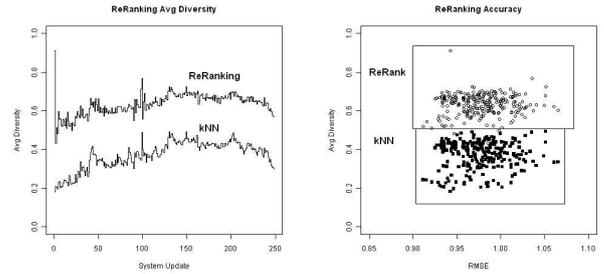
The only overhead imposed by user-based switching is a single value per user that identifies which algorithm was last used to compute recommendations; however, unlike the temporal switching method in the previous section, we are now required to compute both k NN and SVD at every update, albeit for a subset of users. We do not consider this to be an unsurmountable overhead, given that state of the art algorithms already tend to ensemble the results of multiple CF algorithms.

(c) Re-Ranking Frequent Visitors' Lists. An immediate problem with a temporal switching approach is that it requires multiple CF algorithm implementations. In this



(a) Temporal Diversity (b) RMSE vs. Diversity

Figure 8: Temporal Diversity and Accuracy vs. Diversity With User-Based Temporal Switching



(a) Temporal Diversity (b) RMSE vs. Diversity

Figure 9: Temporal Diversity and Accuracy vs. Diversity When Re-Ranking Frequent Visitors' Lists

section, we provide a means of diversifying recommendations to any desired degree of diversity when only a single CF algorithm is used.

One of the observations we made above is that users who have very regular sessions with the recommender system have low top- N temporal diversity. One way of improving overall average temporal diversity thus entails catering to the diversity needs of this group. To do so, we take advantage of the fact that they are *regular* visitors, and only re-rank their top- N recommendations.

The re-ranking works in a very straightforward manner: given a list that we wish to diversify with depth N (e.g., $N = 10$), we select M , with $N < M$ (e.g., $M = 20$). Then, in order to introduce diversity d into the top- N , we replace $(d \times N)$ items in the top- N with randomly selected items from positions $[(N + 1) \dots M]$. In the case of $d = 1$, all elements in the first $[1 \dots N]$ positions are replaced with elements from positions $[(N + 1) \dots M]$. This is the method that we used to diversify the recommendations in the user survey S2 (Section 2); in that case, $N = 10$ and $M = 100$ (the 100 all time box office hits).

In our experiments, we opted to re-rank the top-10 results for any users who had previously visited the system less than two weeks before (recall that our system is updated weekly). The temporal diversity results, shown in Figure 9(a), clearly improve the overall average. Furthermore, the accuracy (Figure 9(b)) remains the same: the diversity has simply been shifted in the positive direction. However, how does this not hurt accuracy? There are three points to keep in mind: (a) we are only reranking the lists for frequent visitors, others' recommendations are untouched; (b) the items

in the top- N are there due to both high prediction value and high confidence (there is a good chance the user will like those items); and (c) we do not promote items that it is unlikely the user would like would not like (by only re-ranking the top- M).

How do these techniques affect recommendation novelty? If we aggregate the temporal results of Figure 3(c), we find that the baseline top-10 recommends, on average, 13.53 ± 2.86 items over time; the SVD top-10 suggests 26.17 ± 12.51 items over time, and the k NN top-10 recommends the highest number of items over time: 79.86 ± 59.33 . This ensures that k NN will also produce the highest number of *new* recommendations. Weekly switching slightly lowers k NN's average, to 75.36 ± 53.98 ; repeatedly visiting the SVD recommendations reduces the number of total items that can be recommended. However, user based switching maintains the average number of recommended items over time at 79.86 ± 55.12 ; it essentially highly promotes temporal diversity without impacting the number of new items that enter the top-10 list over time. However, re-ranking bolsters both the average and standard deviation to 97.93 ± 78.82 ; re-ranking thus seems like a promising approach to solving the related problem of temporal novelty in recommendation.

5. RELATED WORK AND CONCLUSION

Diversity is a theme that extends beyond recommender systems; for example, Radlinski and Dumais examine how it can be used in the context of personalised search [8]. In other cases, diversifying search results is done in order to reduce the risk of query misinterpretation [9]. Similarly, diversity relates to user satisfaction; more specifically, to users' impatience with duplicate results [10]. We have observed similar 'impatience' in our survey: users who completed the survey with no diversity began to rate recommendations lower as they saw that they were not changing.

