

## Predicting User Engagement in Twitter with Collaborative Ranking

Questa è la versione preprint della seguente opera:

*Original*

Predicting User Engagement in Twitter with Collaborative Ranking / Diaz-Aviles, Ernesto; Thanh Lam, Hoang; Pinelli, Fabio; Braghin, Stefano; Gkoufas, Yiannis; Berlingerio, Michele; Calabrese, Francesco. - (2014), pp. 41-46. ( RecSysChallenge '14: Proceedings of the 2014 Recommender Systems Challenge Foster City, USA October 10, 2014) [10.1145/2668067.2668072].

*Availability:*

This version is available at: 20.500.11771/24179

*Publisher:*

Association for Computing Machinery

*Published*

DOI:10.1145/2668067.2668072

*Terms of use:*

This publication is made accessible in accordance with the terms for deposit in the institutional repository, as defined by the IMT School for Advanced Studies Lucca's Open Access Policy. ([https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib\\_0.pdf](https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib_0.pdf)).

Si prega di consultare le pagine informative dell'editore relative alle politiche di autoarchiviazione.

(Article begins on next page)

# Predicting User Engagement in Twitter with Collaborative Ranking

Ernesto Diaz-Aviles, Hoang Thanh Lam, Fabio Pinelli, Stefano Braghin, Yiannis Gkoufas, Michele Berlingerio, and Francesco Calabrese

IBM Research – Ireland

{e.diaz-aviles, t.l.hoang, fabiopin, stefanob, yiannisg, mberling, fcalabre}@ie.ibm.com

## ABSTRACT

Collaborative Filtering (CF) is a core component of popular web-based services such as Amazon, YouTube, Netflix, and Twitter. Most applications use CF to recommend a small set of items to the user. For instance, YouTube presents to a user a list of top-n videos she would likely watch next based on her rating and viewing history. Current methods of CF evaluation have been focused on assessing the quality of a predicted rating or the ranking performance for top-n recommended items. However, restricting the recommender system evaluation to these two aspects is rather limiting and neglects other dimensions that could better characterize a well-perceived recommendation. In this paper, instead of optimizing rating or top-n recommendation, we focus on the task of predicting which items generate the highest user engagement. In particular, we use Twitter as our testbed and cast the problem as a Collaborative Ranking task where the rich features extracted from the metadata of the tweets help to complement the transaction information limited to user ids, item ids, ratings and timestamps. We learn a scoring function that directly optimizes the user engagement in terms of nDCG@10 on the predicted ranking. Experiments conducted on an extended version of the MovieTweetings dataset, released as part of the RecSys Challenge 2014, show the effectiveness of our approach.

## Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—*Information filtering*; I.2.6 [Artificial Intelligence]: Learning—*Parameter learning*

## General Terms

Algorithms; Experimentation; Measurement.

## Keywords

Recommender Systems; User Engagement; Learning to Rank; Twitter.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*RecSysChallenge'14*, October 10, 2014, Foster City, CA, USA.

Copyright 2014 ACM Copyright is held by the owner/author(s).

Publication rights licensed to ACM. ACM 978-1-4503-3188-3/14/10...\$15.00  
<http://dx.doi.org/10.1145/2668067.2668072>.

## 1. INTRODUCTION

Collaborative Filtering (CF) is an important method that helps users to deal with information overload by making personalized predictions of items of interest based on the preference information of many users. CF is also important to web-based services such as Netflix, YouTube, Amazon, and Twitter, which are able to improve customer satisfaction and increase revenue based on personalized recommendations.

Two fundamental tasks of CF are rating prediction and top-n recommendation. In particular, the latter can be cast as a ranking problem whose goal is to define an ordering among items (ie.g., Web pages, documents, news articles, Web sites, CDs, books, or movies) placing relevant ones in higher positions of the retrieved list. Despite the advances in recent years of supervised learning-to-rank methods [19], with very few exceptions (e.g., [10, 4]) their application to CF has gained relatively little attention. A potential reason is that learning-to-rank approaches usually rely upon pre-engineered features that characterize the items to be ranked. Such a set of features can be difficult to compute or maintain, and kNN methods [27] or matrix factorization methods, both for ranking (e.g., [9, 24]) and regression (e.g., [15, 17]), have been the preferred approaches for CF when only transaction information in the form of user-item pairs is available.

