

Taxonomy-Informed Latent Factor Models for Implicit Feedback

Andriy Mnih

*Gatsby Computational Neuroscience Unit
University College London*

AMNIH@GATSBY.UCL.AC.UK

Editor: G. Dror, Y. Koren and M. Weimer

Abstract

We describe a latent-factor-model-based approach to the Track 2 task of KDD Cup 2011, which required learning to discriminate between highly rated and unrated items from a large dataset of music ratings. We take the pairwise ranking route, training our models to rank the highly rated items above the unrated items that are sampled from the same distribution. Using the item relationship information from the provided taxonomy to constrain item representations results in improved predictive performance. Providing the model with features summarizing the user’s rating history as it relates to the item being ranked leads to further gains, producing the best single model result on Track 2.

Keywords: Collaborative Filtering, Implicit Feedback, Machine Learning, Ranking

1. Introduction

Collaborative filtering deals with inferring complete user preferences from large but incomplete collections of preference expressions (Marlin, 2004). Users can express their preferences using explicit feedback, such as item ratings, or implicit feedback, such as item views or purchases (Hu et al., 2008). Explicit feedback is more reliable than implicit feedback but is typically available in much smaller quantities because it requires an effort on the part of the user to explicitly express their preferences through special actions, such as rating items. Implicit feedback, on the other hand, is much easier to obtain in large quantities as it is a byproduct of the natural use of the system. This wide availability makes modelling implicit feedback an important task which until recently has received surprisingly little attention.

What makes implicit feedback fundamentally different from two-valued explicit feedback (e.g. like/dislike) is that users provide only positive examples, that is just the items they are interested in (Hu et al., 2008). In other words, though users can express their interest in an item by selecting it, they have no direct way to express their disinterest. Note that not having expressed interest in an item is not the same as having expressed disinterest. After all, the set of items a user has expressed interest in is almost certainly incomplete for a number of reasons, such as budgetary and time constraints as well as incomplete knowledge of the inventory. As a result, the set of items not selected by a user typically contains both interesting and uninteresting items. It is this uncertainty about the unobserved items that makes implicit feedback more difficult to model than its explicit counterpart.

The simplest way to deal with this uncertainty is to disregard it and treat all items not selected by the user as negative examples. If we view the dataset as a binary user-item

matrix with 1s indicating the items selected by users, this strategy amounts to replacing the unobserved entries with 0s. Finding a low-rank approximation to the resulting dense matrix (Hu et al., 2008; Pan et al., 2008; Pan and Scholz, 2009) is the most popular approach to modelling implicit feedback. Such models are typically trained using alternating least squares in time linear in the number of positive examples but cubic in the rank of the approximation matrix, which can be interpreted as the number of real-valued latent factors describing each user or item.

A recently introduced approach, called Bayesian Personalized Ranking (BPR), makes more realistic assumptions about the unobserved items (Rendle et al., 2009) by treating them as merely less interesting than the observed items, as opposed to completely uninteresting. The model is trained to rank the positive items above the negative ones in positive-negative item pairs, with the negative items sampled from the uniform distribution. Both users and items are represented using latent factor vectors which are learned using stochastic gradient ascent, ensuring excellent scalability. In this paper we adapt BPR to the Track 2 task of KDD Cup 2011, modifying the training algorithm to use a popularity-based distribution of negative examples and extending the model to take advantage of the item taxonomy.

2. Task Description

The Track 2 task deals with music recommendation using a dataset of item ratings collected from the Yahoo! Music service. The dataset consists of ratings assigned to music items of four types (tracks, albums, artists, and genres) by about 249K users of the site. There are about 64M ratings on a scale from 0 to 100 for over 296K items, about 224K of which are tracks. See Dror et al. (2011) for a more detailed description of the data.

The task is to learn to discriminate between the highly rated (ratings 80 and above) and the unrated tracks for each user. The test set contains six tracks per user, three of which are highly rated and three are unrated. Note that because this rated/unrated breakdown is known, it is possible to discriminate between items in the test set even with a model that ranks items instead of classifying them.