It is certainly possible to envisage a finer grained notion of diversity that takes semantic data into account—by measuring, for example, the extent that the same genre or category of items are being recommended. To that end, diversity may also be measured within a *single* top- N list, rather than a pair or sequence of recommendations; such a metric may, for example, take into account the number of highly related items (such as a movie and its sequels, or multiple albums by the same artist) that are being simultaneously recommended. For example, Smyth and McClave [11] apply strategies to improve recommender systems based on case-based reasoning; diversity, in this case, is viewed as the complement of similarity. Zhang and Hurley [12] also focus on intra-list diversity, and optimize the trade off between users' preferences and the diversity of the top- N results. In this work, we focus on the temporal dimension (inter-list diversity) and whether the exact same items are being offered to users more than once; we do not take semantic relationships between the recommended items into account nor improve the diversity of individual top- N lists. However, both lines of research are not in conflict: ideally, one would like a recommender system that offers diverse results that *change* over time to suit each users' tastes. All of the work we have done here diversifies recommendations based on rating data: additional data (such where users click, or what reviews they read) may further inform this process.

When investigating how recommendations change over time, we found that state of the art CF algorithms generally pro-

duce low temporal diversity; they repeatedly recommend the same top- N items to users. We then defined a metric to measure temporal diversity, based on the set theoretic difference of two sequential top- N lists, and performed a fine-grained analysis of the factors that may influence diversity. We found that, while users with large profiles suffer from lower diversity, those who rate a lot of content in one session are likely to see very diverse results the next time. We also observed that diversity will naturally improve if a lot of time passes between user sessions. We then designed and evaluated three methods of improving temporal diversity without extensively penalising recommendation accuracy. Two were based on *switching* CF algorithm over time; users are first given recommendations produced with (for example) a k NN approach, and then offered the results of an SVD algorithm. The last method was based on re-ranking the results of frequent visitors to the system.

In future work, we plan on extending the evaluation methodology that we have applied here in order to examine how *novel* items find their way into recommendations, and how user rating patterns can be used to improve recommender system's resilience to attack.

6. REFERENCES

- [1] J. Herlocker, J. Konstan, L. Terveen, and J. Riedl. Evaluating Collaborative Filtering Recommender Systems. In *ACM TOIS*, volume 22, pages 5–53, 2004.
- [2] M. Mull. Characteristics of High-Volume Recommender Systems. In *Proceedings of Recommenders '06*, Bilbao, Spain, September 2006.
- [3] Y. Koren. Collaborative Filtering With Temporal Dynamics. In *ACM KDD*, Paris, France, June 2009.
- [4] N. Lathia, S. Hailes, and L. Capra. Temporal Collaborative Filtering With Adaptive Neighbourhoods. In *ACM SIGIR*, Boston, USA, 2009.
- [5] G. Adomavicius and A. Tuzhilin. Towards the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE TKDE*, 17(6), June 2005.
- [6] S.M. McNee, S.K. Lam, C. Guetzlaff, J.A. Konstan, and J. Riedl. Confidence Displays and Training in Recommender Systems. In *ACM CHI*, 2003.
- [7] R. Burke. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*, 12(4), 2002.
- [8] F. Radlinski and S. Dumais. Improving Personalized Web Search Using Result Diversification. In *ACM SIGIR*, Seattle, USA, 2006.
- [9] R. Agrawal, S. Gollapudi, A. Halverson, and S. Jeong. Diversifying Search Results. In *ACM WSDM*, Barcelona, Spain, 2009.
- [10] D. Hawking, T. Rowlands, and P. Thomas. C-TEST: Supporting Novelty and Diversity in Testfiles for Search Evaluation. In *ACM SIGIR Workshop on Redundancy, Diversity and Interdependent Document Relevance*, Boston, USA, 2009.
- [11] B. Smyth and P. McClave. Similarity vs. Diversity. In *4th International Conference on Case-Based Reasoning*, Vancouver, Canada, 2001.
- [12] M. Zhang and N. Hurley. Avoiding Monotony: Improving the Diversity of Recommendation Lists. In *ACM RecSys*, Lausanne, Switzerland, 2008.