The core contribution of this work is a collaborative ranking approach for user engagement prediction in Twitter. We make use of the rich Twitter metadata to extract features that characterize user-item-tweet interactions. Such features help us to create a training dataset amenable to a learning-to-rank task, which allows us to leverage one of the many methods capable to directly optimize Information Retrieval (IR) metrics such as nDCG@10.

## Collaborative Ranking Framework

We cast the problem of user engagement prediction in Twitter as a Learning to Rank task [19]. In learning (training), a collection of users and their corresponding tweets are given. Furthermore, the *engagement* (i.e., relevance judgments or labels) of the tweet with respect to the user are also provided. The engagement for a user-item-tweet triple,  $(u, i, d)$ , is defined as the expected user interaction, which is expressed in terms of sum of retweets and favorite counts:

$$\text{engagement}(u, i, d) = \text{retweets}(u, i, d) + \text{favorites}(u, i, d) .$$

Based on the value of engagement it is possible to represent ranks of tweets per user. The objective of learning is to construct a ranking model, e.g., a ranking function, that achieves the best result on test data in the sense of optimization of a performance measure, e.g., nDCG@10 [3].

In the recommendation phase (test), given a unseen user, the learned ranking function is applied, returning a ranked list of tweets in descending order of their relevance scores. Suppose that  $U = \{u_1, \dots, u_{|U|}\}$  is the set of users. Users express their preferences over items using tweets, let  $I$  be the set of such items, e.g., the user can comment about a movie or video or give an explicit rating in a tweet, in this case the movie item  $i$  belong to set  $I$ . Furthermore, let  $D = \{d_1, \dots, d_{|D|}\}$  be the set of tweets, then the training set is created as a set of user-item-tweet triples,  $(u, i, d) \in U \times I \times D$ , upon which a relevance judgment (e.g., a label) indicating the user engagement for  $u$  over item  $i$  expressed in tweet  $d$  is associated to the triple.

Suppose that  $Y = \{y_1, \dots, y_{|Y|}\}$  is the set of labels and  $y_{uid} \in Y$  denotes the label of user-item-document triple. A feature vector  $\phi(u, i, d)$  is created from each triple  $(u, i, d)$ . The training set is denoted by  $S = \{\phi(u, i, d)_k, y_{uid_k}\}_{k=1}^m$ , where  $m$  is the number of training examples, i.e., the number of observed triples. The ranking model is a real valued function of features:

$$f(u, i, d, \Theta) = h(\phi(u, i, d); \theta) + b_{uid} \quad (1)$$

where  $\Theta = \{\theta, b_{uid}\}$  is the set of free parameters to be learned. Here  $h(\phi(u, i, d); \theta)$  is a scoring function on the user-item-tweet feature vectors  $\phi(u, i, d)$ , and  $b_{uid}$  is a bias term.

In ranking for user  $u$ , the model associates a personalized score to each of the tweets  $d$  as their degree of relevance with respect to the user  $u$  using  $f(u, i, d, \Theta)$ , and sorts the tweets based on their scores, e.g., in decreasing order.

## 2. OUR APPROACH

The main idea behind our approach is to transform the user engagement prediction problem into a collaborative ranking one and then learn the ranking function using a learning-to-rank method [19]. The collaborative ranking task can be placed into the learning-to-rank framework by noting that the users correspond to queries and tweets to documents. For each user the observed engagement indicates the relevance of the corresponding tweets with respect to that user and can be used to train the ranking function. The basic steps of our approach are summarized in Procedure 1.

---

### Procedure 1 Collaborative Ranking for User Engagement

---

**Input:** Training set  $S = \{\phi(u, i, d)_k, y_{uid_k}\}_{k=1}^m$

**Output:** Ranking function  $f(u, i, d, \Theta)$

- 1: Scale and normalize feature vectors  $\phi(u, i, d)$
- 2: Remove user outliers
- 3: Learn a ranking function

$$f(u, d, i, \Theta) = h(\phi(u, i, d); \theta) + b_{uid}$$

by optimizing nDCG@10 directly

- 4: **return**  $f(u, i, d, \Theta)$
- 

In the rest of the section, we present our feature extraction approach followed by the details on how the ranking function is learned.