This task is most naturally viewed as an implicit feedback modelling problem because the goal is to learn to discriminate between the observed positive examples (the highly rated tracks) and the unobserved negative examples (the unrated tracks). However, there are two important ways in which the task differs from the classic implicit feedback problem. First, models are evaluated at discriminating between positive items and negative items that come from the same distribution. This makes it impossible to use item popularity for discrimination between positive and negative examples, which is not the case in the classic setting where negative examples are effectively uniformly distributed. Second, in addition to the identities of the positive items, which are also available in the classic setting, we have access to the item taxonomy and rating values. We will show how to take advantage of these additional sources of information to improve model performance.

3. Basic Latent Factor Model

We represent both users and items using D -dimensional real-valued vectors, which will be learned automatically from the data. We will refer to entries in these vectors as *latent factors* to emphasize that they are learned. Though latent factors are sometimes called *features* in the collaborative filtering literature (e.g. [Salakhutdinov and Mnih \(2008\)](#)), we will use the term *features* to refer only to the precomputed quantities given to the learning algorithm as (fixed) inputs. For example, the number of items rated can be one of the features associated with a user. Though such hand-engineered features can often be used to boost system performance, designing effective features is a difficult and time-consuming task, which is why our approach relies mostly on learned representations.

3.1. Scoring Function

Our approach to the problem is based on learning a scoring function for each user to capture their preferences. Since the task of interest is discrimination between positive and negative items, we would like to learn functions that assign higher scores to positive items than to negative ones. Once the scoring function for a user has been learned, we can rank any set of items by ordering them according to their scores.

We restrict our attention to functions of the form

$$s_u(i) = U_u^\top V_i + b_i, \quad (1)$$

where U_u and V_i are the latent factor vectors for user u and item i respectively, and b_i is the item-specific bias term that models the popularity of the item. We do not include a user-specific bias term because it would not affect the induced ranking. In spite of their simple parametric form, such scoring functions are very expressive because the factor vectors they operate on are learned from the data and can be as high-dimensional as necessary.

In order to learn scoring functions we need to introduce an objective function that links item scores to observations. Unfortunately, the classification error rate, which is the official Track 2 performance metric, is discontinuous which makes optimizing it difficult. We take the probabilistic approach instead and use the (penalized) log-likelihood as the objective function, which is much easier to optimize because it is smooth.

3.2. Classification and Ranking Objective Functions

Our first model classifies items for each user into positive and negative by computing the probability that the item is positive based on its score. The probability of an item being positive is obtained by applying the logistic function to its score:

$$P(i \in \mathcal{P}_u) = \frac{1}{1 + \exp(-s_u(i))}, \quad (2)$$

where \mathcal{P}_u is the set of positive items for user u . This model can be seen as a logistic regression classifier with the item factor vector serving as the input and the user factor vector as the weights.

While this model is appealing because of its simplicity, it is trying to solve a harder problem than necessary by estimating the probability of an item being positive. After all,

in collaborative filtering we are typically interested in ranking items based on some criterion, such as potential interest to the user, not classifying them. Since ranking is in a sense an easier task than classification because there is no need to decide on a score threshold for separating the positive and the negative items, we also tried learning scoring functions by training a probabilistic ranking model. We take the pairwise ranking approach followed by BPR (Rendle et al., 2009) and model the probability that user u prefers item i to item j as the logistic function of the difference between the corresponding item scores:

$$P(i >_u j) = \frac{1}{1 + \exp(-(s_u(i) - s_u(j)))}. \quad (3)$$

The model is trained on positive item i / negative item j pairs. We describe how the negative items are generated in Sec. 4.3.

Predictions on the test set are made by ranking each user’s items based on their scores and labelling the three top-scoring ones as positive and the remaining three as negative.

4. Incorporating Taxonomy Information

One distinguishing feature of the Yahoo! Music dataset is that items come in several types: tracks, albums, artists, and genres. We are also given a taxonomy that links items of different types, providing the album, the artist, and the list of genres for each track, as well as the artist and the list of genres for each album.

