Optimal Ranking for Video Recommendation

Zeno Gantner, Christoph Freudenthaler, Steffen Rendle, and Lars Schmidt-Thieme

Machine Learning Lab, University of Hildesheim
Marienburger Platz 22, 31141 Hildesheim, Germany
{gantner,freudenthaler,srendle,schmidt-thieme}@ismll.de

Abstract. Item recommendation from implicit feedback is the task of predicting a personalized ranking on a set of items (e.g. movies, products, video clips) from user feedback like clicks or product purchases. We evaluate the performance of a matrix factorization model optimized for the new ranking criterion BPR-Opt on data from a BBC video web application. The experimental results indicate that our approach is superior to state-of-the-art models not directly optimized for personalized ranking.

1 Introduction

Recommendations are an important feature of many websites. For example, online shops like Amazon provide customers with personalized product offers. In media applications, personalization is attractive both for content providers, who can increase sales or views, and for customers, who can find interesting content more easily. In this paper, we focus on item recommendation from implicit data. The task of item recommendation is to create a user-specific ranking for a set of items. User preferences are learned from the users' past interaction with the system, e.g. their buying/viewing history.

BPR was recently [3] proposed as a generic optimization method for item prediction from implicit feedback. To our knowledge, models optimized for BPR-Opt are among the most competitive methods for item prediction from implicit feedback. BPR has also been adapted to tag prediction; a factorization model using BPR won the ECML PKDD Discovery Challenge 2009 [4].

Its advantages make BPR an obvious choice for video recommendation from click data. In this paper, we evaluate the performance of a matrix factorization model optimized for BPR-Opt on data from a video web application ran by the British Broadcasting Corporation (BBC).

2 Bayesian Personalized Ranking (BPR)

We briefly introduce BPR, which consists of the objective criterion BPR-Opt, for which we sketch the underlying idea, and the learning algorithm LearnBPR. For more details refer to the long version of Rendle et al. [3].

2.1 Motivation for BPR-Opt

Implicit feedback often consists only of positive observations. The non-observed user-item pairs – e.g. a user has not viewed a video – are a mix of real negative feedback (not interested in viewing it at all) and missing values (may want to view in the future).

Let U be the set of all users and I the set of all items. The known implicit feedback be $S \subseteq U \times I$. The task is now to provide a personalized total ranking $>_u \subset I^2$ for each user.

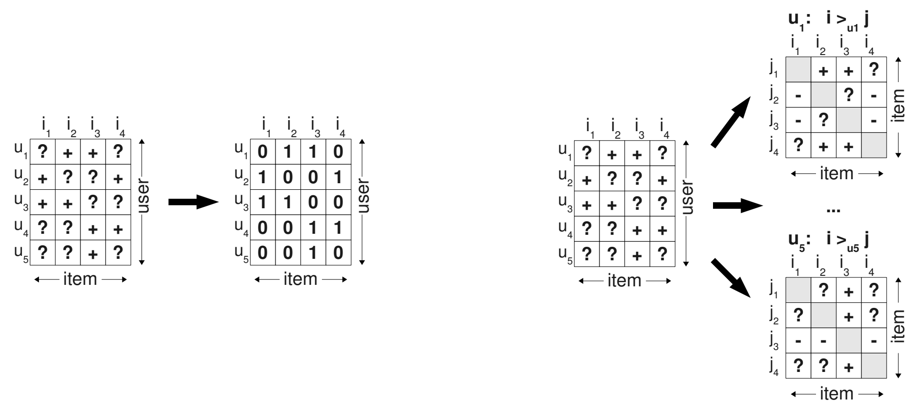

Fig. 1. *Left:* Usually, negative data is generated by filling the matrix with 0 values. *Right:* BPR creates user-specific preferences $i >_u j$ between item pairs. On the right side, $+$ means a user prefers item i over item j; $-$ means they prefer j over i.

Item recommenders [1, 2] typically create the training data from S by giving pairs $(u, i) \in S$ a positive class label and all other combinations a negative one (see Fig. 1). Then a model is fitted to this data, which means is optimized to predict the value 1 for elements in S and 0 for the rest. The problem of this approach is that all elements to be ranked in the future are presented as negative examples to the learning algorithm.

We propose a different approach by using item pairs as training data and optimize for correctly ranked item pairs instead of scoring single items. This represents the problem better than just replacing missing values with negative ones. From S we try to reconstruct $>_u$ for each user. If an item has been viewed by a user, then we assume that the user prefers this item over all other non-observed items. See Fig. 1 for an example. For items that have both been seen by a user, we cannot infer any preference. The same is true for two items that a user has not seen yet.

$$D_S := \{(u, i, j) | i \in I_u^+ \wedge j \in I \setminus I_u^+\} \subseteq U \times I \times I$$

The interpretation of $(u, i, j) \in D_S$ is that user u prefers i over j. Our approach has two advantages: (1) The training data D_S consists of both positive and negative pairs and missing values. The missing values between two non-observed items are exactly the item pairs that have to be ranked in the future. (2) The training data is created for the actual objective of ranking, i.e. the observed subset D_S of $>_u$ is used for training.