### Feature Extraction

We propose to use the user-item-tweet interactions and the corresponding metadata of the tweets as main source to extract the feature vectors. The dataset used in this work corresponds to the *MovieTweatings Extended*, released as part of the RecSys Challenge [25, 20], 2014. The dataset is an extended version of the MovieTweatings dataset [12]. The data includes the ratings extracted from users of the Internet Movie Database (IMDb) iOS app<sup>1</sup> that rate movies

<sup>1</sup>IMDb iOS app: <http://www.imdb.com/apps/ios/>

and share the rating on Twitter. MovieTweatings dataset is built by extracting information of the tweets on a daily basis by using Twitter API [12]. The original MovieTweatings contains only rating information, i.e., (user, item, rating, timestamps) tuples. The extended version released for the challenge also includes all metadata of tweets as provided by Twitter API.<sup>2</sup>

We extract the following features  $F$  to form the  $\phi(u, i, d)$  feature vectors for each user-item-tweet triple.

F1: User rating,  $F1 = r_{uid}$ , given to the movie item  $i$  as expressed in tweet  $d$ .

F2: Deviation of user rating with respect to the median of previous user ratings, i.e.,  $F2 = r_{uid} - \tilde{r}_u$

F3: Average user engagement from her history, i.e.,

$$F3 = \overline{engagement}(u)^{0.5}$$

F4: Boolean indicator that takes the value of 1 if the average user engagement is greater than 0, and 0 otherwise, i.e.,

$$F4 = \mathbb{1}[\overline{engagement}(u) > 0],$$

where  $\mathbb{1}[a]$  is one if  $a$  is true.

F5: Average rating per user,  $F5 = \bar{r}_u$

F6: Ratio of number of user friends to the number of her followers, i.e.,

$$F6 = \left( \frac{\#friends(u)}{\#followers(u)} \right)^{0.5}$$

F7: User tweet count:

$$F7 = \#tweets(u)^{0.5}$$

F8: Average user engagement for a given movie item  $i$ , i.e.,

$$F8 = \overline{engagement}(i)^{0.5}$$

F9: Boolean indicator that takes the value of 1 if the average user engagement for item  $i$  is greater than 0, and 0 otherwise, i.e.,

$$F9 = \mathbb{1}[\overline{engagement}(i) > 0]$$

F10: Average rating per item,  $F10 = \bar{r}_i$

F11: Average ratio of number of user friends to the number of her followers aggregated over the movie item  $i$ :

$$F11 = \left( \frac{1}{K} \sum_{u \in U_i} \frac{\#friends(u)}{\#followers(u)} \right)^{0.5}$$

where  $U_i := \{u \in U | (u, j, d) \in S \wedge j = i\}$  and  $K = |U_i|$ .

F12: Average of user tweet counts aggregated per item  $i$ :

$$F12 = \left( \frac{1}{K} \sum_{u \in U_i} \#tweets(u) \right)^{0.5}$$

F13: User mentions: boolean indicator that takes the value of 1 if the tweet contains any mention, i.e., any Twitter update that contains “@username” anywhere in the body of the Tweet, and 0 otherwise, i.e.,

$$F13 = \mathbb{1}[has\_mention(d)]$$

<sup>2</sup>Twitter API: <https://dev.twitter.com/>

F14: Tweet is a retweet. From the tweet metadata field `retweeted_status` we can infer if the current tweet  $d$  is actually a retweet and use a boolean indicator to capture this information:

$$F14 = \mathbb{1}[is\_retweet(d)]$$

F15: The same field `retweeted_status` for  $d$  also includes the tweet id (`tweet_id(d_o)`) of the original tweet, if such tweet  $d_o$  is present in the dataset we know that it received a non-zero engagement,  $F15$  represents this additional information for  $d_o$ :

$$F15_{\phi(u,i,d_o)} = \mathbb{1}[is\_retweet\_of(d, d_o)] \quad ,$$

where `is_retweet_of(d, d_o)` is true if  $d$  is  $d_o$ 's retweet.