We used taxonomy information to parameterize track factor vectors in a more sophisticated way and to generate features that relate items to the user’s rating history in a more direct manner than latent factors allow.

4.1. Taxonomy-Based Parameterization

Learning latent factor models that generalize well can be difficult because of the large number of free parameters that have to be estimated from the data. For example, a model from Sec. 3 with 100-dimensional latent vectors has over 54M parameters, which is of the same order as the number of positive examples in the training set. If we know that some of the parameters should take on similar values, we can improve generalization by incorporating this information into the model. We apply this principle to encourage tracks from the same album or by the same artist to have similar representations by parameterizing latent factor vectors in a hierarchical manner:

$$V_i^H = V_i + w_{al}^r V_{album(i)} + w_{ar}^r V_{artist(i)}, \quad (4)$$

where $album(i)/artist(i)$ refer to the album/artist for track i .¹ In other words, a track is now represented by a linear combination of its old unconstrained representation V_i and the representations of its album and artist. The contributions of the album and artist representations are modulated by w_{al}^r and w_{ar}^r , which are learned along with all other parameters. We also define new track biases in an analogous manner:

$$b_i^H = b_i + w_{al}^b b_{album(i)} + w_{ar}^b b_{artist(i)}, \quad (5)$$

1. If the album or artist for a track is unknown we drop the corresponding term from Eq. 4.

and use V_i^H and b_i^H in place of V_i and b_i in Eq. 1.

This kind of hierarchical parameterization encourages the latent vectors (and biases) of tracks from the same album and/or by the same artist to have similar representations. It also encourages track representations to be similar to those of its album and artist. As a result, updating the latent vector V_i^H for a track will affect not only the track’s representation but also the representations of its album and artist, as well as the representations of all other tracks sharing the album or the artist. Similarly, updating a representation for an album will also affect representations for all tracks off that album. Thus this kind of parameterization allows pooling of information among all types of items, except for genres. We also experimented with including a contribution from the genre representations in Eq. 4 but found that it did not improve model performance.

4.2. Taxonomy-based Features

There is another, more direct, way to use the item relationship information provided by the taxonomy. It involves identifying the items related to the current track (i.e. its album, artist, and tracks sharing the album or the artist) and then computing a set of features that summarize the user’s rating history for these items. The resulting feature vector $f(u, i)$ is then used as an additional input to the scoring function:

$$s_u(i) = U_u^\top V_i + b_i + (W_u + W_i + W)^\top f(u, i). \quad (6)$$

Here W_u , W_i , and W are the user-specific, the item-specific, and the global feature weights respectively, which are learned jointly with all other model parameters.

Though this approach is conceptually simple, its effectiveness is heavily dependent on the choice of features. This sort of feature-based approach has been very successful in natural language processing (Manning and Schütze, 1999) and computer vision (Lowe, 1999). Designing useful features however requires a reasonably good understanding of the problem and extensive experimentation.

We experimented with a number of features and found that while many of them reduced the training error, only a few of them helped on the validation set. We also observed that some features helped models with low-dimensional latent factor vectors considerably, but either hurt or had no effect on models with higher-dimensional factor vectors.

Our feature-based models used (some) of the following features:

1. The rating given to the track’s album by the user (divided by 100)
2. Is the album’s rating 80 or higher?
3. The rating given to the track’s artist by the user (divided by 100)
4. Is the artist’s rating 80 or higher?
5. The fraction of the track’s genres rated by the user (either directly or by rating an album or a track of that genre)
6. Has the user rated (directly or indirectly) any of the track’s genres?
7. The fraction of user’s genre ratings the track’s genres account for

8. Has the user rated any other track from the album?
9. Has the user rated any other track by this artist?
10. The mean rating of the track’s genres that have been rated by the user (divided by 100)

In general, we found that only the first four of these features improved model performance in all cases. As we show in Sec. 6, using too many features can easily lead to overfitting in models with a large number of latent factors.