2.2 Learning Algorithm and Application to Matrix Factorization

To optimize for BPR-Opt, [3] we use LearnBPR, a stochastic gradient-descent algorithm (see Fig. 2). Using bootstrap sampling instead of full cycles through the data is especially useful as the number of examples is very large, and for convergence often a fraction of a full cycle is sufficient.

procedure LearnBPR(D_S, Θ)
  initialize Θ
  repeat
    draw (u, i, j) from D_S
    $\Theta \leftarrow \Theta + \alpha \left(\frac{e^{-\hat{x}_{uij}}}{1+e^{-\hat{x}_{uij}}} \cdot \frac{\partial}{\partial \Theta} \hat{x}_{uij} - \lambda_\Theta \cdot \Theta \right)$
  until convergence
  return $\hat{\Theta}$
end procedure

Fig. 2. Optimizing for BPR-Opt with bootstrapping-based gradient descent.

Because we have triples $(u, i, j) \in D_S$, we decompose the estimator $\hat{x}_{uij}$: $\hat{x}_{uij} := \hat{x}_{ui} - \hat{x}_{uj}$. Now we can apply any model that predicts $\hat{x}_{ul}, l \in I$.

Matrix factorization (MF) models are known to outperform [5] many other models for the related task of rating prediction. They are also state-of-the-art for item prediction. MF approximates the target matrix X by the product of two low-rank matrices $W : |U| \times k$ and $H : |I| \times k$ by $\hat{X} := WH^t$. For estimating whether a user prefers one item over another, we optimize the parameters $\Theta = (W, H)$ for the BPR-Opt criterion using our algorithm. To apply LearnBPR to MF, only the gradient of $\hat{x}_{uij}$ wrt. every model parameter has to be derived.

3 Evaluation

We compare learning the MF model with LearnBPR to weighted regularized matrix factorization (WR-MF) [1, 2]. We also report results for a baseline method that ranks the items by global frequency.

The BBC data was collected from one of BBC's online services during a ten day period in 2009. It contains 2,867,128 viewing events generated by 189,228 anonymous users on 5,125 different video clips. For comparability, we report results on a subset of the *Netflix* data.[1]

We randomly removed one user-item pair per user to create the training set. and repeated all experiments 10 times by drawing train/test splits in each round.

[1] The results are taken from [3]. The subset contains 10,000 users, 5000 items, and 565,738 rating actions, where each user and item have at least 10 ratings.

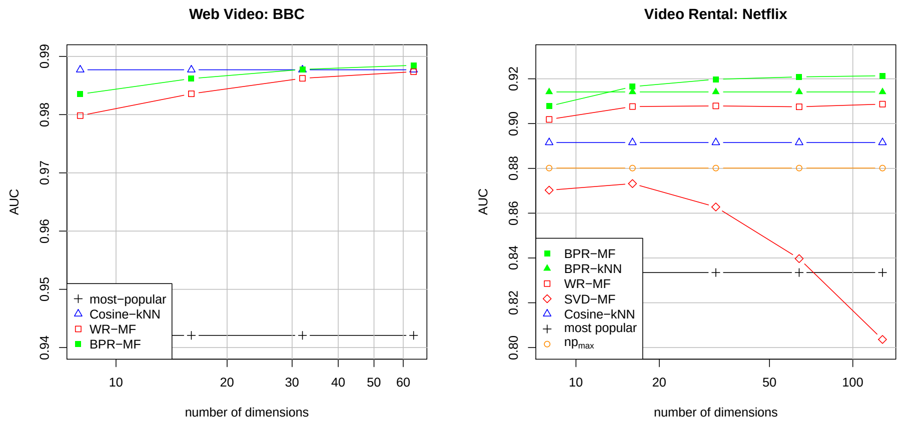

Fig. 3. Area under the ROC curve (AUC) prediction quality for BBC and Netflix.

You can see in Fig. 3 that on Netflix, BPR-MF outperforms all other methods in prediction quality on both datasets.

4 Conclusion and Future Work

We evaluated the performance of a matrix factorization model optimized for the recently proposed generic criterion BPR-Opt on data from a BBC video web application. The results indicate that the method outperforms other state-of-the-art methods. This is justified by the theoretical analysis in [3]. In the future, we will conduct a large-scale study on additional real-world datasets from the MyMedia project field trials.

Acknowledgements We thank the BBC for providing the video dataset, and Chris Newell for fruitful discussions. This work was co-funded by the EC FP7 project MyMedia under the grant agreement no. 215006.

References

1. Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM 2008*, 2008.
2. R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. M. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *ICDM 2008*, 2008.
3. S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *UAI 2009*, 2009.
4. S. Rendle and L. Schmidt-Thieme. Factor models for tag recommendation in bibsonomy. In *Discovery Challenge at ECML PKDD 2009*, 2009.
5. J. D. M. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *ICML 2005*, 2005.