F16: The frequency of observed engagement (i.e., retweet count) extracted per item from the `retweeted_status` field:

$$F16 = \sum_{d \in is\_retweet(D)} \mathbb{1}[engagement(i)] \quad .$$

## Learning the Ranking Function

Given the user-item-tweet features  $\phi(u, i, d)$  extracted based on user-item-tweet interactions and the corresponding metadata of the tweets our goal is to use the observed engagement for each triple to optimize the parameters of the ranking function for the target information retrieval metric, in our case, nDCG@10. In this work we choose to use an ensemble of MART and LambdaMART learning-to-rank methods [5]. In particular, LambdaMART was chosen due to its the ability to directly optimize the Information Retrieval (IR) metric nDCG@10 and the performance exhibited in recent ranking competitions [6, 7].

LambdaMART uses as base learners Multiple Additive Regression Trees or MART, also known as Gradient Boosted Regression Trees (GBRT) [13, 14]. LambdaMART is a pairwise learning-to-rank method capable to optimize arbitrary IR metrics by guiding the learning process using so-called  $\lambda$ -gradients, which reflect small changes in the IR metric while iterating over the training set. We omit a detailed description of the model, since it is out of the scope of this report, and refer the reader to [5] and [6] for a comprehensive analysis of LambdaMART.

## 3. EXPERIMENTS

In this section, we first present some key statistics of the dataset used in the experiments, followed by the experimental evaluation of our approach.

### Dataset

The dataset used in our experiments corresponds to the *MovieTweetings Extended*, released as part of the RecSys Challenge [12, 25, 20], 2014. For this challenge, the dataset is chronologically split in three subsets: training, test, and evaluation set, whose statistics are summarized in Table 1. Note that the number of *unique* users and items (considering the ones present in training and test) are  $|U| = 23,555$  and  $|I| = 14,316$ , respectively.

Figure 1 shows a boxplot of the distribution of the number of interactions (i.e., user-item-tweet triples) per user in the training set. Note that the number of interactions is very low for most users, leading to a median of 2. We can also observe many outliers, e.g., users with over 300 interactions, far from them mean of 7.71.

The training set consists of the first 80% tweets and is used as input for our collaborative ranking approach. The test and evaluation sets both contain 10% of the remaining tweets. The test split is

Split	Users	Items	Tweets	Tweets creation
Training	22,079	13,618	170,285	2013-02-28 14:43 - 2014-01-08 22:06
Test	5,717	4,226	21,285	2014-01-08 22:06 - 2014-02-11 15:49
Evaluation	5,514	4,559	21,287	2014-02-11 15:49 - 2014-03-24 9:57
Total	24,924	15,142	212,857	2013-02-28 14:43 - 2014-03-24 9:57

Table 1: MovieTweetings Extended dataset statistics.




Figure 1: Boxplot of the distribution of the number of interactions (i.e., user-item-tweet triples) per user in the training set. There is a minimum number of 1 and a maximum of 514 interactions per user, with mean value of 7.71 and a median of 2.

used to measure the predictive performance of our method, and the final evaluation split is reserved by the challenge organizers for the final assessment.

## Evaluation Methodology

### Metric

The evaluation for the task of predicting which user-item-tweet triples generate the highest user engagement is measure by the IR metric Normalized discounted cumulative gain with a cut level equal to 10 (nDCG@10), nDCG measures the performance of a recommendation system based on the graded relevance of the ranked tweets based on their predicted engagement. It varies from 0.0 to 1.0, with 1.0 representing the ideal ranking of the tweets. This metric is commonly used in information retrieval and to evaluate the performance of web search engines [3].

Formally, the discounted cumulative gain (DCG) is a measure of ranking quality defined as:

$$DCG@K = \sum_{k=1}^K \frac{2^{rel_k} - 1}{\log_2(k + 1)} \quad (2)$$

if we define  $IDCG@K$  as the maximum possible (ideal) DCG for a given set of user-item-tweet triples and engagement, we can normalize the DCG@K to be between 0.0 to 1.0 and obtain the nDCG@K :

$$nDCG@K = \frac{DCG@K}{IDCG@K} \quad (3)$$

where  $K$  is the maximum number of tweets that are considered in the ranked list for the evaluation.