4.3. Training

The models we described need access to both positive and negative examples for training. We used the user-item pairs from the training set as positive examples. For each user/positive-item pair we generated a negative item by sampling from the empirical distribution of items in the training set. We did not ensure that the sampled negative item had not been rated by the user because we found that the required bookkeeping increased the running time considerably without having much of an effect on the predictive performance.

Since the models are evaluated based on their ability to rank three positive items above three negative items, optimal performance is achieved when each of the positive items is assigned a score higher than the scores of the three randomly-sampled negative items. We encouraged our models to rank positive items above the highest-scoring of the three randomly-chosen negative items by sampling three negative item candidates and keeping the one with the highest score under the model.

We trained all our models using stochastic gradient ascent. The algorithm iterated through the user/positive-item pairs in the training set in a random order, generating a negative item for each such pair before updating the model parameters. Training a model with 100-dimensional latent vectors as well as features took about half a day on a several-year-old Xeon machine. Training models without features was considerably faster.

Model parameters were regularized using L2 weight cost (see [Salakhutdinov and Srebro \(2010\)](#)) as well as early stopping on the validation set.

5. Implementation Details

Except when explicitly stated otherwise, all reported results were obtained by training on user-item pairs with ratings 50 and higher, with the contribution from pairs with ratings between 50 and 79 weighted by a factor of 0.25. The empirical distribution used for generating negative items was based on pairs with ratings 50 and higher. Features described in Sec. 4.2 were computed based on pairs with ratings 10 and above, as we discovered that including pairs with ratings between 10 and 49 to compute features improved predictive accuracy.

We generated negative samples from the empirical distribution of items in the training set by picking a training case uniformly at random and taking the associated item as the sample. This approach is both simpler and more efficient than generating samples based on the precomputed empirical item probabilities in time logarithmic in the number of items.

In order to compute features from Sec. 4.2 efficiently, we stored all the necessary information about user ratings in hash maps² before starting training. Hash maps allowed us to compute features for a user-item pair in time that is effectively constant in the number of training set ratings for the user.

Table 1: The effect of the training objective

Objective function	Test ERR (%)
Classification (Eq. 2)	5.182
Ranking (Eq. 3)	4.961

Table 2: The effect of the sampling distribution

Distribution of negative items	Test ERR (%)
Uniform	8.889
Popularity-based	4.961

Table 3: The effect of the taxonomy-based parameterization and features

Features used	Taxonomy-based parameterization	Test ERR (%)
—	No	5.673
—	Yes	4.961
Best 4 (1-4)	No	3.469
Best 4 (1-4)	Yes	3.314

6. Results

We created a validation set in order to monitor model performance during training as well as to allow quick model evaluation without using the KDD Cup website. For each user present in the test set, we randomly selected three highly-rated tracks from the training set as positive examples and moved them to the validation set.³ The negative examples were generated by sampling from the empirical distribution of tracks rated 80 or higher in the training set. The scores on the resulting validation set were highly correlated with the corresponding test scores, being about 0.45 higher. In other words, a model with a validation error rate (ERR) of 5.45% has a test ERR of about 5%.

Unless stated otherwise, the models used 100-dimensional factor vectors and the taxonomy-based parameterization from Sec. 4.1, but did not use features. The feature-based models were trained using the best four features (features 1-4) from Sec. 4.2. Early stopping on the validation set was performed to prevent overfitting to the test set. All the scores we report were computed on the part of the test set used to calculate the leaderboard scores.

2. We also tried using binary trees but found them slower, though more memory-efficient, than hash maps.

3. Users with fewer than three such tracks were not included in the validation set.

Table 4: The effect of fixing the album/artist contribution weights w_{al}^r / w_{ar}^r

Weights learned or fixed?	Test ERR (%)
Learned	3.314
Fixed to 1	3.847

First, we looked at the effects of the objective function and the distribution of negative items on model performance. As can be seen from Table 1, the pairwise ranking approach outperforms the classification approach, though not drastically. The effect of sampling negative items from the empirical distribution of positive items instead of the uniform distribution used by BPR is much more dramatic. As shown in Table 2 it reduces the error rate by a factor of 1.8. As a result, we performed all the remaining experiments using the ranking objective with the negative items sampled from the empirical distribution.