For each user in the test set we compute the corresponding nDCG@10 and then average over all users. For the evaluation, we use the tool provided by the organizers, which is based on RiVal, a toolkit for recommender systems evaluation.<sup>3</sup> All results reported in this section corresponds to nDCG@10 obtain in the test split only.

<sup>3</sup>RiVal: <https://github.com/recommenders/rival>

## Baselines

We compare the performance of our collaborative ranking approach for user engagement prediction, or *CRUE*, to the baselines recommender systems described next.

- **FM** a Factorization Machine [23]. We use the FM reference implementation provided by libFM<sup>4</sup> and train the model for a ranking task using Markov Chain Monte Carlo (MCMC), and found that a standard deviation of 0.1, 16 factors, and 1000 iterations, lead to good results. We include the indicator variables corresponding to users, items, and rating, and the engagement as label, which is equivalent to a pairwise tensor factorization for ranking [23].
- **recRating** predicts the rating associated to the triple user-item-tweet as the user engagement value.
- **recHEI** predicts as future engagement the average value of the user engagement that a given item received in the past.
- **recRandom** predicts a random value for the user engagement.

Observe that recRating and recHEI correspond to features F1 and F8, respectively, as defined in Section 2.

### CRUE Setup

As discussed in Section 2, CRUE uses a linear ensemble of MART and LambdaMART to learn the ranking function for user engagement prediction. LambdaMART uses as base learners Multiple Additive Regression Trees or MART, also known as Gradient Boosted Regression Trees (GBRT) [13, 14], whose main parameters to adjust is the number of estimators, the number of trees in the forest, and the size of the random subsets of features to consider when splitting a node.

In order to reduce noise in the training dataset (cf., Figure 1), we drop the user interactions of users with less than 4 and over 200 interactions, and use the renaming users for training. We observe more stable results after this pruning.

We set the number of trees using an additional 80%-20% partition of the training set, keeping the 20% as a validation split to determine the optimal number of trees. Based on the same validation procedure, we found that a number of leaves for each tree equal to 10 and a learning rate (or shrinkage) equal to 0.1 gave us good results.

We perform early stopping in the training when we observed that no additional improvement, in terms of nDCG@10, was achieved in the validation split after 50 consecutive rounds. We set the maximum number of features for splitting equal to the number of features in the training set.

To implement CRUE we use a set of Python scripts that rely upon NumPy, SciPy<sup>5</sup> and scikit-learn [22]. We use RankLib’s implementation of LambdaMART to learn the ranking function and directly optimize our goal is to reduce noise and wize for nDCG@10 [18].

## Results

Table 2 summarizes the best results obtained by each method in our experiments. We found that the explicit feedback expressed as rating is quite powerful predictor for user engagement, this is the only source of information for the simple recommender recRating that achieves a nDCG@10 of 0.8182, outperforming the factorization model FM. Figure 2 shows a scatter plot depicting the global

<sup>4</sup>libFM: [www.libfm.org](http://www.libfm.org)

<sup>5</sup>NumPy and SciPy: <http://www.scipy.org/>




Figure 2: Scatter plot in (rating, engagement) coordinates, where the area of the circle depicts the frequency observed for the corresponding pair, i.e., the larger the radius, the higher the frequency.

Recommender	nDCG@10
CRUE (our approach)	<b>0.8701</b>
FM	0.8023
recRating	0.8182
recHEI	0.8031
recRandom	0.7532

Table 2: Results.

rating-engagement relationship in the training set, we can observe that higher ratings tend to receive higher engagement.

The historical user engagement that an item received in the past is also a very good predictor for the task at hand, recHEI also outperforms the FM, but is not as good as recRating.

We also experiment training a FM with additional contextual features, e.g., the ones extracted from tweet metadata, as we did for CRUE, but the performance obtained in our experiments was below to the one of using the user-item-rating triples alone. We are still engineering additional features to train a new FM in order to include the model in a potential ensemble together with CRUE.

CRUE outperforms the baselines considered here and, at the time of writing, our approach places our team (SUD) within the top-5 in the RecSys Challenge leaderboard [26].