We investigated the benefits of the taxonomy-based parameterization and taxonomy-based features by training four models, one for each possible parameterization / features combination. The scores shown in Table 3 indicate that the features help considerably more than the hierarchical parameterization. The two uses of taxonomy information are complementary, as using them both results in the best-performing model.

The importance of learning the weights for the contributions of the album and artist factor vectors to the hierarchically parameterized track factor vectors is demonstrated in Table 4: fixing these weights to 1 instead of learning them increases test ERR by over 0.5.

To explore the interplay between the features used and the latent vector dimensionality we trained models with different numbers of latent factors using only the best four features or all features. Table 5 shows the results for models with 100-dimensional latent vectors and models with no latent factors. Using more features seems to hurt models with many latent factors and help models without (or with few) latent factors, which suggests that the models with a lot of features and latent factors are overfitting. We have also observed that models without features benefit from higher-dimensional latent vectors (e.g. 200D and higher) much more than the models with features do, which supports the overfitting theory.

To determine the contribution of different types of feature weights, we trained three models, each of which used only one type of feature weights (global, user-specific, or item-specific). The results shown in Table 6 suggest that the user-specific feature weights do most of the work, though the item-specific and the global weights also help.

Table 5: The interplay between the number of latent factors and the usefulness of features

Features used	Number of latent factors	Test ERR (%)
All	0	6.709
Best 4 (1-4)	0	11.386
All	100	3.524
Best 4 (1-4)	100	3.314

Table 6: The contribution of different feature weights

Feature weights used	Test ERR (%)
All	3.314
Global only	3.632
Item only	3.591
User only	3.445

Table 7: Summary of results. “Hier. param.” indicates whether the taxonomy-based hierarchical representation was used for item factor vectors. “Min. rating” refers to the minimal value a rating needed to have in order to be included when computing features. “Finetuned” indicates further training with reduced learning rates. “Fold. valid.” indicates whether the validation set was folded into the training set before training.

Model	Model type	# of factors	Hier. param.	Features used	Min. rating	Fine-tuned	Fold. valid.	Test ERR
1	Ranker	500	Y	–	10	Y	N	4.010
2	Ranker	1000	Y	–	10	Y	N	3.728
3	Ranker	2000	Y	–	10	Y	N	3.562
4	Ranker	1000	Y	1-10	10	Y	N	3.358
5	Ranker	100	Y	1-6,10	10	Y	N	3.281
6	Ranker	100	Y	1-5	10	Y	N	3.210
7	Classifier	100	Y	1-5	10	Y	N	3.348
8	Classifier	500	Y	1-5	10	Y	N	3.557
9	Blend of 1-8	N/A	N/A	N/A	10	N/A	N	2.934
10	Classifier	100	Y	–	10	N	N	5.182
11	Ranker	100	Y	–	10	N	N	4.961
12	Ranker	100	N	–	10	N	N	5.673
13	Ranker	100	N	1-4	10	N	N	3.469
14	Ranker	100	Y	1-4	10	N	N	3.314
15	Ranker	100	Y	1-10	10	N	N	3.524
16	Ranker	0	Y	1-4	10	N	N	11.386
17	Ranker	0	Y	1-10	10	N	N	6.709
18	Ranker	100	Y	1-4	10	N	Y	3.191
19	Ranker	100	Y	1-4	0	N	Y	3.104
20	Ranker	100	Y	1-4	0	Y	Y	2.911

7. Our Best Submission

Our best submission combined predictions from the following models:

1. Ranker without features with 500 latent factors
2. Ranker without features with 1000 latent factors
3. Ranker without features with 2000 latent factors
4. Ranker with features 1-10 with 1000 latent factors
5. Ranker with features 1-6,10 with 100 latent factors
6. Ranker with features 1-5 with 100 latent factors
7. Classifier with features 1-5 with 100 latent factors
8. Classifier with features 1-5 with 500 latent factors

The following learning rates were used: 10^{-2} for user representations and item biases, 3×10^{-2} for item representations, 10^{-4} for artist/album contribution weights (w_{ar}^r / w_{al}^r), 10^{-2} for user/item feature weights, and 10^{-4} for global feature weights. The weight cost for user/item representations was 10^{-4} in models without features and 10^{-3} in models with features. The weight cost values for user/item-specific feature weights and global feature weights were 3×10^{-3} and 10^{-3} respectively. User and item representations were initialized from the uniform distributions over the interval $[-10^{-2}, 10^{-2}]$. All other parameters were initialized to zero.

Each model was trained using the above parameters until its validation set ERR started increasing. Then the model was finetuned by reducing all learning rates by a factor of 10 (except for the artist/album contribution weight learning rate, which was set to 0) and training until the validation ERR started increasing again. The finetuning process was repeated one more time for models 1-3 after removing the observations with ratings lower than 80 from the training set. The feature weights for model 4 were learned before other model parameters and were kept fixed while learning the remaining parameters.

Model scores were combined linearly using the weights learned by minimizing the ranking loss (Eq. 3) on the validation set. The resulting ensemble achieved the score of 2.9337 on the leaderboard.

8. Further Improvements

Model performance can be improved further by making a number of tweaks to the training procedure. Our starting point is model 14 (100-dimensional latent vectors and the best four features) from Table 7. The first improvement involves folding in our validation set into the training set. In order to avoid biasing the test ERR by performing early stopping on the test set, we simply make the same number of passes through the training set as we did when training model 14. Using the complete training set reduces the test ERR from 3.314 to 3.191. Computing the features on all ratings during training and testing, instead of just the ones with values 10 and higher, reduces the score down to 3.104. Finally, finetuning the

model by reducing the learning rates as described in Sec. 7 results in our best-performing model with the test set ERR of 2.911. Note that this single model performs better than our best ensemble submission which placed fourth on Track 2. Using more than 100 latent factors is likely to result in even better performance as long as the regularization parameter values are chosen appropriately.

9. Discussion

We have described several improvements to the Bayesian Personalized Ranking model (Rendle et al., 2009) that dramatically improve its performance on Track 2 of KDD Cup 2011. They include generating negative examples from the popularity-based item distribution and taking advantage of the item taxonomy by introducing a hierarchical parameterization of item representations. We also used the taxonomy to provide the model with features summarizing the user’s rating history as it relates to a particular item.

Some of the proposed improvements have been independently discovered by other KDD Cup 2011 participants. Several of the top teams on Track 2 used BPR-like models and sampled negative examples from the popularity-based distribution (Toscher, 2012; Gantner et al., 2012; McKenzie et al., 2012). A taxonomy-based parameterization was also used by some of the teams (Toscher, 2012; Laiy et al., 2012), though they used fixed album/artist contribution weights instead of learning them like we do. Manually designed taxonomy-based features were also quite popular (e.g. Zhang et al. (2012)), with one team taking this approach to the extreme and achieving competitive results without using any latent factors (Xie et al., 2012). To the best of our knowledge, however, no other team has proposed a latent factor model that used both the taxonomy-based parameterization and features. This combination yielded our best-performing model with the test ERR of 2.911, which at the time of this writing is the best score achieved by a single model on Track 2 by a considerable margin.

There is a number of directions we did not have time to explore. Since BPR is based on the pairwise ranking approach Burges et al. (2005) introduced in the “learning to rank” literature, it might be worthwhile to consider the more recent ranking methods developed in that field. LambdaRank (Burges et al., 2006) as well as the list-wise ranking methods (Xia et al., 2008; Cao et al., 2007) seem especially promising.