### Some general thoughts.

- During the competition we concentrated first in directly optimizing the IR metric used in the challenge, namely nDCG@10. The results of our collaborative ranking approach, CRUE, are the ones reported in this paper. Our goal towards the end of the challenge is to continue trying different models, which can explain the variability of the data.
- We found, in initial experiments, that factorization models that use the engagement as label, either for regression or classification, do not perform as good as the rating (recRating) or the item average engagement (recHEI).
- We observe that smoothing of the features, given the long tail distribution of most of them, works well. We smooth them using square root of the values, as shown in Section 2.

- Normalization of the features using their z-scores also helps us to significantly improve ranking performance.
- Removing user outliers also helps.
- An initial exploration on the individual nDCG@10 performance per user indicates that for some users are rather *easy* to predict the engagement that their tweets will receive. Training models on *hard* users only, e.g., those whose nDCG@10 is less than a threshold, is equally effective than training using all of the user available in the training set. This idea goes towards selective sampling and active learning, which have been successfully applied in the context of CF (e.g., [9]).
- In our experiments, we found that traditional kNN and Matrix factorization models for CF did not achieve a performance better than using the rating as predictor for user engagement. We are still trying to improve the performance of these approaches to include them in an ensemble together with CRUE.
- We also collected complementary information from IMDb in order to enrich the movie profiles, e.g., we collected the year of release of the movie, the description, tags, as well as the director and actors. We found that in some cases the year of release improve marginally our predictions. Note that the results reported in this paper are based only on the dataset provided.
- It would be very interesting to consider a different evaluation protocol that better reflects the user impact in the engagement that her tweets will receive when compared to other users as well. In other words, assess the user engagement that triples user-item-tweet from user  $u$  will receive over other user  $v$ 's triples.
- In independent experiments, we have identified that the time of the retweets can be very useful to predict engagement at different levels, e.g., userwise or topicwise. Having access to this additional feature could also help us boost the results.
- If an application requires the user engagement to be predicted on-the-fly, it can leverage the explicit feedback (rating) and the historic user engagement per item for an effective and fast prediction.

#### 4. RELATED WORK

The popularity of Twitter, easy access to data and the unique characteristic of velocity and real-time information propagation, have made it an attractive domain recently. For instance, in [21] the authors make an effort on characterizing retweeters and their retweeting behavior, which is related to our goal of predicting user engagement. However, existing work does not consider a collaborative ranking approach to this task, as the one introduced in this paper.

Our approach is also related to context-aware recommender systems [2], which use additional contextual information such as time, location, or the company of other people to improve recommendation performance. Our approach leverages features extracted from the rich metadata available for the tweets and follow a collaborative ranking approach, which is different to the state-of-the-art approaches for context aware recommender systems that primarily rely on factorization methods [23, 16].

Existing collaborative ranking approaches include the ones proposed in [10] and [4], these methods also cast the recommendation

problem as a learning-to-rank task, but in the absence of rich features, they use factors output by a matrix factorization algorithm for CF as feature vectors, and then learn a ranking function. We instead use the rich metadata from tweets to build feature vectors that characterize the user-item-tweet interactions. As future work, we plan to also include the latent factors as part of our existing feature vectors to asses to what extent we can improve the personalized ranking. Another interesting collaborative filtering approach is the one introduced in [28], which constructs the feature vectors based on metrics derived from the user neighborhoods.

None of the existing approaches have explored the collaborative ranking setting for user engagement prediction as we do in this work.

#### 5. CONCLUSION AND FUTURE WORK

In this work we showed how our collaborative ranking approach is able to predict user engagement in Twitter. The rich metadata available from Twitter API help us to build feature vectors that characterize the user-item-tweet interactions, and complements the information captured by the explicit feedback (i.e., rating) expressed in the tweet.

Experimental results show that by using our collaborative ranking approach is possible to predict attractive tweets that will engage more people. Such insights can be then leveraged, for example, to make interesting predictions for business success [1] or public safety [11].

We are currently investigating additional features such as characteristics of the user and item social neighborhoods to improve ranking performance. An additional direction for future work is the extension of our model for streaming data scenarios, e.g., [9], where the model parameters are learned online without compromising ranking performance.