Using a linear function to map features to item score contributions is the simplest but probably not the best choice. More powerful mappings such as neural networks might be able to make better use of the features.

Acknowledgments

I would like to thank Yee Whye Teh for helpful discussions and feedback on a draft of this paper.

References

Christopher J. C. Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Gregory N. Hullender. Learning to rank using gradient descent. In *ICML*, pages 89–96, 2005.

- Christopher J. C. Burges, Robert Ragno, and Quoc Viet Le. Learning to rank with nonsmooth cost functions. In *NIPS*, pages 193–200, 2006.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to listwise approach. In *ICML '07*, pages 129–136, 2007.
- Gideon Dror, Noam Koenigstein, Yehuda Koren, and Markus Weimer. The Yahoo! Music Dataset and KDD-Cup'11. In *KDD-Cup Workshop*, 2011.
- Zeno Gantner, Lucas Drumond, Christoph Freudenthaler, and Lars Schmidt-Thieme. Bayesian personalized ranking for non-uniformly sampled items. In *JMLR W&CP Volume 18*, 2012.
- Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, pages 263–272, 2008.
- S. Laiy, L. Xiang, R. Diao, Y. Liu, H. Gu, L. Xu, H. Li, D. Wang, K. Liu, J. Zhao, and C. Pan. Hybrid recommendation models for binary user preference prediction problem. In *JMLR W&CP Volume 18*, 2012.
- D.G. Lowe. Object recognition from local scale-invariant features. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 2, pages 1150–1157 vol.2, 1999.
- Christopher D. Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT Press, Cambridge, MA, USA, 1999. ISBN 0-262-13360-1.
- Benjamin Marlin. Collaborative filtering: A machine learning perspective. Master's thesis, University of Toronto, 2004.
- T. G. McKenzie, C.-S. Ferng, Y.-N. Chen, C.-L. Li, C.-H. Tsai, K.-W. Wu, Y.-H. Chang, C.-Y. Li, W.-S. Lin, S.-H. Yu, C.-Y. Lin, P.-W. Wang, C.-M. Ni, W.-L. Su, T.-T. Kuo, C.-T. Tsai, P.-L. Chen, R.-B. Chiu, K.-C. Chou, Y.-C. Chou, C.-C. Wang, C.-H. Wu, H.-T. Lin, C.-J. Lin, and S.-D. Lin. Novel models and ensemble techniques to discriminate favorite items from unrated ones for personalized music recommendation. In *JMLR W&CP Volume 18*, 2012.
- Rong Pan and Martin Scholz. Mind the gaps: weighting the unknown in large-scale one-class collaborative filtering. In *KDD*, pages 667–676, 2009.
- Rong Pan, Yunhong Zhou, Bin Cao, Nathan Nan Liu, Rajan M. Lukose, Martin Scholz, and Qiang Yang. One-class collaborative filtering. In *ICDM*, pages 502–511, 2008.
- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Schmidt-Thieme Lars. BPR: Bayesian personalized ranking from implicit feedback. In *UAI '09*, pages 452–461, 2009.
- Ruslan Salakhutdinov and Andriy Mnih. Probabilistic matrix factorization. In *Advances in Neural Information Processing Systems*, volume 20, 2008.

- Ruslan Salakhutdinov and Nathan Srebro. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. In *NIPS*, 2010.
- Michael Jahrer Andreas Toscher. Collaborative filtering ensemble for ranking. In *JMLR W&CP Volume 18*, 2012.
- Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. Listwise approach to learning to rank: theory and algorithm. In *ICML '08*, pages 1192–1199, 2008.
- Jianjun Xie, David Lisuk, Scott Leishman, Liang Tian, Seongjoon Koo, and Matthias Blume. Feature engineering in users music preference prediction. In *JMLR W&CP Volume 18*, 2012.
- Hang Zhang, Eric Riedl, Valery Petrushin, Siddharth Pal, and Jacob Spoelstra. Committee based prediction system for recommendation: KDD Cup 2011, Track2. In *JMLR W&CP Volume 18*, 2012.