#### 6. REFERENCES

- [1] F. Abel, E. Diaz-Aviles, N. Henze, D. Krause, and P. Siehdnel. Analyzing the blogosphere for predicting the success of music and movie products. *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 0:276–280, 2010.
- [2] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 217–253. Springer US, 2011.
- [3] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, 2nd edition, 2011.
- [4] S. Balakrishnan and S. Chopra. Collaborative ranking. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*, WSDM '12, pages 143–152, New York, NY, USA, 2012. ACM.
- [5] C. J. Burges. From ranknet to lambdarank to lambdamart: An overview. Technical Report MSR-TR-2010-82, June 2010.
- [6] C. J. C. Burges, K. M. Svore, P. N. Bennett, A. Pastusiak, and Q. Wu. Learning to rank using an ensemble of lambda-gradient models. In Chapelle et al. [8], pages 25–35.
- [7] O. Chapelle and Y. Chang. Yahoo! learning to rank challenge overview. In Chapelle et al. [8], pages 1–24.
- [8] O. Chapelle, Y. Chang, and T.-Y. Liu, editors. *Proceedings of the Yahoo! Learning to Rank Challenge, held at ICML 2010, Haifa, Israel, June 25, 2010*, volume 14 of *JMLR Proceedings*. JMLR.org, 2011.
- [9] E. Diaz-Aviles, L. Drummond, L. Schmidt-Thieme, and W. Nejdl. Real-time top-n recommendation in social streams.

- In *Proceedings of the Sixth ACM Conference on Recommender Systems*, RecSys '12, pages 59–66, New York, NY, USA, 2012. ACM.
- [10] E. Diaz-Aviles, M. Georgescu, and W. Nejdl. Swarming to rank for recommender systems. In *Proceedings of the Sixth ACM Conference on Recommender Systems*, RecSys '12, pages 229–232, New York, NY, USA, 2012. ACM.
- [11] E. Diaz-Aviles, A. Stewart, E. Velasco, K. Denecke, and W. Nejdl. Epidemic intelligence for the crowd, by the crowd. 2012.
- [12] S. Doods, T. De Pessemier, and L. Martens. Movietweetings: a movie rating dataset collected from twitter. In *Workshop on Crowdsourcing and Human Computation for Recommender Systems*, CrowdRec at RecSys 2013, 2013.
- [13] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 10 2001.
- [14] J. H. Friedman. Stochastic gradient boosting. *Comput. Stat. Data Anal.*, 38(4):367–378, Feb. 2002.
- [15] S. Funk. Netflix Update: Try This at Home. <http://sifter.org/~simon/journal/20061211.html>, 2006. [Online; accessed 2014-08].
- [16] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the Fourth ACM Conference on Recommender Systems*, RecSys '10, pages 79–86, New York, NY, USA, 2010. ACM.
- [17] Y. Koren. Factorization meets the neighborhood: A multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 426–434, New York, NY, USA, 2008. ACM.
- [18] Lemur Project. Ranklib. <http://www.lemurproject.org/>, 2014. [Online; accessed 2014-08].
- [19] T.-Y. Liu. *Learning to Rank for Information Retrieval*. Springer, 2011.
- [20] D. Loiacono, A. Lommatzsch, and R. Turrin. Recsys challenge 2014: Learning to rank. 2014.
- [21] S. Macskassy and M. Michelson. Why do people retweet? anti-homophily wins the day! 2011.
- [22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [23] S. Rendle. Factorization machines with libFM. *ACM Trans. Intell. Syst. Technol.*, 3(3):57:1–57:22, May 2012.
- [24] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI '09, pages 452–461, Arlington, Virginia, United States, 2009. AUAI Press.
- [25] A. Said, S. Doods, B. Loni, and D. Tikk. Recommender systems challenge 2014. In *Proceedings of the eighth ACM conference on Recommender systems*, RecSys '14, New York, NY, USA, 2014. ACM.
- [26] A. Said, S. Doods, B. Loni, and D. Tikk. RecSys Challenge 2014: User Engagement as Evaluation. <http://2014.recsyschallenge.com/>, 2014. [Online; accessed 2014-08].
- [27] 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, New York, NY, USA, 2001. ACM.
- [28] M. Volkovs and R. S. Zemel. Collaborative ranking with 17 parameters. In F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2294–2302. Curran Associates, Inc., 2012.